Global Policy Spillovers:

How Environmental Policies Propagate through Product Attributes*

Koichiro Ito[1], James M. Sallee[2], and Andrew Smith[3]

[1]University of Chicago and NBER

[2]University of California, Berkeley and NBER

[3]University of Wisconsin-Madison

January 8, 2026

Abstract

In a globalized economy, a country's domestic policies can generate global spillover effects through products designed and manufactured by multinational firms. Standard economic analyses often abstract from this channel, potentially leading to an understatement of the impacts of many economic policies. We study this phenomenon in the context of environmental regulation in the automobile market. We find that Japan's fuel economy subsidy led to significant improvements in the fuel economy of vehicles sold in the U.S. market, thereby generating global environmental benefits. We then develop a model of multinational automobile markets to examine how cross-market linkages in revenues and costs give rise to such global spillovers. Using the estimated model, we conduct counterfactual policy simulations to quantify the environmental benefits and welfare effects of these global policy spillovers.

*Ito: Harris School of Public Policy, University of Chicago, 1307 East 60th St., Chicago, IL 60637 (e-mail: ito@uchicago.edu). Sallee: Haas School of Business, UC Berkeley, 2220 Piedmont Ave, Berkeley, CA 94720 (e-mail: sallee@berkeley.edu). Smith: Department of Economics, University of Wisconsin-Madison (e-mail: jasmith52@wisc.edu). We would like to thank Caleb Halvorson-Fried, Linnea Holy, and David Xu for excellent research assistance, and Meghan Busse, Michael Greenstone, Ryan Kellogg, and seminar participants at the Coase Project conference, Carnegie Mellon University, Northwestern University, Michigan State University, Stanford University, Tulane University, University of California, Berkeley and University of California, Davis for their helpful comments.

1 Introduction

How should we evaluate policies in one jurisdiction that affect products sold in many places? Standard economic analysis typically focuses on domestic outcomes when evaluating policy, but in a world where multinational firms design global products, policies in one country can propagate to markets worldwide by inducing changes in a product's attributes. This channel, which we call *attribute propagation*, represents a potentially significant but understudied mechanism through which domestic policies generate international effects. Accounting for attribute propagation has the potential to fundamentally alter the evaluation of environmental policies, safety regulations, and antitrust remedies, as conventional economic analyses may substantially understate the full impacts of many economic policies.[1]

In this paper, we study attribute propagation in the context of automobile environmental policy, highlighting a linkage between the Japanese and U.S. auto markets. Our goal is to develop empirical methods to measure and quantify attribute propagation, investigate its underlying mechanisms, and provide tools for more comprehensive policy evaluation in globalized markets.

Our analysis begins with a difference-in-differences (DID) design that exploits variation generated by a Japanese fuel-economy subsidy introduced in 2009. The policy created strong incentives for firms to improve the fuel economy of models sold in Japan. We leverage the fact that while many models sold in Japan are also marketed abroad (e.g., in both Japan and the United States), multinational Japanese automakers also produce vehicles for foreign markets that are not sold in Japan. This feature allows us to construct treatment and control groups within the U.S. automobile market, where the former may be affected by the Japanese subsidy through attribute propagation. We show that the treatment and control groups have similar observable characteristics and exhibit parallel pre-trends in fuel economy prior to the policy's introduction.

The DID analysis provides statistical evidence that the Japanese fuel-economy subsidy propagated to the U.S. market. Specifically, the subsidy led to an 8.65% improvement in the fuel economy of related vehicles sold in the United States. We also find that the direct effect of the policy in the Japanese market was a 25.2% improvement in fuel economy. Taken together, these results imply an incomplete "pass-through" of fuel-economy gains from Japan to the U.S. market.

Importantly, these estimates do not imply that the policy's environmental impact in the United States is 0.34 (= 8.65 / 25.2) of its impact in Japan, because environmental externalities depend not only on fuel-

[1] Although our focus is on attribute propagation generated by economic policy, the same mechanism can also transmit changes in preferences or costs in one location to welfare outcomes in other locations.

economy improvements but also on sales volumes and vehicle miles traveled. To evaluate the policy's environmental impact in each country, we quantify the *spillover multiplier of environmental impacts* as the ratio of the Japanese policy's effect on CO_2 reductions in Japan and the United States to its domestic effect. We find that this spillover multiplier equals 6.72, implying that the Japanese policy's environmental impact abroad substantially exceeds its domestic impact.

A limitation of the DID design is that it cannot capture potential indirect equilibrium effects: vehicles that are not sold in Japan may adjust their fuel economy in response to changes in the fuel economy of competing vehicles. To account for these equilibrium effects, the second part of our paper develops a structural model of multinational vehicle markets with global spillovers. The model features two markets served by a mix of multinational products and products sold in a single location, in which firms choose product attributes (quality) and prices under Nash–Bertrand competition.

Our model links the two markets through firms' cost functions. In particular, the cost function includes a fixed cost of improving fuel economy, and for multinational products this fixed cost is lower when the attribute is more similar across versions sold in the two markets. Within this framework, we characterize the channels through which a subsidy to the attribute in one market affects equilibrium outcomes in the other market—directly by altering the attribute choice for multinational products, and indirectly by changing the equilibrium attribute choices of local products, as well as all equilibrium prices.

We then estimate the model following the tradition of Berry, Levinsohn and Pakes (1995) and the approach of Fan (2013) and Barwick, Kwon and Li (2024*a*) for identifying the slope of the fixed cost of adjusting product attributes. We estimate the automobile demand system separately for Japan and the United States to allow for differences in consumer preferences across the two markets. We then recover marginal costs and marginal fixed costs using the first-order conditions implied by automakers' profit maximization. We find statistical evidence of economies of scope in fuel economy across markets for a given vehicle model: deviating fuel economy in one market from that in the other incurs additional fixed costs of improving fuel economy.

Finally, we use the model and estimated parameters to conduct a counterfactual simulation in which we remove the fuel-economy subsidy policy and compute an equilibrium. Our simulation results confirm the headline finding on attribute propagation: most of the environmental benefits of the policy arise in the U.S. market. In the United States, vehicles not directly affected by the subsidy improve fuel economy in response to competitors' adjustments, generating positive indirect equilibrium effects. These indirect equilibrium effects further increase the spillover ratio of the policy's environmental impacts relative to the ratio computed using the DID estimates.

Overall, our findings indicate that incorporating indirect equilibrium effects underscores the importance of global spillovers. Although the Japanese subsidy reduces CO_2 emissions in both Japan and the United States, the resulting emissions reductions in the United States can be substantially larger than the domestic reductions in Japan. Consequently, abstracting from global spillover effects could lead to a substantial understatement of the policy's environmental impacts.

The automobile sector and environmental policy provide a natural setting in which to study such spillovers, but attribute propagation can arise in many other contexts, triggered by factors including safety regulations, antitrust rules, or differences in consumer preferences. For example, the European Union's 2022 directive requiring USB-C charging ports led Apple to adopt USB-C globally for subsequent iPhone models, rather than maintaining separate Lightning and USB-C versions. Similarly, pharmaceutical firms frequently comply with the stringent guidelines of the International Council for Harmonisation worldwide in order to access major markets, even when selling in countries with less demanding regulatory requirements. In aviation, manufacturers such as Boeing and Airbus design aircraft to meet both Federal Aviation Administration and European Aviation Safety Agency standards simultaneously, and then market these designs globally. Likewise, the European Union's REACH directive has become a de facto global standard, with firms adopting EU chemical regulations across their entire product lines to ensure access to the European market.

Each of these examples illustrates how policies in one jurisdiction can propagate through firms' global product-design decisions, generating spillovers that extend well beyond the regulating authority's borders. Our findings suggest that evaluating only domestic effects in such settings is likely to miss a substantial share of the total policy impact.

Related literature and our contributions—Our study contributes to literatures on cross-jurisdictional policy effects by identifying and quantifying a novel channel: attribute propagation through multinational firms' global product design decisions. There are other channels through which policy in one location can affect other jurisdictions. First, policy can cause shifts in the location of production. In environmental economics, this relates to the pollution haven hypothesis, which posits that tightening environmental regulation in one location can shift dirty activities to other markets (Levinson and Taylor, 2008; Copeland, 2008). A related literature considers how policy can induce trade in products (Davis and Kahn, 2010) and end of life treatment (Tanaka, Teshima and Verhoogen, 2022), rather than production per se. Second, policies can also affect outcomes in other jurisdictions when they induce the adoption of similar *policies*, which is sometimes referred to as the "California effect" (Vogel, 1995). Third is a cost channel: if local policy accelerates learning by doing or enables scale economies, this can affect additional markets through prices. This mechanism has been documented in electric vehicle batteries (Barwick, Kwon, Li and Zahur, 2024*b*; Head, Mayer, Melitz and Yang, 2025), semiconductors (Goldberg, Juhász, Lane, Lo Forte and Thurk, 2024) and clean energy technologies (Gerarden, 2023). Our paper examines a complementary but distinct mechanism in which regulatory pressure leads firms to modify product attributes—in our case, fuel economy in automobiles—with these design changes then carrying over to products sold in unregulated markets.

Attribute propagation is one manifestation of what is sometimes called the "Brussels effect" (Bradford, 2020), and there are some well-known examples, like the "50-state emissions" phenomenon in which automakers design all vehicles to meet California's stricter standards rather than maintain separate production lines. Even so, there appears to be very little prior empirical work, and only Sabal (2024), and the related work in Castro-Vincenzi, Menaguale, Morales and Sabal (2024), seem to address the same issue we do here, though with a different methodology and a different set of questions. Sabal (2024) develops a model of the global car industry where policy in one location can influence other market entries through a form of attribute propagation, but this occurs entirely through product entry. In order to achieve computational tractability, Sabal (2024) uses a highly aggregated unit of observation for a vehicle and holds attributes fixed within each vehicle type. In our context, we see no evidence of differential product entry in response to the Japanese policy, but instead quantify attribute propagation through modifications to existing models when looking at more granular definitions of a vehicle. Moreover, Sabal (2024) does not consider environmental outcomes or quantify environmental spillovers. As such, we view the two papers as complementary—we answer different questions using different modeling assumptions.

Our paper establishes the importance of attribute propagation in one particular setting, and it provides a methodology that others could follow. Our methodology could be used to study not just policy shocks, but also how preference shocks, or simply preference differences, in one location spillover into other locations through embodied attributes. In terms of policy analysis, ignoring attribute propagation would substantially understate the environmental benefits of the Japanese policy. We do not know if the Japanese government thought about these spillover effects when designing the policy, but we do believe this motive does sometime animates policy. The State of California, for example, often pursues an environmental leadership role where it hopes to accelerate the deployment of a low carbon product, like electric vehicles or clean trucks, that would be sold outside the state.

2 Background and Data

2.1 Japanese government's subsidy for fuel-efficient cars

Starting in 2009, the Japanese government provided a subsidy for consumers purchasing a new car with fuel economy in excess of the fuel economy target. There are three unique features of this policy that will help our empirical analysis. First, the subsidy ranged approximately between $700 and $1,500, which was a significant amount for consumers (about 5 to 10% of an average new car price).

Second, this subsidy was based on each vehicle's own fuel economy, rather than a corporate average fuel economy. A model was qualified for the subsidy if the model's fuel economy was above the fuel economy target given the model's weight. This allows us to exploit variation at the model level rather than at the corporate level.

Third, because the fuel economy target was designed as a step function of weight, the policy created variation across models in the difficulty of meeting the fuel economy target. For example, if a model's combination of fuel economy and weight before the policy was enacted (i.e., the pre-policy period) was already close to the target, the model would be able to qualify for the subsidy with small changes in its product attributes. In contrast, if a model was located far from the target function, it would need relatively larger changes in its product attributes in order to qualify for the consumer subsidy.

Fourth, many models sold in the US market were also possibly affected by Japan's policy because they were sold in both countries, whereas there were many other models in the US market that were not sold in Japan. This provides another source of variation to conduct difference-in-differences (DID) estimation. we take advantage of these four features to analyze data with three empirical methods.

2.2 Potential global spillover effects

Our question is whether Japan's fuel-economy subsidy policy generated global spillover effects through products sold by multinational firms. Before we begin with empirical analysis, it is helpful to see key descriptive statistics in the Japanese and global car markets to hypothesize which automakers are more likely to generate global spillover effects than others.

Figure 1 shows the market shares for new car sales in the Japanese car market in 2012 (Panel A). The top 10 were all Japanese automakers, and about 80% of new car sales were from Japanese firms. In contrast, while American firms sold a variety of cars in Japan, their sales quantities were extremely small, which are in part of "other" firms in the figure. These statistics imply that qualifying for the fuel-economy subsidy was likely to be important for Japanese firms but not for American firms—it is unlikely to make sense for American firms to spend a fixed cost to change their car designs to qualify for the subsidy as their sales quantities were low.

Panel B of Figure 1 illustrates the extent to which the Japanese market is important for each automaker relative to their worldwide sales. This figure shows again that Japan is a key market for Japanese firms but not for American firms. In addition, this figure suggests that for most of the major Japanese firms, such as Toyota, Honda, Nissan, Suzuki, Subaru, and Mitsubishi, Japan is a major market but Japan's share relative to these firms' worldwide sales are around 15-20%, which implies that these firms have high sales quantities in the rest of the world.

Overall, these descriptive statistics suggest that the global spillover effects can be heterogeneous among automakers, depending on their market share in Japan and the rest of the world. We provide our empirical analysis in Section 3.1 to test this hypothesis.

2.3 Data

We use three primary datasets. The first dataset records car specifications data from Japan and the United States. The Japanese and U.S. data sources for the specifications datasets are Car Sensor and Wards Auto data, respectively. The second dataset is monthly car sales quantity data from Marklines. The third dataset is monthly car production data from Marklines. Our datasets covers all car models sold in Japan and the United States between the 2003 and 2019 model years.

3 Difference-in-Differences Analysis

In this section, we use a difference-in-differences (DID) analysis to investigate whether the fuel-economy subsidy in Japan described in Section 2.1 generated global spillover effects through international car markets. Our main focus is on the US car market, although we find similar results in other countries, which we show in the Appendix.[2]

We begin by presenting evidence of the Japanese subsidy's spillover effects on fuel economy in the US car market in Section 3.1. We then investigate potential threats to the identification in Section 3.2, the underlying drivers of these spillover effects in Section 3.3, followed by an analysis of potential spillovers on other product attributes in Section 3.4 and on product entry and exit in Section 3.5.

3.1 Global Spillover Effects on Fuel Economy

During our sample period from 2003 to 2018, we observe that many vehicle models sold in the U.S. market by Japanese automakers were also offered in Japan, while many others were not. We exploit this variation to construct treatment and control groups for evaluating the impact of the Japanese subsidy policy, which was introduced in the 2010 model year. Using model year 2009 as the pre-policy baseline, we define the treated group as consisting of models sold in both the U.S. and Japan, and the control group as models sold in the U.S. but not in Japan.

Table 1 reports summary statistics of vehicle characteristics in 2009 for the treated and control groups for Japanese vehicles sold in the U.S. market. While the average vehicle price is slightly higher for the treated group, the difference is not statistically significant. Other characteristics are similar on average, and the differences are statistically insignificant across the two groups.

Our hypothesis is that, if the Japanese subsidy policy generated an international spillover effect on fuel economy, we would expect to observe an improvement in fuel economy for the treated group relative to the control group in the U.S. market. Our DID design is likely to yield a lower bound of the international spillover effect if there was also a within-firm technological spillover—where innovations adopted for models affected by the subsidy may be shared across untreated models within the same firm. If this was the case, it could lead to improvements in fuel economy even among the control group. Such within-firm spillovers would attenuate the estimated difference between the treated and control groups, thus making our estimates a lower bound of the international spillover effect.

Figure 2 presents the time trends of sales-weighted average log fuel economy in the U.S. market for vehicles sold by Japanese automakers. The figure indicates that the treated models—defined as those sold in both the U.S. and Japanese markets—and the control models—those sold in the U.S. but not in Japan—

[2]Table A.4 reports the DID results for Germany and India, which are similar to the findings for the United States presented in this section.

exhibited similar trends during the pre-policy period, spanning from 2003 to 2009. These parallel pre-trends include a decline in average fuel economy in 2008, which reflects a market-wide phenomenon driven by falling gasoline prices during that year.[3]

Following the introduction of the subsidy policy in 2009, fuel economy began to diverge between the treated and control groups, with the treated group exhibiting an increase of approximately 0.1 log points (roughly 10%) relative to the control group. The figure further suggests that the improvement in fuel economy for the treated group did not occur entirely in the immediate aftermath of the policy's implementation. Rather, the response appears to involve both short-run and medium- to long-run adjustments. This pattern is consistent with the typical product development cycle in the automobile industry, where major specification changes to a vehicle model occur only every few years. As such, some of the automakers' responses to the subsidy policy likely materialized with a delay.

While the graphical analysis provides a visual representation of the raw data trends, it does not account for potential confounding factors. To obtain DID estimates with controls, we estimate the following equation using ordinary least squares (OLS). The dataset comprises all vehicles sold by Japanese automakers in the U.S. market from 2003 to 2018, at the model-year (t) and model-by-trim (j) level:

$$\ln e_{jt} = \alpha D_{jt} + \theta_j + \lambda_t + \epsilon_{jt}, \tag{1}$$

where e_{jt} denotes the fuel economy, measured in miles per gallon (MPG), for vehicle trim i in model-year t. The treatment indicator D_{jt} equals 1 if the model is also sold in Japan and the model-year t is after the introduction of the Japanese subsidy. The specification includes model fixed effects (θ_j) to control for time-invariant heterogeneity across vehicle models and year fixed effects (λ_t) to account for common shocks over time. In Section 3.2, we also include year fixed effects interacted with a truck–car indicator and year fixed effects interacted with firm indicators to assess robustness. Standard errors are clustered at the model level to address serial correlation.

Table 2 reports the OLS estimates of Equation (1) for Japanese vehicles sold in the U.S. market. The difference-in-differences estimates in column 1 indicate that the Japanese subsidy policy generated an international spillover effect on fuel economy in the *U.S. market*, increasing fuel economy by 0.073 log points (7.57 percent).

[3] It is well established that lower gasoline prices tend to reduce average fuel economy, as consumers typically respond to contemporaneous fuel prices when making vehicle purchase decisions.

In Table A.1, we replicate the same DID estimation for *American vehicles* in the U.S. market. As discussed in Section 2.2, although American automakers do sell a range of models in both Japan and the U.S., their sales volumes in the Japanese market are very low. Consequently, we hypothesize that American automakers have limited incentive to respond to the Japanese subsidy. The empirical results in Table A.1 support this prediction: we find economically and statistically insignificant effects on fuel economy for vehicles produced by American automakers.

3.2 Potential Threats to Identification

The identification assumption underlying the DID estimation is the parallel trends assumption—namely, that in the absence of the Japanese subsidy policy, trends in fuel economy would have evolved similarly for the treated and control groups. A potential threat to this assumption is the presence of a confounding factor in the U.S. market that varies over time and differentially affects the treated and control groups.

Several things happened in the US auto market around this time. The Cash for Clunkers program was initiated as a stimulus program. Transactions were eligible for this subsidy only if they met minimum fuel economy requirements. This might have created an incentive to improve fuel economy. However, this would have affected both our treatment and control group, and the program was so short-lived (two months) that automakers had limited ability to respond by modifying and certifying a new configuration.

US fuel-economy standards (Corporate Average Fuel Economy, or CAFE) also changed in this time period. The law was changed in 2009, but changes in the standard did not take effect until 2012. Nevertheless, because there is some banking and borrowing allowed, automakers could have responded to the new law earlier than 2012. The 2009 reform to CAFE also introduced credit trading across firms. The Japanese automakers, particularly in the passenger car segment, were operating well above the standard, so changes in the standard would have little direct effect on them. However, trading meant that they could potentially monetize improvements in fuel economy by selling credits to other automakers, after 2009.

A tighter future fuel-economy standard and the availability of credit trading would have increased the incentives to improve fuel-economy today, but this should have a similar effect on treatment and control for Japanese vehicles. Even so, our identification strategy could be threatened if the effect of these reforms was different across the two sets of vehicles. The main reason we can think of why this would be true was if the control vehicles were skewed towards the truck segment (which includes SUVs and vans), where the Japanese companies had far less head room relative to the standard. In other words, prior to the advent of credit trading, fuel-economy standards were creating some shadow price on fuel economy improvements among Japanese trucks, but not among Japanese passenger cars.

To address this concern, we include interactions between time fixed effects interacted with car/truck indicator to allow for vehicle-type-specific time trends (columns 2 and 3). In column 3, we also add time fixed effects interacted firms indicators to allow for automaker-specific time trends. The results of these specifications are presented in column 2 and 3 of Table 2, and they suggest that the estimated treatment effect remains largely unchanged, providing further support for the validity of our identification strategy.

3.3 What Drives the Spillover Effects?

Beyond estimating the average treatment effect, we explore two potential sources of heterogeneity in the spillover effects. First, our data in the pre-subsidy period reveal that some vehicle models exhibited similar fuel economy within a model across the U.S. and Japanese markets, while others showed considerable differences in fuel economy between the two countries, despite being sold under the same model name. This suggests that the degree of product differentiation within a model across the two markets was heterogeneous in the baseline period.

We hypothesize that this pre-existing cross-market product differentiation may influence the magnitude of the international spillover effect. Specifically, less-differentiated models (i.e., those with similar fuel economy across markets) were likely designed and manufactured for the two markets. In contrast, more-differentiated models (i.e., those with greater fuel economy differences) were likely tailored separately for each market. Based on this reasoning, we expect that the spillover effects are larger for models with less pre-existing product differentiation.

We empirically test this prediction in column of Table 2 by interacting the treatment variable, D_{jt} in equation (1), with a measure of cross-market product differentiation. For each vehicle model in 2009, we calculate the average fuel economy separately for the U.S. and Japanese markets, and then compute the absolute value of the log difference between the two. We find that the interaction term is negative and statistically significant, indicating that models with higher levels of pre-existing product differentiation across markets exhibit smaller spillover effects. This finding is consistent with our theoretical prediction.

3.4 Potential Impacts on Other Product Attributes

The analysis thus far has focused on the spillover effects of the Japanese subsidy policy on fuel economy in the U.S. market. We can extend our DID framework in equation (1) to examine whether the policy also had spillover effects on other product attributes. In Table A.2, we apply our primary DID specification—reported in column 3 of Table 2 to estimate treatment effects on additional vehicle characteristics in the U.S. market. We do not find statistically significant effects on product attributes other than fuel economy, suggesting that automakers' responses were primarily by improving fuel economy.

3.5 Potential Impacts on Product Entry and Exit

Recent studies in the international trade literature emphasize that firms, including automakers, may respond to policy shocks through adjustments in product entry and exit (Sabal, 2024). In our context, however, such a response is less likely because the Japanese subsidy policy targeted a single product attribute: fuel economy. As a result, automakers may have found it more cost-effective to adjust the fuel economy of existing models rather than engage in more costly product entry or exit decisions.

Nevertheless, we empirically test whether the Japanese subsidy policy affected product entry and exit behavior. For each model-year, we identify product entries and exits to calculate net entry counts, which we then plot separately for the treated and control groups in Figure A.1. The net entries are similar between the two groups and similar before and after the introduction of the Japanese subsidy policy, and thus, it suggests little evidence that the subsidy policy differentially influenced net entry between the treatment and control groups.

We also provide statistical evidence in Tables A.3 by estimating DID regressions. The DID estimates indicate that the Japanese subsidy policy did not have statistically significant effects on net entry, entry, or exit, further supporting the view that automakers were likely to focus on adjusting the fuel economy of existing models in response to this policy.

3.6 Direct Policy Effects in the Japanese Market

A natural question is how large the spillover effect is relative to the policy's direct effect in the Japanese market. For the treated vehicles in Equation (1), we observe fuel economy in *Japan*. We can therefore estimate the DID specification using fuel economy in Japan as the outcome variable for the treated vehicles. The identifying assumption is that the fuel economy trend of the control group provides a valid counterfactual trend for the treated vehicles in the absence of the policy.

We have two possible predictions. One is *complete spillover*, in which automakers make identical fuel economy improvements in Japan and the United States. Under this scenario, the direct treatment effect on fuel economy in Japan would be similar to the spillover effect on fuel economy in the United States in Section 3.1, where we find an improvement of 0.083 log points (8.65 percent). Another possibility is *partial spillover*, where the direct effect on fuel economy in Japan would be larger than the spillover effect in the United States.

Table 3 provides evidence consistent with partial spillover. Column 3 shows that the direct effect on fuel economy in Japan is a 0.225 log point (25.2 percent) improvement on average, which exceeds the spillover effect estimated for the U.S. market (an 8.65 percent improvement). The resulting spillover-to-direct ratio in terms of fuel economy improvements is therefore 0.34 (= 8.65 / 25.2). However, this ratio alone does not capture the relative environmental impacts, as it does not account for differences in reduced externalities, which we examine in the next section.

3.7 Policy Implications: Spillover Multiplier of Environmental Impacts

The DID estimation results indicate that the Japanese policy's direct effect was an 25.2 percent improvement in fuel economy in Japan, whereas the spillover effect in the U.S. market was a 8.65 percent improvement.

Importantly, these estimates do not imply that the policy's environmental impact in the United States is 0.34 (= 8.65 / 25.2) of its impact in Japan, as environmental externalities depend not only on fuel economy improvements but also on sales volumes and vehicle miles traveled. First, the U.S. automobile market is substantially larger than that of Japan. Therefore, more vehicles in the US can be affected by the policy. Second, U.S. drivers travel substantially more miles per vehicle than Japanese drivers, resulting in more gasoline consumption per vehicle. The average annual miles traveled per vehicle are 14,489 in the United States and 4,181 in Japan. As a result, a given improvement in fuel economy may generate a larger reduction in environmental externalities in the United States.

To evaluate the policy's environmental impact in each country, we make the following steps. Our DID estimates provide the percentage change in fuel economy (miles per gallon) for affected vehicles. We convert these estimates into changes in gasoline consumption per mile (gallons per mile) and multiply them by the average annual miles driven per vehicle in each country to obtain the reduction in gasoline consumption.[4] We then translate the reduction in gasoline consumption into metric tons of CO_2 emissions avoided.

We define the *spillover multiplier* of environmental impacts (ρ) as follows:

$$\begin{aligned}
\rho &\equiv \frac{\text{Japanese policy's environmental impacts in Japan and the U.S.}}{\text{Japanese policy's environmental impacts in Japan}} \\
&= 1 + \frac{\Delta\text{Externality per vehicle}_{US} \times Q_{US}}{\Delta\text{Externality per vehicle}_{JP} \times Q_{JP}} \\
&= 1 + \frac{0.43 \text{ tons of CO}_2 \times 5,395,182}{0.20 \text{ tons of CO}_2 \times 2,074,181} \\
&= 1 + \frac{2,351,186 \text{ tons of CO}_2 \text{ per year}}{411,203 \text{ tons of CO}_2 \text{ per year}} \\
&= 6.72,
\end{aligned} \tag{2}$$

where ΔExternality per vehicle denotes the policy-induced change in CO_2 emissions per vehicle, and Q denotes the quantity of affected vehicles. ΔExternality per vehicle is larger in the United States because average annual miles traveled per vehicle are higher. Moreover, Q is larger in the United States because the affected models—those sold in both countries—have higher total sales in the U.S. market.

The spillover ratio of environmental impacts (ρ) is greater than one, implying that the Japanese policy's environmental impact abroad *exceeds* its domestic impact. This finding carries important policy implications. Standard analyses of environmental policies often focus exclusively on domestic effects, potentially leading to a substantial understatement of overall policy impacts.

There is one important limitation in our calculation of ρ based on the DID estimates. The calculation above focuses only on vehicles directly affected by the spillover effects of the fuel economy subsidy. What may be missing are potential equilibrium effects: vehicles that are not directly affected by the global spillover may nevertheless adjust their fuel economy in response to competitors' changes. In the next section, we develop a model of multinational automobile markets with global policy spillovers to incorporate these equilibrium effects into the calculation of ρ.

[4] For now, we abstract from a potential rebound effect—drivers may increase vehicle usage when fuel economy improves. This extension can be incorporated straightforwardly. Furthermore, if proportional rebound effects are identical across the two countries, they cancel out in Equation 2.

4 A Model of Multinational Car Markets with Global Policy Spillovers

Our model has two goals. First, we aim to model and estimate a potential mechanism of global policy spillovers. To do so, we extend a standard differentiated-product market model Berry, Levinsohn and Pakes (1995) to incorporate firms selling products in multinational markets and incorporate firms' endogenous attribute choices. This allows us to estimate potential cross-market links in revenues and costs.

The second objective of our model is to investigate the welfare implications of the global policy spillover effect. The equilibrium model allows us to examine the welfare implications of the global spillover effects by quantifying consumer surplus, producer surplus, and environmental externalities.

4.1 Demand

We follow Berry, Levinsohn and Pakes (1995) to model a consumer's new car purchase with a random utility model. We estimate demand in Japan and the United States separately, allowing the demand systems different between the markests.

We use p_{jc} to denote price for product j in market c and x_{jc} for a vector of product characteristics for product j in market c. Conditional indirect utility of consumer i who purchases product j can be written by: $u_{ijc} = \beta_i x_{jc} + \alpha_i p_{jc} + \xi_{jc} + \epsilon_{ijc}$, where ξ_{jc} is unobserved factors at the market-product level and ϵ_{ijc} is unobserved factors at the market-product-consumer level (type-I extreme value). The market share for product j in country c is:

$$s_{jc} = \int \frac{\exp(\beta_i x_{jc} + \alpha_i p_{jc} + \xi_{jc})}{\sum_{j'=0}^{J} \exp(\beta_i x_{j'c} + \alpha_i p_{j'c} + \xi_{j'c})} f(\mu_i) d\mu_i, \tag{3}$$

where $f(\mu_i)$ is the distribution of random-coefficients. The outside option is not to buy product $j = 1, ..., J$. This market share is usually unobservable from a dataset. A typical approach is to assume that s_{0c} is the number of consumers (households) in market c that did not buy any product j.

We begin by estimating demand using the standard logit model without random coefficients. In this specification, the preference parameters (α and β) do not vary across consumers, allowing Equation (3) to be written in linear form as $\ln s_j - \ln s_0 = \beta x_{jc} + \alpha p_{jc} + \xi_{jc}$. An advantage of this approach is that it can be consistently estimated using linear instrumental variables methods with valid instruments. A key limitation, however, is that the standard logit model imposes restrictive substitution patterns through the Independence of Irrelevant Alternatives (IIA) assumption.

To address the issue with the IIA, we use a random-coefficient logit approach for our main specification. We allow heterogeneity in β and α with log-normal distributions. An advantage of this approach is that it allows for flexible substitution patterns, less restrictive price elasticities, and heterogeneous consumer tastes. A key challenge is that nonlinear GMM estimation requires numerical simulation and does not guarantee convergence to a unique global optimum; therefore, careful implementation is necessary to obtain globally optimal estimates (Knittel and Metaxoglou, 2013; Conlon and Gortmaker, 2020).

4.2 Supply and Equilibrium

We describe the operating profit of multinational, multi-product firm f in each market as follows:

$$\begin{aligned} \text{Japan: } \pi_f &= \sum_{j \in J_f} \left[(p_j - c_j(e_j, x_j)) \cdot q_j \left(p_j - \tau_j(e_j), e_j, x_j \right) \right] \\ \text{US: } \tilde{\pi}_f &= \sum_{j \in \tilde{J}_f} \left[(\tilde{p}_j - \tilde{c}_j(\tilde{e}_j, \tilde{x}_j)) \cdot \tilde{q}_j \left(\tilde{p}_j, \tilde{e}_j, \tilde{x}_j \right) \right] \end{aligned} \tag{4}$$

where J_f denotes the set of cars sold by firm f; p_j is the price of car j; c_j is its marginal cost; e_j denotes fuel economy; x_j is a vector of other product attributes; q_j represents demand; and $\tau_j(e_j)$ is the fuel-economy subsidy in Japan. We use tildes to denote the corresponding variables in the U.S. market.

Firm f maximizes the joint profit from the two markets, with respect to prices and fuel economy:

$$\max_{p, e, \tilde{p}, \tilde{e}} \pi_f = \pi(p, e, x) + \tilde{\pi}(\tilde{p}, \tilde{e}, \tilde{x}) - \sum_{j \in J_f} FC(e_j, \tilde{e}_j) \tag{5}$$

where p, e, $\tilde{p}$, and $\tilde{e}$ are vectors of prices and fuel economy for all products. The function $FC(e_j, \tilde{e}_j)$ denotes the fixed cost of changing fuel economy, which we allow to depend on both e_j and $\tilde{e}_j$. In Section 5.3, we describe how we model this fixed cost to allow for cross-market complementarity.

Equation 5 implies that, in equilibrium, four first-order conditions—with respect to $(\tilde{p}_j, \tilde{e}_j, p_j, e_j)$—

must be satisfied for each product j.

$$\tilde{q}_j + \sum_{k \in \tilde{J}_f} \left[(\tilde{p}_k - \tilde{c}_k) \frac{\partial \tilde{q}_k}{\partial \tilde{p}_j} \right] = 0, \tag{6}$$

$$-\frac{\partial \tilde{c}_j}{\partial \tilde{e}_j} \tilde{q}_j + \sum_{k \in \tilde{J}_f} \left[(\tilde{p}_k - \tilde{c}_k) \frac{\partial \tilde{q}_k}{\partial \tilde{e}_j} \right] = \frac{\partial FC(e_j, \tilde{e}_j)}{\partial \tilde{e}_j}, \tag{7}$$

$$q_j + \sum_{k \in J_f} \left[(p_k - c_k) \frac{\partial q_k}{\partial p_j} \right] = 0, \tag{8}$$

$$-\frac{\partial c_j}{\partial e_j} q_j + (p_j - c_j) \left(\frac{\partial q_j}{\partial e_j} - \frac{\partial q_j}{\partial (p_j - \tau_j)} \frac{\partial \tau_j}{\partial e_j} \right) + \sum_{k \neq j \in J_f} \left[(p_k - c_k) \frac{\partial q_k}{\partial e_j} \right] = \frac{\partial FC(e_j, \tilde{e}_j)}{\partial e_j}. \tag{9}$$

Equations (6) and (8) are the first-order conditions with respect to prices, which are standard in the literature on differentiated product markets. For each firm f in each market, these conditions yield a system of J_j equations in J_j unknown marginal costs, allowing us to recover marginal costs given demand estimates.

Equations (7) and (9) are the first-order conditions with respect to fuel economy.[5] The left-hand side of Equation (7) represents the net marginal revenue from an increase in fuel economy. $\sum_{k \in \tilde{J}_f} (\tilde{p}_k - \tilde{c}_k) \frac{\partial \tilde{q}_k}{\partial \tilde{e}_j}$ captures the marginal revenue, while $-\frac{\partial \tilde{c}_j}{\partial \tilde{e}_j} \tilde{q}_j$ reflects the increase in marginal cost. This first-order condition therefore implies that firms equate the net marginal revenue with respect to fuel economy—the left-hand side—with the *marginal* fixed cost—the right-hand side—when endogenously choosing the optimal level of fuel economy.

Each element of Equations (7) can be obtained from data or the estimated demand and marginal cost functions. Once we obtain these, we can estimate a function of marginal fixed cost, $\frac{\partial FC(e_j, \tilde{e}_j)}{\partial \tilde{e}_j}$, with a parametric assumption. We will discuss this estimation strategy in Section 5.3.

Compared to Equation (7) for the U.S. market, Equation (9) for the Japanese market includes an additional term, $-\frac{\partial q_j}{\partial (p_j - \tau_j)} \cdot \frac{\partial \tau_j}{\partial e_j}$, for j. This term captures the marginal effect of fuel economy e_j on the subsidy $\tau_j(e_j)$ and the subsidy's effect on demand q_j. This term reflects how a change in e_j influences the level of the subsidy, thereby indirectly affecting consumer demand in Japan.

In Section 5, we use our data and policy-induced variation to estimate the model. We then use the estimated model to conduct counterfactual simulations in Section 6. In these simulations, the four first-order conditions in Equations 6–9 play a central role. To illustrate the mechanism, consider a change in the subsidy τ_j and the resulting new equilibrium. The direct effect of the subsidy change enters only the

[5] Our approach follows Fan (2013) in modeling endogenous product attributes, as in her analysis of newspapers.

fuel-economy first-order condition in Japan (Equation 9). However, any induced change in fuel economy in Japan can affect the optimal level of fuel economy in the U.S. (Equation 7) through the marginal fixed cost, provided that e_j and $\tilde{e}_j$ are complements in the marginal fixed cost function. These adjustments in optimal fuel economy then feed into the price first-order conditions in Equations 6 and 8, leading firms to choose new equilibrium levels of both prices and fuel economy.

4.3 Endogeneity and instruments

The standard BLP estimation considers that firms endogenously choose only p_j, taking other attributes exogenous. This approach is not appropriate in our context because we allow firms endogenously choose p_j and e_j.

Following Ito and Sallee (2018), we use a unique feature of the Japanese subsidy to create an instrumental variable for e_j. To be qualified for the subsidy, e_j needed to be above the target. As shown in Figure A.2, the fuel-economy target was a non-linear step function. This created variation in easiness/difficulties to qualify for the subsidy so that it created a policy-induced change in e_j in policy period. Recall that the subsidy was introduced in 2009. We create $\Delta e_j = e_j^{\text{target}} - e_{j,2008}$ as an instrument for e_j.

Figure A.2 visually shows the policy induced variation. We construct panel data of car models by linking cars sold in 2008 (before the policy change) and 2012 (three years after the policy introduction). Each dot in the figure shows a car's starting values of fuel economy and weight in 2008. For the cars that qualified for the new subsidy in 2012, we also show vectors connecting each car's starting position in 2008 to its final position in 2012.

This figure provides several useful results. First, many of the cars that gained the subsidy were redesigned in a way that they were just above the subsidy cutoff. Second, cars that started closer to the new standard were more likely to get the subsidy; that is, the "distance" to the subsidy cutoff explains most of the variation in which cars obtained the subsidy.

5 Estimation of the Model

In this section, we estimate the model presented in Section 4. We begin by estimating demand in Section 5.1, followed by the estimation of marginal costs in Section 5.2, and marginal fixed costs in Section 5.3. We then use the estimated model to conduct counterfactual policy simulations in Section 6.

5.1 Demand Estimation Results

Table 4 reports the demand estimation results for Japan and the United States, with the standard logit estimates shown in columns 1–2 and the random-coefficients logit estimates shown in columns 3–4.

In both markets, consumers value fuel economy, horsepower, and lower prices. The parameter σ represents the standard deviation of the log-normal random coefficient on price and indicates substantial unobserved heterogeneity in price sensitivity. Figure A.3 illustrates the implied distribution of the price coefficient to visualize this heterogeneity in price elasticity.

Overall, the demand estimates suggest that preference parameters are broadly similar across the two markets. One notable exception is the coefficient on fuel economy. Our results indicate that consumers in the United States place slightly greater value on fuel economy than consumers in Japan, potentially because U.S. drivers travel substantially more miles than Japanese drivers, as discussed in Section 3.7.

5.2 Marginal Cost Estimation Results

As discussed in Section 4.2, the first order conditions with respect to prices in Equations (6) and (8) yield a system of J_f equations in J_f unknown marginal costs for firm f, allowing us to recover the marginal costs (c_j) given the estimated demand system.

We regress the recovered marginal costs c_j on product attributes to estimate the marginal cost function. We estimate it separately for each market to allow for cross-market heterogeneity.

Table 5 reports the marginal cost estimation results for each market, with and without firm fixed effects. The estimated coefficient on horsepower is similar across the two countries, while the coefficient on fuel economy is larger in the United States. This pattern suggests that a unit increase in fuel economy leads, on average, to a larger increase in marginal cost in the United States than in Japan.

5.3 Marginal Fixed Cost Estimation Results

Our approach builds on the estimation of marginal fixed costs with respect to endogenous product attributes in Fan (2013) and Barwick, Kwon and Li (2024*a*). Our approach extends this method to incorporate cross-market complementarity in the marginal fixed cost.

Firm f's first order conditions with respect to fuel economy—Equations (7) and (9)—provides an estimate of the *margins* fixed cost with respect to an improvement in fuel economy, $\frac{\partial FC(e_j,\tilde{e}_j)}{\partial \tilde{e}_j}$. The left-hand sides of these equations can be calculated by data and estimates from the demand and marginal cost estimation. Then we can estimate the marginal fixed cost function with a parametric functional form assumption.

Consider that firms face the following fixed cost function for improving fuel economy:

$$FC(e_j, \tilde{e}_j) = \gamma + \gamma_0 \bar{e}_j + \gamma_1 \bar{e}_j^2 + \gamma_2 (e_j - \tilde{e}_j)^2, \tag{10}$$

where e_j and $\tilde{e}_j$ denote fuel economy in Japan and the United States, respectively, and $\bar{e}_j$ is the global average fuel economy: $\bar{e}_j = \frac{1}{2}(e_j + \tilde{e}_j)$. This specification implies that firms incur a quadratic cost when improving the global average level of fuel economy. In addition, the final term captures a potential economy of scope between fuel economy choices in Japan and the United States: γ_2 represents the cost of deviating fuel economy across the two markets.

Recall that Equations (7) and (9) characterize the *marginal* fixed costs, rather than the fixed cost itself. We therefore take derivatives with respect to fuel economy in Japan (e_j) and in the United States ($\tilde{e}_j$):

$$\begin{aligned}
\frac{\partial FC(e_j, \tilde{e}_j)}{\partial e_j} &= \frac{1}{2}\gamma_0 + \gamma_1 \bar{e}_j + 2\gamma_2 (e_j - \tilde{e}_j), \\
\frac{\partial FC(e_j, \tilde{e}_j)}{\partial \tilde{e}_j} &= \frac{1}{2}\gamma_0 + \gamma_1 \bar{e}_j - 2\gamma_2 (e_j - \tilde{e}_j).
\end{aligned} \tag{11}$$

We estimate these equations to recover the parameters γ_0, γ_1, and γ_2, controlling for firm fixed effects to capture unobserved heterogeneity across firms. The key parameter of interest is γ_2, which governs the degree of scope economies in fuel-economy choices across the Japanese and U.S. markets by capturing the cost of divergence in fuel economy for a given model. We use the policy-induced instruments discussed in Section 4.3 as instrumental variables for e_j and $\tilde{e}_j$ to address the potential endogeneity of fuel economy choices.

Table 6 reports the marginal fixed-cost estimation results. The positive and statistically significant γ_2 provides evidence of cross-market complementarity in firms' fuel-economy choices—the divination of fuel economy for given model between the two markets incurs an additional marginal fixed cost of improving fuel economy. To investigate the mechanism underlying the spillover effect, we use production-location data to construct an indicator variable, $1\{\text{Produced in each market}\}$, which equals one if vehicle model j has production location in both Japan and North America, and zero otherwise. The interaction term, $\gamma_2 \times 1\{\text{Produced in each market}\}$, indicates that the cross-market complementarity captured by γ_2 is close to zero for models that have production facilities in both markets. This finding implies that the estimated complementarity effect is driven primarily by models produced in a common plant and transported to both markets.

6 Counterfactual Policy Simulation

To investigate the welfare impact of the global policy spillover, we use our structural model in Section 4 and parameter estimates from Section 5 to simulate two scenarios. The first scenario is the actual scenario, where we include the fuel-economy subsidy policy in the Japanese market. The second scenario is a counterfactual scenario, where we remove the fuel-economy subsidy policy and compute an equilibrium.

6.1 Simulation Algorithm

In our policy simulation, we first introduce a counterfactual policy environment (e.g., removing Japan's fuel economy subsidy). Firms then endogenously choose four variables—prices and fuel economy in Japan (p_j, e_j) and in the United States $(\tilde{p}_j, \tilde{e}_j)$—by solving the first-order conditions (FOCs) in Equations (6)–(9), yielding a new equilibrium.

Solving the four first-order conditions simultaneously is computationally intensive due to the large number of products produced by multi-product firms and the presence of nonlinear equilibrium conditions, including a random-coefficients demand system. We therefore solve the first-order conditions using the following iterative procedure.

In the first iteration, we initialize the algorithm using the observed values of p_j, e_j, $\tilde{p}_j$, and $\tilde{e}_j$ from the data. Within the iteration, we treat these values as given and solve the first-order conditions in Equations (6)–(9) separately. For example, we solve Equation (9) with respect to fuel economy in Japan (e_j), holding fixed the other three endogenous variables, p_j, $\tilde{p}_j$, and $\tilde{e}_j$. Similarly, we solve Equation (7) with respect to fuel economy in the United States $(\tilde{e}_j)$, holding fixed p_j, $\tilde{p}_j$, and e_j.

Solving all four first-order conditions in this manner yields an updated set of p_j, e_j, $\tilde{p}_j$, and $\tilde{e}_j$. At the end of the iteration, we update the values of these endogenous variables, as well as demand and cost functions, which depend on them.

In the second and subsequent iterations, we repeat this procedure: each first-order condition is solved separately, taking the remaining endogenous variables from the previous iteration as given. At the end of each iteration, we update prices, fuel economy, demand, and costs, and proceed to the next iteration. We continue until the algorithm converges to a new equilibrium.

6.2 Counterfactual Policy Simulation Results

Table 7 reports the counterfactual policy simulation results based on the structural model described in Section 4 and the parameter estimates from Section 5. Column 1 presents the observed equilibrium with the fuel economy subsidy in place in the Japanese market. Column 2 reports the counterfactual equilibrium in which the subsidy is removed. Columns 3 and 4 report the differences between the two scenarios in levels and percentages, respectively.

We report sales-weighted average fuel economy (MPG) separately for all vehicles, affected vehicles, and unaffected vehicles to distinguish the overall effect, the direct effect, and indirect equilibrium effects. Affected vehicles are those that respond to the subsidy incentive and receive the subsidy in the observed equilibrium. As discussed in Section 3.7, vehicles not directly affected by the subsidy may nevertheless adjust fuel economy in response to competitors' changes, generating indirect equilibrium effects.

For affected vehicles, the simulation results are consistent with our findings based on the DID design in Sections 3.1 and 3.6. On average, the policy increases fuel economy by 21.53% in Japan, while generating a spillover effect of an 8.62% improvement in the U.S. market.

Indirect equilibrium effects are smaller than direct effects at the per-vehicle level. In Japan, the indirect effect on unaffected vehicles is a 3.51% decline in fuel economy, whereas in the United States the indirect effect is a 3.43% increase.

In Japan, these patterns are also reflected in aggregate CO_2 emissions. The overall policy effect is a reduction of 0.56 megaton of CO_2 per year. Most of this reduction comes from the direct effect on affected vehicles (0.60 megaton per year), while the indirect effect is positive but small, due to the large market share of affected vehicles.

In contrast, in the United States, although the per-vehicle indirect effect is smaller than the direct effect, it generates a substantial contribution to total emissions reductions because unaffected vehicles account for a large share of the market. We estimate CO_2 reductions of 2.12 megatons per year from affected vehicles and 3.51 megatons per year from unaffected vehicles, for a total reduction of 5.64 megatons per year.

Accounting for indirect equilibrium effects, the spillover ratio of the policy's environmental impacts (ρ) rises to 11.07 ($= 1+5.64/0.56$). This finding suggests that accounting for indirect equilibrium effects further underscores the importance of global spillovers: while the Japanese subsidy policy reduces CO_2 emissions in both Japan and the United States, the resulting reductions in the United States can be substantially larger than the domestic reductions in Japan, and therefore, abstracting from the global spillover effect could considerably understate policy impacts.

7 Conclusion

In a globalized economy, a country's domestic policies can generate global spillover effects through products designed and manufactured by multinational firms. In this paper, we study this phenomenon in the context of environmental regulation in the automobile market.

We find that Japan's fuel economy subsidy led to significant improvements in the fuel economy of vehicles sold in the U.S. market, thereby generating global environmental benefits. We then develop a model of multinational automobile markets to examine how cross-market linkages in revenues and costs give rise to such global spillovers. Using the estimated model, we conduct counterfactual policy simulations to quantify the environmental benefits and welfare effects of these global policy spillovers.

References

Barwick, Panle Jia, Hyuk-soo Kwon, and Shanjun Li. 2024*a*. "Environmental Externalities, Product Attributes, and Market Power: Implications for Government Subsidies." *Working Paper*.

Barwick, Panle Jia, Hyuk-Soo Kwon, Shanjun Li, and Nahim B Zahur. 2024*b*. "Drive down the cost: Learning by doing and government policies in the global EV battery industry." National Bureau of Economic Research Working Paper 33378.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica*, 63(4): 841–90.

Bradford, Anu. 2020. *The Brussels Effect: How the European Union Rules the World.* Oxford, UK:Oxford University Press.

Castro-Vincenzi, Juan Manuel, Eugenia Menaguale, Eduardo Morales, and Alejandro Sabal. 2024. "Market entry and plant location in multiproduct firms." Working Paper.

Conlon, Christopher, and Jeff Gortmaker. 2020. "Best practices for differentiated products demand estimation with pyblp." *The RAND Journal of Economics*, 51(4): 1108–1161.

Copeland, Brian R. 2008. "The pollution haven hypothesis." In *Handbook on Trade and the Environment*. , ed. Kevin P Gallagher, Chapter 4. Edward Elgar Publishing.

Davis, Lucas W, and Matthew E Kahn. 2010. "International trade in used vehicles: The environmental consequences of NAFTA." *American Economic Journal: Economic Policy*, 2(4): 58–82.

Fan, Ying. 2013. "Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market." *American Economic Review*, 103(5): 1598–1628.

Gerarden, Todd. 2023. "Private innovation and government subsidies with learning-by-doing spillovers." *Management Science*. DOI: 10.1287/mnsc.2022.4662.

Goldberg, Pinelopi K, Réka Juhász, Nathan J Lane, Giulia Lo Forte, and Jeff Thurk. 2024. "Industrial policy in the global semiconductor sector." National Bureau of Economic Research Working Paper 32651.

Head, Keith, Thierry Mayer, Marc Melitz, and Chenying Yang. 2025. "Industrial policies for multi-stage production: The battle for battery-powered vehicles." *Manuscript*.

Ito, Koichiro, and James Sallee. 2018. "The Economics of Attribute-Based Regulation: Theory and Evidence from Fuel-Economy Standards." *Review of Economics and Statistics*, 100(2): 319–336.

Knittel, Christopher R., and Konstantinos Metaxoglou. 2013. "Estimation of Random-Coefficient Demand Models: Two Empiricists' Perspective." *Review of Economics and Statistics*, 96(1): 34–59.

Levinson, Arik, and M Scott Taylor. 2008. "Unmasking the pollution haven effect." *International Economic Review*, 49(1): 223–254.

Sabal, Alejandro. 2024. "Global product entry in the automobile industry." Working Paper, Princeton University.

Tanaka, Shinsuke, Kensuke Teshima, and Eric Verhoogen. 2022. "North-South displacement effects of environmental regulation: The case of battery recycling." *American Economic Review: Insights*, 4(3): 271–288.

Vogel, David. 1995. *Trading up: Consumer and environmental regulation in a global economy.* Cambridge, MA:Harvard University Press.

Figures and Tables

Figure 1: Market Shares

Panel A: Market Share in the Japanese market in 2012

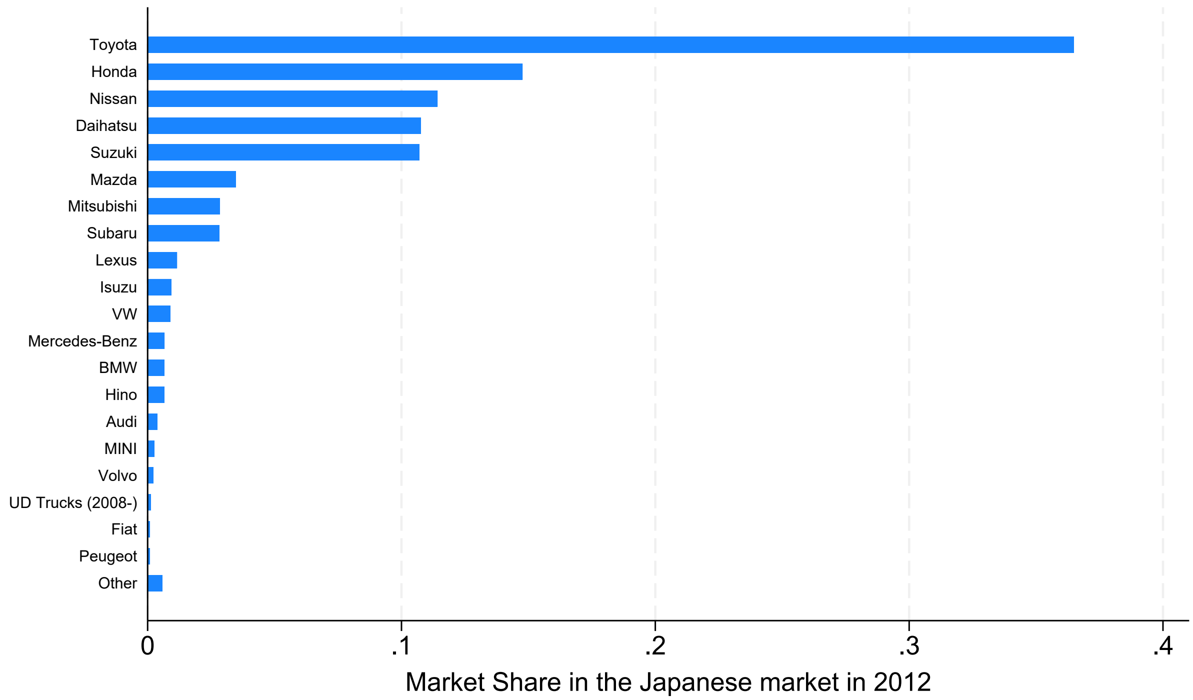

Panel B: Each firm's Japanese market share relative to its worldwide sales

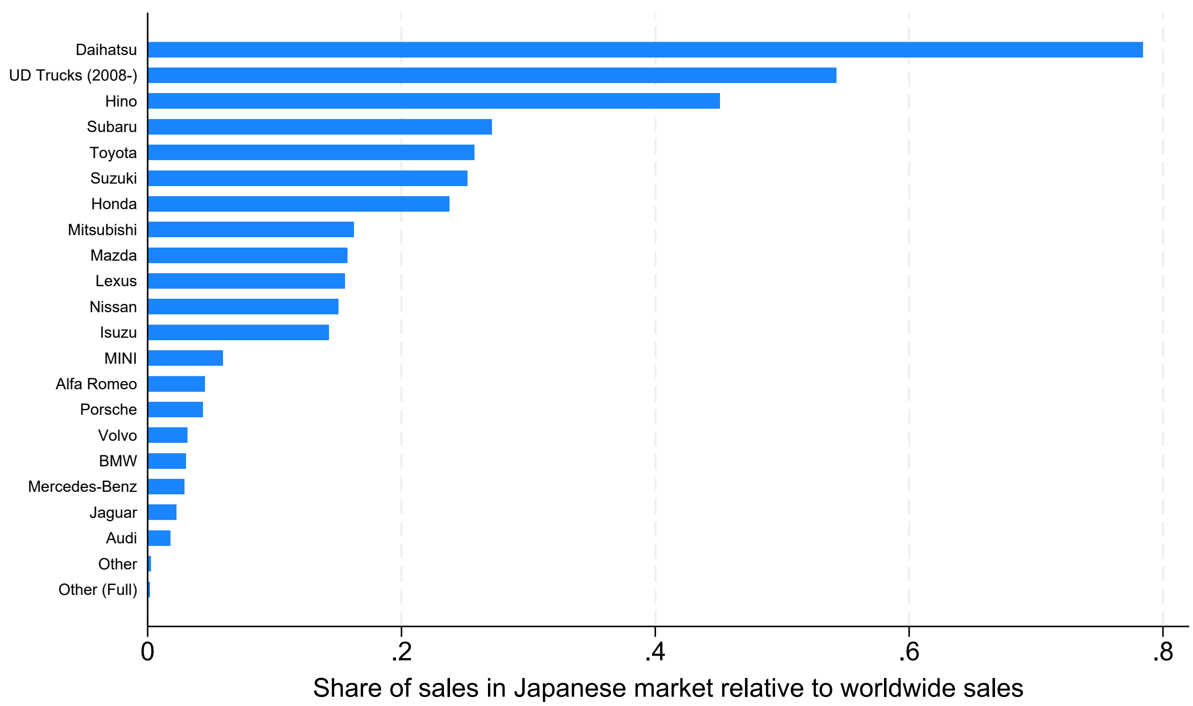

Figure 2: Average Fuel Economy of Japanese Vehicles in the US market

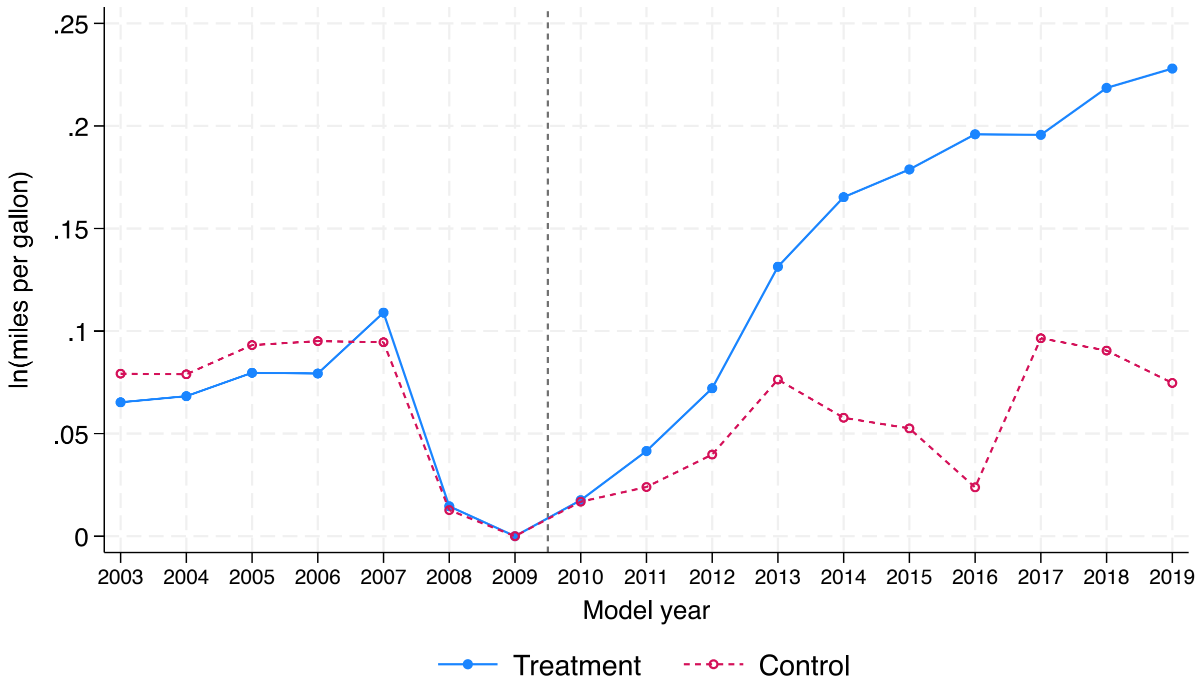

Note: Each dot shows the log of fuel economy in the US auto market for each group, normalized at their 2009 level so that it shows the changes in the log of fuel economy relative to 2009. The treatment group includes Japanese automakers' vehicles that were sold in both the United Sates and Japan. The control group includes Japanese automakers' vehicles that were sold in the United Sates but not in Japan.

Table 1: Baseline Characteristics of Treatment and Control Groups Prior to the Subsidy

	Treated		Control		Difference		
	Mean	S.D.	Mean	S.D.	Mean	S.E.	p-value
Price (1,000 USD)	35.0	10.4	32.4	9.6	2.7	1.9	0.17
Miles per gallon	18.2	3.2	18.9	3.7	-0.7	1.1	0.55
Horsepower	288.3	65.3	268.1	79.9	20.2	21.7	0.35
Length (feet)	18.3	2.6	17.7	2.3	0.6	1.0	0.57
Width (feet)	6.6	0.6	6.4	0.6	0.2	0.2	0.34
Height (feet)	6.0	0.6	5.9	0.8	1.5	0.1	0.1
Wheelbase (feet)	11.3	2.0	10.8	1.8	0.6	0.8	0.46
Footprint (square feet)	75.3	17.7	69.6	16.2	5.7	7.0	0.42
Weight (1,000 lbs)	5.5	1.5	4.7	1.2	0.8	0.6	0.23

Notes: This table reports summary statistics of vehicle characteristics in 2009, one year prior to the introduction of the Japanese fuel economy subsidy, for Japanese vehicles sold in the U.S. market. The treated group consists of models sold in both the U.S. and Japan, while the control group includes models sold in the U.S. but not in Japan. *S.D.* denotes the standard deviation, *S.E.* refers to the standard error of the difference in means, and the *p*-value corresponds to a test of the null hypothesis that the difference in means between the two groups is equal to zero.

Table 2: Global Spillover Effects of the Japanese Fuel-Economy Subsidy on the US Market

Dependent variable: log fuel economy

	(1)	(2)	(3)	(4)
Treated × Post	0.073	0.090	0.083	0.151
	(0.024)	(0.022)	(0.026)	(0.038)
Treated × Post × Differentiation				-0.300
				(0.090)
N	9,098	9,098	9,098	7,159
Model FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Year × Truck FE	No	Yes	Yes	Yes
Year × Make FE	No	No	Yes	Yes

Note: This table shows the OLS regression results of equation (1). The dependent variable is the log of fuel economy (miles per gallon) at the make-model-trim level between model years 2003 and 2019. All regressions are weighted by the average annual sales. Standard errors are clustered at the model level. The data include all vehicles sold by Japanese automakers in the US automobile market.

Table 3: Direct Effects of the Japanese Fuel-Economy Subsidy in Japan

	Dependent variable: log fuel economy		
	(1)	(2)	(3)
Treated × Post	0.203	0.276	0.225
	(0.081)	(0.125)	(0.115)
N	12,812	12,810	12,810
Model FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Year × Truck FE	No	No	Yes
Year × Make FE	No	Yes	Yes

Note: This table shows the OLS regression results in Section 3.6. The dependent variable is the log of fuel economy (miles per gallon) at the make-model-trim level between model years 2003 and 2019. All regressions are weighted by the average annual sales. Standard errors are clustered at the model level. The data include all vehicles sold by Japanese automakers in the US automobile market.

Table 4: Demand Estimation Results

	Standard logit		Random-coefficient logit	
	Japan	US	Japan	US
Mean Price/Income (USD)	-1.28	-4.81	-11.57	-14.42
	(0.496)	(0.428)	(3.37)	(2.79)
Fuel economy (mpg)	0.079	0.117	0.058	0.133
	(0.010)	(0.020)	(0.013)	(0.020)
Horsepower	0.012	0.015	0.018	0.016
	(0.004)	(0.003)	(0.005)	(0.003)
sigma			0.688	0.580
			(0.248)	(0.168)
Observations	2142	1469	2142	1469

Note: This table shows the demand estimation results of our structural model in Section 5.1. We estimate random-coefficients on prices with the log-normal distribution and report their means and standard deviations σ. Figure A.3 visualizes the distribton of the price coefficients in each market.

Table 5: Marginal cost estimation results

	Japan		US	
	(1)	(2)	(3)	(4)
Fuel economy (MPG)	348.27	348.77	1645.76	1434.95
	(51.19)	(33.13)	(238.88)	(235.89)
Horsepower	334.54	313.58	280.37	261.27
	(18.32)	(19.71)	(26.31)	(32.70)
Constant	-38095.77		-79591.67	
	(4408.34)		(11593.30)	
Firm FE	No	Yes	No	Yes
Observations	2142	2142	1469	1469

Note: This table shows the marginal cost estimation results described in Section 5.2.

Table 6: Marginal fixed cost estimation results

	(1)	(2)	(3)	(4)
γ_1	0.117	0.121	0.081	0.173
	(0.017)	(0.017)	(0.013)	(0.018)
γ_2	0.074	0.029	0.055	0.053
	(0.009)	(0.007)	(0.006)	(0.010)
$\gamma_2 \times 1\{\text{Produced in each market}\}$			-0.056	-0.045
			(0.019)	(0.010)
Firm FE	No	Yes	No	Yes
Observations	3611	3611	3611	3611

Note: This table shows the marginal fixed cost estimation results described in Section 5.3.

Table 7: Counterfactual Simulation Results

	Actual	Counterfactual (no subsidy in Japan)	Difference	Difference (%)
Japan				
Fuel economy (MPG): All vehicles	49.97	43.51	6.46	14.85
Fuel economy (MPG): Affected vehicles	56.49	46.48	10.01	21.53
Fuel economy (MPG): Others	35.71	37.01	-1.30	-3.51
Total CO2 (Mt): All vehicles	4.17	4.72	-0.56	-11.83
Total CO2 (Mt): Affected vehicles	2.43	3.03	-0.60	-19.71
Total CO2 (Mt): Others	1.73	1.70	0.04	2.25
USA				
Fuel economy (MPG): All vehicles	25.29	24.45	0.84	3.43
Fuel economy (MPG): Affected vehicles	28.38	26.13	2.25	8.62
Fuel economy (MPG): Others	24.35	23.94	0.41	1.71
Total CO2 (Mt): All vehicles	65.33	70.97	-5.64	-7.94
Total CO2 (Mt): Affected vehicles	13.79	15.91	-2.12	-13.34
Total CO2 (Mt): Others	51.55	55.06	-3.51	-6.38

Note: This table shows the counterfactual simulation results described in Section 6.

Appendix Figures

Figure A.1: Net Product Entry

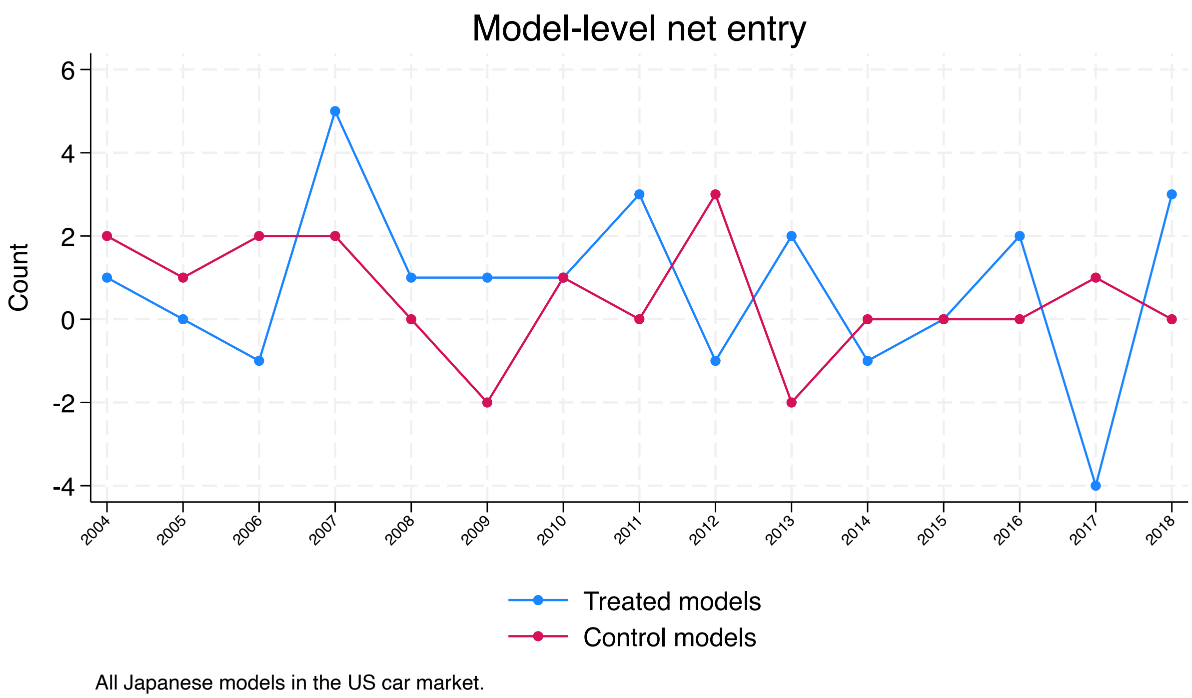

Note: This figure shows the product entry and exit discusse in Section 3.5. For each model-year, we identify product entries and exits to calculate net entry counts, which we then plot separately for the treated and control groups.

Figure A.2: Subsidy take-up

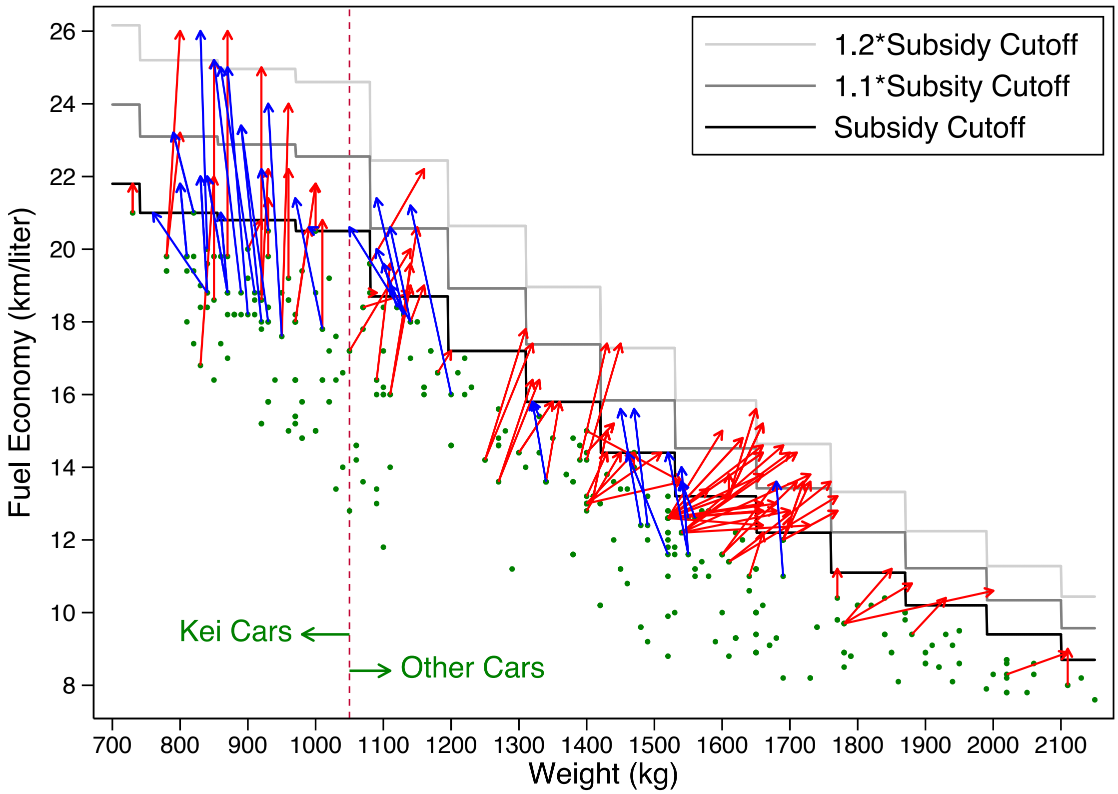

Note: This figure shows the policy induced variation. We construct panel data of car models by linking cars sold in 2008 (before the policy change) and 2012 (three years after the policy introduction). Each dot in the figure shows a car's starting values of fuel economy and weight in 2008. For the cars that qualified for the new subsidy in 2012, we also show vectors connecting each car's starting position in 2008 to its final position in 2012.

Figure A.3: Log-Normal Distribution of Price/Income Coefficients

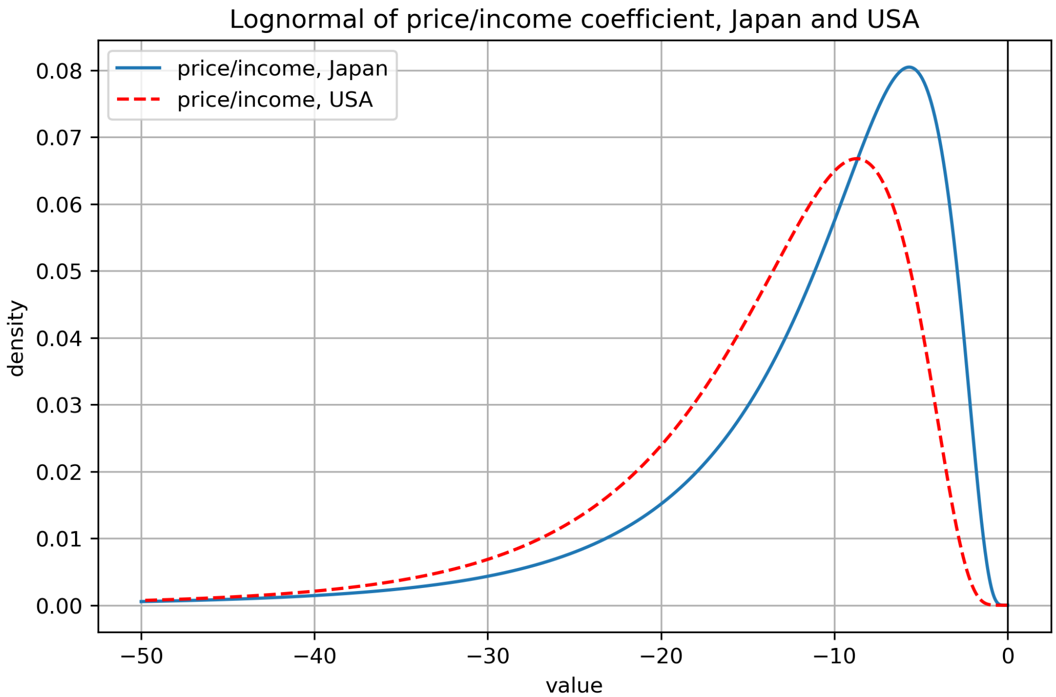

Note: This figure shows the distributions of price coefficients estimated by the random-coefficient logit model in Section 5.1.

Appendix Tables

Table A.1: International Spillover Effects of Japan's Fuel Economy Subsidy on the US Automobile Market (American cars in the US market)

American cars in the US market. Dependent variable is the log of each attribute.

	(1) MPG	(2) Horsepower	(3) Price	(4) Wheelbase	(5) Footprint	(6) Weight
Treated × Post	-0.025	0.051	0.016	-0.003	0.000	-0.007
	(0.028)	(0.035)	(0.023)	(0.007)	(0.009)	(0.013)
N	21,567	21,752	21,661	21,762	21,762	21,719

Note: This table shows the OLS regression results of equation (1). The dependent variable is the log of each car attribute at the make-model-trim level between model years 2003 and 2019. All regressions are weighted by the average annual sales. Standard errors are clustered at the model level. Panel A includes vehicles sold by Japanese automakers (Honda, Isuzu, Lexus, Mazda, Mitsubishi, Nissan, Subaru, Suzuki, and Toyota). Panel B includes vehicles sold by American automakers (Buick, Cadillac, Chevrolet, Chrysler, Dodge, Fisker, Ford, GMC, Hummer, Jeep, Lincoln, Mercury, Oldsmobile, Pontiac, Saturn, Tesla, and Wheego). Year by Car or Truck FE

Table A.2: Potential Spillover Effects on Other Product Attributes

Dependent variable is the log of each vehicle characteristic in the US market.

	(1) MPG	(2) Horsepower	(3) Price	(4) Wheelbase	(5) Footprint	(6) Weight
Treated × Post	0.083	-0.062	0.001	-0.010	-0.001	-0.006
	(0.026)	(0.045)	(0.019)	(0.007)	(0.008)	(0.016)
N	9,098	9,134	9,124	9,134	9,134	9,120

Note: This table shows the OLS regression results of equation (1). The dependent variable is the log of each car attribute at the make-model-trim level between model years 2003 and 2019. All regressions are weighted by the average annual sales. Standard errors are clustered at the model level. Panel A includes vehicles sold by Japanese automakers (Honda, Isuzu, Lexus, Mazda, Mitsubishi, Nissan, Subaru, Suzuki, and Toyota). Panel B includes vehicles sold by American automakers (Buick, Cadillac, Chevrolet, Chrysler, Dodge, Fisker, Ford, GMC, Hummer, Jeep, Lincoln, Mercury, Oldsmobile, Pontiac, Saturn, Tesla, and Wheego). Year by Car or Truck FE, Model FE, and Year by Make FE

Table A.3: Difference-in-differences Estimation on Entry, Exit, Net Entry

	(1) Entry Ratio	(2) Exit Ratio	(3) Net Entry Ratio
Treated × Post	0.063	0.038	0.010
	(0.054)	(0.056)	(0.064)
N	32	32	30
Year FE	Yes	Yes	Yes
T-C group FE	Yes	Yes	Yes

Note: For each model-year, we identify product entries and exits to calculate net entry counts. We then compute the entry, exit, and net entry ratios for each group by dividing the respective counts of entry, exit, and net entry by the total number of models in that group for the given year.

Table A.4: Additional Evidence from Other Countries: Global Spillover Effects on Fuel Economy

Panel A: Germany

	(1)	(2)	(3)	(4)
Treated × Post	0.083	0.076	0.078	0.076
	(0.035)	(0.031)	(0.024)	(0.020)
Treated	-0.263	-0.263		
	(0.114)	(0.115)		
Post	0.061		0.047	
	(0.022)		(0.014)	
N	547	547	543	543
Year FE	No	Yes	No	Yes
Model FE	No	No	Yes	Yes

Panel B: India

	(1)	(2)	(3)	(4)
Treated × Post	0.173	0.144	0.285	0.272
	(0.135)	(0.142)	(0.056)	(0.060)
Treated	-0.016	-0.016		
	(0.139)	(0.143)		
Post	0.115		-0.006	
	(0.123)		(0.009)	
N	147	147	145	145
Year FE	No	Yes	No	Yes
Model FE	No	No	Yes	Yes

Note: These tables shows our DID estimation in Equation (1) using data from Germany and India.