Can Green Car Taxes Restore Efficiency? Evidence from the Japanese New Car Market

Yoshifumi Konishi, Meng Zhao

Abstract: To quantify the economic impacts of Japan's feebate policy, a random-coefficients logit model is estimated for quarterly automobile sales between 2007 and 2012 from the Japanese new car market. For identification of the structural parameters, we exploit the policy-induced variation in effective car prices and the location of product-specific vehicle taxes as instruments. The estimated demand system allows us to simulate counterfactual Bertrand-Nash equilibria in response to alternative policy scenarios. Our results indicate that Japan's feebate policy induced a sizable increase in economic surplus, yet only a small improvement in sales-weighted average fuel efficiency, relative to the no-policy counterfactual. We also design an optimal feebate policy, which maximizes total economic surplus subject to a tax revenue constraint, by explicitly accounting for market power, product attributes, and carbon dioxide emissions rates. The policy is predicted to induce sizable improvements in both economic surplus and average fuel efficiency over Japan's feebate policy without requiring any decrease in tax revenues.

JEL Codes: H23, H30, L62, Q53

Keywords: Automobile demand, Discrete choice models, Imperfect competition, Optimal vehicle taxation, Product differentiation, Random-coefficients logit, Vehicle emissions

CARBON DIOXIDE (CO_2) EMISSIONS from motor vehicles continue to present a daunting challenge to policy practitioners. In first-best settings, an efficient gasoline tax can fully restore economic efficiency both at the extensive margin (i.e., car ownership) and the intensive margin (i.e., car utilization) (Innes 1996; Fullerton and West

Yoshifumi Konishi (corresponding author) is at Sophia University, Faculty of Liberal Arts, 7-1 Kioi-cho, Chiyoda-ku, Tokyo 102-8554, Japan (y-konishi@sophia.ac.jp). Meng Zhao is at Gakushuin University, Faculty of International Social Science. We are grateful for all the helpful comments from two anonymous referees and the coeditor (Don Fullerton) and from seminar participants at Kobe University, Sophia University, Waseda University, University of Tokyo,

Received May 18, 2015; Accepted August 29, 2016; Published online January 13, 2017.

JAERE, volume 4, number 1. 2333-5955/2017/0401-0002$10.00 http://dx.doi.org/10.1086/689701

2002).[1] Yet, a number of real-world complexities make the gasoline tax a less attractive policy instrument. Such complexities include the imperfectly competitive nature of the automobile market, the regressive property of the gasoline tax, and the optimization failures often reported on consumer choices in adopting new technologies such as hybrid and electric vehicles due to risks, information costs, and/or switching costs. Due partly to these complications, many countries have started to explore other complementary policies to control vehicle CO_2 emissions. One of such policies is known as a "feebate" policy. A feebate is a fiscal instrument, that imposes a "fee" on purchase of high-emission, fuel-inefficient vehicles and gives a "rebate" on purchase of low-emission, fuel-efficient vehicles. It is known for the potential to avoid (some of) the demerits of the gasoline tax yet it can be readily implemented on top of existing vehicle taxation and other incentive systems without undermining their intended effects (Anderson et al. 2011). Hence, variants of feebates have been recently explored in several countries.[2]

While there is a growing interest among environmental economists in quantifying the economic impacts of feebates, doing so is inherently complicated precisely due to the imperfectly competitive nature of the automobile market. Automobile industries are often oligopolistic with a small number of automakers competing in multi-product pricing. On one hand, the markup pricing tends to underprovide the goods relative to the perfectly competitive equilibrium. On the other hand, the negative externality as-

and the Fifth World Congress of Environmental and Resource Economists. The study was in part supported by financial support from the Japan Society for the Promotion of Science (JSPS), Grant-in-Aid for Young Scientists (grant 25780171). We also thank Yukiko Omata for excellent research assistance in collecting and coding the data and two undergraduate assistants, Saya Kuwata and Tomoka Sumikawa, for data entry. An earlier version of the paper was circulated as Tokyo Center for Economic Research Working paper E-82-2014. This is a substantially improved version of that paper.

1. The result holds even in the presence of heterogeneous consumers, but only for CO_2 emissions. The result does not hold for other vehicle pollutants such as carbon monoxide (CO), nitrogen oxides (NO_x), and reactive hydrocarbons (HC) because emissions per unit of fuel consumption may vary substantially due to vehicle characteristics. To address these pollutants, optimal coordination of vehicle and fuel taxes is necessary (see Fullerton and West [2002] for detailed discussions).

2. Examples include France, Germany, Japan, Sweden, and the United States. The feebate policy is only one of the complementary policy instruments used in these countries, however. Many developed countries impose sufficiently high gasoline taxes to account for the negative externality cost of fuel consumption (Ley and Boccardo 2010). Other instruments include, but are not limited to, emissions standards and corporate average fuel economy (CAFE) standards. Feebates and CAFE with flexible credit trading can be equivalent in theory, yet the feebates may be favored over the CAFE on the ground that the former tend to be additive while the latter tends to be incompatible with other preexisting incentive policies. See Anderson et al. (2011) for stimulating discussions on this point.

sociated with vehicle emissions implies that the market equilibrium tends to overprovide the goods relative to the social optimum. Which of the effects dominates depends on heterogeneous consumer preferences that generate product-level demand systems. In principle, automakers have incentives to underprovide (overprovide) fuel-efficient cars in product segments where consumers value them less (more) than in others (Fischer 2010). Hence, corrective instruments must be tailored and coordinated across firms and products, properly accounting for imperfect competition, externality, and consumer heterogeneity generating product-level demand. What makes the feebate attractive in this context is that it is by design a combination of product-specific fees and rebates and, therefore, can be designed, at least in theory, to correct for these types of market failures. Yet, feebates in most practical settings are tied closely to emissions and fuel efficiency characteristics of products, not the measures of product-level market power. An important question then is whether it is possible to design the feebate scheme that would correct for imperfect competition while achieving the original goal of encouraging the purchase of low-emission, fuel-efficient vehicles.

The primary objective of the paper is to investigate this question empirically. Doing so requires us to measure the degree of market power at the product level. Our strategy is to estimate the product-level demand and use the estimated demand jointly with the markup pricing rules implied by the Bertrand-Nash equilibrium to recover marginal costs for all products. With the estimates of demand parameters and marginal costs, we are then able to simulate counterfactual Bertrand-Nash pricing equilibria corresponding to alternative feebate policies. This ability to simulate counterfactual equilibria is exploited in designing the optimal feebate system, which would maximize total economic surplus (the sum of consumer surplus, producer surplus, and tax revenues in net of subsidy expenditures) subject to a tax revenue constraint. In the optimization algorithm, we linearize the tax/subsidy rates around observed markups and carbon emissions rates. This linearization can, therefore, explicitly account for the degree of product-level market power and environmental attractiveness simultaneously, while making it easy for policy practitioners to adopt the scheme in practical settings.

To implement this general strategy, we employ the random-coefficient discrete-choice model also known as the Berry-Levinsohn-Pakes (BLP) estimator. The BLP estimator was developed in Berry (1994) and Berry, Levinsohn, and Pakes (1995) and has been successfully applied in a number of empirical studies since then (e.g., Berry, Levinsohn, and Pakes 1999; Nevo 2001; Petrin 2002; Villas-Boas 2007; Crawford and Yorukoglu 2012). The approach makes use of market-level data only and deals with endogeneity of prices, yet it allows for heterogeneity in consumer tastes for product characteristics and, hence, generates rich and realistic substitution patterns.[3] For estima-

3. Its main drawback has been the computational burden and numerical accuracy, as it requires running a nested fixed point (NFP) algorithm as an inner-loop subroutine for the gen-

tion of the model, we make use of detailed market-level data on sales by car model and quantifiable car characteristics in the Japanese new car market between 2007 and 2012. We focus on the Japanese new car market, as it offers an attractive empirical setup for pursuing our objective. The market is characterized by an oligopolistic industry with nine domestic automakers. The Japanese preexisting taxation system consists of both a gasoline tax and a suite of vehicle taxes based on car characteristics such as size, weight, and displacement levels. Most importantly, the Japanese government started a series of subsidy and tax incentive programs for low-emission, fuel-efficient cars, called Ecocar Subsidy (ES) and Ecocar Tax Credits (ETC), since 2009. Their unique features created large policy-induced variations in the effective prices of cars across car models and over time, in a manner analogous to feebates.

Implementation of the BLP estimator requires a set of instruments for identification. To that end, we exploit the quasi-experimental nature of Japan's ES/ETC policy. Earlier studies often used the "location" of observed product characteristics in the product space as instruments (called "BLP instruments" hereafter), arguing that product location is at least predetermined prior to the determination of consumer demand. Though this may be a valid assumption in some contexts, there is a concern in our context that the location of observed product attributes may be highly correlated with unobservable product attributes such as nonprice sales promotions or brand images (e.g., Toyota Prius's brand image may come from its high fuel efficiency). We circumvent this concern by constructing variables that represent the location of product-specific vehicle taxes in the characteristics space. The vehicle taxes in Japan are indeed a function of observed product characteristics (i.e., prices, displacement levels, size, and vehicle weight). Hence, they are correlated with prices. Yet, the frequent changes in the location of these taxes are likely to remove much of the causal link with respect to the unobserved product characteristics such as style and brand images, which presumably stay more or less constant over time. In section 3, we document the problems we encountered with BLP instruments, offer more detailed arguments, and report on evidence in support of our instruments. We then report the results of our estimation with our instrumental variables (IVs) in section 5.

Our study is closely related to an ample body of literature that has empirically examined the economic impacts of fiscal instruments on the demand for car ownership and utilization (e.g., Goldberg 1998; West 2004; Bento et al. 2009; Feng, Fullerton, and Gan 2013; Klier and Linn 2013, 2015; D'Haultoeuille, Givord, and Boutin 2014) as well as studies that have applied the BLP estimator in a variety of empirical contexts (e.g., Berry et al. 1999; Nevo 2001; Petrin 2002; Villas-Boas 2007; Crawford and

---

eralized method-of-moments (GMM) estimation. To circumvent some of the computational problems, we take advantage of recent advances in the study of the BLP estimator (Dube et al. 2012; Knittel and Metaxoglou 2014).

Yorukoglu 2012). Of these, ours is probably most closely related to D'Haultoeuille et al. (2014), who conducted an empirical study investigating the impact of the feebate policy in France. Exploiting rich household-level data and estimating both car ownership and utilization, they show that the policy is estimated to increase carbon dioxide emissions primarily due to its scale effects: that is, increases in driving distance, car sales, and overall stock of cars. Furthermore, to overcome the price endogeneity, they posit the price differential to be a linear function of fees/rebates for identification of their model. This identifying assumption is essentially the same as ours, for it means that the unobserved errors (and hence the changes, too) are causally unrelated to the location of changes in car taxes (i.e., the location of fees/rebates).

Our work also contributes to an extensive literature on the theory of optimal taxation in second-best settings. On one hand, the public finance literature has it that optimal subsidies can fully restore economic efficiency, provided that the regulatory authority has perfect information about the degree of market power and access to lump-sum taxation elsewhere (see Auerbach and Hines [2003] and related studies cited therein). On the other hand, the environmental economics literature has it that without imperfect competition, an efficient gasoline tax can fully restore economic efficiency in addressing vehicle CO_2 emissions (Innes 1996; Fullerton and West 2002). With imperfect competition in the goods market, however, this negative consumption externality and consumers' valuation of it can interact, in an intricate manner, with the degree of market power (Fischer 2010). Hence, full economic efficiency cannot be restored without additional policy instruments. The issue is further complicated because in many practical settings, the government's ability to restore economic efficiency is often constrained by the need to raise tax revenues. Hence, a feebate scheme must be optimized under some tax revenue constraint à la Ramsey (1927). One important contribution of the paper is that we offer an approach to designing an optimal feebate policy in these second-best settings that is relatively easy to implement in practice. In the appendix (available online), we elaborate more on the motivation, economic intuition, and key design issues for our study using a simple two-product monopoly setup.

1. INSTITUTIONAL BACKGROUND

Under the Japanese vehicle taxation system, consumers pay three types of car taxes at the time of new car purchase and during ownership. First, automobile acquisition tax is a prefectural ad valorem tax, which charges 5% of the sales value at the time of car purchase. Second, vehicle weight tax is a national tax collected at the time of car inspections every 1–3 years and was set at ¥12,600 (or ¥10,000) per ton of vehicle weight before (or after) April of 2010. Third, annual automobile tax is another prefectural tax imposed on car ownership, which ranges from ¥0 to ¥111,000 depending on displacement level. There is a special class of cars called *Kei-cars* or "minicars" sold in Japan: that is, extremely small vehicles with displacement level of 660 cubic centimeters or less. These minicars are exempt from the annual automobile tax. The last two taxes are taxes on ownership, but consumers also pay them at the time of car registration.[4]

Prior to 2009, these car taxes were only tied to vehicle weights, displacement levels, and sales values of cars and, hence, were not explicitly linked to either fuel efficiency or emissions performance. In 2009, partly backed up by then Prime Minister Aso's Green New Deal, the Japanese government started to implement a series of policy experiments on the taxation of automobiles. The policy roughly consists of the Ecocar Tax Credits (ETC) program and the Ecocar Subsidy (ES) program. The ETC offered a variety of tax incentives based on fuel efficiency and emissions performance. For example, models exceeding the 2010 fuel efficiency standard by 15% (but less than 25%) and receiving a four-star rating on the 2005 emissions standard would receive a 50% cut on vehicle weight tax, a 50% cut on acquisition tax, and a 25% cut on annual automobile tax.[5] The ETC program was originally scheduled to continue until March 31, 2012 (April 30, 2012, for vehicle weight tax) but was extended (in March 2012) to April 2015. The ES program, on the other hand, offered a cash rebate of ¥100,000 (¥50,000) for purchase of a passenger car (minicar) if it achieves 15% above the 2010 fuel efficiency standard and the four-star rating on the 2005 emissions standard.[6] Initially, the ES program was scheduled to last until March 31, 2010. However, it was extended to September 30, 2010, as part of the 2010 economic stimulus package. Furthermore, the second phase of the ES program was reimplemented on December 20, 2011, and continued until January 31, 2013. The eligibility requirements in the second phase were made stricter than those in the first phase. Consequently, the policy period can be further subdivided into three distinct periods: (i) April 2009–September 2010, in which ETC and the first phase of ES were in place; (ii) October 2010–December 2011, in which only ETC was in effect; and (iii) January 2012–December 2012, in which ETC and the second phase of ES were in effect. Table 1 summarizes the eligibility requirements for different ES and ETC programs.

An attractive feature of the ES/ETC policy for our analysis is that its frequent changes provide important policy-induced variations in the effective car tax rates over time and across car models. Importantly, because these ES/ETC programs were tightly

4. On top of these car taxes, the consumers also need to pay the 5% ad valorem sales tax, which did not change throughout the study period.

5. To be more precise, the tax incentive on the automobile tax started in April 2004 before the ETC program, and its eligibility requirements have been changing over time. The text refers to the requirements for cars sold in fiscal year 2009.

6. The cash rebate is increased to ¥250,000 (¥125,000) for purchase of a passenger car (minicar) if it replaces old cars aged 13 years or more and meets the 2010 fuel efficiency standard. Because an average year of car ownership in Japan is substantially less than 13 years, we ignore this complication in our analysis.

Table 1. Model Eligibility Requirements for ES and ETC

	2005 Emissions Standard 4 stars	2010 Fuel Efficiency Standard			
		115% or Above	Incentives	125% or Above	Incentives
Passenger cars:					
ES_1	✓	✓	JP¥100,000	✓	JP¥100,000
ES_2	✓	. . .	. . .	✓	JP¥100,000
ETC (vehicle weight tax)	✓	✓	50% tax cut	✓	75% tax cut
ETC (acquisition tax)	✓	✓	50% tax cut	✓	75% tax cut
ETC (auto tax)	✓	✓	25% tax cut	✓	50% tax cut
Minicars:					
ES_1	✓	✓	JP¥50,000	✓	JP¥50,000
ES_2	✓	. . .	. . .	✓	JP¥70,000
ETC (vehicle weight tax)	✓	✓	50% tax cut	✓	75% tax cut
ETC (acquisition tax)	✓	✓	50% tax cut	✓	75% tax cut
ETC (auto tax)	. . .	. . .	. . .	. . .	. . .

Note. Under the first ES policy, the subsidy amount would increase to JP¥250,000 for passenger cars and JP¥125,000 for minicars if consumers replace their owned cars aged 13 years or above. ES_1 and ES_2 stand for the first and the second phases of the ES program, respectively. The eligibility requirements for tax credits vary over the study period. The requirements in this table refer to those in 2009. ES = Ecocar subsidy; ETC = Ecocar tax credits.

linked to fuel efficiency, it allowed the car taxes to be closely linked to the carbon emissions rates of the vehicles. Figure 1*A* shows the scatter plots of the car taxes against the corresponding carbon emissions rates for all car models sold during the pre-policy period (January 2007–March 2009) and during the policy period (April 2009–December 2012). The figure demonstrates that the linkage between the car taxes and the emissions performance of the cars became much tighter during the policy period than during the pre-policy period. This is also confirmed with figure 1*B*, which plots the kernel densities of car taxes corresponding to four different policy periods. Prior to the policies, dispersion in car taxes is relatively small, with the mode of the distribution around ¥180,000. During the policy period, the distribution of car taxes is shifted to the left, made more disperse, with some of the car models receiving even negative tax rates due to the ES program. Importantly, these distributions changed substantially not just between the pre-policy and policy periods but also across the three distinct policy periods.

Figure 2*B* shows the trend in average tax rates (incorporating the subsidy and the tax credits), which confirms that the changes in the distribution of car taxes also translate into a substantial intertemporal variation in vehicle tax rates. The average tax

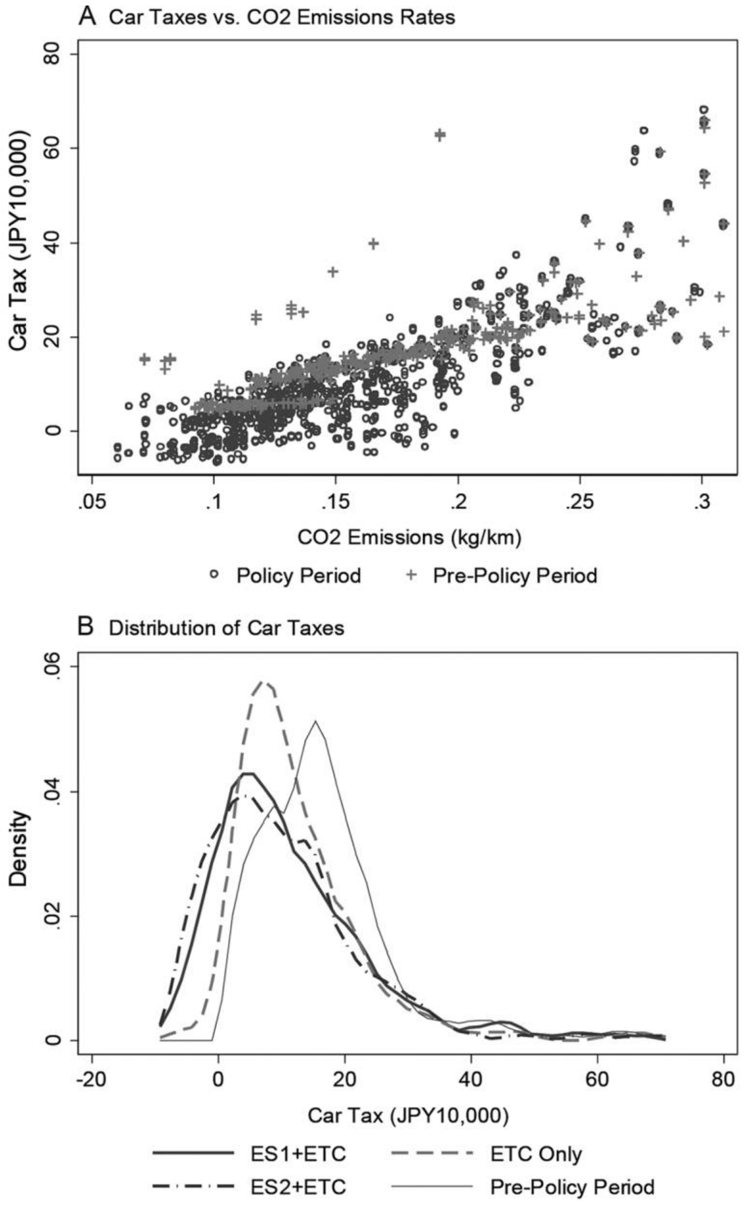

Figure 1. Regulatory changes in car taxes in Japan. Note that CO_2 emissions for each model = Average CO_2 emissions per liters of gasoline/mileage per liter of gasoline. Average CO_2 emissions per liter of gasoline are taken from EPA (2014). Kernel density estimation used the Epanechnikov kernel and the bandwidth of 2.5. "Pre-policy" period = from the first quarter of 2007 to the first quarter of 2009; ES1 + ETC period = from the second quarter of 2009 to the third quarter of 2010; ECT Only period = the fourth quarter of 2010 to the fourth quarter of 2011; and ES2 + ETC period = all quarters in 2012.

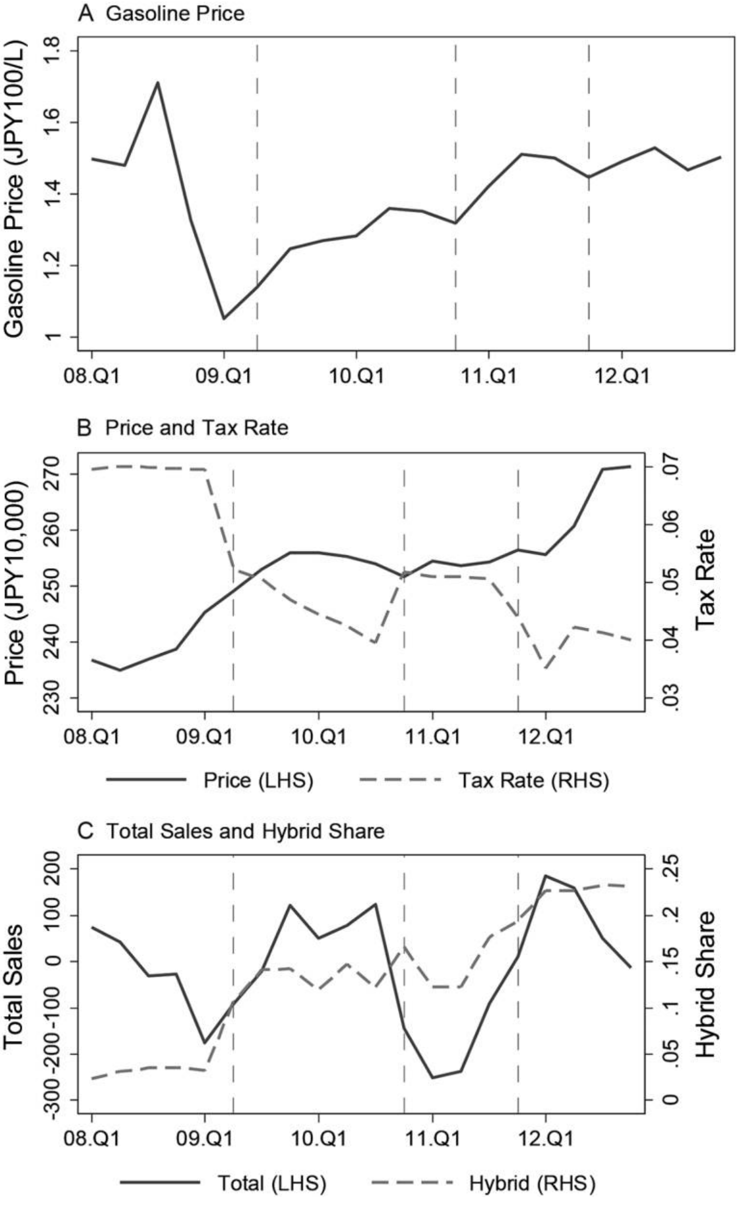

Figure 2. Trends in gasoline price, car prices, car tax rates, new car sales, and hybrid shares

rate sharply dropped during the first policy subperiod.[7] It then increased slightly during the second policy subperiod due to the temporary suspension of the ecocar subsidy and then decreased again during the third subperiod when the second phase of the ES was implemented. The policy's impacts on the sales mix and the total sales are less clear-cut. A casual look at the sales patterns over time suggests that these changes in tax rates may appear to have induced substantial behavioral changes in terms of both aggregate consumption and substitution patterns across models. First, figure 2C shows that the share of hybrid cars in total car sales increased dramatically during the first policy subperiod, and the trend continued throughout the policy period.[8] Second, total sales quantity (detrended by regressing it on quarter dummies) also jumped dramatically during the first policy subperiod and then dropped sharply after the ES was ceased. However, there are clear confounders during the study period. The gasoline price (deflated using consumer price index) also increased substantially during the same period (see fig. 2*A*), which must have also induced consumers to buy fuel-efficient cars. The impact on total sales is also somewhat ambiguous because the Japanese economy went through two substantial macroeconomic shocks during the study period (the financial crisis, known as the Lehman Shock, and the 2011 Tohoku earthquake). The effects of these two macroeconomic shocks appear particularly evident during 2008/Q3–2009/Q1 and during 2011/Q1–2011/Q2. Hence, to get at the causal impacts of the policy, we need to estimate the automobile demand controlling for these time-varying factors.

The Japanese government also mandates corporate average fuel economy (CAFE) standards in a manner similar to that of the US CAFE. The standards were changed in 2007 from the 2010 standards to the 2015 standards because many firms already met the 2010 standards by 2005. Furthermore, though the fuel economy standard is set for each segment (by car weight), each firm is only expected to meet the overall CAFE standards. Hence, firms faced the same 2015 standards throughout our study period (2007–12), which became binding only at the firm level and after 2015.[9] In contrast, ecocar subsidies and tax credits were tied to different standards for different car segments over different time periods. This distinction helps us to isolate the effects of these policies from those of the CAFE standards.

7. The average tax rate was calculated as a simple unweighted average over all car models sold during each time period.

8. In Japan, diesel-based cars represent a tiny fraction of total sales. Instead, hybrid cars such as Toyota Prius and Honda Civic Hybrid are more closely equated with "eco-friendly" cars in the minds of Japanese consumers.

9. The economic impact of the CAFE standards may still materialize through a firm's product strategy. This pathway is not addressed in the paper since our model does not endogenize a firm's product choices.

2. EMPIRICAL MODEL

2.1. Consumer

Let us first start with a generic empirical framework, building upon the extensive literature on discrete choice models of automobile demand. In each quarterly market t, the indirect utility of consumer i choosing alternative j depends on both observable and unobservable product and consumer characteristics:

$$u_{ij}(\theta) = \delta_j(\theta) + \mu_{ij}(\theta) + \epsilon_{ij},$$

where $\boldsymbol{\theta}$ is a vector of parameters to be estimated. Throughout our presentation, we omit the identifier t since the model structure stays the same for all t.[10] The first term δ_j only depends on product characteristics (either observable or unobservable) and is common to all individuals. The second term μ_{ij} depends on consumer attributes and observed product characteristics and captures heterogeneity in consumer tastes for observed product characteristics. The last component ϵ_{ij} is the mean-zero random utility and is assumed to be independently and identically distributed (i.i.d.).

Much of the recent advance in the literature centers on how to incorporate the term μ_{ij} in the estimation of automobile demand. If this term is not included, the only consumer-level heterogeneity comes from the i.i.d. error ϵ_{ij}, and hence, the choice probability for any consumer only depends on observable product characteristics δ_j. The most unappealing implication of the omission is the unrealistic substitution pattern à la McFadden's red bus/blue bus problem. When consumer-level data are available, the observed choice probabilities of new purchasers can be directly linked to their household and product attributes. Goldberg (1998) and Bento et al. (2009) follow this approach. When only market-level data are available, however, we cannot directly link these two. The Berry et al. approach is to assume that the consumer-level taste variation arises from some known distribution such as multivariate normal (Berry et al. 1995) and χ^2 distributions (Petrin 2002). Then the observed market shares are matched with the model's predicted choice probabilities to consistently estimate the parameters of the term μ_{ij}.[11]

Following Berry et al. (1995, 1999) and Nevo (2000, 2001), we specify the utility as follows:

$$u_{ij} = \alpha_i(y_i - p_j^e) + \mathbf{x}_j\boldsymbol{\beta}_i + \xi_j + \epsilon_{ij} = V_{ij} + \epsilon_{ij}, \tag{1}$$

10. In the empirical specifications in sec. 5, we do include quarter dummies, year dummies, and GDP growth rates (without random coefficients) to allow the utility relative to the outside option to vary over time due to some time-varying factors à la Berry et al. (1999).

11. To further improve the precision of the BLP estimators, Nevo (2001) and Petrin (2002) independently offered methods to link the aggregate-level demongraphics of consumers to the characteristics of the products. We do not follow Nevo or Petrin's approach in this paper, since in our data we do not have enough variation in, or enough information on, aggregate-level consumer demographics across quarterly markets to implement their approaches.

where $p_j^e = (1 + \tau_j)p_j$ is the effective (i.e., tax-inclusive) price of car j, $\mathbf{x}_j$ the K-dimensional vector of observable characteristics of car j, ξ_j the unobservable characteristics of car j, y_i the income of individual i, and $(\alpha_i, \boldsymbol{\beta}_\text{i})$ is a vector of "random coefficients" to be estimated and assumed to vary over individuals, which are specified as:

$$\begin{pmatrix} \alpha_i \\ \beta_i \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix} + \sum \circ \nu_i, \tag{2}$$

where $\sum = (\sigma_p, \sigma_1, \ldots, \sigma_K)'$ is a $(K + 1)$-dimensional vector of parameters and $\text{v}_i = (\nu_{ip}, \nu_{i1}, \ldots, \nu_{iK})'$ is a $(K + 1)$-dimensional vector of unobservable characteristics of individual i. The number of dimension K is equal to the number of variables in x_j. We assume that $\mathbf{v}_i$ follows an i.i.d. standard multivariate normal $N(\mathbf{0}, \mathbf{I})$, following Berry et al. (1995, 1999), except for the price coefficient α_i. A normal distribution for α_i can be problematic because it would allow the price coefficient to become positive for some individuals.[12] Hence, following Train's suggestion (2003, 142), we experimented a constant price coefficient ($\alpha_i = \alpha$) as well as a lognormal distribution for ν_i. We decided to use a lognormal distribution based on the model's performance in terms of statistical significance and estimated elasticities. This distributional assumption implies that the marginal utility from k-th product characteristic (or its logarithm in the case of price) has a mean β_k (α) and a standard deviation σ_k (σ_p). For this reason, β_k is often called a mean parameter and σ_k is called a standard deviation parameter in the literature.

Note that the term ξ_j represents product attributes that are observed by consumers and firms but are unobservable or unquantifiable by the researcher. One way to interpret the term ξ_j is that it measures brand images, style, and prestige. Another way to interpret it is that it represents the measurement errors in observed market prices such as product-specific sales promotions and marketing strategies. Either way, it is likely to be correlated with p—for example, consumer demand is higher for products with better brand images, and measurement errors with respect to prices are likely to be related to sales promotions and sales channels. Hence, if uncontrolled, it is likely to bias our parameter estimates. We take the estimation strategy proposed by Berry et al. (1995, 1999) to take care of this endogeneity, which we shall turn to in section 3.

The discrete choice model is closed with the inclusion of an outside option. In each period, the consumer is assumed to buy at most one car and may choose to buy one of the J_t car models or not buy any car ($j = 0$). In the latter case, she may choose to use public transportation or continue to use a car she already owns. The inclusion of the outside option allows us to estimate the impact of a homogeneous decrease or

12. We appreciate our referee for pointing this out.

increase in the effective prices of all car models on quantities purchased. Given our specification in (1), the indirect utility from the outside option is given by

$$u_{i0} = \alpha_i y_i + \sigma_0 v_{i0} + \epsilon_{i0} .$$

Note that the term v_{i0} still needs to be included, despite there being no observable attributes for the outside option, to account for the possibility that the idiosyncratic variance for this option may be larger than that for the "inside" goods (Berry et al. 1995; Nevo 2000).[13] Then assuming that ϵ_{ij} are i.i.d. with a Type I extreme value distribution, the market share of car model j is given by

$$s_j = \int \frac{\exp(\mathbf{x}_j\beta_i - \alpha_i p_j^e + \xi_j)}{1 + \sum_{r=1}^{J_t} \exp(\mathbf{x}_r\beta_i - \alpha_i p_r^e + \xi_r)} dP(\mathbf{v}), \tag{3}$$

where $P(\cdot)$ is the population distribution of the individual attributes v. The integral is only with respect to v because y vanishes in our linear income specification.

2.2. Producer

There are F firms in all markets and each firm produces a subset of the products $\mathcal{J}_f$. In each quarterly market t, the profits of firm f are given by:

$$\sum_{j \in \mathcal{J}_f} (p_j - mc_j) M s_j(\mathbf{p}^e) - FC_f,$$

where s_j is the market share of car model j as defined in (3), $\mathbf{p}^e$ is the vector of effective, tax-inclusive prices defined as $p_j^e = (1 + \tau_j)p_j$, mc_j is the marginal cost of each car model j, M is the market size of the new car market, and FC_f is the fixed cost of production.

Assuming that firms compete in the Bertrand manner and the unique pure-strategy Bertrand-Nash equilibrium exists (as in Berry et al. 1995, 1999; and Nevo 2000, 2001), the price of each product j satisfies the following first-order condition:

$$s_j(\mathbf{p}^e) + (1 + \tau_j) \sum_{k \in \mathcal{J}_f} (p_k - mc_k) \frac{\partial s_k}{\partial p_j} = 0 . \tag{4}$$

13. In practice, however, the coefficient σ_0 cannot be identified since it cannot be separated out from the standard deviation coefficient on the constant term. Hence, the standard practice is to set σ_0 to equal zero. Because we assume a linear income effect in (1), the term $\alpha_i y_i$ eventually vanishes. Thus, setting σ_0 to equal zero is equivalent to normalizing the indirect utility from the outside option to zero (see Nevo [2000, 2001] for a detailed discussion on this point). With this normalization, the idiosyncratic differences in tastes for the outside option are subsumed in the standard deviation parameter on the constant term. Hence, if we expect different consumers to behave differently with respect to the outside option, say, due to differences in access to public transportation or in the ownership of cars, then we should expect the standard deviation parameter on the constant term to be statistically significant.

For each market, this set of J equations determines the optimal markup for each product. These markups can be solved explicitly à la Nevo (2001). Let us define the matrix Ω such that each element of Ω is defined as $\Omega_{jk} = O_{jk} * D_{jk}$, where O_{jk} is the matrix describing the ownership structure:

$$O_{jk} = \begin{cases} 1 & \text{if } \exists f : \{j, k\} \in \mathcal{J}_f \\ 0 & \text{otherwise} \end{cases},$$

and D_{jk} is the matrix of share derivatives with respect to prices, multiplied by -1: $D_{jk} = -\partial s_k / \partial p_j$. Then the first-order condition implies:

$$\mathbf{mc} = \mathbf{p} - \Omega^{-1} \mathbf{s}^e(\mathbf{p}^e), \tag{5}$$

where $\mathbf{s}^e$ is a vector of market shares adjusted for tax rates: that is, the jth element of $\mathbf{s}^e$ is $s_j^e = s_j/(1 + \tau_j)$.

Once we obtain the consistent estimates of demand parameters, we can estimate the price-cost margins or the marginal costs using (5), which can then be used to simulate the policy-induced effects on counterfactual Bertrand-Nash equilibria. One could impose further structures on the supply relationship, and the cost parameters could then be jointly estimated with the demand parameters. For example, Berry et al. (1995, 1999) consider the estimates of mc_j's obtained from (5) as a log-linear function of cost shifters such as observed product attributes, wages, and unobservable product attributes and estimate the cost parameters jointly with the demand-side parameters. Such a strategy would improve the efficiency of the estimates, but at the cost of imposing more structures and increasing the computational burden. As we do not directly make use of the cost-side parameters in our simulation analysis, we shall take Nevo's approach to avoid undue complexity.

3. EMPIRICAL STRATEGY

For estimation of the model, we closely follow the generalized method of moments (GMM) method proposed in Berry et al. (1995). Suppose we have data on a set of exogenous instruments $\boldsymbol{z}$ such that the unobserved product attributes are mean independent of $\boldsymbol{z}$:

$$E[\boldsymbol{\xi}(\boldsymbol{\theta}) \mid \boldsymbol{z}] = 0\,. \tag{6}$$

This gives us a set of population moment restrictions. Then the GMM estimates of the parameters are:

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}}\, \boldsymbol{\xi}(\boldsymbol{\theta})' \boldsymbol{z} W^{-1} \boldsymbol{z}' \boldsymbol{\xi}(\boldsymbol{\theta}), \tag{7}$$

where W is a consistent estimate of $E[\mathbf{z}'\xi\xi'\mathbf{z}]$, which is used to weigh moments in accordance to their variance. Implementation of this GMM estimator is not easy, as $\boldsymbol{\xi}$ is by assumption unobservable to researchers and needs to be estimated empirically. Our estimation is done by carefully modifying the Matlab code supplied at Nevo's website.[14]

The question is, what variables would qualify as $\mathbf{z}$ for the moment condition (6)? The common identifying assumption, used in Berry et al. (1995, 1999) and subsequent studies, is that the "location" of observed product attributes for each car model in the characteristics space is exogenous, or at least predetermined prior to the determination of a consumer's valuation of unobserved product-specific attributes. More specifically, Berry et al. used the observed product characteristics, the values of the characteristics summed over all products produced by each firm, and the values of the characteristics summed over all products produced by other firms (BLP instruments). However, there is a growing concern in the literature about the validity of BLP instruments—the location of observed product attributes may indeed correlate with brand images and may be closely related to the average cost only rather than the marginal cost of production. In our case, this concern is even more severe. For example, Toyota's well-known compact-car/hybrid-car strategies suggest that the location of observed attributes such as size and fuel efficiency for their best-selling car models such as Prius and Vitz (known as Yaris in the United States and Europe) may be causally correlated with unobserved brand images consumers have about these products. In addition, in Japan, some car models are sold exclusively through certain sales channels. For example, Toyota Camry and Vitz, two flagship models, are sold only through stores under the franchises of the Corolla and the Netz, respectively. Because we only observe regular market prices, $\boldsymbol{\xi}$ can also include product-specific or franchise-specific sales promotions or marketing campaigns, information on which is not readily available to us. Some of the location variables, such as those for size and fuel efficiency, may then be causally related to these unobservable sales promotions.[15] Indeed, our earlier attempt to estimate the model with BLP instruments resulted in large GMM objective values and price coefficients that are highly sensitive to the random draws of $\boldsymbol{\nu}$.

The aforementioned concern suggests that we need an alternative set of instruments that are correlated with prices (or observable product attributes) yet are uncor-

14. A technical note describing the estimation algorithm in more detail is available in the appendix.

15. One may argue (correctly) that if we believe $\boldsymbol{\xi}$ represents unquantifiable brand images or measurement errors in observed prices, simply including product-level fixed effects in the set of covariates $\mathbf{x}$ might just take care of the concern. The problem with this approach is that if we include brand dummies in the regression, the matrix of $\mathbf{z}'\mathbf{z}$ will be essentially singular, as they do not vary across products and over time. Hence it cannot be inverted. See Nevo (2001) for more detailed discussions on this and related issues.

related, at least causally, with the unobserved product attributes ξ. Two alternatives have been proposed in the literature. The first type is called "Hausman instruments" and uses prices of the same products across different markets. The second type exploits variation in input prices (such as wages) across products, makers, and/or markets. The problems with either set of instruments are well documented in the literature. For the former type of instruments, the identifying assumption is that geographic variation in prices of the same products comes from some supply-side variation across markets, yet the problem is that some of the variation may indeed come from geographic variation in consumer demand for these products. Hence, the exclusion restriction may not hold. For the latter type, either the exclusion restriction may not hold, or when it does, it may be a weak instrument. Input prices may indeed affect the demand for products through changes in consumers' real incomes. When input prices change due purely to exogenous factors, that change most likely affects many products and firms simultaneously, and hence it is difficult to obtain product-specific variation. See Bresnahan (n.d.) and Byrne et al. (2015) for detailed discussions on these points. In our case, neither set of instruments is available because we do not have access to regional market price data, nor do we observe product-specific variation in input prices.

We instead exploit the unique quasi-experimental setup in the Japanese new car market and construct the "tax-location" variables in a manner analogous to Berry et al.: that is, tax amounts, the sums of own-firm tax amounts, and the sums of rival-firm tax amounts.[16] It is straightforward to show that changes in taxes/subsidies can work as equilibrium shifters in almost the same way as cost shifters like BLP and other instruments. As discussed in section 1, the series of green tax policies generated exogenous variations in vehicle taxes across products and over time. Because these taxes are functions of the observable product characteristics (price, weight, and displacement level) and because firms choose markups accounting for the effects of taxes on consumer demand (see eq. [5]), they would surely be correlated with prices. On the other hand, the ES/ETC policy caused the effective tax rates to change four times over the study period, which shifted the location of the vehicle taxes in the characteristics space four times, while the unobserved product characteristics such as style and brand images presumably stayed largely constant. Hence, the location of the vehicle taxes is unlikely to be causally related to the unobserved characteristics.[17] Note, however, that we are not arguing for "perfect" instruments here that would make no correlation between the vehicle taxes and

16. In our earlier trials, we experimented with different tax-location variables, such as tax amounts only, tax rates only, or the location of tax rates. We use the location of tax amounts because they seem to perform the best.

17. We discuss our identifying assumptions more rigorously in a technical note available in the appendix.

the unobserved product attributes ξ. The question here is, like in Berry et al. (1995) and others that followed, which instruments would be better equilibrium shifters that are more likely to be causally unrelated to these unobservables than the other instruments. Unfortunately, we cannot directly test the exogeneity of our instruments. We instead report some indirect evidence below that suggests that our instruments are likely to be more exogenous than BLP instruments.

Figure 3 exhibits scatter plots of our instruments against BLP instruments before and during the policy period (*A*) and kernel densities of our instruments and BLP instruments across different policy periods (*B*). To economize our space, we only report BLP instruments using horsepower divided by weight (HP/weight), but we observe essentially the same patterns with car size as well. We call the values of HP/weight summed over all products produced by each firm "hpw_owfirm" and those summed over all products by other firms "hpw_otherfirm." Analogously, "tax_owfirm" and "tax_otherfirm." First, figure 3*A*1 and *A*2 demonstrate that tax_owfirm have almost a one-to-one relationship to hpw_owfirm, and similarly, tax_otherfirm to hpw_otherfirm, during the pre-policy period. However, during the policy period, substantially more variation is generated in tax_owfirm at each level of hpw_owfirm. This is even more so in the case of tax_otherfirm against hpw_otherfirm. This is indeed the kind of variation that helps the identification of the demand parameters—for the car model in the same location on the product space, its tax location shifted four times during the study period. The question is, whether this variation is caused due to factors related to product attributes or not. Though we cannot directly investigate this question, figure 3*B*1–*B*4 seem to indicate it is very unlikely. The distributions of hpw_owfirm and hpw_otherfirm have gradually shifted over time possibly reflecting technological advances, yet the shifts in the distributions of tax_owfirm and tax_otherfirm do not show any detectable pattern corresponding to those of hpw_owfirm and hpw_otherfirm. In fact, about 37% of the variation in our tax-location variables comes from purely intertemporal variations. Hence, we conclude that our IVs would be better, if not perfect, than the conventional IVs.

There are two additional issues that may raise a question into the validity of our IVs. First, some may argue that automakers might have had significant influence over the design of the policy in favor of particular products (e.g., hybrid cars) or particular automakers in accordance with their brand images. To take care of this concern, we include minicar and hybrid car dummies as well as maker dummies. Moreover, because the tax rates changed four times over time and across products during the study period, the frequent changes should minimize the causal link, if any at all, between the unobservable attributes and the tax rates. Second, Ito and Sallee (2014) provide evidence that Japanese automakers had the tendency to increase their vehicle weights in response to the weight-based fuel economy regulations. If this were indeed true, we may need to be concerned about the endogeneity not only of prices but also of all

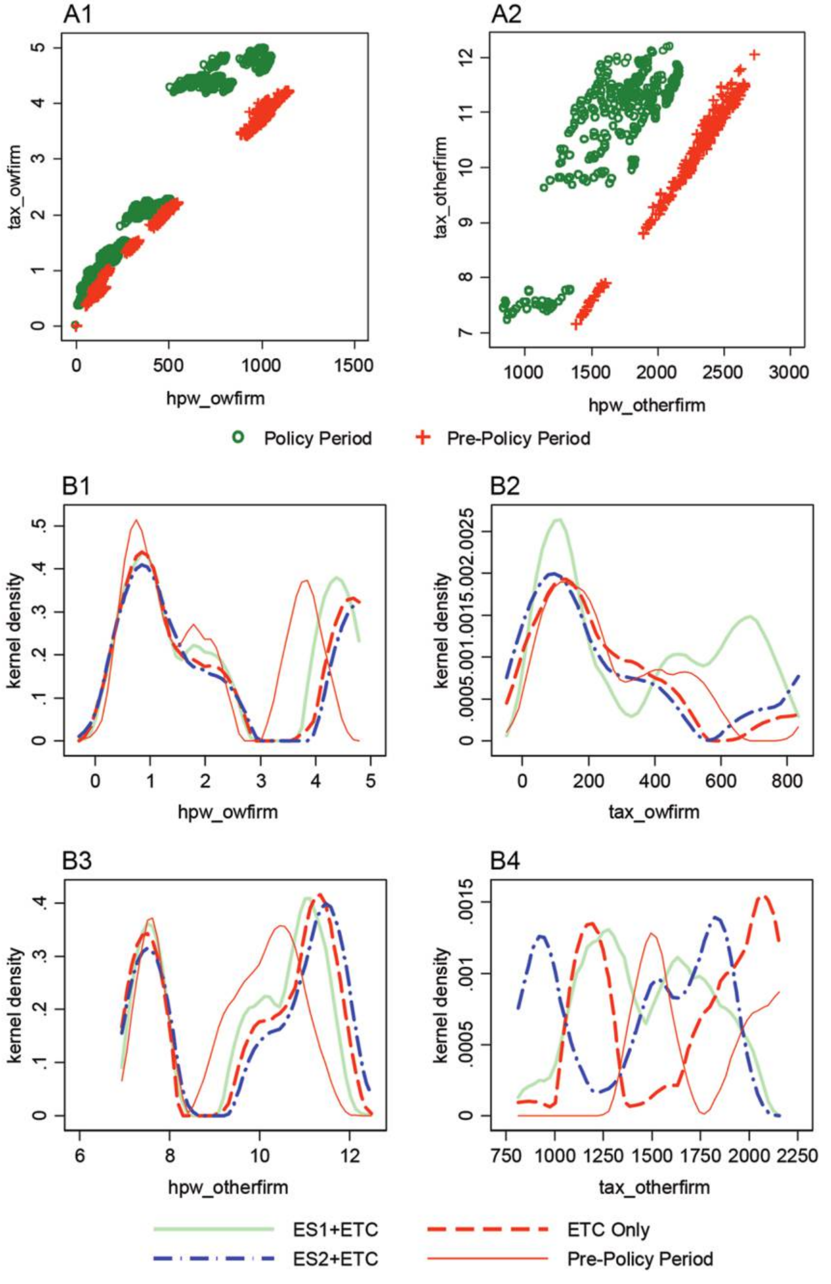

Figure 3. Variation in BLP instruments versus our instruments

the other product attributes.[18] This is the area of active research in the empirical industrial organization literature (see Crawford [2012] for detailed discussions and other related studies on the state of the research). Dealing with this issue is, therefore, outside the scope of our paper.

4. DATA

Our data analysis covers the period from January 2007 to December 2012 with the pre-policy period (before April 2009) as the control period.[19] We chose this study period because detailed sales data on minicars and hybrid cars are available only after 2007. Furthermore, as discussed in section 1, the Japanese government changed the fuel economy standard in 2007 and maintained it throughout the study period, which helps us identify the impact of the ES/ETC policy. We obtained the data on product characteristics and listed prices for all domestic passenger car models from Carsensor .Net, one of the largest used car information services in Japan.

Our price variable is defined as the tax-inclusive price $p^e = (1 + \tau)p$, where p is the market price. Ideally, we would use transaction prices for p, which include other incentives such as options, sales promotions, trade-in values, and preferential interest rates on loans. Unfortunately, we do not observe transaction prices because they usually vary at the individual level and, thus, are hard to come by even in detailed household or price surveys. In Japan, even the most comprehensive price-data services such as Carsenso.net, Goo-net.com, and Kakaku.com do not offer such data on new cars. Hence, we follow Berry et al. (1995, 1999) and Petrin (2002) and use list prices instead and deflate them by the consumer price index. The list prices are generally less variable over time than market prices. However, we still do observe significant changes in list prices over time. The F-test of time dummies (or policy dummies) from regressing prices on all product attributes and these dummies is statistically significant at the 0.01 level. The use of list prices also creates some measurement errors in the price variable. For example, Copeland, Dunn, and Hall (2011) show that transaction prices decline almost monotonically for each vintage within the same model year primarily because sales dealers offer incentives in response to inventory buildups. In case of Japan, sales dealers instead offer incentives in bonus seasons (June/July and December)

18. The concern would be the severest with the MPY variable since the tax incentives are closely tied to fuel efficiency ratings during the policy period. Yet, it is generally difficult for automakers to upgrade fuel efficiency levels given other product attributes within the three-year term. Automakers with technological advantages such as Toyota and Honda might have been able to do so. Therefore, inclusion of maker dummies, we hope, would reduce the extent of the bias if any.

19. Due to space limitation, the descriptive statistics of key variables by quarter and by car segment are reported in the appendix.

and the end of a fiscal year (March). We control for this type of measurement error by including quarter dummies.[20]

To make our analysis comparable to previous studies, we consider the following major product attributes: the ratio of horsepower to car weight (HP/weight), mileage per yen (MPY), car size (Size), and a dummy indicating whether the model has automatic or continuously variable transmission (AT/CVT).[21] Information on displacement, emissions performance, and fuel efficiency was also used to determine the ES and ETC eligibility and to calculate MPY, which is the mileage per liter of gasoline divided by the price of gasoline per liter. We treat the same model produced in different time periods as different models: that is, Honda Accord 2009 versus Honda Accord 2010, as they could be very different due to the rapid technological upgrading. We also make use of some macroeconomic data, such as GDP growth rate, CPI, total number of households, and gasoline prices. The GDP and CPI data are taken from the statistics published by the Cabinet Office of the Japanese government. The data on the number of households are based on the estimates from the Institute of Population and Social Security. The monthly prices of gasoline are from the Institute of Energy Economics in Japan.

The quarterly sales data are obtained from the Japan Automobile Dealers Association (JADA) and the Japan Automotive Products Association (JAPA). Since we have only the total sales for each model and, in many cases, there are many variants (or "grades") of each model, we obtain the corresponding product attributes and prices by taking the averages over all the variants of the same model marketed in the same time period.[22] We confirmed the validity of our treatment in two ways. First, we were able to obtain detailed used-car sales data by grade for a small fraction of the car models. We used the data to verify that the majority of sales are concentrated around the variant of the model that has close proximity to the mean attributes. Second, we estimated

20. There is a related issue concerning the salience of taxation. Li, Linn, and Muehlegger (2014) show that consumers are more responsive to gasoline tax changes than those of tax-inclusive retail gasoline prices because the former are usually permanent and more salient. A similar behavioral response could be possible in the case of car prices and taxes. We did not explore this possibility, leaving that for future research, as we had to rely on list prices and the variation in tax-inclusive prices for identification of the model parameters.

21. Berry et al. (1995, 1999) used a dummy indicating whether the model has air conditioning or not as a default. For our data, this resulted in virtually no variation across models. We thus replaced this variable with the auto transmission dummy. Recently, small-sized cars and hybrid cars increasingly use continuously variable transmission (CVT) to improve fuel efficiency.

22. Although aggregation of choice sets in this manner has been a common practice, it has also been a concern to researchers. Bunch and Brownstone (2013), for example, employ a maximum likelihood approach to addressing the measurement error bias arising from the choice aggregation. To our knowledge, however, the applicability of their approach to the BLP framework has yet to be confirmed. Dealing with this issue is, therefore, left for future research.

the IV logit model using the maximum, minimum, and median as alternatives, and our major results are quite robust to the different choices.

5. ESTIMATION RESULTS

We report the results of our full random-coefficient (RC) logit model in table 2. For all models, we include the same set of variables: constants, effective prices, HP/weight, MPY, size, AT/CVT, GDP growth rate, minicar-hybrid dummies, year-quarter dummies, and maker dummies. Choice of these variables is based on the results of a simple IV logit (appendix). Column 2 reports the result with the conventional BLP instruments whereas column 3 displays the result with our "tax-location" IVs. Column 1 reports the result of the simple IV logit with ours for comparison.

There are in general two ways to explain the effect of each product attribute. For example, a large-sized car might be popular, either because an average consumer places a high value for the large-sized car (i.e., the effect of the mean utility) or because there is a large variance in consumers' tastes for the large-sized car (i.e., the effect of the distribution of the random utility). Statistical significance on mean parameters would get at the significance of the former effects whereas that on standard deviation parameters would get at the latter. If we expect that any of these variables has significant influence on purchase decision, we should observe that at least one of these on each variable is significant. Moreover, for all variables except for the price, we assume a standard normal distribution. Because the normal distribution is symmetric around the mean zero, the signs of the standard deviation parameters should not matter. Hence, the estimates are reported in absolute values. For the price variable, however, we assume a lognormal distribution, for which the sign of the parameter does matter. Hence, we report the raw values for the price coefficients.

With BLP instruments, we observe that virtually all of the mean parameters are significant at the conventional significance levels, except on AT/CVT. However, the sign of the mean parameter on HP/weight was negative. With the insignificant standard deviation parameter, this would imply that virtually all consumers are inclined to buy a car when it is less powerful, holding other characteristics of the car constant—a behavioral response very hard to believe. Moreover, all of the standard deviation parameters (except on constant) are insignificant. A more serious concern is that its sign on the standard deviation parameter of the price variable is positive. Despite its statistical insignificance, this positive standard deviation parameter tends to make the price elasticities of some car models smaller or even positive. We also note that the signs and significance of the estimated parameters are quite sensitive to both the size and seed of random draws of v. These are the problems we did not encounter with our tax-location IVs, and thus, this is another reason why we think our instruments are better, at least in our empirical context.

With our tax-location IVs, the results are more encouraging. All of the mean parameters are significant at the conventional levels with signs in line with our expectation,

Table 2. Estimation Results: Full Random-Coefficients Logit

	IV Logit (1)	RC Logit (2)		RC Logit (3)	
		Mean	SD	Mean	SD
Constant	−22.6680***	−21.1570***	11.8160***	−33.8222***	11.0310***
	(.9927)	(1.3978)	(4.8117)	(2.6268)	(.5164)
Price	−.0142***	−.0157***	.0063	−.0054***	−.0067***
	(.0038)	(.0056)	(.0186)	(.0006)	(.0026)
HP/weight	33.8305***	−42.7740***	4.6993	15.0663***	8.3313***
	(11.2386)	(16.5780)	(21.7300)	(2.2992)	(.4387)
MPY	.0582	.3318***	.1139	.1432***	.1212
	(.0671)	(.0977)	(.3177)	(.0352)	(.1460)
Size	.0010***	.0006***	.0002	.0015***	.0004
	(.0001)	(.0002)	(.0005)	(.0003)	(.0006)
AT/CVT	1.3489***	−.3901	1.2720	.6450***	1.3666***
	(.2812)	(.4037)	(13.2730)	(.1638)	(.3600)
Minicar dummy	✓	✓		✓	
Hybrid dummy	✓	✓		✓	
Maker dummies	✓	✓		✓	
Year dummies	✓	✓		✓	
Quarter dummies	✓	✓		✓	
Macroeconomic variables	✓	✓		✓	
Location IVs used	Taxes	Characteristics		Taxes	
Observations	3,703	3,703		3,703	
GMM objective	. . .	267.2		33.6	

Note. In parentheses are standard errors. Inner-loop tolerance for NFP = 1E-14. Outer-loop tolerance for GMM = 1E-3.
* Significant at .10.
** Significant at .05.
*** Significant at .01.

suggesting that consumers, on average, prefer lower price, higher HP/weight, higher MPY, larger size, and AT/CVT. Moreover, many of the standard deviation parameters are significantly different from zero. The statistically significant standard deviation parameters on price, HP/weight, and AT/CVT suggest that there is indeed a variation in consumer tastes for these attributes. In particular, the large standard deviation parameter on HP/weight implies that some consumers who have very strong preferences for acceleration will still buy a car with a high HP/weight rating when its price gets higher, whereas others who do not have such preferences will not.[23] The result makes sense in the context of Japan. Some areas in Japan have very steep hills, for which some consumers may prefer more powerful cars. Yet, in Japan, public roads are notoriously narrow in urban areas so that a majority of urban consumers may not need powerful cars for daily operations. On the other hand, the standard deviation parameters on MPY and size are not statistically significant. Recall that we included minicar and hybrid car dummies. Within the hybrid or minicar segment, there should be much less variation in MPY and size, which might have removed much of the taste variation on these attributes.

One well-documented advantage of the RC logit model over simpler logit models is that it gives richer and more realistic own- and cross-price elasticities of demand (Nevo 2000, 2001). With simple logit models, own- and cross-price elasticities depend only on the constant parameter, own and cross prices, and observed market shares, which results in (i) nearly constant own-price elasticities and (ii) counterintuitive substitution patterns that do not take into account similarities between car models. With the RC logit, the own- and cross-price elasticities are instead given by:

$$\varepsilon_{jk} = \frac{\partial s_j p_k}{\partial p_k s_j} = \begin{cases} -\dfrac{p_j}{s_j}\displaystyle\int \alpha_i s_{ij}(1 - s_{ij})dP(v_i) & \text{if} \quad j = k \\ \dfrac{p_k}{s_j}\displaystyle\int \alpha_i s_{ij} s_{ik} dP(v_i) & \text{if} \quad j \neq k \end{cases}, \tag{8}$$

where s_{ij} is the choice probability for car model j by individual i. In this expression, each individual has different price elasticities, which are averaged out to yield mean elasticities.

Table 3 displays the sales and product characteristics, the estimated own-price elasticities as well as price-cost markups for 15 top-selling car models, based on our preferred model 3 for year 2012. Our estimated price elasticities for these car models range from −2.0 to −3.2, and the sales-weighted average own-price elasticity for all car models in 2012 was −2.66. The estimated elasticities seem roughly comparable

23. We follow Murdock (2006) for this intuitive interpretation of parameter estimates.

Table 3. Product Characteristics, Estimated Elasticities, and Implied Price-Cost Markups of Top-Selling Car Models for 2012

Brand Name	Sales	Price	HP/Weight	MPY	Size	AT/CVT	Own-Price Elasticities	Price-Cost Markups
Toyota Prius	317,675	256	.068	21.9	7,881	1.000	−2.821	.359
Toyota Aqua Hybrid	266,567	163	.068	25.7	7,135	1.000	−2.165	.468
Daihatsu Mira	218,295	107	.066	18.1	6,398	1.000	−2.033	.499
Honda N-Box	211,155	143	.060	14.0	6,656	1.000	−2.386	.423
Daihatsu Tanto	170,609	134	.054	16.3	6,620	1.000	−2.316	.435
Daihatsu Move	146,016	128	.064	17.1	6,500	1.000	−2.276	.443
Honda Fit Hybrid	116,212	174	.078	19.3	7,342	.927	−2.496	.404
Suzuki Alto	112,002	106	.070	14.5	6,398	.775	−2.228	.458
Toyota Vitz	105,611	138	.091	14.5	7,095	.935	−2.435	.426
Honda Fit	93,041	144	.099	12.8	7,275	.699	−2.442	.416
Nissan Note	85,330	141	.088	13.1	7,274	.905	−2.449	.419
Nissan Moco	66,460	120	.068	14.7	6,495	1.000	−2.338	.438
Honda StepWgn	63,707	269	.094	9.8	8,207	1.000	−3.259	.311
Suzuki Palette	60,136	128	.059	12.7	6,610	1.000	−2.464	.410
Mazda Demio	57,820	123	.095	14.7	7,070	.702	−2.391	.428

Note. A hybrid version of the same car model is treated as a separate model, so the sales and other product characteristics of the nonhybrid model exclude those of the hybrid model. Price = average list price in JP¥10,000; HP/weight = HP/weight in kw/kg; MPY = mileage in km per JP¥; Size = the sum of length, width, and height in millimeters; AT/CVT = the fraction of the car grades that have automatic (AT) or continuously variable (CVT) transmission. All quantities are simple averages, except for sales, which is the sum of sales for 2012.

to the reported elasticities in Berry et al. (1995), which range from –3 to –4.5.[24] Partly because our elasticity estimates are slightly smaller than BLP estimates, our estimates of the price-cost margins were slightly higher than those of Berry et al. In our case, estimated markups for these top-selling models range from 0.31 to 0.49, whereas Berry et al.'s sales-weighted average markup was 0.23. We believe this difference can be partly attributed to Japanese automakers' cost structures, which are known to have a high ratio of fixed costs to variable costs. Using Toyota's operating profit margin for automobile sales in 2012, and assuming that roughly one-third of the costs are fixed costs, we arrive at an accounting-based estimate of its price-cost markup of 0.42, a number very much in line with our estimates. Hence, we conclude that consumer demand for these popular car models in Japan is more inelastic than US consumer demand, which allows Japanese automakers to maintain high markups and cover their large fixed costs. Though not reported, we also estimate cross-price elasticities, elasticities with respect to product attributes, and substitution probabilities to the outside option and confirmed that all of them are roughly consistent with Berry et al. (1995). These elasticity estimates are available in the appendix.

6. POLICY EVALUATION

6.1. Construction of Counterfactuals

We now use the estimated product-level demand and marginal costs to run several counterfactuals for policy evaluation. To isolate the pure impact of each policy, we first simulate a no-policy counterfactual in which vehicle taxation was maintained at the pre-policy level during the 2009–12 policy period. Quantification of each policy's economic impacts is reported in relative terms to this no-policy counterfactual. We do this because we need some benchmark against which we calculate compensating variation (see [10] below). We also ran two additional counterfactual simulations. The first is called the ETC only policy, in which only the ETC program was implemented. The second is the second-best optimal feebate (SBF) policy, the details of which are discussed below.

One advantage of our structural estimation approach is that it enables us to fully simulate firms' equilibrium responses in pricing strategies to any changes in tax incentives. Given the estimated demand and marginal costs, we can re-solve (5) for a vector of new Bertrand-Nash equilibrium prices for a given tax/subsidy policy $\boldsymbol{\tau}$. The problem, however, is that doing so requires solving 24 systems of nonlinear equations, each having dimension of roughly 150 with no apparent bounds on the solution space. This

---

24. In Bento et al. (2009), price elasticities range from –0.88 to –1.97, which are much smaller in absolute terms than ours. However, our estimates must be more comparable to Berry et al. than to Bento et al. because car models are aggregated in the latter, which should make the estimated demand more inelastic. We thank our reviewer for pointing this out.

is computationally quite challenging. We therefore use the following approximation à la Nevo (1997) and Knittel and Metaxoglou (2014) instead:

$$\mathbf{p}^{\text{new}}(\boldsymbol{\tau}^{\text{new}}) = \widehat{\mathbf{mc}} + \Omega^{-1}(\mathbf{p}^{\text{old}}, \boldsymbol{\tau}^{\text{new}})\mathbf{s}^{e}(\mathbf{p}^{\text{old}}, \boldsymbol{\tau}^{\text{new}}). \tag{9}$$

From here on, the hat indicates estimates based on the RC logit model 3.

To quantify the policy impact on consumer surplus, we use, as in previous studies, the compensating variation measure of the changes in effective prices à la Small and Rosen (1981):

$$\begin{aligned} CS(\mathbf{p}^{m}(\boldsymbol{\tau}^{m}), \boldsymbol{\tau}^{m}; \hat{\boldsymbol{\theta}}) &\simeq CV(\mathbf{p}^{e,m}; \mathbf{p}^{e,0}, \hat{\boldsymbol{\theta}}) \\ &= \sum_{i} \frac{\ln\left(\sum_j \exp(V_{ij}(\mathbf{p}^{e,m})) - \ln(\sum_j \exp(V_{ij}(\mathbf{p}^{e,0})\right)}{\hat{\alpha}_i}, \end{aligned} \tag{10}$$

where $\mathbf{p}^{e,m}$ is a vector of tax-inclusive prices under policy m and $\mathbf{p}^{e,0}$ that of the no-policy benchmark. Note that this compensating variation measure does not include the negative externality cost of vehicle emissions. We do not quantify the monetary value of the vehicle emissions because with lack of data on vehicle miles traveled, we cannot fully quantify the impact on vehicle emissions.

Producer surplus and tax revenues are computed as:

$$PS(\mathbf{p}^{m}(\boldsymbol{\tau}^{m}), \boldsymbol{\tau}^{m}; \hat{\boldsymbol{\theta}}) = \sum_{f \in \mathcal{F}} \sum_{j \in \mathcal{J}_f} (p_j^m - \widehat{mc}_j) M s_j(\mathbf{p}^{e,m}), \tag{11}$$

$$TR(\mathbf{p}^{m}(\boldsymbol{\tau}^{m}), \boldsymbol{\tau}^{m}; \hat{\boldsymbol{\theta}}) = \sum_{f \in \mathcal{F}} \sum_{j \in \mathcal{J}_f} p_j^m \tau_j^m M s_j(\mathbf{p}^{e,m}). \tag{12}$$

Total surplus is then calculated as $TS = CS + PS + TR$, which we use as a metric for our policy evaluation. Making inferences about the policy impacts also requires us to obtain the standard errors of the estimated impacts. Doing so in our context is not easy. We could linearize the policy impacts in the parameters and use the delta method. However, as the policy impacts are highly nonlinear in the parameters, this approach may not be appropriate. Berry et al. (1999) instead use a Monte Carlo approach, taking draws from the estimated asymptotic normal distribution of the parameters. We took 500 draws of parameters and calculated the standard deviations of the policy impacts as the estimates of the standard errors.

Following Ramsey (1927) and other subsequent studies, the second-best feebate policy $\boldsymbol{\tau}^{sbf}$ should be the solution to:

$$\max_{\boldsymbol{\tau}} TS(\mathbf{p}(\boldsymbol{\tau}), \boldsymbol{\tau}; \hat{\theta}) \text{ subject to } TR(\mathbf{p}(\boldsymbol{\tau}), \boldsymbol{\tau}; \hat{\theta}) \geq \overline{TR}, \tag{13}$$

for some revenue target $\overline{TR}$. We may naturally choose to set $\overline{TR} = 0$ because correcting for imperfect competition generally requires tax shifting, which would result in deadweight losses in other sectors. However, doing so obscures our policy evaluation because the preexisting tax scheme generates large tax revenues and, therefore, setting the revenue target at zero would generate too large gains in consumer surplus and producer surplus at the expense of large losses in tax revenues. On the other hand, setting $\overline{TR}$ at the no-policy level would be too restrictive as it would effectively eliminate use of subsidies. Hence, as an intermediate case, we set $\overline{TR}$ at the ES/ETC level.

One complication in solving (13) is that *CS*, *PS*, and *TR* are again systems of highly nonlinear equations, each having a dimension of the number of products (recall how we define $s(\cdot)$ and $p(\cdot)$). Hence, solving this optimization program is computationally quite demanding, if not infeasible. We instead consider linearizing the feebate scheme in product attributes. That is,

$$\boldsymbol{\tau}^{sbf} = \mathbf{y}'\gamma,$$

where $\boldsymbol{\gamma}$ is a vector of parameters and $\mathbf{y}$ is a vector of product attributes over which tax rates are varied. We plug this in (13) and solve for the optimum $\boldsymbol{\gamma}^{*}$. This linearization gives us multiple benefits. First, it helps us reduce the dimension of the search for optimization substantially while allowing for differential tax/subsidy rates across car models. Second, it makes the tax/subsidy rates a linear function of quantifiable attributes. Hence, its implementation by the regulatory authority is quite simple in practical settings.

There is a question as to what variables are to be included in $\mathbf{y}$. To allow for sufficient variation, we include all key product attributes in $\mathbf{x}$, with the following exceptions. First, the theory of optimal taxation tells us that corrective subsidies for imperfect competition must be closely linked to markup levels. Hence, we include the estimated markups in place of the price variable. To fix optimal feebates, we use the estimated markups under the observed market conditions (i.e., the ES/ETC policy). Second, we replace MPY with $1/MPG$ to keep tax rates invariant with changes in gasoline prices while allowing the tax/subsidy rates to vary with carbon emissions rates. Indeed, this is consistent with the idea that the fee/rebate must be proportional to fuel consumption per mile instead of *MPG* (Anderson et al. 2011). Third, we also include $\boldsymbol{\tau}^{es/etc}$, the tax rates under the ES/ETC policy, in $\mathbf{y}$. Doing this enables us to assess how $\boldsymbol{\tau}^{sbf}$ differs from $\boldsymbol{\tau}^{es/etc}$. If $\boldsymbol{\tau}^{es/etc}$ is indeed close enough to the second-best feebate policy, (13) should return 1 as a coefficient on $\boldsymbol{\tau}^{es/etc}$ and 0 on other attributes. On the other hand, if $\boldsymbol{\tau}^{sbf}$ differs substantially from $\boldsymbol{\tau}^{es/etc}$, the coefficient on $\boldsymbol{\tau}^{es/etc}$ should be different from 1 and those on other attributes significantly different from 0.[25]

25. We also experimented with square terms and other product attributes, but inclusion of these additional variables did not significantly improve the objective value. Hence, we only report the result from the parsimonious set of variables without square terms.

6.2. Properties of the Second-Best Feebates

We first investigate the properties of these two feebate systems. Figure 4 displays scatter plots of tax rates $(1 + \tau)$ under the ES/ETC and the SBF against observed markup levels, carbon emissions rates, HP/weight, and size. The SBF tax rates indeed have a negative relationship to the observed markup levels. This is in line with our expectation because all else equal, there is a large welfare gain from subsidizing cars with high markups. Interestingly, the SBF tax rates have roughly positive relationships to all the other product characteristics. This is particularly evident on car size. As a result, car models that are larger with higher emissions rates are heavily taxed while those that are smaller with lower emissions rates are heavily subsidized. We emphasize here that we obtained the results despite the fact that the optimization program (13) does not explicitly take into account the environmental benefits. Importantly, the SBF generates substantial variation within the same car type. That is, holding some car attributes, say, emissions rates at 0.15 or HP/weight at 0.7, tax rates can range

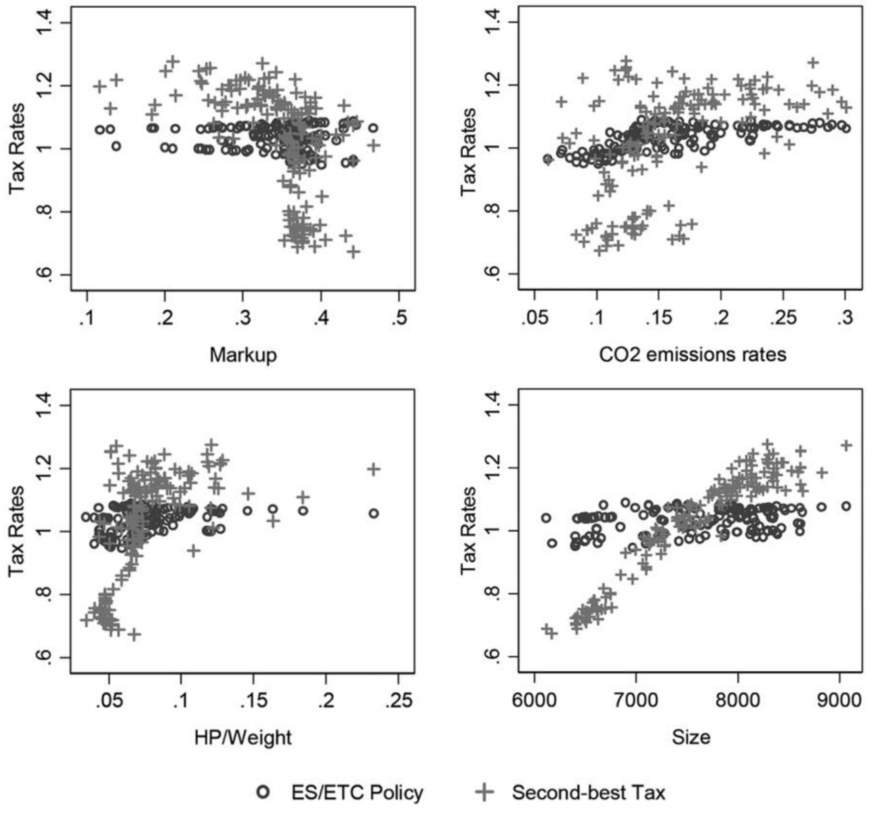

Figure 4. SBF versus ES/ETC tax rates $(1 + \tau)$ for year 2012

from −30% to +30%. This indeed signifies the importance of accounting for product-level demand in designing optimal feebates. Intuitively, even within the same car segment, different car models might have different levels of market power due to differences in substitutability, maker reputation, and brand image. The SBF properly accounts for all these as our demand estimation does so through random coefficients, maker dummies, and instruments.

Next, we investigate how these design features translate into firms' strategic responses in pricing equilibrium. Table 4 reports tax rates $(1 + \tau)$, equilibrium prices p, markups mu, tax-inclusive prices p^e, and sales quantities q for five top-selling car models in 2012 under alternative policy counterfactuals. All these top-selling models received substantial tax reductions under the ES/ETC relative to the preexisting tax system. In response, firms producing these models were able to increase their market prices yet to keep their tax-inclusive prices at lower levels than under no policy, which resulted in higher sales quantities than would have been. These equilibrium responses are indeed consistent with economic theory.

There is indeed an important interplay between the SBF scheme and the strategic pricing responses. The SBF mandates a very high tax rate (15%) to Toyota Prius, a moderate subsidy (−2%) to Toyota Aqua Hybrid, and large subsidies (−26%, −23%, and −24%) to Daihatsu Mira, Honda N-Box, and Daihatsu Tanto. All these cars are highly fuel-efficient cars, but the last three cars are very small-sized minicars, whereas the first two are regular-sized hybrid cars. Toyota Prius is slightly larger in size with a smaller markup than Toyota Aqua Hybrid. These differences in product attributes roughly explain these highly differentiated tax rates. Firms make intricate responses to these tax rates. Toyota would decrease Prius's price only by 3.4% relative to the ES/ETC policy for a tax rate increase of 16 percentage points, whereas it would increase Aqua's price by 2.9% for a subsidy decrease of only 1 ppt. Toyota's Prius and Aqua are indeed close substitutes to Daihatsu Mira and Tanto and Honda N-Box, as all these cars are highly fuel efficient. Because these minicars were able to raise prices substantially due to large subsidies, Toyota was able to maintain relatively high prices and markups for these models. Yet, this Toyota's pricing would come with the expense of large losses in sales for these models and gains for their competitors.

6.3. Economic Impacts of Alternative Feebates

Figure 5 demonstrates the impacts of alternative policies on the sales shares of hybrid cars and minicars. As expected, the ES/ETC policy increased the share of hybrid cars relative to the ETC-only counterfactual, and even more so relative to the no-policy counterfactual. Interestingly, however, the sales share of minicars was lower under the ES/ETC policy than under no policy. We believe this occurred because, although minicars were equally eligible for the same tax incentives, the size of these tax incentives relative to the preexisting tax system was smaller for minicars than for hybrid cars (see sec. 1). In contrast, the SBF policy is predicted to decrease the share of hybrid cars relative to

Table 4. Simulated Impacts of Alternative Policies on Prices, Markups, and Sales for Top-Selling Car Models in 2012

	No Policy (Est.)	ETC Only (Est.)	ES/ETC (Obs.)	SBF (Est.)
Toyota Prius:				
Price, p	249.9	252.7	255.9	247.1
Tax rates, $(1+\tau)$	1.07	1.03	.99	1.15
Tax-incl. price, p^e	268.0	260.8	254.0	285.0
Markup, mu	.34	.35	.36	.34
Sales, q	293.7	307.8	318.0	231.1
Toyota Aqua Hybrid:				
Price, p	157.2	159.2	162.6	167.3
Tax rates, $(1+\tau)$	1.08	1.04	.97	.98
Tax-incl. price, p^e	170.2	165.1	158.4	163.9
Markup, mu	.45	.46	.47	.48
Sales, q	249.0	256.7	266.8	246.1
Daihatsu Mira:				
Price, p	104.0	104.4	106.6	121.9
Tax rates, $(1+\tau)$	1.05	1.03	.96	.74
Tax-incl. price, p^e	108.9	107.2	102.5	90.6
Markup, mu	.49	.49	.50	.56
Sales, q	215.1	212.2	218.4	256.1
Honda N-Box:				
Price, p	141.0	141.5	143.1	158.6
Tax rates, $(1+\tau)$	1.04	1.02	.98	.77
Tax-inclp Price, p^e	147.0	144.9	140.9	122.4
Markup, mu	.41	.42	.42	.48
Sales, q	211.2	209.5	211.1	281.4
Daihatsu Tanto:				
Price, p	131.3	132.1	134.0	148.6
Tax rates, $(1+\tau)$	1.04	1.02	.96	.76
Tax-incl. price, p^e	137.0	134.3	129.2	112.8
Markup, mu	.42	.43	.43	.49
Sales, q	165.3	165.4	170.7	219.4

the no-policy counterfactual while increasing the share of minicars. These heterogeneous impacts on sales in different car segments had ambiguous impacts on the time path of sales-weighted average fuel efficiency. In the first three years of the policy period, the SBF seems to outperform the ES/ETC policy, yet in the last year, the difference between the two largely disappears.

How do these impacts on the time path of sales in different car segments translate into economic efficiency? We investigate this in table 5, which reports the simulated

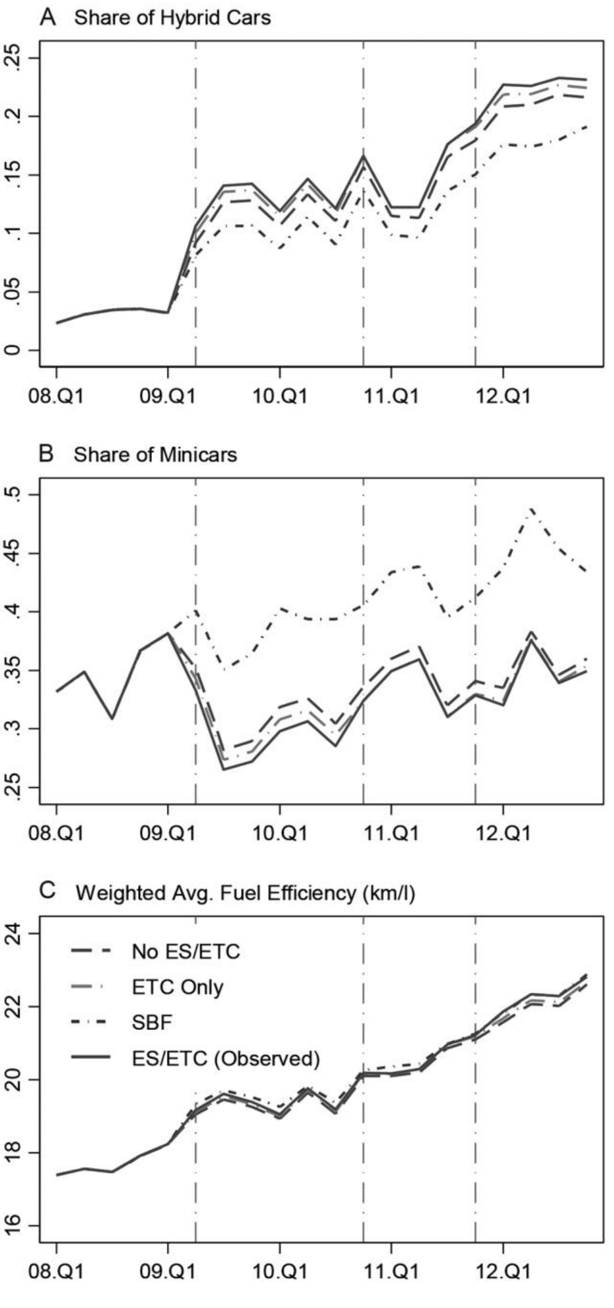

Figure 5. The impacts of alternative feebate policies on the time path of key variables

Table 5. Simulated Impacts of the ES/ETC and SBF Policies on Compensating Variation, Industry Profits, Tax Revenues, Vehicle CO_2 Emissions, and Sales-Weighted Average Fuel Efficiency

	2009		2010		2011		2012		Average	
	Coefficient	SE	Coefficient	SE	Coefficient	SE	Coefficient	SE	Coefficient	SE
Compensating variation (billion ¥):										
ES/ETC	146.80	(389.6)	240.16	(544.0)	114.24	(305.4)	285.20	(328.5)	196.60	(390.0)
ETC only	72.05	(301.7)	127.06	(450.1)	106.52	(298.1)	131.14	(235.5)	109.19	(320.2)
SBF	114.48	(1,666.7)	130.70	(1,842.0)	16.10	(1,610.5)	150.25	(1,712.3)	102.88	(1,706.3)
Industry profits (billion ¥):										
ES/ETC	74.25	(93.8)	116.15	(97.3)	67.04	(56.4)	149.98	(75.2)	101.85	(80.1)
ETC only	33.89	(55.8)	59.71	(62.1)	61.10	(53.3)	60.70	(36.5)	53.85	(51.6)
SBF	218.78	(222.1)	273.04	(251.7)	208.93	(254.7)	340.15	(241.8)	260.22	(238.8)
Tax revenues (billion ¥):										
ES/ETC	−147.78	(372.1)	−238.98	(519.7)	−112.09	(282.1)	−286.91	(307.7)	−196.44	(368.2)
ETC only	−71.17	(290.4)	−124.11	(432.6)	−104.14	(275.6)	−128.66	(221.6)	−107.02	(303.6)
SBF	−193.57	(1,783.0)	−226.20	(1,945.9)	−90.05	(1,746.1)	−286.62	(1,872.4)	−199.11	(1,835.5)
Total surplus (billion ¥):										
ES/ETC	73.26	(105.4)	117.33	(111.9)	69.18	(74.6)	148.27	(88.7)	102.01	(94.5)
ETC only	34.76	(64.2)	62.66	(73.9)	63.49	(71.0)	63.17	(47.2)	56.02	(63.8)
SBF	139.69	(129.0)	177.54	(175.8)	134.98	(162.8)	203.78	(134.9)	164.00	(141.4)
Vehicle carbon emissions (1,000 tons of CO_2):										
ES/ETC	11.21	(71.3)	36.69	(103.2)	3.14	(117.5)	4.18	(113.1)	13.80	(100.3)
ETC only	4.47	(56.0)	18.81	(87.0)	4.44	(114.7)	4.75	(87.9)	8.12	(85.8)
SBF	−54.68	(68.1)	−33.72	(79.7)	−63.41	(170.0)	−17.92	(203.1)	−42.43	(121.5)
Sales-weighted average fuel efficiency (km/L):										
ES/ETC	17.29	(2.7)	17.58	(2.6)	19.10	(3.9)	19.54	(4.0)	19.41	(.8)
ETC only	17.29	(2.7)	17.59	(2.6)	19.05	(3.9)	19.50	(4.0)	19.37	(.8)
No policy	17.29	(2.7)	17.59	(2.6)	19.00	(3.8)	19.43	(4.0)	19.32	(.8)
SBF	17.87	(2.4)	18.12	(2.3)	19.32	(3.6)	19.67	(3.8)	19.68	(.8)

Note. ES = Ecocar subsidy; ETC = Ecocar tax credits.; SBF = second-best optimal feebate.

impacts of three policy scenarios on compensating variation, industry profits for domestic automakers, tax revenues, sales-weighted average fuel efficiency, and vehicle CO_2 emissions.[26] All values, except average fuel efficiency, are in terms of changes relative to the no-policy counterfactual.

The ES/ETC policy indeed had a positive impact on both consumer welfare and industry profits, with increases of ¥196.6 billion and ¥101.9 billion annually relative to no policy. The increase in compensating variation and producer surplus more than offset the decrease in tax revenues and resulted in an increase in economic surplus of ¥102 billion per year. The ES/ETC policy performed better than the ETC-only policy. This occurred presumably because the ecocar subsidy worked as a corrective instrument for imperfect competition. In line with our expectation, the SBF policy indeed induced a sizable gain in economic surplus over the ES/ETC policy. We emphasize here that the SBF policy did not have to increase public expenditures more than what the ES/ETC policy did in order to obtain this gain in economic surplus. Interestingly, the gain from the SBF comes mainly from that in producer surplus, accounting more than 70% of the total gain. This is in sharp contrast to the ES/ETC policy, the gain from which largely comes from the gain in consumer surplus.[27]

As for environmental benefits, the ES/ETC policy indeed led to a small increase in average fuel efficiency from 19.32 (km/L) to 19.41 (km/L) (or equivalently, 45.45 mpg

---

26. The vehicle emissions reported here are only approximate and essentially measure how much of CO_2 emissions would be emitted, on average, from the cars sold during each period t annually. We would ideally estimate the demand for driving jointly with the automobile demand by combining the market-level and household-level data on car ownership and utilization. Bento et al. (2009) jointly estimate the two types of demand using the household-level data only. To our knowledge, no household-level data that are sufficiently comprehensive enough to allow researchers to make accurate inferences about economic behaviors are available in Japan during our study period. Hence, we instead adopt the following measure of expected aggregate emissions, in a manner analogous to Fullerton and Gan (2005) and Klier and Linn (2015):

$$E_t \simeq \sum_{j \in J_t} M_t s_{jt} \left(\frac{EPG \times VMT}{MPG_{jt}} \right),$$

where MPG_{jt} is the expected miles per gallon of gasoline for car model j, VMT is the expected annual vehicle miles traveled, and EPG is the average CO_2 emissions per unit of gasoline. We assume EPG is constant and use the EPA estimate of 8.887 kilograms per gallon (EPA 2014). For VMT, we use the average annual driving distance of 10,575 km in Japan (MLITT 2012).

27. As in Berry et al. (1999), the standard errors of the policy impacts are generally larger than the size of the policy impacts. We see particularly large standard errors of the simulated policy impacts for the SBF policy, particularly on consumer surplus and tax revenues. This is somewhat anticipated because the SBF by design must be solved for each draw of parameters, yet the simulation fixes the SBF scheme for all draws. Hence, the SBF can be very far from optimum for some

to 45.66 mpg). Klier and Linn (2013) report the estimated impact of a \$1 increase in fuel price per gallon on fuel efficiency in the United States and Europe to be between 0.15 and 1.30 mpg (see their table 8). We, therefore, conclude that the estimated impact of the ES/ETC was not too large. In contrast, somewhat unexpectedly, the SBF policy resulted in a larger increase of 0.36 (km/L) to 19.68 (km/L) (or 0.84 mpg to 46.28 mpg). Interestingly, the improvement in average fuel efficiency did not translate into a reduction in annual gasoline-consumption-related CO_2 emissions under the ES/ETC policy, whereas it did under the SBF policy. The ES/ETC policy indeed increased annual vehicle CO_2 emissions relative to no policy by 13,800 tons. This occurred because the ES/ETC policy increased the likelihood of car purchase. In contrast, the SBF policy led to a small reduction in vehicle CO_2 emissions. A somewhat discouraging observation is that the Japanese government's decision to add the ES policy on top of the ETC policy seemed to have induced a further increase in vehicle emissions rather than decreasing them. In sum, our results suggest that the SBF policy indeed outperforms the ES/ETC policy in terms of both economic efficiency and environmental metrics and that the welfare gain is quite large.

7. CONCLUDING REMARKS

To quantify the economic impact of alternative feebate policies, a random-coefficients logit model was estimated for quarterly automobile sales data in Japan between 2007 and 2012. We exploited the unique quasi-experimental setup created through a series of green car tax policies in the Japanese new car market in two ways. First, we constructed a new set of instrumental variables, arguing that the location of vehicle taxes over the product space is more exogenous than the conventional product-location IVs. Second, we took advantage of the large and persistent variation in the effective prices of cars that varied across models and over time in identifying the price elasticities. The estimated product-level demand was then used to simulate counterfactual Bertrand-Nash equilibria under alternative policy scenarios. We also proposed an approach for designing an optimal feebate scheme utilizing the product-level demand system and solving for product-specific tax/subsidy rates as a function of markups and product attributes. Our approach is relatively simple to implement in practical settings and is general enough to be applicable in other markets characterized by oligopolistic industries with multi-product firms and consumption externality.

We found evidence that indicates (i) our tax-location IVs are valid, (ii) Japan's feebate policy implemented since 2009 led to a sizable increase in economic surplus

draws (we did observe such incidences). Hence, we are not too concerned with the size of standard errors for the SBF policy. We could, of course, solve for a different SBF policy for each draw of parameters. We did not do so for two reasons. First, it makes the interpretation of the standard errors difficult. Second, it takes us approximately 10 hours to solve for the optimum on each draw. Hence, solving that for 500 draws would be too impractical.

relative to the preexisting tax system, and (iii) the second-best feebate scheme induced even larger improvements in both economic surplus and sales-weighted average fuel efficiency over Japan's feebate policy, without the need for any further decrease in tax revenues. The optimal feebate scheme also exhibits a number of characteristics that are significantly different from those of Japan's feebate policy. Our results suggest that the optimal tax/subsidy rates must be substantially more varied across car models, have a roughly negative relationship to the estimated markup levels, but have positive relationships to carbon dioxides emissions rates, HP/weight, and vehicle size.

While our study offers several advantages over the previous studies, it also has several important limitations. First, due to data limitation, we could not estimate the car ownership and utilization decisions jointly. Recent studies have shown that (i) combining the market-level data with household-level data (Berry et al. 2004; Petrin 2002) and (ii) imposing cross-equation restrictions by imposing the Roy's identity for the demand for car utilization (Bento et al. 2009) would improve the consistency and efficiency of the estimates. Second, we did not investigate the effects of the feebates on used car and scrap markets. In theory, the policy must have had two counteracting effects. On one hand, the policy would induce consumers to buy new, fuel-efficient cars and, therefore, may facilitate retirement of old, fuel-inefficient cars. On the other hand, the policy would also induce consumers to buy used cars because it would increase the supply of used cars, thereby decreasing the prices of used cars. Hence, it seems largely an empirical question whether inclusion of used car and scrap markets would increase or decrease the estimated impacts on aggregate emissions. Third, there is a growing interest in endogenizing product attributes in the BLP-type estimation of product-level demand, a complication we did not explore in the paper that may have important welfare implications in our context (see Crawford, Shcherbakov, and Shum [2015] for discussions). Addressing these important limitations would define new and important agendas for future research.

REFERENCES

Anderson, Soren T., Ian W. H. Parry, James M. Salee, and Carolyn Fischer. 2011. Automobile fuel economy standards: Impacts, efficiency, and alternatives. *Review of Environmental Economics and Policy* 5 (1): 89–108.

Auerbach, Alan J., and James Hines. 2003. Perfect taxation with imperfect competition. In *Public finance and public policy in the new century*, ed. S. Cnossen and H. Sinn. Cambridge: MIT Press.

Berry, Steven. 1994. Estimating dicrete-choice models of product differentiation. *RAND Journal of Economics* 25 (2): 242–62.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995 Automobile prices in market equilibrium. *Econometrica* 63 (4): 841–90.

———. 1999. Voluntary export restraints on automobiles: Evaluating a trade policy. *American Economic Review* 89 (3): 400–430.

———. 2004. Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of Political Economy* 112 (1): 68–105.

Bento, Antonio, Lawrence Goulder, Mark Jacobsen, and Roger von Haefen. 2009. Distributional and efficiency impacts of increased US gasoline taxes. *American Economic Review* 99 (3): 667–99.

Bresnahan, Timothy F. n.d. The Apple-Cinnamon Cheerios war: Valuing new goods, identifying market power, and economic measurement. Unpublished manuscript, http://www.stanford.edu/126tbres/.

Bunch, David S., and David Brownstone. 2013. Estimating demand for high fuel economy vehicles. Paper presented at International Choice Modelling Conference, Sydney, July.

Byrne, David P., Susumu Imai, Neelam Jain, Vasilis Sarafidis, and Masayuki Hirukawa. 2015. Identification and estimation of differentiated products models using market size and cost data. Working paper, https://sites.google.com/site/susumuimai1/home.

Copeland, Adam, Wendy Dunn, and George Hall. 2011. Inventories and the automobile market. *RAND Journal of Economics* 42 (1): 121–49.

Crawford, Gregory S. 2012. Accommodating endogenous product choices: A progress report. *International Journal of Industrial Organization* 30:315–20.

Crawford, Gregory S., Oleksandr Shcherbakov, and Matthew Shum. 2015. The welfare effects of endogenous quality choice in cable television markets. University of Zurich, Department of Economics Working paper 202.

Crawford, Gregory S., and Ali Yurukoglu. 2012. The welfare effects of bundling in multichannel television markets. *American Economic Review* 102:301–17.

D'Haultfœuille, Xavier, Pauline Givord, and Xavier Boutin. 2014. The environmental effect of green taxation: The case of the French bonus/malus. *Economic Journal* 124 (578): 444–80.

Dube, Jean-Pierre, Jeremy T. Fox, and Che-Lin Su. 2012. Improving the numerical performance of static and dynamic aggregate discrete choice random coefficients demand estimation. *Econometrica* 80 (5): 2231–67.

EPA (US Environmental Protection Agency). 2014. Fact sheet: Greenhouse gas emissions from a typical passenger vehicle. Office of Transportation and Air Quality, EPA-420-F-14-040a, https://www.epa.gov/sites/production/files/2016-02/documents/420f14040a.pdf.

Feng, Ye, Don Fullerton, and Li Gan. 2013. Vehicle choices, miles driven, and pollution policies. *Journal of Regulatory Economics* 44:4–29.

Fischer, Carolyn. 2010. Imperfect competition, consumer behavior, and the provision of fuel efficiency in light-duty vehicles. Resources for the Future Discussion paper 10–60, Washington, DC.

Fullerton, Don, and Li Gan. 2005. Cost-effective policies to reduce vehicle emissions. *American Economic Review* 95 (2): 300–304.

Fullerton, Don, and Sarah E. West. 2002. Can taxes on cars and gasoline mimic an unavailable tax on emissions? *Journal of Environmental Economics and Management* 43:135–57.

Goldberg, Pinelopi Koujianou. 1998. The effects of the corporate average fuel efficiency standards in the US. *Journal of Industrial Economics* 46 (1): 1–33.

Innes, Robert. 1996. Regulating automobile pollution under certainty, competition, and imperfect information. *Journal of Environmental Economics and Management* 31:219–39.

Ito, Koichiro, and James M. Sallee. 2014. The economics of attribute-based regulation: Theory and evidence from fuel-economy standards. NBER Working paper no. 20500, National Bureau of Economic Research, Cambridge, MA.

Klier, Thomas, and Joshua Linn. 2013. Fuel prices and new vehicle fuel economy: Comparing the United States and Western Europe. *Journal of Environmental Economics and Management* 66:280–300.

———. 2015. Using vehicle taxes to reduce carbon dioxide emissions rates of new passenger vehicles: Evidence from France, Germany, and Sweden. *American Economic Journal: Economic Policy* 7 (1): 212–42.

Knittel, C., and K. Metaxoglou. 2014. Estimation of random coefficients demand models: Two empiricists' perspectives. *Review of Economics and Statistics* 96 (1): 34–59.

Ley, Eduardo, and Jessica Boccardo. 2010. The taxation of motor fuel: International comparison. World Bank Policy Research Working paper 5212, Washington, DC.

Li, Shanjun, Joshua Linn, and Erich Muehlegger. 2014. Gasoline tax and consumer behavior. *American Economic Journal: Economic Policy* 6 (4): 302–42.

MLITT (Ministry of Land, Infrastructure, Transport and Tourism). 2012. Annual statistical report on motor vehicle transport. Transport Research and Statistics Office. http://www.mlit.go.jp/k-toukei/search/pdf/06/06201500a00000.pdf.

Murdock, Jennifer. 2006. Handling unobserved site characteristics in random utility models of recreation demand. *Journal of Environmental Economics and Management* 51:1–25.

Nevo, Aviv. 2000. A practitioner's guide to estimation of random-coefficients logit models of demand. *Journal of Economics and Management Strategy* 9 (4): 513–48.

———. 2001. Measuring market power in the ready-to-eat cereal industry. *Econometrica* 69 (2): 307–42.

Petrin, Amil. 2002. Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy* 110 (4): 705–29.

Ramsey, Frank P. 1927. A contribution to the theory of taxation. *Economic Journal* 37 (145): 47–61.

Small, Kenneth A., and Harvey S. Rosen. 1981. Applied welfare economics with discrete choice models. *Econometrica* 49 (1): 105–30.

Train, Kenneth E. 2003. *Discrete choice methods with simulation*. New York: Cambridge University Press.

Villas-Boas, Sofia B. 2007. Vertical relationships between manufacturers and retailers: Inference with limited data. *Review of Economic Studies* 74 (2): 625–52.

West, Sarah E. 2004. Distributional effects of alternative vehicle pollution control policies. *Journal of Public Economics* 88:735– 57.

Online Appendix

Appendix A. Corrective Taxation Revisited

To motivate our study, we present a simple, static model of a monopolist selling two competing car models, characterized by a single negative consumption externality (i.e., carbon emissions), which depends on product attributes (i.e., fuel efficiency) as well as car utilization (i.e., driving). In the absence of any tax/subsidy, the monopolist faces two demand curves $D_j(P_j, P_{-j})$ for $j = 1, 2$. To pin down the social optimum, let the monopoly charge perfectly competitive prices $P_j^s = MC_j$ for the moment. Given the constant marginal costs (which we also assume in the subsequent analyses), $D_j(P_j, MC_{-j})$ (black solid lines in Figure A1 below) represent the private demand when prices of competing products are set at the socially optimal level. Now, in the presence of negative externality from gasoline consumption, the socially efficient demand curve must be lower than this private demand curve. This social demand is labeled $D_j^s(P_j, MC_{-j})$ (green solid lines). The social optimum quantity should then be Q_j^s given by the intersection of D_j^s and MC_j. If the government imposes an efficient gasoline tax, consumers should, in theory, fully incorporate the negative externality in purchasing their cars. Hence, the private demand curve for cars should match the social demand in the presence of the efficient gasoline tax. The logic is intact even in the presence of consumer heterogeneity [see eq. (20) in Fullerton and West (2002)]. In the presence of imperfect competition, however, the firm will exert market power and charge prices higher than marginal costs. Hence, provided that $P_{-j}^* > MC_{-j}$, the private demand for each car model given the gasoline tax must lie above D_j^s. The private demand curves in the presence of the gasoline tax and the markup pricing $D_j^g(P_j, P_{-j}^*)$ are represented by the red dash-dot lines in the diagram. The monopolist then chooses a price over marginal cost, following the standard markup pricing rule, which is represented by the intersection of red dash line and MC curve. Applying the same logic to the substitute, we obtain the equilibrium prices (P_1^*, P_2^*). It is worth emphasizing here that *the socially optimal quantity Q_j^s is no longer optimal in this imperfect competition setup*. Because the equilibrium quantity in the substitute good market Q_{-j}^* is less than the optimal quantity Q_{-j}^s, it raises the private benefit (in net of negative consumption externality) of consuming good j, and hence, the optimal quantity of good j given (P_{-j}^*, Q_{-j}^*) is $Q_j^s(Q_{-j}^*) > Q_j^s$. The same logic applies to the substitute good. The resulting deadweight losses are drawn as red shaded triangles.

Hence, the 'optimal' taxation (or feebate) scheme under this imperfect competition setup must minimize the sum of these deadweight losses (or equivalently, maximize the sum of consumer surplus and producer surplus). The diagram below illustrates a myriad of difficulties in designing such an optimal feebate scheme. To see this point, first hold Q_1^* and give a production subsidy to car model 2. This would make the price of car model 2 lower and the quantity higher. For the sake of exposition, let's say the resulting price and quantity are (P_2', Q_2'). As a result, the demand for car model 1 would decrease. Hence, holding Q_1^* constant, the subsidy to car model 2 can decrease deadweight loss from car model 1. Alternatively, of course, one might give a subsidy to car model 1. Holding the price and quantity of car model 2, this would lower the price of car model 1, increase the quantity, and reduce deadweight loss. Hence, a subsidy to one product and a subsidy to the competing product can be considered substitute policies. However, the logic assumes no strategic response by the monopolist. With monopoly pricing, however, given the production subsidy to car model 2, the monopolist now faces a lower demand for car model 1, labeled $D_1^g(P_1, P_2')$. The monopolist should then choose a markup based upon this demand curve, which would result in a lower quantity Q_1' with somewhat ambiguous impact on the equilibrium price. The resulting deadweight loss is drawn as a blue unshaded triangle. The exact size of the deadweight loss is ambiguous as it depends on own- and cross-price elasticities, or more precisely, on how different consumers substitute car model 1 for car model 2 or to the outside option. Furthermore, what complicates the task even more is that this change in the price (and quantity) of car model 1 also has the feedback effect on car model 2. More generally, in a *n*-firm, multi-product setting, firms choose prices of their multiple products strategically in response to other competitors' prices. Provided that firms compete in the Bertrand manner, any tax/subsidy in one or more products can have perverse impacts on prices of all firms' products via shifting in the (pure strategy) Bertrand-Nash equilibrium. Hence, the government must coordinate taxes/subsidies across car models, fully taking into account strategic responses in firms' multi-product pricing. Our empirical model is suited for this as it allows us to fully model firms' equilibrium pricing to changes in vehicle taxation.

Figure A1. Feebate Policy in the Presence of Gasoline Tax, Imperfect Competition with Multi-product Pricing, and Negative Externality from Carbon Emissions

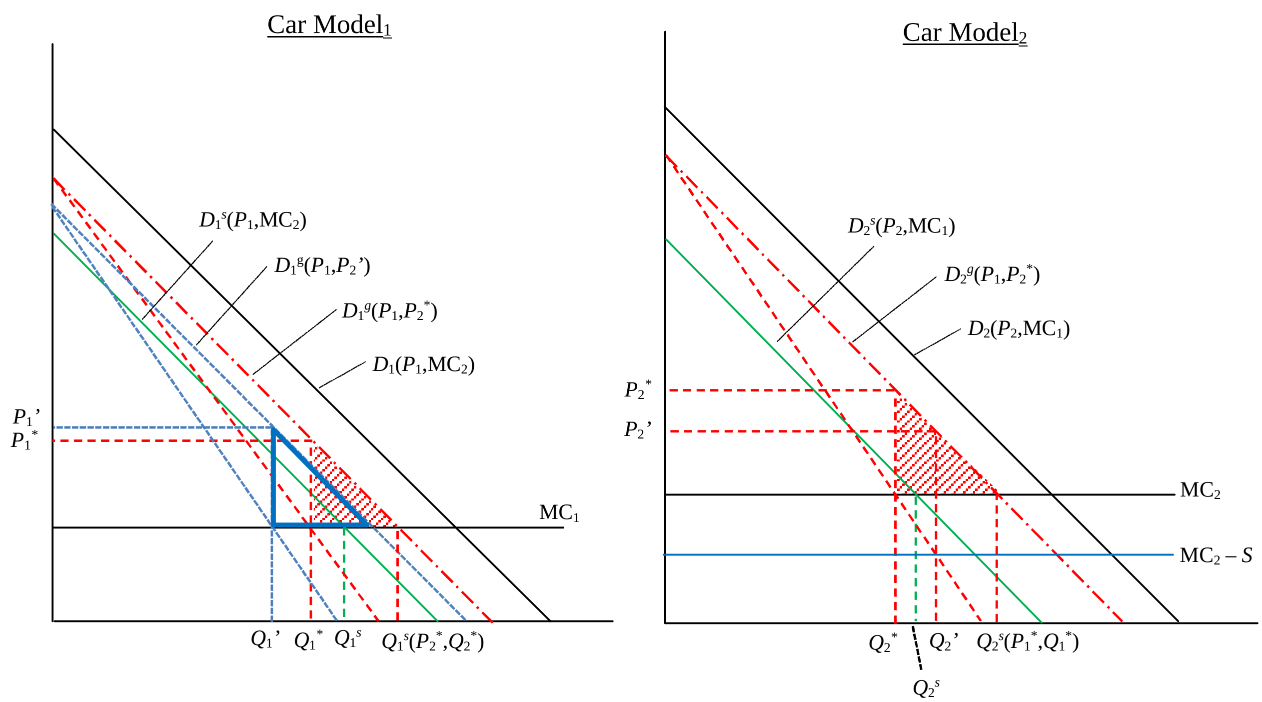

We emphasize here that our paper focuses on the *second-best* feebate scheme, and we evaluate Japan's ES/ETC policy against this second-best counterfactual, instead of the first-best. Ours are in second-best settings for two reasons. First, it is relatively well-established that under imperfect competition with multi-product pricing, the regulator can achieve the first-best optimum if she has perfect information about the degree of market power and access to specific subsidies in the market *and* lump-sum taxes outside the market. To see this, consider giving production subsidies to both products just enough to absorb markups from both products simultaneously. The problem, of course, is that doing so requires tax revenues collected elsewhere without generating deadweight loss. On the contrary, in many countries, vehicle taxes are often a source of substantial tax revenues. Hence, it is of more policy relevance to consider the optimal feebate scheme in the manner analogous to Ramsey's optimal taxation, i.e., with a constraint on tax revenues. Given the tax revenue constraint, the optimal policy should naturally look like a feebate policy, a mixture of taxes and subsidies across different car models. Unlike typical feebates, however, this (second-best) optimal policy should relate the tax/subsidy to markup levels (rather than, or in addition to, fuel efficiency or environmental characteristics of cars). Second, we lack one important piece of information to pin down the first-best optimum — the degree of negative externality for each car model, which depends on consumers' heterogenous demand for driving. Consumers who demand cars with high fuel efficiency ratings may naturally drive more, and hence, the negative externality from each unit of car for such car models can be high despite the fact that they use less gasoline per unit of driving. Without access to household-level data, we are unable to pin down either the optimal gasoline tax (which may not be necessarily set to equal the marginal external cost per ton of CO_2 emissions in our second-best setup) or the optimal quantities of different car models. It is important to note, however, that the second-best feebate scheme found this way is still fully efficient *conditional on any* level of the gasoline tax (given the tax revenue target, of course). Hence, the general approach presented here also works when the efficient gasoline tax is available as well as when it is not. In the meantime, because the original goal of Japan's feebate policy is to promote adoption of eco-friendly cars, and hence, we do characterize the environmental properties (such as average emissions and weighted average fuel efficiency) of the second-best policy as well as the ES/ETC policy.

B. Technical Notes on Empirical Strategy

B-1. Model Specification: We make two clarifications on our model specification in eq. (1) and (2). First, our utility specification slightly diverges from that of BLP (1995; 1999) and excludes the nonlinear income effect. In this sense, ours is similar to that in Nevo (2000; 2001). If we are to include the nonlinear income effect, we would either take $\log(y_i - p^e_j)$ or make α_i inversely proportional to income $\alpha_i = \alpha/y_i$ in (1). We chose our specification because our earlier attempts to estimate such a model resulted in either insignificant or positive price coefficients. Second, we also diverge from Nevo in that we do not interact the random-utility terms with observable demographic variables. We chose to do so for two reasons. First, identification of interaction parameters would require variation in the distribution of demographic variables over different markets, yet we found there was very little variation during the study period. In contrast, Nevo was able to use variation across cities as an additional source of variation. Second, we had to estimate the model with a much larger number of products (150-160 car models per market) than Nevo's study (25 brands). Thus, we concluded that little variation in the distribution of demographic variables compared with a larger number of car models would result in inefficient estimates of the parameters. Indeed, our trial runs with different sets of demographic variables resulted in non-convergence of the estimation algorithm.

B-2. Estimation Algorithm: Implementation of the GMM estimator (7) requires an estimate of $\boldsymbol{\xi}$ because $\boldsymbol{\xi}$ is by assumption unobservable to researchers. BLP (1995) proposed a nested fixed point (NFP) algorithm to numerically solve for ξ. A key here is to recognize that ξ can be considered as an unobservable error in the mean utility δ. Rearrange terms in (1), we then obtain explicit expressions for the mean utility δ_j and the idiosyncratic utility μ_{ij}:

$$\delta_j(\boldsymbol{\theta}) = -\alpha p^e_j + \mathbf{x}_j\boldsymbol{\beta} + \xi_j; \quad \mu_{ij}(\boldsymbol{\theta}) = -\sigma_p \nu_{ip} p^e_j + \textstyle\sum_{k=1}^{K} \sigma_k \nu_{ik} x_{jk}.$$

As long as we have a consistent estimate of δ_j, we can obtain the consistent estimate of ξ_j by simply running a linear regression of the estimate of δ_j on product attributes. Let S be a vector of observed market shares and s be the market share function defined by eq. (3). Then the value of the mean utility term δ can be solved numerically by the contraction mapping:

$$\delta^{h+1} = \delta^h + \ln(S) - \ln(s(\delta^h|\boldsymbol{\theta})) \quad \text{for } h = 1, ...H.$$

BLP (1995) also offers a proof of the convergence of this NFP algorithm.

Thus in essence, the estimation is done by repetition of the two-step procedure. First, given the initial guess of the parameters $\hat{\boldsymbol{\theta}}_0$, run the NFP algorithm to get the estimate of $\hat{\delta}_0$ and obtain the estimate of the error $\hat{\xi}_0$ (this is the "inner loop" of the estimation). Second, solve the optimization program (7) to get the estimate of $\boldsymbol{\theta}$. We repeat the process until the optimization routine achieves desired tolerance. Our estimation is done by carefully modifying the Matlab code supplied at Nevo's website. The modifications include, but are not limited to: allowing the set of products in each market to vary, modifying the inner loop tolerance, replacing the minimization routine, replacing the mean-distance procedure, supplying the code for calculation of own- and cross-price elasticities for both 'inside' and 'outside' goods, and supplying the code for calculation of price-cost margins.

Recently, studies have found important problems with implementation of the NFP algorithm and the resulting estimates [see Dube *et al.* (2012) and Knittel and Metaxoglou (2014) for a more detailed review of such issues]. In particular, Dube *et al.* showed that use of loose tolerance criteria for the inner-loop algorithm to ease the computational burden may result in (i) failure of the optimization program to converge or (ii) the optimization finding parameter estimates that are not even local optima. Indeed, our earlier attempt to directly use Nevo's code revealed both of these problems. To overcome these problems, we adjusted Nevo's original code and used inner-loop tolerance of 1E-14 as suggested by Dube *et al*. We also replaced Matlab's optimization routine "fminu" with Zeina's KNITRO program, which turned out to be substantially more robust and efficient than 'fminu.' Dube *et al.* also suggested an alternative algorithm known as a mathematical program with equilibrium constraints (MPEC). We use Nevo's code for our estimation because it was easier for us to flexibly adjust for different specification runs. Our earlier trial with both codes revealed that the estimates were largely similar. The standard errors of the GMM estimator are also computed using Nevo's code, which follows the standard asymptotic variance-covariance formula discussed in Newey and McFadden (1994).

We also note that in Table 2 in the manuscript, we used different draws on $\boldsymbol{\nu}$ for models (II) and (III). We did this because model (II) failed to produce a negative price coefficient for a large number of estimation runs with different draws of $\boldsymbol{\nu}$. This would mean that consumers demand more of a product when its price is higher — a sign that the model fails to properly control for unobservables, which would cause prices to be higher for products with higher consumer demand. The failure rate was surprisingly high. In contrast, with model (III), virtually all estimation runs successfully produced a negative price coefficient. Hence, Table 2 of the manuscript displays the estimates of model (II) that appear to be the best out of all estimation runs in terms of sign and significance of coefficients. Even then, the value of the GMM objective is substantially lower with our 'tax-location' IVs than with the conventional IVs.

B-3. Identification Condition: In this section, we offer a more elaborate discussion on our identification condition. As in BLP (1995; 1999), Nevo (2000; 2001), and Petrin (2002), we assume a static model in which consumers make no intertemporal substitution in purchase of cars and firms choose prices in each period, given their product choices, without coordinating prices over time. More recently, the literature has started to explore methods that allow the BLP framework (i) to incorporate dynamic pricing [e.g., Gowrisankaran and Rysman (2012) and Copeland (2014)] and (ii) to endogenize firms' product choices [e.g., Crawford (2012) and Crawford *et al.* (2015)]. We abstract away from these important complications. Our identifying condition, therefore, implicitly depends on this modeling assumption. Below then, we discuss our identifying assumptions that do not directly come from the modeling assumption itself.

Let us first rewrite our utility specification (1) so that the dependence of the indirect utility on period t is more explicit:

$$u_{ijt} = \alpha_i(y_{it} - p^e_{jt}) + \mathbf{x}_{jt}\boldsymbol{\beta}_i + \zeta_t + \mathbf{1}(j \in s)\eta_s + \xi_{jt} + \epsilon_{ijt},$$

where ζ_t is the unobservable utility term that represents the homegenous influence of time-specific valuation (say, due to unobserved sales incentives offered each season or changes in the value of outside options), η_s represents that of segment-specific valuation (say, due to unobservable segment-specific valuation), and $\mathbf{1}(j \in s)$ is an indicator variable which equals 1 if j is in segment s. We can directly control for ζ_t by including quarter and year dummies and for η_s by including hybrid/minicar dummies. Thus, we only need identifying assumptions to control for unobservable product attributes after controlling for ζ_t and η_s, which homogenously influence valuation of all products.

The identifying assumption is, therefore,

$$E[\xi_{jt}|\mathbf{z}_{jt}(\boldsymbol{\tau}_t), \mathbf{x}_{jt}, \zeta_t, \eta_s] = 0 \quad \text{for all } j, t,$$

where $\mathbf{z}_{jt}$ is a vector of tax-location variables for car model j in period t as a function of car taxes $\boldsymbol{\tau}_t$. As in Nevo (2000; 2001), it is presumable that the unobserved product characteristics have both time-invariant and time-varying components:

$$\xi_{jt} \equiv \xi_j + \Delta\xi_{jt}.$$

Therefore, we can rewrite the identification condition in two parts:

$$E[\xi_j|\mathbf{z}_{jt}(\boldsymbol{\tau}_t), \mathbf{x}_{jt}, \zeta_t, \eta_s] + E[\Delta\xi_{jt}|\mathbf{z}_{jt}(\boldsymbol{\tau}_t), \mathbf{x}_{jt}, \zeta_t, \eta_s] = 0 \quad \text{for all } j, t.$$

Per our discussion in Section 4 of the main manuscript, we have reasons to believe the location of car taxes in the characteristics space change exogenously. Hence, $E[\xi_j|\mathbf{z}_{jt}(\boldsymbol{\tau}_t), \mathbf{x}_{jt}, \zeta_t, \eta_s] = 0$ should be trivially satisfied since the location of car taxes in *each period* should be unrelated to the time-invariant component of the unobservable product attributes. Our identification, therefore, relies more critically on the second part:

Assumption: Unobserved product attributes stay constant over time ($\Delta\xi_{jt} = 0$) or do not change coincidentally with the location of car taxes: i.e., $E[\Delta\xi_{jt}|\mathbf{z}_{jt}(\boldsymbol{\tau}_t), \mathbf{x}_{jt}, \zeta_t, \eta_s] = 0$.

A similar identifying assumption has been used in D'Haultfœuille *et al.* (2014). By definition, $\Delta\xi_{jt}$ is the time-specific deviation of unobservable product valuation from its mean ($\Delta\xi_{jt} = \xi_{jt} - \xi_j$). Presumably, this term reflects changes in consumer's brand images, preferences for style etc — changes that happen only gradually over time, and hence, it seems quite unlikely that such a change in each period happens coincidentally with the location of car taxes in the same period. On the other hand, if this term also reflects changes in unobservable product-specific sales incentives (our price variable p does not include them), then it might indeed happen that such changes may be correlated with the location of car taxes and may bias our estimates. Unfortunately, we do not have data on transaction prices to address this concern. However, our informal interview with a sales representative from Toyota's dealer told us that they did not change their sales incentive plans in accordance with ES/ETC programs.

Because our model does not account for intertemporal substitution, it would complicate the idenfiticaion of the demand parameters and policy impacts if some consumers had shifted their consumption in anticipation of future policies. In our case, however, the effect seems negligible. The ES/ETC policy was announced in April, 2009 and administered in June, 2009, yet covered cars purchased in April and May, 2009. Moreover, the ES program was initially scheduled to end in March, 2010, but was unexpectedly extended to September, 2010. The second ES period was also similar. It was adopted on December 20, 2011 and started its administration in April, 2012, but covered cars purchased since December 20, 2011. Figure A2 shows the (detrended) trends in monthly new car sales from 2008 to 2010. The sales amount and seasonal pattern were quite stable before and after the first ES program. Although the sales were relatively lower at the beginning of 2009 compared to the same period in the previous years, the trend actually continued until the end of the second quarter of 2009.

Figure A2. Trends in Monthly New Car Sales from 2008 to 2010

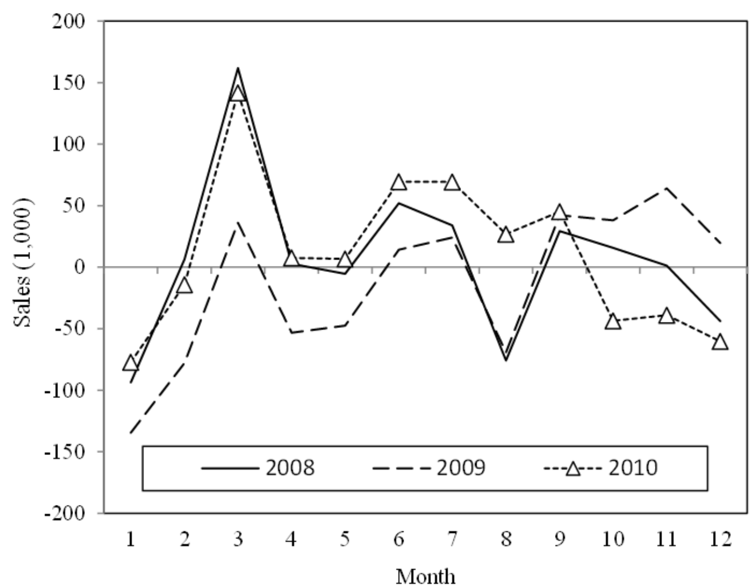

C. Descriptive Statistics

Table A1 shows the trends in the sales, prices and major product attribute variables used in our analysis over the study period. The prices and major product attributes are sales-weighted means. Around 145-159 car models were sold in each quarter. Total quarterly sales series clearly displays a seasonal pattern. Car sales are generally strong in the first and the third quarters, followed by drops in the second and the fourth quarters. There are two reasons for this seasonal pattern. First, working individuals usually receive bonuses in June and December, and each bonus is a lump-sum payment approximately twice of their monthly wages. Second, because March is the end of a fiscal year in Japan, sales subsidiaries offer a variety of sales promotions then. After taking into account the seasonal cycle, the sales generally trend downward over time: i.e., the first quarter sales decreased from 1,177,911 in 2007 to 883,547 in 2009 right after the *Lehman Shock* and further hit the bottom of 807,624 in 2011 due to the *Tohoku* Earthquake. It started to recover quickly after that, with the first quarter sales reaching 1,245,498 in 2012.

Table A1. Sales, Price and Product Characteristics of All Car Models over Time

Quarter	Models	Sales	Price		HP/Weight		MPY		Size		AT/CVT	
			Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
2007.1	153	1,177,911	172	72	0.089	0.023	13.4	3.6	7249	727	0.930	0.141
2007.2	151	807,883	172	81	0.088	0.023	13.3	3.5	7227	732	0.932	0.141
2007.3	145	864,876	175	85	0.089	0.022	12.3	3.2	7272	727	0.932	0.137
2007.4	147	811,305	178	83	0.089	0.024	11.8	3.2	7281	738	0.937	0.134
2008.1	149	1,134,377	176	84	0.088	0.025	11.6	3.0	7251	727	0.930	0.139
2008.2	152	799,539	177	89	0.087	0.024	12.0	3.2	7235	746	0.933	0.138
2008.3	153	862,397	178	84	0.088	0.023	10.2	2.7	7286	736	0.932	0.138
2008.4	155	729,635	175	84	0.086	0.023	13.8	3.6	7219	711	0.937	0.138
2009.1	156	883,547	168	73	0.085	0.022	17.3	4.4	7168	663	0.933	0.139
2009.2	158	663,686	176	80	0.085	0.022	16.8	5.1	7236	655	0.938	0.135
2009.3	156	872,018	179	76	0.086	0.021	15.7	5.4	7314	637	0.939	0.121
2009.4	155	878,585	185	83	0.086	0.021	15.3	5.2	7337	657	0.939	0.115
2010.1	156	1,110,119	185	89	0.087	0.022	14.8	4.9	7318	666	0.921	0.130
2010.2	156	833,896	181	80	0.085	0.021	14.5	4.9	7293	652	0.916	0.138
2010.3	157	1,016,468	184	81	0.086	0.021	14.2	4.7	7332	658	0.920	0.134
2010.4	155	613,634	185	86	0.085	0.023	15.3	5.3	7298	677	0.927	0.143
2011.1	155	807,624	178	82	0.084	0.022	14.2	4.5	7236	666	0.924	0.167
2011.2	156	518,600	175	77	0.083	0.022	13.4	3.9	7219	650	0.918	0.174
2011.3	157	801,895	182	81	0.083	0.022	14.0	4.3	7291	638	0.928	0.169
2011.4	159	768,096	187	84	0.083	0.022	14.7	4.6	7305	679	0.936	0.160
2012.1	159	1,245,498	186	81	0.081	0.022	14.7	4.9	7298	665	0.941	0.154
2012.2	158	916,647	182	84	0.079	0.021	14.6	4.7	7227	667	0.939	0.154
2012.3	155	942,678	184	83	0.081	0.022	15.2	4.9	7264	651	0.936	0.161
2012.4	156	743,547	184	95	0.081	0.023	15.2	5.0	7246	649	0.941	0.157

Note: Price = average retail price in 10,000 JPY; HP/Weight = HP/weight in kw/kg; MPY = mileage in km per JPY; Size = the sum of length, width and height; AT/CVT = the fraction of the car grades that have automatic or continuously variable transmission.

An increasing trend in prices is observed during the study period. The sales-weighted average car prices in 2012 were roughly 8% higher than those in 2007. On the other hand, HP/weight has been fairly constant, but decreased slightly in recent years. MPY first declined from 13.4 km/yen in 2007 to 10.2 km/yen in the third quarter of 2008, and then bounced back and reached 15.2 km/yen in the end of 2012. The downward trend was mainly driven by the increasing price of gasoline, which reached its peak in the third quarter of 2008, while the upward trend reflects the improvement in the fuel efficiency of some car models marketed after 2009. The increasing trend in the MPY is likely to be due to the green vehicle tax policy introduced in the second quarter of 2009. Lastly, the car size and the share of cars equipped with AT/CVT have been roughly constant over time.

Table A2. Sales, Price and Product Characteristics of Hybrid Cars over Time

Quarter	Models	Sales		Price		HP/Weight		MPY		Size		AT/CVT	
		Total	Share	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
2007.1	7	23,628	0.020	275	85	0.067	0.016	22.5	4.8	7,852	316	1	0
2007.2	7	19,125	0.024	328	230	0.073	0.031	21.5	5.5	7,864	327	1	0
2007.3	7	21,309	0.025	362	279	0.077	0.037	19.7	5.5	7,880	338	1	0
2007.4	7	21,187	0.026	316	223	0.071	0.030	19.8	4.7	7,825	307	1	0
2008.1	8	26,634	0.023	321	223	0.072	0.031	19.0	4.7	7,843	318	1	0
2008.2	8	24,644	0.031	329	210	0.078	0.037	19.2	5.0	7,839	286	1	0
2008.3	7	29,963	0.035	323	173	0.079	0.036	16.2	4.2	7,856	286	1	0
2008.4	7	25,832	0.035	299	160	0.071	0.029	22.5	4.8	7,802	266	1	0
2009.1	8	28,426	0.032	245	99	0.068	0.018	28.6	4.1	7,676	221	1	0
2009.2	9	70,575	0.106	239	101	0.074	0.016	27.2	4.4	7,680	230	1	0
2009.3	10	122,798	0.141	233	86	0.074	0.013	26.2	4.2	7,704	204	1	0
2009.4	11	124,989	0.142	245	103	0.075	0.014	25.2	4.4	7,715	204	1	0
2010.1	12	132,374	0.119	264	130	0.079	0.018	24.5	5.1	7,742	217	1	0
2010.2	12	122,369	0.147	242	92	0.077	0.015	23.8	4.3	7,689	205	1	0
2010.3	12	123,358	0.121	248	107	0.077	0.017	23.8	4.5	7,703	219	1	0
2010.4	14	102,012	0.166	227	99	0.076	0.016	24.6	3.8	7,580	306	0.992	0.064
2011.1	14	98,710	0.122	235	103	0.076	0.016	22.6	3.6	7,562	313	0.986	0.082
2011.2	14	63,359	0.122	244	93	0.075	0.016	20.6	2.6	7,623	298	0.986	0.082
2011.3	15	141,087	0.176	254	96	0.073	0.015	20.6	2.4	7,748	235	0.996	0.045
2011.4	19	148,781	0.194	254	86	0.073	0.015	20.9	2.9	7,782	239	0.996	0.044
2012.1	19	282,623	0.227	231	81	0.071	0.012	21.3	3.5	7,636	371	0.996	0.043
2012.2	21	207,197	0.226	233	99	0.073	0.014	20.9	3.8	7,594	406	0.987	0.043
2012.3	24	219,347	0.233	234	98	0.074	0.015	21.6	4.4	7,619	424	0.947	0.195
2012.4	24	171,833	0.231	229	120	0.073	0.015	21.7	4.4	7,557	430	0.949	0.196

Note: Price = average retail price in 10,000 JPY; HP/Weight = HP/weight in kw/kg; MPY = mileage in km per JPY; Size = the sum of length, width and height; AT/CVT = the fraction of the car grades that have automatic or continuously variable transmission.

Similarly, Tables A2 and A3 provide the summary descriptive statistics for the hybrid cars and minicars only, respectively. It is evident that the sales of hybrid cars have been increasing rapidly, especially after the first quarter of 2009. The market share of hybrid cars became seven times larger from 0.03 in 2008 to 0.23 in 2012. The prices of hybrid cars are generally higher than the average car prices. During 2007, hybrid car prices rose quickly probably because of the increasing demand due to the rising gasoline price. Compared to non-hybrid cars, hybrid cars tend to have lower ratio of horsepower to weight and larger size, but much higher fuel efficiency. Minicars account for approximately one third of the Japanese new car sales. They are generally more compact, less powerful, and cheaper. One take-away message from Tables A1-A3 is that the trends in the key product characteristics did not change dramatically by the introduction of green car tax policies, yet the variety, sales composition, and prices appear to have changed during the policy period.

Table A3. Sales, Price and Product Characteristics of Minicars over Time

Quarter	Models	Sales		Price		HP/Weight		MPY		Size		AT/CVT	
		Total	Share	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
2007.1	34	393,824	0.334	116	12	0.066	0.005	16.5	1.5	6,456	412	0.925	0.168
2007.2	35	279,763	0.346	116	12	0.066	0.005	16.1	1.5	6,456	446	0.935	0.167
2007.3	34	268,102	0.310	117	12	0.066	0.005	15.0	1.5	6,455	472	0.948	0.165
2007.4	34	249,394	0.307	117	12	0.066	0.005	14.5	1.4	6,451	516	0.946	0.166
2008.1	35	376,292	0.332	118	12	0.065	0.005	14.1	1.3	6,470	464	0.948	0.170
2008.2	35	278,631	0.348	118	12	0.065	0.005	14.4	1.4	6,474	500	0.950	0.167
2008.3	36	266,074	0.309	118	12	0.065	0.005	12.4	1.2	6,473	490	0.947	0.172
2008.4	36	267,351	0.366	119	12	0.065	0.005	16.3	1.6	6,490	295	0.946	0.172
2009.1	35	337,038	0.381	119	12	0.065	0.005	20.3	2.2	6,506	96	0.946	0.168
2009.2	35	221,059	0.333	119	11	0.065	0.004	18.7	2.0	6,512	97	0.950	0.168
2009.3	33	231,049	0.265	119	11	0.065	0.004	17.1	1.9	6,508	94	0.954	0.151
2009.4	34	239,032	0.272	119	11	0.065	0.005	16.5	1.9	6,511	93	0.938	0.139
2010.1	34	330,702	0.298	119	12	0.064	0.005	16.3	1.6	6,518	93	0.901	0.155
2010.2	34	255,808	0.307	119	13	0.065	0.005	15.2	1.6	6,518	97	0.894	0.162
2010.3	34	289,906	0.285	120	12	0.065	0.005	15.3	1.7	6,518	94	0.897	0.159
2010.4	33	198,791	0.324	121	13	0.064	0.005	15.9	2.1	6,523	97	0.902	0.172
2011.1	33	282,195	0.349	121	12	0.064	0.005	15.2	2.3	6,515	89	0.930	0.159
2011.2	33	186,504	0.360	121	11	0.064	0.005	14.5	2.2	6,516	90	0.924	0.166
2011.3	33	248,801	0.310	121	12	0.062	0.006	14.9	2.5	6,519	94	0.941	0.162
2011.4	33	252,487	0.329	119	13	0.063	0.005	15.9	2.8	6,502	96	0.945	0.157
2012.1	32	398,727	0.320	122	14	0.063	0.005	15.2	2.4	6,520	100	0.949	0.147
2012.2	32	344,514	0.376	124	14	0.062	0.005	14.7	2.1	6,531	103	0.952	0.142
2012.3	29	319,920	0.339	125	14	0.062	0.005	15.3	2.1	6,534	105	0.964	0.122
2012.4	30	259,827	0.349	125	13	0.063	0.005	15.2	2.1	6,532	102	0.962	0.138

Note: Price = average retail price in 10,000 JPY; HP/Weight = HP/weight in kw/kg; MPY = mileage in km per JPY; Size = the sum of length, width and height; AT/CVT = the fraction of the car grades that have automatic or continuously variable transmission.

D. OLS and IV Logit Results

This appendix reports the results from the OLS and IV Logit models. Although these models are known to yield unrealistic substitution patterns [see Nevo (2001) for a thorough discussion on this point], and hence, are not used to make real inferences about policy impacts, the results give us a sense of the performance of different sets of instrumental variables for use in the full RC logit model. Note that in the logit models, the stochastic error term includes the unobserved product attribute ξ_j (so does the random utility term μ_{ij}). Therefore, if the set of instruments are correlated with any of these terms, the price coefficients would be biased and the overidentification tests would likely reject the null hypothesis.

The results are obtained from regressing $\ln(S_{jt}) - \ln(S_{0t})$ on constants, effective prices, HP/weight, MPY, size, AT/CVT, minicar dummy, hybrid car dummy, year dummies, quarter dummies, and maker dummies, along with a macroeconomic variable (seasonally adjusted GDP growth rates) to account for the effects of the *Lehman* crisis and the *Tohoku* earthquake.[1] The first two columns in Table A4 report the results from OLS logit, with and without the minicar and hybrid car dummies. Columns (iii)-(v) display the results of IV logit with different sets of instruments, without the minicar and hybrid car dummies. Columns (v)-(viii) report the same, but with the minicar and hybrid car dummies.

We include quarter fixed effects, because in Japan at least, the car sales has large seasonal effects, particularly in the first quarter and the third quarter (see our discussion in Section C above). This occurs because they correspond to the two bonus seasons and the Japanese automakers put together sales promotions in response. As discussed in the main text, including model fixed effects in the regression is problematic. We thus instead included maker fixed effects to control for maker-specific brand values (Nevo, 2001). Inclusion of maker dummies as well as minicar/hybrid dummies should, in principle, take care of the concern that the Japanese government might have designed the ES/ETC policy as a means to favor (or disfavor) a particular car maker or a particular type of cars (for the latter, minicars and hybrid cars may be considered natural candidates in Japan).

For virtually all models, coefficients on price, HP/weight, MPY, size, and AT/CVT are significant at traditional levels, with signs consistent with our expectation. With the OLS logit, the coefficients on prices are negative and highly significant. Without minicar and hybrid dummies, coefficients on price and HP/weight turn less significant when the traditional 'product-location' IVs are used, whereas with our 'tax-location' IVs, they continue to be highly significant. Comparing models (iii)-(v) with models (vi)-(viii), we see the inclusion of the minicar and hybrid dummies improves the efficiency of estimates and the overidentification test presumably because it allows us to control for unobservables better. Hence, we favor the models with these dummies over those without them. Importantly, when the 'product-location' IVs are used [i.e., model (vi)], the standard errors are generally larger and MPY becomes insignificant, though the overidentification test cannot reject the null at the 30% level. When our 'tax-location' IVs are used instead [i.e., model (vii)], the overidentification test cannot reject the null even at the 50% level, whereas the coefficient on MPY returns highly significant again. We also examine the weak IV problem. Though not reported, all the tax-location IVs are significant at the 1% level in the first stage regression, suggesting that our IVs do not suffer from the weak IV problem. On the other hand, with the traditional IVs, many of the own-firm and rival-firm

[1] Inclusion of the macroeconomic variable follows BLP (1999). As BLP points out, it is somewhat arbitrary to include such variables. However, the effects of these macroeconomic shocks appear to be very significant, and removing these variables may bias the estimates. An alternative would be to exclude observations from these periods. However, these periods also overlap with policy periods that are important for our analysis. Thus, excluding observations from these periods appears at least as arbitrary as inclusion of macroeconomic variables.

IVs are highly insignificant. We take these as evidence that our 'tax-location' IVs are more reliable than the conventional 'product-location' IVs for our full model.

Table A4. Estimation Results of OLS Logit and IV Logit

	OLS Logit		IV Logit					
	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)	(viii)
Constant	-17.1314 ***	-21.0986 ***	-21.4549 ***	-18.4343 ***	-17.3291 ***	-22.6680 ***	-21.7701 ***	-21.8887 ***
	(0.5567)	(0.6717)	(3.7046)	(0.5841)	(1.2240)	(0.9927)	(0.6953)	(0.8040)
Price	-0.0045 ***	-0.0042 ***	-0.0154 *	-0.0078 ***	-0.0050 *	-0.0142 ***	-0.0085 ***	-0.0093 ***
	(0.0002)	(0.0003)	(0.0092)	(0.0004)	(0.0028)	(0.0038)	(0.0004)	(0.0025)
HP/Weight	1.8893 *	4.4348 ***	35.6327	12.0581 ***	3.4322 *	33.8305 ***	17.0132 ***	19.2347 ***
	(1.0528)	(1.1654)	(28.3791)	(1.3978)	(8.5773)	(11.2386)	(1.4580)	(7.4023)
MPY	0.1599 ***	0.2303 ***	0.1564 ***	0.1588 ***	0.1597 ***	0.0582	0.1566 ***	0.1436 ***
	(0.0093)	(0.0130)	(0.0126)	(0.0096)	(0.0093)	(0.0671)	(0.0143)	(0.0449)
Size	0.0003 ***	0.0007 ***	0.0008 **	0.0005 ***	0.0004 ***	0.0010 ***	0.0009 ***	0.0009 ***
	(0.0000)	(0.0001)	(0.0004)	(0.0000)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
AT/CVT	0.9180 ***	0.7066 ***	1.4054 ***	1.0649 ***	0.9403 ***	1.3489 ***	0.9814 ***	1.0299 ***
	(0.1179)	(0.1186)	(0.4376)	(0.1220)	(0.1701)	(0.2812)	(0.1239)	(0.2022)
Minicar Dummy	---	✓	---	---	---	✓	✓	✓
Hybrid Dummy	---	✓	---	---	---	✓	✓	✓
Maker Dummies	✓	✓	✓	✓	✓	✓	✓	✓
Year Dummies	✓	✓	✓	✓	✓	✓	✓	✓
Quarter Dummies	✓	✓	✓	✓	✓	✓	✓	✓
Macroecon. Var.	✓	✓	✓	✓	✓	✓	✓	✓
Location IVs used	---	---	Characteristics	Taxes	Taxes + Characteristics	Characteristics	Taxes	Taxes + Characteristics
Overidentification Tests								
Sargan χ^2	---	---	8.0013	1.2162	16.7479	7.3639	1.2659	14.4740
(p-values)			(0.3325)	(0.5444)	(0.0528)	(0.3920)	(0.5310)	(0.1064)
Basmann χ^2	---	---	7.9536	1.2084	16.6786	7.3148	1.2571	14.3974
(p-values)			(0.3367)	(0.5465)	(0.0540)	(0.3969)	(0.5334)	(0.1089)
# of Obs.	3,703	3,703	3,703	3,703	3,703	3,703	3,703	3,703

E. Cross-Price Elasticities, Attribute Elastisities, and Outside Substitution

Table A5 reports cross-price elasticities for 15 top-selling car models in 2012. Each entry (j, k) represents a percentage change of the market share for brand j with respect to a percentage change of the price of brand k. The estimated elasticities exhibit expected signs and magnitudes and are roughly comparable with those reported in BLP (1995) on U.S. counterparts. There is also substantial variation across brands, unlike with the standard logit model, which would display the identical cross-price elasticity for all entries in each column. Note that with the standard logit, the cross-price elasticity formula is $\varepsilon_{jk} = \alpha p_k s_k$ for all $j \neq k$. Many of the top-selling brands in the table had small cross-price elasticities, the magnitudes that are roughly comparable to those reported in BLP (1995). In BLP (1995), even the brands that had the largest own- and cross-price elasticities exhibited cross-price elasticities that were in the order of $1/100$ or smaller relative to their respective own-price elasticities. Yet, some of the top-selling brands had relatively large cross-price elasticities with respect to each other, particularly to brands with similar characteristics.

Table A5. Estimated Own- and Cross-Price Elasticities for Top-selling Car Models in 2012

	Brand Name	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	Honda StepWgn	-3.259	0.046	0.057	0.046	0.048	0.055	0.035	0.049	0.046	0.039	0.041	0.047	0.034	0.039	0.041
2	Toyota Prius	0.208	-2.821	0.222	0.311	0.229	0.225	0.351	0.235	0.229	0.201	0.247	0.206	0.235	0.250	0.215
3	Honda Fit	0.044	0.038	-2.442	0.049	0.063	0.058	0.043	0.060	0.059	0.070	0.059	0.063	0.064	0.059	0.064
4	Honda Fit Hybrid	0.049	0.074	0.068	-2.496	0.074	0.069	0.088	0.075	0.075	0.078	0.081	0.073	0.086	0.083	0.078
5	Mazda Demio	0.020	0.021	0.034	0.029	-2.391	0.034	0.028	0.037	0.035	0.047	0.038	0.038	0.044	0.038	0.042
6	Nissan Note	0.039	0.037	0.054	0.048	0.059	-2.449	0.044	0.059	0.061	0.068	0.062	0.066	0.065	0.061	0.067
7	Toyota Aqua Hybrid	0.087	0.194	0.137	0.201	0.165	0.150	-2.165	0.167	0.170	0.181	0.199	0.158	0.225	0.211	0.183
8	Toyota Vitz	0.039	0.042	0.063	0.056	0.071	0.066	0.055	-2.435	0.072	0.087	0.074	0.078	0.085	0.075	0.082
9	Honda N-Box	0.081	0.090	0.128	0.122	0.145	0.146	0.119	0.152	-2.386	0.177	0.165	0.177	0.176	0.164	0.178
10	Suzuki Alto	0.026	0.030	0.058	0.048	0.073	0.061	0.049	0.070	0.069	-2.228	0.075	0.077	0.093	0.076	0.086
11	Daihatsu Tanto	0.052	0.070	0.095	0.095	0.112	0.105	0.101	0.115	0.120	0.141	-2.316	0.129	0.145	0.130	0.137
12	Suzuki Palette	0.020	0.020	0.035	0.030	0.040	0.039	0.028	0.042	0.044	0.051	0.045	-2.464	0.049	0.044	0.050
13	Daihatsu Mira	0.043	0.066	0.103	0.100	0.132	0.111	0.116	0.131	0.131	0.181	0.148	0.143	-2.033	0.154	0.163
14	Daihatsu Move	0.040	0.057	0.077	0.078	0.093	0.084	0.087	0.094	0.097	0.119	0.106	0.104	0.125	-2.276	0.113
15	Nissan Moco	0.018	0.021	0.036	0.032	0.043	0.040	0.032	0.044	0.046	0.057	0.048	0.051	0.057	0.049	-2.338

We also report the estimated elasticities with respect to price and other product attributes in Table A6. For comparison, we report the results with the full RC logit as well as the simple IV logit.[2] Note first that with the simple logit, the elasticity of car j with respect to attribute k is given by $\varepsilon_{jk} = \beta_k x_{jk}(1 - s_j)$. Hence, we should expect the elasticity estimates to be highly correlated with the values of the attributes. In contrast, with the RC logit, the elasticity with respect to k-th attribute is given by the formula analogous to (8) in the main text, with p_j replaced by x_{jk} and α_i replaced by β_{ik}. Hence, the elasticities need not monotonically change with the values of the attributes. These patterns are indeed observed in the table. Price elasticities are substantially less variable with our model, whereas they almost monotonically increase with prices in the case of the IV logit. We still do, however, observe highly monotonic relationships between the values of

[2] The RC logit with BLP instruments produced elasticities that seemed highly unrealistic, and hence, are not reported here. The results are available upon request.

the attributes and the corresponding elasticities with the RC logit. For example, the elasticity with respect to MPY declines almost monotonically with respect to MPY ratings — that is, consumers who buy fuel-efficient cars care more about fuel efficiency than those who buy fuel-inefficient cars. Analogous comments apply to HP/weight and car size.

Our estimates are roughly consistent with BLP (1995), except that our elasticity estimates on MPY and car size are much larger than the corresponding estimates from BLP. Our estimates indicates that a 1% increase in MPY ratings would increase the sales of Prius by roughly 5%. Analogously, 1% increase in car size would induce a 7% increase in the sales of Prius. Comparing the estimates of the IV logit with those of the RC logit, the elasticities with respect to size are largely in the same magnitude between the two models, and thus, we are less concerned with them. The difference in elasticity estimates between the IV logit and the RC logit must be attributed to the interaction between the variations in β_{ik} and s_{ij}. Hence, for MPY, the existence of consumers with a high marginal valuation for fuel efficiency *and* with a high purchase probability of a car model with high fuel efficiency tends to magnify the elasticity with respect to MPY for that car model. A question then is, what would be the elasticities for low fuel-efficiency cars that are roughly comparable to gas gazzlers in U.S.? Indeed, the elasticities for such car models are in the range of 1 to 2, much closer to those reported in BLP. Although these values are still larger than those of BLP, we are somewhat suspicious of their reported elasticities because their values turn negative for some car models, implying that consumers decrease their purchase probabilities for these car models when these models improve fuel efficiency — a behavioral prediction is very hard to believe. We suspect there might have been some uncontrolled errors that caused this to happen. We thus conclude that Japanese consumers care a great deal about fuel efficiency and car size, much more than U.S. consumers, and that this cross-country difference may be attributed in part to the fact that the Japanese cars have substantially higher MPY ratings and are much smaller in size than U.S. counterparts.

Table A6. Demand Elasticities with respect to Attributes/Price for Top-selling Car Models for 2012

Brand Name	IV Logit with Our Instruments					RC Logit with Our Instruments				
	Elasticities with respect to:					Elasticities with respect to:				
	Price	HP/Weight	MPY	Size	AT/CVT	Price	HP/Weight	MPY	Size	AT/CVT
Toyota Prius	-3.298	2.097	1.154	7.350	0.981	-2.821	1.438	5.235	7.329	1.849
Toyota Aqua Hybrid	-2.130	2.121	1.381	6.763	0.997	-2.165	1.652	7.365	5.703	2.367
Daihatsu Mira	-1.417	2.099	0.988	6.156	1.012	-2.033	2.140	7.260	3.298	3.603
Honda N-Box	-1.908	1.905	0.763	6.420	1.014	-2.386	1.728	4.857	4.691	3.008
Daihatsu Tanto	-1.808	1.729	0.903	6.465	1.027	-2.316	1.619	5.975	4.421	3.204
Daihatsu Move	-1.736	2.065	0.954	6.396	1.034	-2.276	1.953	6.388	4.279	3.277
Honda Fit Hybrid	-2.382	2.546	1.083	7.289	0.967	-2.496	2.045	5.946	6.473	2.355
Suzuki Alto	-1.459	2.277	0.817	6.360	0.810	-2.228	2.391	6.190	3.219	2.978
Toyota Vitz	-1.893	2.995	0.820	7.066	0.979	-2.435	2.745	5.287	5.116	2.945
Honda Fit	-1.995	3.254	0.726	7.273	0.735	-2.442	2.789	4.320	5.896	2.009
Nissan Moco	-1.952	2.907	0.742	7.289	0.953	-2.449	2.558	4.539	5.641	2.684
Toyota Ractis	-1.665	2.255	0.836	6.545	1.059	-2.338	2.245	5.956	3.880	3.627
Honda Freed	-3.744	3.121	0.557	8.277	1.060	-3.259	2.136	2.504	8.734	1.945
Suzuki Palette	-1.783	1.968	0.727	6.673	1.061	-2.464	1.877	4.916	4.388	3.404
Toyota Passo	-1.718	3.148	0.842	7.142	0.746	-2.391	2.941	5.539	4.950	2.305

The RC logit model also allows us to estimate the substitutability of the inside goods to the outside option, which is calculated, à la BLP (1995), as the estimated percentage of consumers who substitute to the outside good as a percentage of those who substitute away from a car model, given a price increase for that model: i.e., $(ds_0/dp_j)/|ds_j/dp_j| \times 100$. The number essentially indicates, given a small price increase for the car model j, of those who decided not to purchase the car model, what percentage of them would choose not to buy any of the car models. As in BLP (1995), the estimated substitution elasticities vary substantially across car models. We emphasize here that these numbers are again roughly comparable to those in BLP (1995).

Table A7. Outside Substitution Probabilities for Top-selling Car Models in 2012

Our Estimates		BLP (*Econometirca* , 1995)	
Brand Name	Precentage who substitute to the outside good	Brand Name	Precentage who substitute to the outside good
Toyota Prius	6.08	Mazda 323	27.12
Toyota Aqua Hybri	9.32	Nissan Sentra	26.13
Daihatsu Mira	14.67	Ford Escort	28.00
Honda N-Box	8.43	Shevy Cavalier	26.39
Daihatsu Tanto	9.73	Honda Accord	21.84
Daihatsu Move	10.52	Ford Taurus	25.21
Honda Fit Hybrid	6.76	Buick Century	25.40
Suzuki Alto	15.59	Nissan Maxima	21.74
Toyota Vitz	9.87	Acura Legend	20.79
Honda Fit	9.10	Lincoln Town Car	20.31
Nissan Note	8.78	Cadillac Seville	16.74
Nissan Moco	11.87	Lexus LS400	10.09
Honda StepWgn	4.70	BMW 735i	10.10
Suzuki Palette	9.48		
Mazda Demio	12.39		

References

[1] Berry, S., J. Levinsohn, and A. Pakes, Automobile Prices in Market Equilibrium. *Econometrica*, 63 (4), 1995, 841-890.

[2] Berry, S., J. Levinsohn, and A. Pakes, Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy. *American Economic Review*, 89 (3), 1999, 400-430.

[3] Copeland, A., W. Dunn, and G. Hall. Inventories and the Automobile Market. *RAND Journal of Economics*, 42(1), 2011, 121-149.

[4] Copeland, A. Intertemporal Substitution and New Car Purchases. *RAND Journal of Economics*, 45(3), 2014, 624-644.

[5] Crawford, G. S., Accommodating Endogenous Product Choices: A Progress Report. *International Journal of Industrial Organization*, 30, 2012, 315-320.

[6] Crawford, G. S., O. Shcherbakov, M. Shum, The Welfare Effects of Endogenous Quality Choice in Cable Television Markets. *University of Zurich Department of Economics Working Paper 202*, 2015.

[7] D'Haultfœuille, X., P. Givord, and X. Boutin. The Environmental Effect of Green Taxation: The Case of the French Bonus/Malus. *The Economic Journal*, 124 (578), 2014, 444–480.

[8] Dube, Jean-Pierre, Jeremy T. Fox, and Che-Lin Su. 2012. Improving the Numerical Performance of Static and Dynamic Aggregate Discrete Choice Random Coefficients Demand Estimation. *Econometrica* 80 (5): 2231-2267.

[9] Fullerton, D. and S. West, Can Taxes on Cars and Gasoline Mimic an Unavailable Tax on Emissions? *Journal of Environmental Economics and Management*, 43, 2002, 135-157.

[10] Gowrisankaran, G. and M. Rysman. 2012. Dynamics of Consumer Demand for New Durable Goods. *Journal of Political Economy*, 120 (6), 1173-1219.

[11] Knittel, C. and K. Metaxoglou. 2014. Estimation of Random Coefficients Demand Models: Two Empiricists' Perspectives. *The Review of Economics and Statistics* 96 (1): 34-59.

[12] Nevo, A., A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand. *Journal of Economics and Management Strategy*, 9 (4), 2000, 513-548.

[13] Nevo, A., Measuring Market Power in the Ready-to-Eat Cereal Industry. *Econometrica*, 69 (2), 2001, 307-342.

[14] Newey, W. and D. McFadden. 1994. Estimation in Large Samples, in *The Handbook of Econometrics*, Vol. 4, ed. by D. McFadden and R. Engle. Amsterdam: North-Holland.

[15] Petrin, A., Quantifying the Benefits of New Products: The Case of the Minivan. *Journal of Political Economy*, 110 (4), 2002, 705-729.