Predicting the Effects of New Goods, with an Application to Chinese Electric Vehicles in the United States

Hunt Allcott Luke Heeney Alejandro Sabal*

PRELIMINARY AND INCOMPLETE: PLEASE DO NOT CIRCULATE

November 6, 2026

Abstract

We develop a framework to predict demand for existing differentiated products in new markets, weighting the markets where each product is already sold by how closely their preferences correlate with those of the target market. We apply the framework to one of the most popular new goods available globally, which American consumers currently cannot buy: Chinese electric vehicles (EVs). We exploit a new multi-country survey of vehicle buyers and second choices, where Chinese EV buyers outside of China report that if Chinese EVs were not available, 33 percent would have bought gas vehicles, and only 26 percent would have bought from other Asian manufacturers. Using second-choice data from that survey, our model predicts that across 19 countries where Chinese EVs have already entered, in 2025 they increased consumer surplus by $2.49 billion, reduced CO2 emissions by 2.02 million tons, and reduced competing firms' profits by $1.86 billion. We solve an entry game calibrated to 2030 preferences for EVs in the United States and find that Chinese EVs would increase U.S. consumer surplus by $9.66 billion and reduce CO2 emissions by 18.65 million tons, but would reduce American automakers' profits by $2.49 billion and reduce other automakers' American labor demand by 29.0 thousand jobs. Methodologically, we evaluate our prediction framework out-of-sample and show that it yields less biased and variable predictions than existing approaches, though predictions remain sensitive to errors.

JEL Codes: H21, H22, H23, H24, L13, L50, L91, Q58.

Keywords: Electric vehicles, new goods, demand estimation, forecasting.

*Allcott: Stanford and NBER. allcott@stanford.edu. Sabal: University of Chicago. sabalalejandro@uchicago.edu. We thank Jim Sallee, Joe Shapiro, and seminar participants at the University of California Energy Institute for helpful comments.

1 Introduction

One of the key economic policy questions facing the United States is the extent to which the country should decouple its economy from China. Perhaps the most stark example of the tradeoffs underlying this question is in electric vehicles (EVs). Chinese EVs' high quality and low costs could offer large benefits to consumers and the environment, but potentially at the expense of other firms that might have become more competitive over time. While Chinese EVs have exploded in global popularity over the past five years, they remain excluded from the United States market. This is striking: excluding such an important foreign consumer product has no precedent in modern U.S. history.

Evaluating the tradeoffs in this policy decision is an example of a more general economic problem: evaluating the effects of existing goods entering in new markets. Electronics, smartphone apps, consumer packaged goods, pharmaceuticals, and other goods and services often enter in some markets before expanding to others. The entrants, their competitors, and policymakers may want to predict demand, profits, and other outcomes following potential entry in new markets.

A holistic equilibrium prediction requires at least three ingredients: demand, cost, and competition. Applied to Chinese EVs, these three ingredients raise the following questions. On the demand side, how certain can we be about how well Chinese EVs would sell in the United States, and what products are their closest substitutes? Substitution patterns matter for understanding how much Chinese EVs would accelerate adoption of EVs overall, and how much they would cannibalize sales of domestic versus foreign incumbents. On the cost side, how much would it cost to modify and produce Chinese EVs for the United States market? And on the competition side, how would incumbent firms respond with their pricing and product offerings?

Our paper studies these questions with new data and new models of vehicle supply and demand. We first model the effects of Chinese EVs in markets where they have already entered. We then forecast the potential future effects of Chinese entry into the United States market. To understand substitution patterns, we partnered with a large global management consultancy to include "second choice" questions in their major multi-country survey of recent vehicle buyers—i.e., questions asking what car people would have bought if the model they actually bought (or if all Chinese EVs) were not available.

To predict the preference for existing goods in new markets, we apply tools from statistics that account for the fact that preferences in the market of interest may be more correlated with some markets than others.[1] We show that market shares across overlapping products are differentially correlated across markets, which suggests that the performance of Chinese EVs in markets "closer" to the United States might be more informative than their performance in more dissimilar markets. Our approach applies interactive fixed effects (IFE) as in Bai (2009) and Athey et al. (2021) in the context of a mixed-logit demand model, similarly to

[1] We could also apply our procedure to costs, but we focus on preference predictions because we find there to be much less uncertainty on the cost side.

Moon, Shum, and Weidner (2018), to predict preferences for Chinese EVs in new markets. We evaluate our procedure out-of-sample and show that it provides less biased and more precise predictions than standard approaches in this literature, which estimate product fixed effects or preferences for observed characteristics and apply those estimates to new markets. Our procedure also highlights the importance of capturing out-of-sample uncertainty when evaluating the effect of entry into new markets.

We begin by establishing several basic facts about global sales of Chinese EVs. As of the end of 2025, we observe Chinese EV sales in 40 of the largest 50 vehicle markets. In those 40 countries, Chinese EVs went from 0.007 percent of new vehicle sales in 2019 to 3.0 percent in 2025. Chinese firms now make more EVs than any other country's firms. Compared to other EVs of the same body style (cars, trucks, pickups, SUVs) sold in the same country, Chinese EVs have 22.0 percent lower prices.

In the vehicle buyer survey, 33 percent of Chinese EV owners reported that if all Chinese EVs were unavailable, they would have bought a gasoline vehicle instead. This suggests that the availability of Chinese EVs meaningfully speeds overall EV adoption, instead of simply cannibalizing other EV sales. Only 26 percent of Chinese EV owners reported that they would have bought from another Asian firm if all Chinese EVs were unavailable, suggesting substantial Chinese-EV cannibalization of U.S. and European firms.

To simulate the effects of Chinese EV entry, we estimate country-specific supply and demand systems for 2015–2025, following the framework of Berry, Levinsohn, and Pakes (1995, 2004) and Petrin (2002). Our demand system includes random coefficients on key characteristics (including powertrain, body style, and brand origin) identified from second-choice data. We recover the mean utilities for each product offered in each market and model mean utility as depending on observed characteristics and a product-market-specific residual unobserved to the researcher, ξ_{jmt}. To capture residual correlation in preferences across markets, we explicitly model this residual as depending on a free product-market component and a low-rank IFE component $\lambda_j' F_{mt}$, similarly to Moon, Shum, and Weidner (2018), though in an unbalanced rather than a balanced panel. We assume static Nash-Bertrand pricing to infer marginal costs and simulate prices.

For the 19 countries where Chinese EVs have already entered and complete data are available, we simulate a counterfactual scenario without Chinese EVs. Globally, our model predicts that Chinese EVs increased consumer surplus by $2.49 billion in 2025, with the gains primarily driven by higher-income car buyers. We predict that Chinese EVs increased global EV sales by 0.141 million units. Given each country's average CO2 emission rate for electricity generation, this would reduce CO2 emissions by 2.02 million tons over the vehicles' lifetimes. The model predicts that through both diverted sales and more competitive pricing, Chinese EVs reduced non-Chinese firms' profits by $1.86 billion.

We validate our IFE approach for predicting demand in new markets in two steps. First, we leverage the 19 markets in which Chinese EVs have entered by 2025 and we hold out the mean utilities of the Chinese EVs in each of these countries separately. Then, we compare how IFE and OLS (with and without product fixed effects) predict product-specific mean utilities for Chinese EVs in each of these countries. We find

that IFE achieves a mean absolute error of 1.12 relative to 1.55 and 1.66 under OLS with and without product fixed effects, respectively. Moreover, IFE dominates the other approaches in terms of bias and spread. Across the 19 countries, the bias in total-surplus predictions from removing Chinese EVs is 0% when replacing Chinese-EV mean utilities with those predicted with IFE, versus 12% and 28% under OLS with and without product effects. While we find that all approaches suffer from large standard deviations in the out-of-sample error, we find that IFE also dominates in this dimension.

Second, we randomly hold out in 20 iterations half of the products available in the United States in each year in our sample. In each of the iterations, we re-estimate mean utilities using our IFE approach, as well as OLS with and without product fixed effects, and evaluate such procedures for predicting preferences specifically in the United States. We find that IFE also dominates for these purposes, achieving a mean-absolute error of 0.86, while OLS achieves a mean absolute error of 1.14 without product effects and, perhaps surprisingly, an even higher error of 1.25 with product fixed effects. We find that product effects perform poorly because U.S. demand residuals are weakly correlated with the residuals of overlapping products in the average country, despite being systematically correlated with residuals in nearby countries like Canada. Taking stock, our validation exercises demonstrate that IFE performs better than OLS, and that adding product effects may actually worsen rather than improve out-of-sample predictions.

In the final part of the paper, we turn to the forward-looking question of the long-run effects of Chinese entry in the United States market. In recent years, there has been bipartisan agreement among U.S. policymakers to impose prohibitive import tariffs and otherwise discourage entry, due to static protectionism (i.e., protecting American jobs and profits), dynamic infant industry arguments (i.e., that U.S. automakers will eventually grow and be competitive with the Chinese if given more time for learning-by-doing), and other concerns. However, political winds shift, and policy experts and media alike have begun considering the potential benefits of allowing Chinese EVs to enter Western markets (e.g., Wagner and Walsh (2024), Wagner and Wei (2024), The Economist (2024), Washington Post Editorial Board (2026)).

To predict the effect of potential Chinese EV entry in the United States market, we predict Chinese-EV demand using our IFE approach and proceed in three steps. First, we calibrate the United States EV demand shifter to match the BloombergNEF (2026) 2030 forecast of a 17% BEV inside share. Second, we borrow product-market fixed-cost estimates from Sabal (2026), and truncate fixed-cost draws to reflect observed offering and non-offering decisions in 2025: products offered in the United States must have had sufficiently low fixed costs to rationalize their entry, while potential unoffered products must have had fixed costs larger than their marginal contribution to profits. Finally, we solve a complete-information entry game using the solution method developed in Castro-Vincenzi et al. (2026).

Our model predicts that a policy removing prohibitive tariffs would lead to a Chinese-EV market share of 5.6 percent. Chinese EV entry would also lead to $9.66 billion gains in consumer surplus and 18.65 million tons in reduced CO2 emissions. However, we predict that 33.2 percent of diverted sales would be

from American automakers, and 63.7 percent of diverted sales would be from models manufactured in the United States. The model predicts that Chinese EV entry would reduce American automakers' profits by $2.49 billion and would reduce non-Chinese automakers' American labor demand by 29.0 thousand jobs. Plotting these losses on a map, 19 percent of them would be in swing states where the presidential vote share difference was less than five percent in 2024. We also find that Chinese EV entry predominantly benefits richer households which we estimate are less price-sensitive.

To speak to concerns about Chinese-EV entry crowding out U.S. incumbents' incentives to develop new EVs, we allow incumbents to introduce new products in 2030: twins of gas vehicles already offered in 2025. Overall, we do not find strong evidence for Chinese EVs crowding out incumbents on the extensive margin. In the absence of Chinese-EV entry, we find that incumbents would offer a net 2.1 more Cars and 1.2 more SUVs; with Chinese EV entry, incumbents would offer on net 0.3 fewer SUVs. Overall, we find that most of the product exit by 2030 is due to GVs exiting in response to greater relative preference for EVs, and that Chinese entry would not substantially affect incumbent entry.

To evaluate the importance of accounting for uncertainty in predictions about Chinese-EV demand in the United States, we solve the 2030 entry game at multiple draws of the out-of-sample errors obtained in our validation exercise, as well as multiple fixed-cost realizations. We find that the central 95% interval for the predicted U.S. Chinese-EV share that accounts for both sources of uncertainty is [2.0, 11.8] percent. Our point prediction of 5.6 percent, which only averages across fixed-cost draws, lies roughly at the median. Despite the uncertainty in the predictions, the central 95% interval of total surplus gains is [$3.6, $34.9] billion, which shows large gains even in the most pessimistic scenarios. We find that most of the uncertainty in the effect arises from uncertainty about Chinese-EV demand, rather than uncertainty about unobserved (to the researcher) fixed-cost realizations. Our results highlight that, even though economic models are informative and useful to study the entry of new goods in new markets, out-of-sample errors are quantitatively important.

Our analyses should be interpreted with several caveats. First, our model predictions depend on estimated substitution patterns. To ensure they are realistic, we allow random coefficients that capture the most important margins of substitution, and we identify the variance of random coefficients using second choice data from thousands of vehicle buyers. Second, although the vehicle buyer surveys over-sampled EV buyers, the subsample of Chinese EV buyers is still relatively small—only 1,075 people. However, the sample is large enough that the key diversion ratios are relatively statistically precise, and reweighting on observables does not change the key results, reducing concerns about sample selection bias. Third, we do not capture potentially important dynamic forces, such as learning-by-doing.

Our primary applied contribution is to evaluate the past and potential future effects of the expansion of Chinese EVs, one of the most important new products so far in the 21st century. In doing this, we make metholodogical contributions by developing a portable methodology to forecast the effects of introducing existing goods in new markets. Our work builds on several areas of related literature. First, our work

connects to a recent literature on trade protections affecting electric vehicles and, more generally, automobiles. This includes Allcott et al. (2024), Barwick et al. (2025), Castro-Vincenzi et al. (2025), Barwick et al. (2026a), Barwick et al. (2026b), Head et al. (2026), and Sabal (2026). Second, our work connects to a classic literature on the welfare effects of new goods. This includes a trade literature dating at least to Feenstra (1994) and Broda and Weinstein (2006), as well as an industrial organization literature dating at least to Hausman (1996), Petrin (2002), Nevo (2003), and Eizenberg (2014). Our methodological approach focuses more specifically on prospectively forecasting the effects of existing goods entering new markets, as opposed to retrospectively evaluating the effects of new goods.

Sections 2–7, respectively, present the data, descriptive facts, model, estimation, welfare effects of Chinese EV entry, and conclusion.

2 Data

We define "electric vehicles" (EVs) to include both battery electric vehicles (which use no gasoline) and plug-in hybrid vehicles. We define "gasoline vehicles" (GVs) to include both gas hybrids (without a plug) and gas-only vehicles. We refer to each country-year as a "market." Within each market, we define the choice set to be all light-duty EVs and GVs excluding vans and any models with price over $150,000.

We primarily use two types of data: market-level data and microdata. Our market-level data include quantities sold, prices, and characteristics for all new light-duty vehicle models sold in each market from 2015–2025. The characteristics we collect are footprint, horsepower, energy intensity (gallons per mile for GVs, and kWh per 100km for EVs), and body style (car, SUV, or pickup truck). Our sample of markets covers the largest 50 markets for LDVs in 2025, excluding 12 where data are not available.[2] Quantity data are from Marklines. Price and characteristic data from 24 countries (U.S., European countries except for Hungary, Japan, Malaysia, South Africa, and the UAE) are from Teoalida. For the remaining countries, we collected price and characteristic data through online searches. In total, we have 70,937 model-country-year observations. To our knowledge, this is the most comprehensive collection of market-level vehicle data in the academic literature.

To model realistic substitution patterns, we use survey microdata from vehicle buyers. While many previous papers have used second-choice surveys from the United States and China, we are not aware of second-choice surveys covering other countries that are available for research. Moreover, to understand the impacts of excluding all Chinese vehicles from the market, we were interested in a unique type of second-choice question to understand what consumers would have bought if all Chinese vehicles were not available. We thus needed to collect new data.

To do this, we partnered with a leading management consulting firm to add second-choice questions

[2]Table A3 reports our sample's country and year coverage.

to a vehicle owner survey that they delivered in January and February 2026. The survey includes 21,000 total responses from China, Germany, Japan, the U.K., and the United States, including an oversample of EV owners.[3] The survey includes household income and other demographics, as well as the brand, model, powertrain, vehicle age at purchase (new or used), and purchase year for each model owned. We added a standard second-choice question for the most recent vehicle purchase: *Imagine that the [brand & model of most recent purchase] were not available. What car would you have bought instead?* For Chinese brand owners, if their second choice was also a Chinese brand, we added an additional second-choice question: *Imagine that all Chinese brands were not available. What car would you have bought instead?*

For demand estimation, we define the "brand home" country as the country where the brand originated. For describing effects on profits, we define the "firm owner" country as the country where the owner is located. In our estimation samples, brand origins fall into Asia excluding China, China, Europe, and the United States. We classify a brand as Chinese if its brand home or its firm owner is in China, unless the brand has been sold in the United States at any point in our 2004–2025 U.S. sales data. For example, SAIC owns the MG brand, and Geely owns Volvo and Polestar. All three have their brand home in Europe and a Chinese firm owner, but only MG is a Chinese brand: Volvo and Polestar are already sold in the United States. When we refer to "Chinese EVs" in the descriptive facts and model below, we refer to EVs of Chinese brands.

We compute each vehicle model's lifetime uninternalized CO2 externality in each country, given the model-specific energy intensity. We assume that all vehicles are driven 150,000 miles over their lifetime. For plug-in hybrid vehicles, we assume that 63 percent of miles traveled are on the gasoline engine and 37 percent are on the electric motor, following Plötz et al. (2020). We are not aware of a reliable source of marginal CO2 emission rates from electricity generation for all countries in our sample. We approximate that using country-specific average CO2 emission rates from Our World in Data, sourced from Ember (Our World in Data 2026). We assume a social cost of carbon (SCC) of $253 per metric ton in 2025 dollars.[4] Each country's CO2 marginal damage is that SCC net of the applicable carbon price.[5] Vehicles impose other externalities, such as local air pollution and accidents, but we do not attempt to quantify these given that our sample includes many countries where marginal damages may vary.

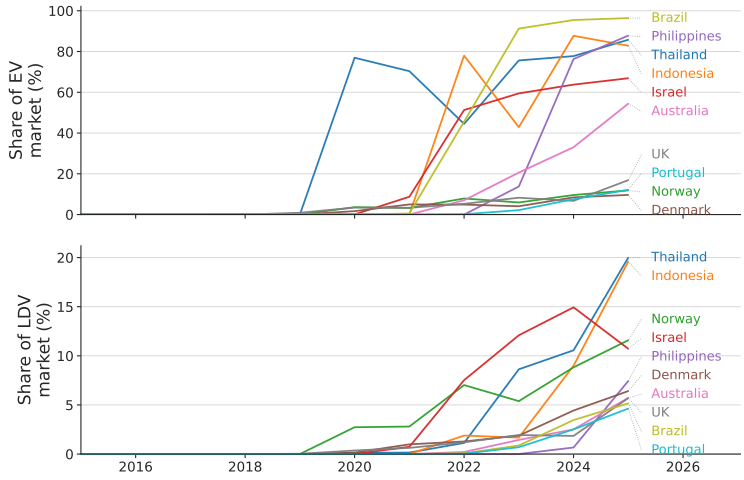
We convert all currencies to 2025 U.S. dollars. Appendix A presents summary statistics and more details on data preparation.

[3] Table A2 reports detailed summary statistics of the survey data.

[4] We convert the SCC assumption used in our EV-trade calculations from its July–August 2023 dollar basis using CPI-U. For the EPA's methodology for valuing CO2 emissions, see U.S. EPA (2023).

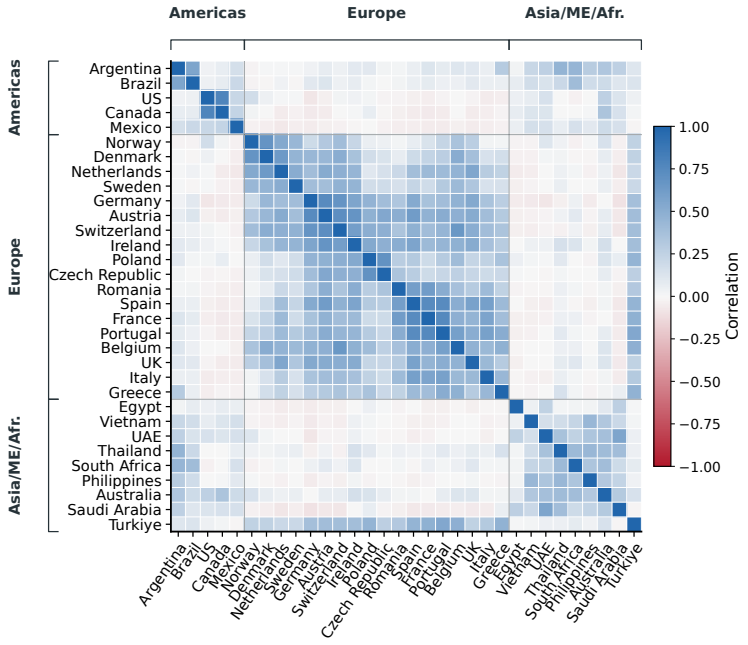
[5] We use historical carbon prices through 2025 from the World Bank's May 2026 release, in 2025 dollars, with separate coverage-weighted electricity and road-fuel prices (World Bank 2026). Welfare also includes manufacturing emissions; the outside option uses a gasoline-vehicle use-phase emissions proxy.

Figure 1: Chinese EV Market Share by Country



Notes: This figure presents the top 10 countries, excluding China, by 2025 Chinese EV share of light-duty vehicle sales.

Figure 2: Market Share Correlations Across Countries

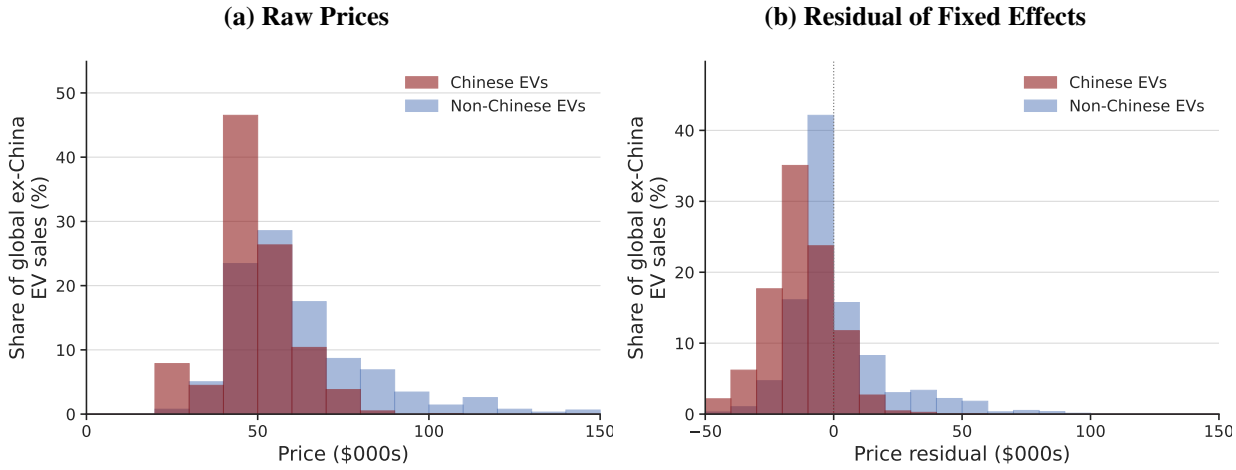


Notes: This figure shows the correlation in 2025 product market shares between each country pair in our sample. Products that are not available in a country are assigned zero market share.

3 Descriptive Facts

In this section, we present key descriptive facts that motivate our research question and methodological approach.

Figure 3: Chinese versus Non-Chinese EV Prices in 2025



Notes: This figure presents the distribution of prices for Chinese and non-Chinese EV models across all countries where Chinese EVs were available in 2025. Panel (b) is residual of country, powertrain, and body type fixed effects. Country-model observations are weighted by sales.

First, Chinese EV sales are growing rapidly worldwide. Appendix Figure A1 presents sales by brand over time. Between 2021 and 2025, Chinese EV sales outside of China rose from 45,000 to 1,005,000, now representing 15.0 percent of EV sales outside of China and 66.0 percent of EV sales globally.[6]

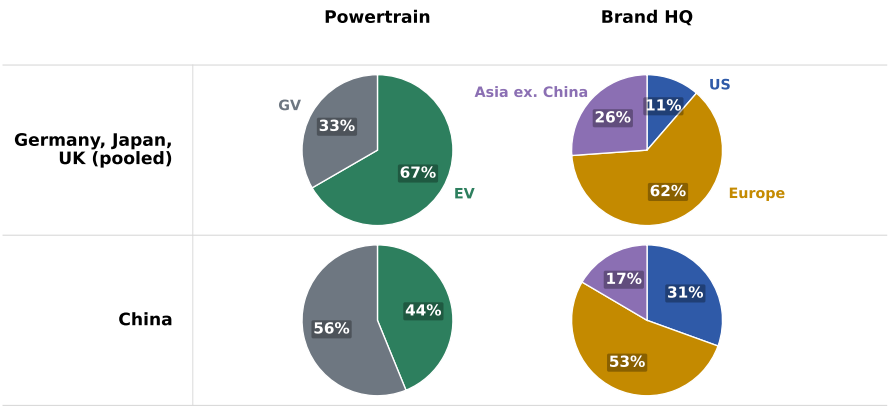
Second, the market shares of Chinese EV brands vary considerably across countries. Figure 1 presents Chinese EV market share by year, in the 10 countries where they comprise the highest share of LDV sales. The top panel presents Chinese EVs' share of EV sales in each country, while the bottom panel presents the share of the full LDV market. The top panel shows that in 2025 Chinese EV brands' EV market share ranged from 9.6 percent in Denmark to 96.4 percent in Brazil. This broad range is important because it suggests that there may be significant uncertainty when forecasting Chinese-EV sales in any additional country.

Third, some countries have more correlated market shares. Figure 2 presents the correlations of vehicle market shares between each country pair in our sample. In this figure, models that are not available in a country are assigned zero market share. Appendix Figure A3 presents correlations after coding unavailable models as missing; this shows that some but not all of the observed correlations are due to the extensive margin of product availability. These correlations are important because they suggest that predicting demand in one country could be aided by exploiting demand patterns in certain similar countries.

Fourth, Chinese EVs are relatively cheap. Figure 3 presents the distribution of prices for Chinese versus non-Chinese EVs across all countries where Chinese EVs were available in 2025, weighting country-model observations by sales. Panel (a) presents the raw prices, while Panel (b) presents prices residual of country,

[6] Appendix Figure A2 plots the evolution of EV sales by brand, including China.

Figure 4: Alternative Choices if Chinese EVs Were Not Available



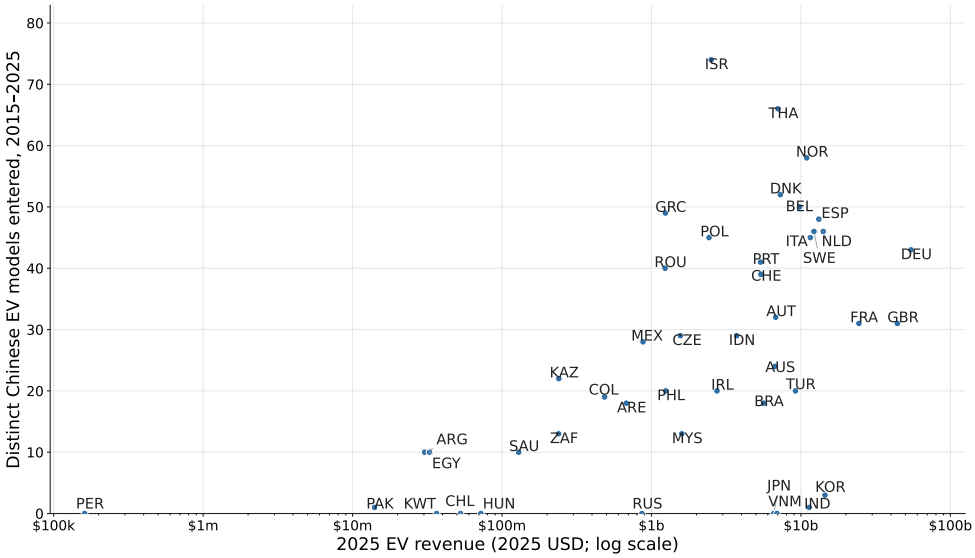
Notes: This figure restricts to consumers whose most recent vehicle purchase was a Chinese EV, defined as a Chinese brand with electric or plug-in hybrid powertrain. It is based on the follow-up survey question for Chinese-brand owners: *Imagine that all Chinese brands were not available. What car would you have bought instead?* Sample sizes for the China powertrain, China brand-HQ, pooled powertrain, and pooled brand-HQ pies are 438, 629, 81, and 88, respectively.

powertrain, and body type fixed effects. Chinese EVs are considerably cheaper than non-Chinese EVs; in Panel (a), the raw price means are $50,100 and $62,700 for Chinese and non-Chinese EVs, respectively. However, EVs in general are more expensive that other LDVs. These facts are important because they help determine the distributional effects of Chinese EV entry across income groups.

Fifth, our new kind of second choice data directly describe the effects of Chinese EV entry, without needing to resort to a demand model. Figure 4 presents the distributions of powertrains and brand homes that current Chinese EV owners reported they would buy if all Chinese brands were unavailable. Among Chinese EV owners in Germany, Japan, and the U.K., 33 percent reported that they would have bought a gasoline vehicle, and only 26 percent reported that they would have bought another Asian brand. Among Chinese EV owners in China, 56 percent reported that they would have bought a gasoline vehicle, and only 17 percent reported that they would have bought another Asian brand. These results are important because they suggest that the entry of Chinese EVs in these countries has (i) substantially increased EV sales overall and (ii) substantially cannibalized sales of American and especially European brands, not other Asian brands. Of course, these substitution patterns could be different for consumers in the United States and other countries, so a formal demand model remains useful.

Sixth, there is a larger count of models entered in larger country markets, consistent with fixed costs. Figure 5 shows that there is a clear positive association between total EV revenues earned and the number of distinct Chinese EV models offered in a country between 2015 and 2025.

Figure 5: Chinese EV Models Entered versus EV Market Size



Notes: Each point represents a country among the top 50 automotive markets, excluding China, Taiwan, the United States, and Canada. The vertical axis counts distinct Chinese-brand EV models with positive sales in any year during 2015–2025. EVs include battery electric and plug-in hybrid vehicles. The horizontal axis reports total 2025 EV sales revenue across all brands, in 2025 U.S. dollars on a logarithmic scale. Revenue is calculated as sales times prices. Country labels are ISO three-letter codes.

4 Structural Model

We use a fairly standard static model of differentiated product supply and demand in the spirit of Berry, Levinsohn, and Pakes (1995), applied simultaneously to many country-level markets. Our key methodological consideration is the need to predict demand in new markets. We first lay out the standard model, then introduce the prediction approach. Finally, we specify a static product entry model.

4.1 Demand and Supply

We index consumers by i, vehicle models by j, country markets by m, and years by t. Define p_{jmt} as the vehicle price, δ_{jmt} as a vertical demand shifter, and Z_{jmt} as the vector of vehicle characteristics with random coefficients: an inside-good indicator, an EV indicator, car and pickup indicators, and indicators for Chinese, European, and U.S. brand origin. SUV and Asia ex-China are the omitted body-style and origin categories. Define σ as a vector of random coefficient standard deviations, ν_i as a vector of standard normal random coefficients, and ε_{ijmt} as a logit error.

Define y_{imt} as income for consumer i in country m. We draw y_{imt} from a lognormal distribution that matches country m's mean income and Gini coefficient.

The utility function is

$$u_{ijmt} = \overbrace{\delta_{jmt}}^{\text{demand shifter}} - \overbrace{(\alpha/y_{imt})\, p_{jmt}}^{\text{price response}} + \overbrace{\sigma \cdot \nu_i \cdot Z_{jmt}}^{\text{random coefficients}} + \beta_y \log(y_{imt}) + \varepsilon_{ijmt}. \tag{1}$$

The random coefficient standard deviations are assumed to hold globally, since we do not have second-choice data from all countries. By contrast, the demand shifters δ_{jmt} vary by product, market, and year, exactly rationalizing observed demand. The coefficient β_y shifts the utility of all inside goods with log income.

Define $D_{jmt}(\boldsymbol{p}_{mt})$ as aggregate demand for vehicle j in market m in year t, as a function of the price vector $\boldsymbol{p}_{mt}$. In Appendix C.1, we provide the standard formulas for choice probabilities and market shares that arise from our specification.

4.2 Supply

Firms price vehicles in a static Nash-Bertrand equilibrium. We assume that firms do not consider dynamic incentives such as learning-by-doing. Each firm f offers a set of vehicles $\mathcal{J}_{fm}$ in country m at constant marginal costs c_{jmt}. Firm f's profit function in country m and year t is,

$$\pi_{fmt} = \sum_{j \in \mathcal{J}_{fm}} (p_{jmt} - c_{jmt}) \cdot D_{jmt}(\boldsymbol{p}_{mt}). \tag{2}$$

Following Berry, Levinsohn, and Pakes (1995), we infer marginal costs from each firm's first-order condition. We model log marginal costs as

$$\log c_{jmt} = \boldsymbol{W}_{jmt}\phi + \omega_{jmt}, \tag{3}$$

where $\boldsymbol{W}_{jmt}$ includes product, country, year, body-style, and powertrain fixed effects; log size and log horsepower; log fuel efficiency for gasoline and non-plug-in hybrid vehicles; log electricity consumption for battery electric vehicles; log electric range for plug-in hybrid vehicles; a home market indicator; and Chinese EV indicators interacted with calendar year. We estimate this projection by OLS.[7]

4.3 Predicting Demand in New Markets

Predicting the demand for goods in new markets, or for new goods, is challenging. As in Petrin (2002), we rely on micro-moments to pin down the distribution of heterogeneous consumer tastes, and to obtain accurate welfare estimates. In our setting, we are primarily interested in predicting the demand for a subset of goods (Chinese electric vehicles) in a set of markets (the United States) in which those goods have not

[7] We do not model ω_{jmt} as depending on an interactive fixed-effects component. Appendix C.5 provides further details about the projection and extrapolation of costs to potential entrants.

been offered. The standard approach in the literature (e.g., Petrin 2002; Wollmann 2018; Castro-Vincenzi et al. 2025; Fan and Yang 2025; Sabal 2026) has been to project δ_{jmt} on observables $\boldsymbol{X}$, and use the predicted values of such a regression model to obtain an estimate of $\mathbb{E}[\delta_{jmt}|\boldsymbol{X}]$ in a market (m,t) in which j is not offered. We argue that a limitation of this approach is that it does not account for residual correlations in preferences across markets, which one should expect given the correlations in market shares observed in Figure 2. We demonstrate that such correlations are quantitatively important and accounting for them greatly improves out-of-sample demand predictions.

Our approach is to explicitly model the residual using interactive fixed effects (IFE). Inspired by Bai (2009) and Moon, Shum, and Weidner (2018), we specify,

$$\delta_{jmt} = \boldsymbol{X}_{jmt}\theta + \xi_{jmt} \tag{4}$$

$$= \boldsymbol{X}_{jmt}\theta + \underbrace{\lambda_j' F_{mt}}_{\Gamma_{jmt}} + \zeta_{jmt}, \tag{5}$$

where $\mathbb{E}[\zeta_{jmt}|\boldsymbol{X},\Gamma] = 0$, and λ_j and F_{mt} are R-dimensional factor loadings and common market factors, respectively.[8] In our baseline specification, we include brand and country-year fixed effects; region-specific effects for (log) size, body type, and powertrain; and (log) horsepower, (log) fuel efficiency (for non-electric products), (log) kWh/100km (for electric products), (log) electric range (plug-in hybrid products), a home market dummy indicating whether the brand has local headquarters, and a Chinese EV trend term.[9] This specification generalizes a two-way fixed effects model by allowing for latent factors for which consumers have heterogeneous preferences across markets. Technically, these terms pick up correlations in ξ_{jmt} across markets. Indeed, $Cov(\xi_{jm\cdot}, \xi_{jm'\cdot}) \propto F_{m\cdot}' F_{m'\cdot}$, i.e., the alignment in the common factors captures the covariance in preferences across any pair of countries. This is in contrast to standard approach, where $\hat{\delta}_{jmt} = \boldsymbol{X}_{jmt}\hat{\theta}_{OLS}$, which only allows for correlations across countries on observable characteristics.

Intuitively, when predicting demand for Chinese electric vehicles in the United States, it makes sense not only to use estimated preferences for observed product characteristics (e.g., U.S. preference for the size of Chinese vehicles), but also to leverage the fact that we recover ξ_{jmt} for Chinese vehicles in some markets which might have highly correlated preferences to the United States. The IFE approach implements this intuition by imposing a factor (low-rank) structure on a component of ξ. Relative to Moon, Shum, and Weidner (2018), our contribution is to treat interactive fixed effects as the object of interest in a new-good prediction problem, rather than treating them as controls to deal with potential endogeneity in prices.[10]

An additional advantage of this approach is that by accounting for richer unobserved factors affecting demand, we reduce the concern that ξ_{jmt} is selected at entry. Under the assumed factor structure, we obtain

[8] In, Section 5, we discuss how R is chosen.

[9] We define 8 global regions: North America (ex Mexico), Latin America, Central Europe, Eastern Europe, Southern Europe, Northern Europe, Developed Asia-Pacific (Japan, Australia), and Emerging markets.

[10] Under our demand specification, prices coefficients only enter non-linearly.

consistent predictions provided firms cannot predict ζ_{jmt} at entry, but they are allowed to know $\lambda_j' F_{mt}$.

In contrast to Moon, Shum, and Weidner (2018), we have an unbalanced panel, with some products offered in some markets and not in others. Thus, in our setting the IFE will play a dual role: allow for rich unobserved factors that reduce selection concerns, while improving our imputation of demand in missing markets, as in Athey et al. (2021). In Section 5, we discuss the implementation and identification in detail.

4.4 Entry Game

Prior to the pricing stage, we assume that firms choose offerings according to a static complete-information pure-strategy Nash equilibrium in each country. To offer a product in a country, firms must incur a market offering fixed cost. As in Sabal (2026), we assume that product-market entry fixed costs follow a log-normal distribution, $\log FC_{jm} \sim_{iid} \mathcal{N}(\mu_{FC}, \sigma^2_{FC})$. Thus, each multi-product firm with potential products $\bar{\mathcal{J}}_f$ chooses which products to offer in country m by best-responding to rival firms' offerings in that country,

$$\mathcal{J}^*_{fm} = \arg \max_{\mathcal{J}_{fm} \subseteq \bar{\mathcal{J}}_f} \sum_{j \in \bar{\mathcal{J}}_f} \mathbb{1}\{j \in \mathcal{J}_{fm}\}[\mathbb{E}[\pi_{jm}(\mathcal{J}_{fm}, \mathcal{J}^*_{-f,m})] - FC_{jm}], \tag{6}$$

where $\mathbb{1}\{\cdot\}$ is the indicator function. At entry, firms have a symmetric information set. In particular, firms know the realization of fixed costs FC_{jm} for all potential products in all markets. Entry occurs prior to firms obtaining knowledge of the realization of demand and marginal cost shocks $(\zeta_{jm}, \omega_{jm})$. Hence, the expectations in equation (6) are with respect to the joint distribution of such shocks. We assume that firms know this distribution; in practice, we take draws from the empirical distribution of $(\zeta_{jm}, \omega_{jm})$ when solving the entry model. We describe further details about the implementation and the set of potential entrants in Section 6.

5 Model Estimation and Results

5.1 Estimation Details and Moments

We estimate demand by GMM following Berry, Levinsohn, and Pakes (1995, 2004) and Petrin (2002). We first estimate the nonlinear preference parameters, which we denote by γ, using the 2025 U.S., U.K., and German markets, where we observe second-choice survey data. For each candidate coefficient vector γ, we invert shares to recover mean utilities and, at our estimated $\widehat{\gamma}$, we solve for marginal costs from firms' pricing first-order conditions for products where we directly observe prices.

5.1.1 Moments and Survey Data

We match 8 moments to estimate the nonlinear parameters and discipline substitution patterns in our model.

Table 1: Nonlinear Demand Estimates

Coefficient	(1) Estimate	(2) Standard error
Price coefficient log-intercept	2.649	0.003
Inside good random coefficient	5.008	0.018
EV random coefficient	3.053	0.197
Car random coefficient	2.557	0.179
Pickup random coefficient	1.670	0.452
US origin random coefficient	1.949	0.230
Europe origin random coefficient	1.756	0.245
China origin random coefficient	1.541	0.277

Notes: The random-coefficient rows report elements of σ. All coefficients are common across countries. Standard errors are conditional local standard errors at the saved estimate; the two elasticity calibration targets are treated as fixed. Fixed parameters are omitted.

The 2 calibration moments inform price responses using existing estimates from the literature. We target a sales-weighted mean own-price elasticity of -5.36, following Grieco, Murry, and Yurukoglu (2024), and an elasticity of total new-vehicle demand to a common price change of -1.00, following Berry, Levinsohn, and Pakes (2004).

The remaining 6 moments come from our vehicle-buyer survey. Of these, 6 measure the second choice behavior of buyers: we match the probability that a buyer's second choice retains the first choice's EV powertrain, car or pickup body style, or U.S., European, or Chinese brand origin. The final moment is mean household income among new-vehicle buyers and governs the interaction between income and preference for the inside good.

The empirical moments are constructed from our survey data, pooling respondents from the United States, United Kingdom, and Germany. All second choice moments are constructed using 2025 new- and used-vehicle buyer responses (with valid first and second choices), except for the Chinese-origin moment, which uses all survey years because Chinese-brand first choices are uncommon in our data. The income moment uses 2025 new-vehicle buyers and converts income to 2025 GDP-PPP international dollars.

As detailed in Appendix C.2, we simulate each moment and stack the difference between the simulated values and their empirical counterparts into the vector $g(\gamma)$. We estimate γ by one-step GMM with identity weighting:

$$\widehat{\gamma} = \arg\min_{\gamma} g(\gamma)' g(\gamma). \tag{7}$$

Appendix C.4 describes the share inversion and numerical integration.

5.1.2 Nonlinear Estimates

Table 1 reports the nonlinear estimates. The table reports price sensitivity, the income interaction with inside-good utility, and random coefficients on vehicle characteristics and brand origin. All estimates except the inside-good intercept are statistically significant. We find substantial heterogeneity in preference for Chinese EVs, EVs, and different body types. We match the 8 empirical moments perfectly, as seen in Appendix Table A5.

Appendix Table A7 reports model price elasticities and compares Chinese-EV owners' counterfactual second choices with the survey. Overall, we match the diversion of Chinese EV owners to their second choices decently well both in Germany and the U.K. and in China. In the German and U.K. survey samples, all three model shares lie within about one survey standard error of the survey (for example, 64.4 versus 70.8 percent of Chinese EV owners switching to another EV), and in China the model's EV diversion is within 2.3 percentage points of the survey (48.0 versus 50.3 percent). The main discrepancy is that in China the model understates diversion to U.S. and European brands (27.1 versus 24.3 percent and 55.4 versus 42.1 percent), and so overstates diversion to other non-Chinese brands.

5.1.3 Mean Utilities and Marginal Costs

Holding $\widehat{\gamma}$ fixed, we recover mean utilities for the remaining products and marginal costs where the required price and ownership data are available. The following subsection uses these mean utilities to predict demand for products not observed in a market. Appendix A9 reports the marginal-cost estimates.

5.2 IFE and Matrix Completion

We jointly estimate θ and impute the missing entries in matrix $\Delta = (\delta_{jmt})_{j\in\mathcal{J},m,t}$ using a procedure akin to Mazumder, Hastie, and Tibshirani (2010) and Athey et al. (2021) applied to inverted mean utilities. Our estimator solves,

$$\min_{\Gamma,\theta} \frac{1}{2}||P_{\mathcal{O}}(\Delta - \boldsymbol{X}\theta - \Gamma)||_F^2 + \psi||\Gamma||_* \tag{8}$$

where $||\cdot||_F$ denotes the Frobenius norm, $||\cdot||_*$ denotes the nuclear norm, $\psi > 0$ is a rank penalty, and $P_{\mathcal{O}}$ is an operator that sets to 0 all (j, m, t) entries that are unobserved in our sample, i.e., not in the set $\mathcal{O}$ of observed entries. We also define $P_{\mathcal{O}}^{\perp}$ as the operator that sets to 0 all entries observed in our sample and selects only the unobserved entries. Note that this estimator nests OLS under $\psi = \infty$. If $\psi = 0$, then clearly $\theta = 0$ and $\Gamma_{jmt} = \delta_{jmt}$ are a solution but the model would have no predictive power out of sample. The estimator chooses Γ to improve model fit (the first term), while penalizing rank (the second term).

Appendix C.6 describes the soft-impute algorithm that solves this program, shows that the estimator is a penalized version of the interactive fixed effects model of Section 4, and discusses identification. Appendix

A9 reports the estimates of θ.

We follow Athey et al. (2021) and select ψ by cross-validation. We partition the observed product-market-year entries of Δ into $K = 4$ random sets; for each candidate ψ and each set, we re-estimate (θ, Γ) on the remaining entries and compute the root mean-squared error of the resulting predictions for the held-out entries. We select the ψ that minimizes this error, averaged across sets. The parameter ψ is then held fixed in all of our subsequent analysis.

5.2.1 Imputation Example

We provide an illustrative example of the imputation procedure. In reality, the common factors F and loadings Λ are estimated jointly, but the imputation procedure is most transparent when F and R are known and only Λ is estimated. In that case our estimator collapses to a product-by-product ridge regression,

$$\hat{\lambda}_j = \arg\min_{\lambda \in \mathbb{R}^R} \sum_{m \in \mathcal{M}_j} (r_{jm} - \lambda' F_m)^2 + \psi \|\lambda\|^2, \tag{9}$$

where $F_m \in \mathbb{R}^R$ is market m's common factor vector (we omit time subscripts for notational simplicity), $\mathcal{M}_j = \{m : (j, m) \in \mathcal{O}\}$ is the set of markets in which product j is offered, and $r_{jm} = \delta_{jm} - \boldsymbol{X}_{jm}\hat{\theta}$ is its demand residual, i.e., its over- or under-performance relative to observables. The main result is that the solution to the problem in equation (9) yields an imputation for product j in any market m^* where product j is unoffered that is equal to a weighted average of the product's *own* residuals in the markets where it is offered (a.k.a., anchor markets),

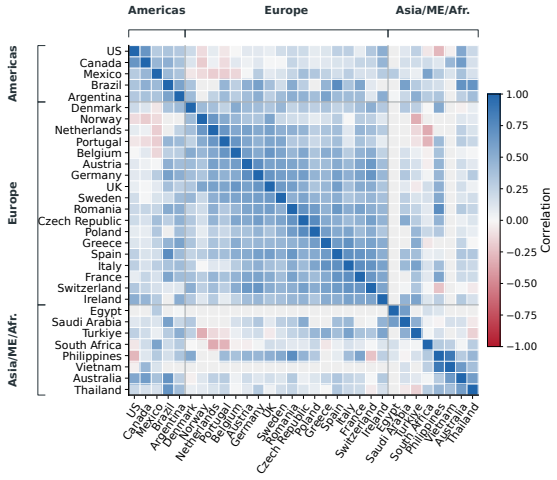
$$\hat{\Gamma}_{jm^*} = \hat{\lambda}_j' F_{m^*} = w' r_{j\cdot}, \qquad w = (F_{\mathcal{M}_j} F_{\mathcal{M}_j}' + \psi I)^{-1} F_{\mathcal{M}_j} F_{m^*}, \tag{10}$$

where $F_{\mathcal{M}_j}$ is the $|\mathcal{M}_j| \times R$ matrix that stacks the factor vectors F_m' of the $|\mathcal{M}_j|$ markets in $\mathcal{M}_j$, $r_{j\cdot}$ is the corresponding $|\mathcal{M}_j|$-vector of residuals r_{jm}, and w assigns one weight to each of these markets. A market receives weight to the extent that its latent tastes, as summarized by its factors, align with those of the target market. Each market's weight also depends on the presence of other markets, so that markets whose factor content is redundant (already informed by other markets in which the product is offered) are not given as much weight. The penalty ψ shrinks the weights toward zero, and the shrinkage loses importance as the product is observed in more markets.

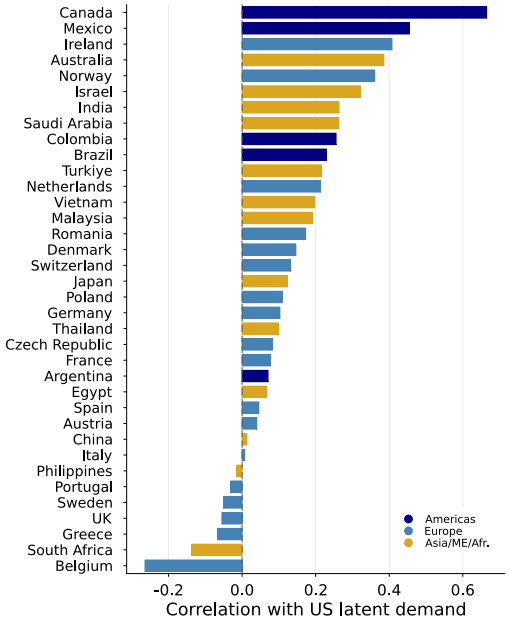
This simpler characterization of the imputation procedure makes both the virtues and the limitations of the approach transparent. Our approach transports latent tastes across countries through a formal regression that weights each anchor country by how its unobserved tastes align with the target country's, rather than through an arbitrary average. Of course, the method has limitations. First, for products offered in few markets, predictions are shrunk toward zero, so our estimates of Chinese EV appeal in the United States are conservative for products offered in few countries whose latent appeal is improving rapidly. Second,

Figure 6: Preference Correlation Across Countries

(a) Heat Map with All Countries



(b) United States Only



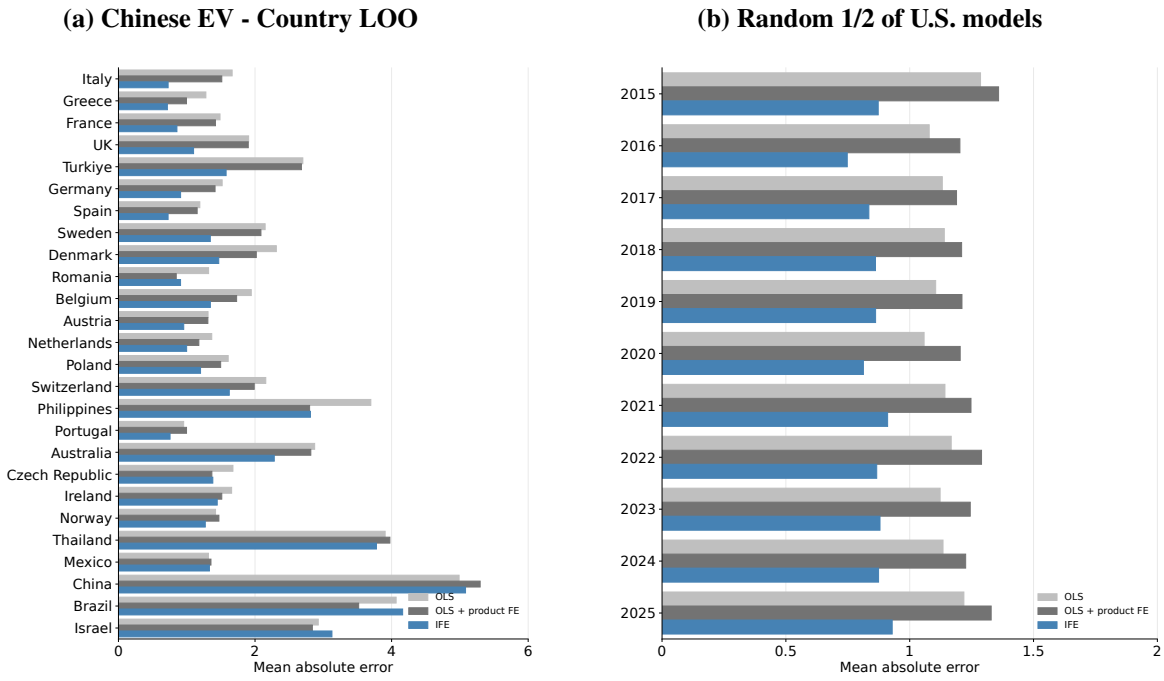
Notes: Panel (a) shows the correlation in 2025 ξ between each country pair in our sample, using products offered in both countries. Panel (b) shows each country's correlation with the United States in the latent-demand component Γ, averaged across sample years and reconstructed from the IFE fit.

there is a span restriction: the procedure can project appeal only along taste dimensions that are reflected in the markets where the product is observed. Any preference component invisible in every anchor market is imputed as zero, however large it may truly be. The method is therefore most reliable for products offered in many countries with heterogeneous tastes, and may underestimate appeal elsewhere.

5.3 Validation of the IFE Approach

Panel (a) of Figure 6 motivates the reason to use interactive fixed effects. Even after controlling for brand, country-year fixed effects, and observable product characteristics, there is residual correlation in preferences across markets. Importantly, the intensity of this correlation varies substantially across different pairs and groups of countries; for example, we observe greater similarity in preference within North American or European countries. The results in Figure 6 are the structural counterpart of those in Figure 2 and suggest that there are latent product characteristics for which there are heterogeneous preferences across countries. Panel (b) shows that the interactive fixed effects capture stronger similarity in preferences between the United States and countries that are both geographically close, such as Canada or Mexico, or economically similar,

Figure 7: Out-of-sample MAE Reduction



Notes: Both panels compare the mean absolute error (MAE) of out-of-sample predictions of product–market–year mean utilities δ_{jmt} under the observables-only projection (OLS), OLS with product fixed effects estimated from each product's training observations in other markets (OLS + product FE), and IFE. Panel (a) is a leave-one-country-out exercise for Chinese EVs: for each country in which at least one Chinese EV is sold, all Chinese-EV observations in that country in all years are removed, all three estimators are re-estimated, and the removed observations are predicted. The panel shows the countries with at least five held-out product–year observations, sorted by the MAE reduction from using IFE. Panel (b) is a random-half holdout for the United States: in each of 20 iterations, all U.S. observations of a random half of the products sold in the United States are removed, all three estimators are re-estimated, and the removed U.S. observations are predicted. Bars pool the iterations by year. The nuclear-norm penalty is fixed at the value selected by cross-validation on the full sample and is not re-tuned within the holdouts.

such as Ireland, Norway, Switzerland, or Australia. An econometric approach that uses this information to make out-of-sample demand predictions should intuitively outperform the common approach that disregards it, by properly weighting latent demand from similar countries more than from more dissimilar countries.

In this section, we conduct two exercises to validate our IFE approach. In the first exercise, we exclude the δ entries in some country in which Chinese EVs are offered, implement both our IFE approach (IFE) and OLS (with and without product FE) on the sample without such entries, and compare the predictions from the three approaches to the recovered Chinese-EV δ values by reporting the mean absolute error (MAE). This procedure allows us to assess the out-of-sample performance of IFE relative to the standard OLS approach by leveraging countries in which Chinese EVs are already offered in our sample. In the second exercise,

we randomly exclude half of U.S.-offered products in all years, implement IFE and OLS, and compare the MAE in their predictions. This exercise evaluates our procedure for predicting demand specifically in the United States.

Figure 7 compares out-of-sample MAEs under IFE and OLS with and without product FEs. Panel (a) reports the comparison by country for Chinese EVs. Three main takeaways arise. First, IFE dominates OLS in 22 out of the 26 reported countries and, in those countries in which it does not achieve a lower MAE, the methods' performance do not differ substantially.[11] As reported in Table A10 of the Appendix, the MAE of IFE is 1.47, while it is 1.82 for OLS with product FE and 1.95 for OLS. Second, all methods tend to perform poorly when preference for Chinese EVs tends to be unrelated to preference for other products in similar countries. For instance, preference for Chinese EVs in China is difficult to predict in China because of a large home bias that is impossible to infer without observing Chinese demand.[12] Finally, while product FEs tend to improve prediction of Chinese EV demand across many markets, the reduction in MAE relative to OLS is small in the majority of countries, particularly in Europe. This small improvement in performance appears because product FEs only help to the extent that a product's appeal in a target country matches its average appeal elsewhere.

Panel (b) of Figure 7 reports how the methods differ in their ability to predict U.S. preferences across years. Two main takeaways arise. First, IFE dominates OLS with or without product FEs, and achieves a substantial reduction in MAE. Second, and perhaps more strikingly, OLS with product FEs is the worst method for predicting U.S. demand. Table A10 reports a MAE of 0.86 for IFE, 1.25 for OLS with product FE and 1.14 for OLS, averaged across all years. To shed light on this finding, Panel B of Table A10 reports how the OLS residuals for products offered in the United States correlate with the residuals for those same products in other markets. The correlation between U.S. residuals and the average residuals for the same product (across years) across all other markets is only 0.30, while it is 0.48 if we weight countries by similarity to the United States (using the full-sample estimates of Γ), and 0.68 for products that are also available in Canada. This shows that product FEs, which capture information about average appeal across markets, do not adequately predict preferences in the United States, given that U.S. tastes do not coincide with those of the average market where U.S.-offered products are available. The higher correlation with countries like Canada further justifies our approach, which weights countries depending on how their preferences correlate with the United States. Finally, Panel C of Table A10 shows that the advantage of IFE and OLS over OLS with product FE is not driven by product FEs for products sold in only a few markets.

[11] In this analysis, we include 7 countries for which we do not have full price data coverage: Australia, Brazil, China, Israel, Mexico, Philippines, Thailand. For these countries, we borrow the price of each product from other countries in which that product is available. In Section 6, we drop these countries to avoid using unreliable marginal-cost estimates in counterfactual exercises.

[12] In addition to home biases and pure preferences, it is possible some of this error can be explained by unobserved policies in China or in countries like Israel that lead to a larger market share for Chinese EVs than in otherwise similar markets.

6 Predicted Effects of Chinese EV Entry

We study the global impact of Chinese-EV entry in two separate exercises. First, we explore the effect of removing Chinese EVs from countries in which they are available in 2025. We leverage this exercise to evaluate the out-of-sample properties of IFE relative to existing approaches, comparing the effects under the recovered mean utilities for such products to the welfare effects with held-out mean utilities. In the second exercise, we study the welfare effects of Chinese-EV entry in the United States in the year 2030. We define the change in global total surplus as the sum of the changes in consumer surplus, environmental externalities, and profits.[13]

6.1 Predicted and Actual Effects of Entry in Existing Countries

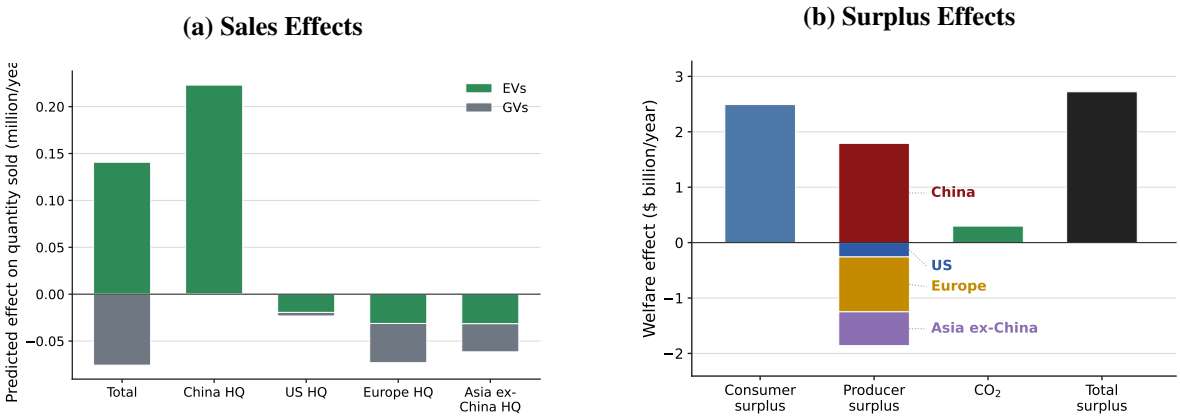
Figure 8 and Column 1 in Table 2 report the total impact of Chinese EVs on welfare outcomes in the 19 countries in our sample in which they are offered in 2025.[14] The availability of Chinese EVs causes the number of new electric vehicles sold to increase by 0.141 million, while the number of new vehicles sold only increases by 0.065 million due to business stealing. Chinese EV sales increase by 0.257 million, larger than the increase in EV sales, which is also due to business stealing between Chinese and non-Chinese EVs. Overall, Chinese EV entry in such countries is positive for welfare. Total surplus rises by $2.72 billion, predominantly due to an increase in consumer surplus of $2.49 billion, a decrease in the negative CO_2 externality, and the fact that the increase in Chinese producer surplus more than offsets the decrease in profits for other sellers. Overall, the main losers from Chinese EV entry globally are European brands, followed by non-Chinese Asian brands. American brands are the least exposed in those markets in which Chinese EVs are already offered in 2025.

To validate the IFE procedure, we compare how this procedure compares with existing demand extrapolation approaches for predicting the effects of Chinese EV entry in each of the countries in which we observe Chinese EVs. We hold out the Chinese EVs offered in each of these 19 countries and implement 4 distinct procedures to predict their mean utilities in the held-out countries: IFE, OLS, OLS with product FE, and a China anchor. IFE is our standard procedure, OLS uses the same observables as described in Section 5, and the China anchor implements OLS (without product FE) but, in the spirit of Barwick et al. (2026b), adds the OLS residual for the product in China in year t' (the closest year to

[13] This definition ignores changes in government revenue.

[14] A limitation of this exercise is that we study the effect of removing Chinese EVs from markets in which they are offered, without accounting for strategic entry and exit responses from non-Chinese incumbents. In this section, we restrict to the 19 countries for which we have matched price data, so as to obtain reliable marginal costs that allow us to perform counterfactual exercises. We also restrict the sample to countries with at least five held-out Chinese-EV observations, which amounts to excluding South Africa (where we only observe 1 Chinese EV). This restriction causes the leave-one-country-one holdout to yield slightly different results, as seen by comparing Table 2 and Table A10 in the Appendix.

Figure 8: Effects of Chinese EV Entry on Sales and Surplus in Existing Countries



Notes: Panel (a) reports the change in total units sold from not removing the Chinese EVs that are available in each of the 19 countries, and the decomposition by brand headquarters country. The green bars report the effects on EVs, and the gray bars on GVs. Panel (b) reports the the overall impact of Chinese-EV availability on each of the components of welfare.

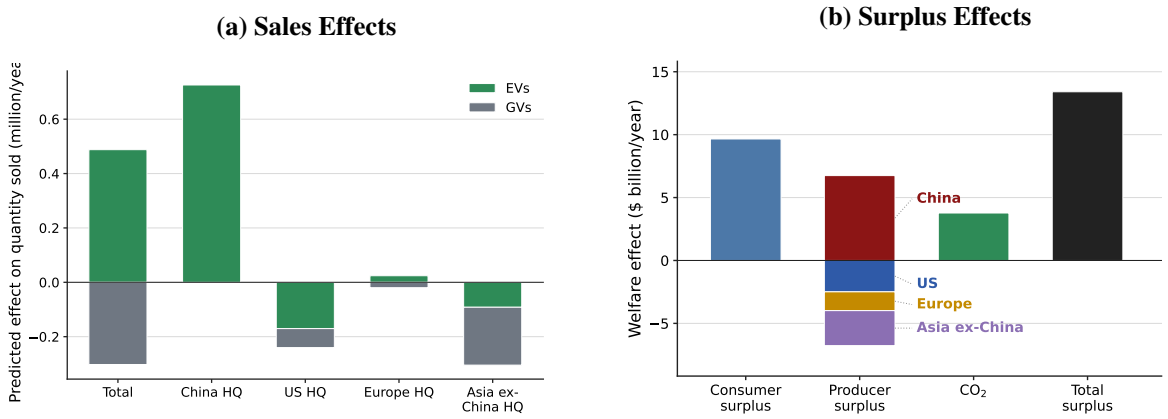
year t in which the product is offered in China), and the average difference in demand residuals across overlapping non-Chinese EV and plug-in hybrid products in China and the target country, $\mathcal{C}_{mt,t'}$, i.e., $\hat{\delta}^{\text{anchor}}_{jmt} = \boldsymbol{X}_{jmt}\hat{\theta} + \hat{\xi}_{j,\text{CN},t'} + \frac{1}{|\mathcal{C}_{mt,t'}|}\sum_{\ell \in \mathcal{C}_{mt,t'}}\left(\hat{\xi}_{\ell mt} - \hat{\xi}_{\ell,\text{CN},t'}\right)$.

Panel C of Table 2 reports how each of these methods' predictions compare to the actual effects that use the recovered mean utilities in each of the 19 countries. Overall, IFE and OLS with product FEs perform best for predicting the aggregate effects across all such countries. Columns 3 and 5 show that OLS and especially the China-anchor approach systematically over-predict the impact of Chinese EVs on new-vehicle sales, EV sales, consumer surplus, the CO_2 externality, and total surplus. Indeed, OLS and the China anchor predict that aggregate surplus rises by \$3.40 billion and \$6.49 billion, respectively, compared to the actual effect of \$2.72 billion. Meanwhile, IFE and OLS with product FEs make more accurate prediction of \$2.55 billion and \$2.99 billion, respectively. As discussed in Section 5, the dominance of OLS with product FE over OLS without product FE is a quantitative result that depends on the target countries, and may not hold for predicting demand in the United States.

Even though OLS with product FEs appears almost as successful as IFE for predicting the total global impact of Chinese EVs, we show that this result hides certain unappealing properties of product FE for predicting the effects in a single country. Panel C of Table 2 shows that predictions are both less biased and less noisy under IFE than under the other methods. Overall, IFE predictions are on average unbiased (0% bias), compared to 12% under OLS with product FE, 28% under OLS, and 164% under the China anchor.[15]

[15]Our bias measure is the log prediction error and weights all countries equally. An advantage of this bias measure is that over-predictions and under-predictions of the same percent magnitude cancel out. In levels, IFE appears to be downward biased because it under-predicts total surplus in larger countries.

Figure 9: Effects of Chinese EV Entry on Sales and Surplus in the United States



Notes: Panel (a) shows the predicted effects of Chinese EV entry in the United States, averaged across simulated fixed-cost draws in both the scenario without Chinese EV entry and in the scenario with Chinese EV entry. Panel (b) plots the expected effect of Chinese EV entry on different components of U.S. welfare.

The bias measure shows that IFE has the advantage of not systematically over- or under-predicting the effects of Chinese EVs across countries, while other methods tend to over-predict the benefits.

Moreover, across all welfare outcomes, we find that IFE's error has smaller spread, which measures the standard deviation of the (log) error. The spread in the total-surplus predictions for IFE is 0.58 , which means that a country's typical prediction of the total surplus effect is within a factor of $(1 + \text{bias})e^{\pm 0.58} = [0.56, 1.79]$ of the actual effect.[16] This shows that our approach, while achieving the lowest biases and spreads, is still subject to considerable prediction errors for each of the 19 countries.

Using information on actual Chinese-EV performance in 19 countries, we have shown that IFE achieves better out-of-sample demand predictions than alternative methods used in the literature. For predicting demand specifically in the United States, we showed that OLS with product FEs may actually perform even worse than OLS without product FEs, and the gains from using IFE might be even higher than in the countries in which Chinese EVs are already offered. Perhaps surprisingly, we find that using China as an anchor country is the worst-performing approach for predicting the global impact of Chinese EVs. This is likely due to systematic home bias and policy premia for Chinese products in China. Having studied the backward-looking global effects of Chinese EVs, we now exploit our IFE approach to predict the impact of potential Chinese-EV entry in the United States.

Table 2: Predicted Equilibrium Effects of Chinese EV Entry in Existing Countries

	(1) Actual	(2) IFE	(3) OLS	(4) OLS + product effects	(5) China anchor
Panel A: Mean Absolute Error of Held-Out Mean Utilities					
Chinese EVs, leave one country out (648 cells)		1.12	1.66	1.55	2.53
US models, random half (28,669 cells)		0.86	1.14	1.25	–
Panel B: Market Effects in 19 Countries					
Δ new-vehicle sales (million/year)	0.065	0.061	0.067	0.065	0.140
Δ EV sales (million/year)	0.141	0.130	0.149	0.144	0.279
Chinese EV sales (million/year)	0.257	0.239	0.276	0.265	0.514
Panel C: Welfare Effects in 19 Countries					
Δ consumer surplus ($ billion/year)	2.49	2.34	2.97	2.72	5.39
Bias across countries		+2%	+23%	+13%	+145%
Spread across countries		0.55	0.65	0.67	0.69
Δ producer surplus ($ billion/year)	-0.07	-0.11	-0.08	-0.10	0.04
Bias across countries†		+3%	0%	+2%	-11%
Spread across countries†		0.08	0.16	0.12	0.23
Δ CO_2 externality ($ billion/year)	-0.30	-0.32	-0.50	-0.37	-1.06
Bias across countries†		+3%	+13%	+5%	+40%
Spread across countries†		0.10	0.20	0.13	0.37
Δ total surplus ($ billion/year)	2.72	2.55	3.40	2.99	6.49
Bias across countries		0%	+28%	+12%	+164%
Spread across countries		0.58	0.69	0.70	0.74

Notes: Each method predicts the mean utilities of Chinese EV and plug-in-hybrid models in a country out of sample. Panel A reports mean absolute errors of held-out mean utilities when all Chinese EV observations in a country (pooled over the 19 countries in Panels B and C) or all U.S. observations of a random half of U.S. models (20 iterations) are removed before estimation. As in Panels B and C, the China anchor uses the OLS prediction for models with no observations in China. Panels B and C report the market and welfare effects, respectively, of the observed 2025 Chinese EV and plug-in-hybrid choice set relative to removing it, under recovered (Actual) or held-out predicted mean utilities. Dollar rows are 2025 USD billions per annual sales cohort. Outcome rows sum effects over the 19 countries. Indented rows in Panel C summarize each method's error in these countries, $e_m = \log(\text{predicted}_m/\text{actual}_m)$, weighting countries equally. Bias is defined as $\exp(\bar{e}) - 1$, the average proportional over- or under-prediction, and spread is the standard deviation of e_m. For producer surplus and the CO_2 externality (†), whose predicted and actual effects differ in sign in some countries, e_m is instead the prediction error as a share of the country's actual total-surplus effect, signed so that positive values overstate the effect, and bias is the mean of e_m.

6.2 Predicted Effects of Entry in the United States

We evaluate the impact of Chinese EV entry in the United States by solving a static entry game under 2030 demand conditions. First, we shift U.S. BEV mean utility by a shifter κ calibrated to match BloombergNEF's 2030 forecast of a 17% BEV inside share (BloombergNEF 2026). Second, we borrow product-market entry fixed-cost estimates from Sabal (2026) and assume that offering a product in the United States is subject to a fixed cost draws from a log-normal distribution with location -4.5 and scale 3.6 (in billion-dollar units). Third, we use the information encoded in 2025 observed entry decisions to draw counterfactual fixed-cost realizations from truncated log-normal distributions. For a product offered in the United States in 2025, we truncate its market entry fixed-cost from above at the expected change in variable profits from withdrawing that product from the United States (at the observed 2025 entry configuration): $\mathbb{E}[\pi_{fm}(\mathcal{J}_{fm}, \mathcal{J}_{-f,m})] - \mathbb{E}[\pi_{fm}(\mathcal{J}_{fm} \setminus \{j\}, \mathcal{J}_{-f,m})]$, where m is the United States.[17] Indeed, the yearly market offering fixed-cost must be smaller to rationalize their offering in 2025. We assume that if Chinese EV products were allowed to be offered in 2030 in the United States, their fixed costs would be drawn from the unconditional (non-truncated) distribution; that is, their 2025 non-entry is assumed to be purely policy-driven. Finally, we allow for additional potential entrants in 2030: EVs sold by brands with 2025 presence in the United States. In addition to the observed EVs offered by these brands in 2025 in countries other than the United States, we allow brands with presence in the United States to offer EV versions of GV body types that they already offer in 2025. Combining existing EVs with potential GV twins yields the set of potential non-Chinese entrants.[18] We truncate their U.S. entry fixed costs from below at $\mathbb{E}[\pi_{fm}(\mathcal{J}_{fm} \cup \{j\}, \mathcal{J}_{-f,m})] - \mathbb{E}[\pi_{fm}(\mathcal{J}_{fm}, \mathcal{J}_{-f,m})]$, to rationalize the fact that these products are not offered in the United States in 2025.[19]

[16] If the log-errors were approximately normal, around two thirds of countries' predictions should lie in this range. Indeed, we find that the predicted change in total surplus for 14 out of 19 countries lies within this range.

[17] The expectation is over realizations of demand and marginal-cost shocks for each offered product, which are assumed to be realized after entry. These shocks are drawn iid from the joint empirical distribution of $(\zeta_{jm}, \omega_{jm})$.

[18] We construct the set of potential non-Chinese entrants in two steps, restricting attention throughout to brands with a 2025 U.S. presence and to the retained body styles (cars, SUVs, and pickups). First, we include the BEV models that these brands offered in 2025 in markets other than the United States but not in the United States; their U.S. mean utilities are the IFE projections and their marginal costs come from the cost projection specified in Appendix C.5. Second, we include "GV twins": for each brand–body-style pair in which the brand offered a GV but no BEV in the 2025 U.S. market, we allow the brand to introduce an electric version of its best-selling gasoline model in that body style (one twin per brand–body-style pair). A twin inherits the donor model's observed characteristics with the powertrain switched to electric (with energy efficiency imputed from comparable BEVs), and its mean utility and marginal cost are again given by the IFE and cost projections. Chinese potential entrants are constructed separately as the Chinese BEV and PHEV models offered outside the United States in 2025.

[19] The set of fixed costs that arise from independently truncating each product's market-entry fixed cost may include values for which the equilibrium of the entry game does not coincide with the observed configuration in 2025, due to multiple equilibria and interdependent market entry decisions for multi-product firms. Still, this rule for selecting draws guarantees that we are drawing fixed costs from a set containing the region in which the observed configuration is an equilibrium.

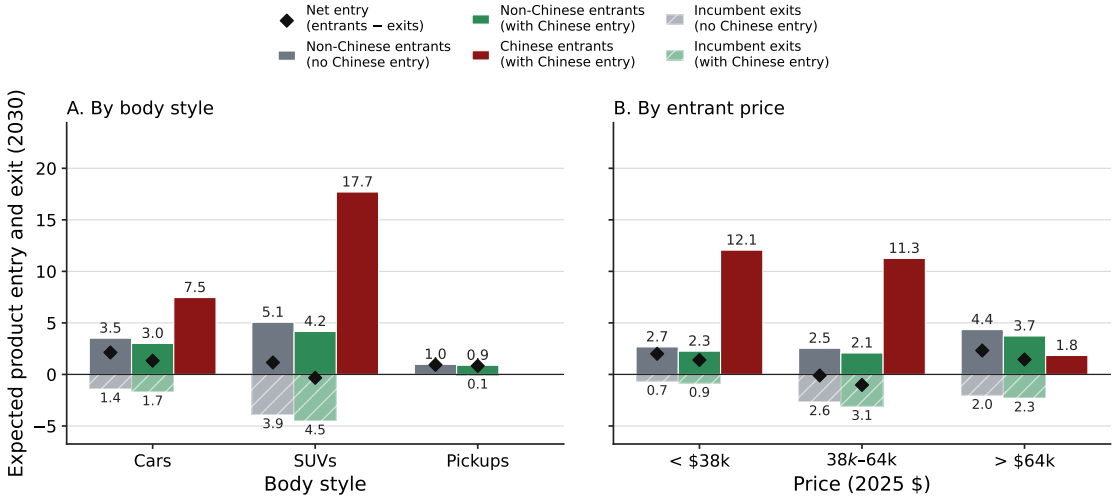
We solve the entry game using the algorithm in Castro-Vincenzi et al. (2026), which extends Jia (2008) to cases with multiple oligopolistic multi-product firms. The algorithm yields binary-decision-specific bounds on equilibrium entry decisions. In practice, we find that the equilibrium is unique under almost all of the fixed-cost realizations.[20]

Figure 9 plots the predicted effects of Chinese EV entry in the United States. Panel (a) reports the effect of Chinese-EV entry on total sales and sales by brand origin. We find that overall, total EV sales rise by 0.49 million units, while GV sales fall by 0.30 million units. EV sales increase exclusively due to the increase in Chinese EV sales, which reach a predicted 0.73 million units. Chinese-EV entry therefore leads to a decline in both the EV and GV sales of non-Chinese brands. Brands with U.S. and European headquarters contract EV sales the most, while non-Chinese Asian brands predominantly reduce their GV sales. Panel (b) plots the welfare effects of Chinese EV entry. Total surplus increases by \$13.41 billion, driven primarily by a \$9.66 billion increase in consumer surplus and a reduction in 18.65 million tons of CO_2 emissions worth \$3.77 billion. The main losers from Chinese EV entry are non-Chinese brands, whose profits fall by \$6.77 billion. Non-Chinese Asian brands see larger losses than American brands, who lose \$2.49 billion in profits, due to greater substitution among Asian brands than between American and Chinese brands.

Figure 10 reports the heterogeneous effect of Chinese-EV entry on different segments of the United States vehicle market. By 2030, increased EV demand in the United States would lead to an increase in EV products offered, even in the absence of Chinese EVs. The black diamonds over the gray bars in Panels A and B show that without Chinese entry, Americans would be offered on net 2.1 more Cars, 1.2 more SUVs, 2.0 more vehicles priced below \$38 thousand, and 2.3 more expensive vehicles, with small effects both for mid-priced vehicles and pickups. With Chinese EV entry, some of this increase in incumbents' product offerings is crowded out: incumbents contract their SUV product offerings relative to 2025 by 0.3, owing to large substitution towards Chinese electric SUVs, of which 17.7 are expected to enter. With Chinese EV entry, U.S. incumbents also reduce their net offering of mid-priced vehicles and Cars. Such heterogeneous crowd-out effects reflect the rich substitution patterns across passenger vehicles and the higher competitiveness of Chinese EVs in segments like SUVs. The results from Figure 10 also suggest that while crowd-out is potentially important, most products that are predicted to exit by 2030 would exit independently of whether Chinese vehicles enter or not due to the increased preference for electric vehicles and therefore the reduced profitability of GVs by 2030.

[20] We find the equilibrium to be unique in 99/100 of the fixed-cost samples with Chinese entry and in 99/100 of the fixed-cost samples without Chinese entry. In these very few cases in which the equilibria are not unique, we select one of the equilbria.

Figure 10: Crowd-Out and Exit in 2030 under Chinese EV Entry



Notes: Expected changes in the set of models offered in the United States by 2030, averaged over simulated fixed-cost draws. Bars above zero count entry by products not offered in 2025: non-Chinese potential entrants without (grey) and with (green) Chinese entry, and Chinese entrants (dark red). Hatched bars below zero count exits among the incumbent models observed in 2025, without (grey) and with (green) Chinese entry. Black diamonds report net non-Chinese product change (entrants minus incumbent exits) within each scenario. Panel B bins products by the projected U.S. entry price for potential entrants and the observed 2025 price for incumbents, at common cutpoints.

6.2.1 Distributional Effects Across Geography and Income

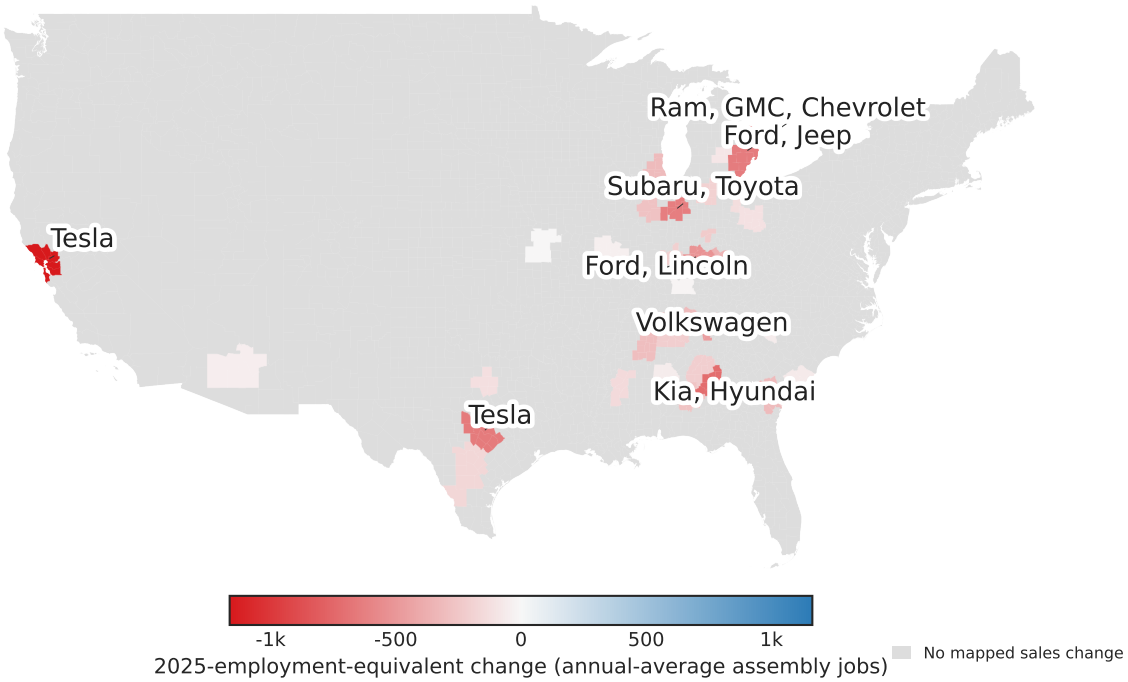
Figure 11 maps the production and employment losses from Chinese EV entry. We assign each model's predicted change in U.S. sales to the United States plants that assemble it and convert units units produced into jobs using the national ratio of assembly employment to vehicles assembled. Appendix Figure A5 shows the corresponding change in vehicles produced. Chinese EV entry reduces sales of U.S.-assembled models by 377 thousand vehicles, or 3.75 percent, which corresponds to 9,449 assembly jobs. Applying the same proportional decline to parts suppliers adds 19,555 jobs, for a total of 29.0 thousand auto-manufacturing jobs. No commuting zone gains employment, but the losses are not concentrated in a single region. They are spread across the 34 commuting zones with affected plants, and the five largest (Tesla's plants in the San Francisco Bay Area and Austin, Detroit, Kia's plant in West Point, Georgia, and Ford's plants in Louisville) account for 40 percent of the losses. By region, 56 percent of the assembly job losses fall in the South, 31 percent in the Midwest, and 13 percent in California and Arizona. Michigan and Georgia, the only states with affected plants in which the 2024 presidential vote-share margin was below five percentage points, account for 19 percent.

In proportional terms, losses are largest at plants that build EVs, consistent with Chinese EVs being closer substitutes for other EVs than for GVs. Production falls by 16.2 percent at Hyundai's EV plant near

Savannah, by 11.7 percent at the Volvo and Polestar plant near Charleston, and by between 6.7 and 7.8 percent at the Tesla, Rivian, and Lucid plants, compared with 1.8 percent in Detroit, whose plants mainly assemble gasoline pickups and SUVs.

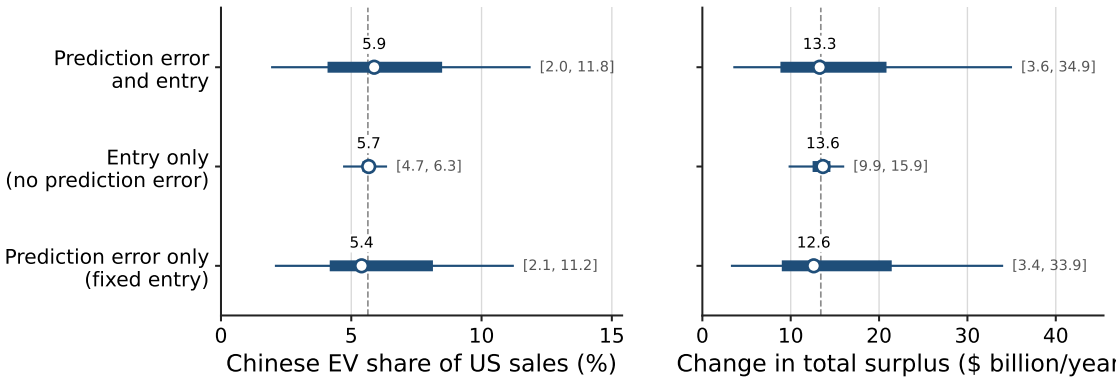
Appendix Figure A6 shows how the $9.66 billion gain in consumer surplus is distributed across household incomes. Households earning more than $100,000 receive 70 percent of the gains, while households earning less than $50,000 receive only 3.5 percent. The gains are tilted toward the top even relative to who buys new vehicles. Households earning more than $150,000 are 33 percent of new-vehicle buyers but receive 44 percent of the gains, and those earning $400,000 or more are 2.7 percent of buyers but receive 5.0 percent. Part of this tilt reflects the price coefficient α/y_{imt}: a given gain in utility is worth more dollars to a higher-income household.

Figure 11: Assembly Employment Changes under Chinese EV Entry in 2030



Notes: Predicted change in annual assembly employment by commuting zone from Chinese EV entry in 2030, averaged over simulated fixed-cost draws. Each model's predicted change in U.S. sales is assigned to the United States plants that assemble it and converted into jobs at the national ratio of assembly employment to vehicles assembled, in 2025 employment equivalents. Labels name the brands with the largest losses in the seven commuting zones with the largest losses. Parts employment, which we scale nationally, is not mapped.

Figure 12: Uncertainty in Predicted Effects of Chinese EV Entry in the United States



Notes: Effects of Chinese EV entry in 2030 on the Chinese EV share of U.S. new-vehicle sales (left) and on total surplus (right), relative to the equilibrium without Chinese EVs. Each row reports, across realizations, the median (dot), the 25th–75th percentiles (thick bar), and the 2.5th–97.5th percentiles (thin bar, endpoints in brackets). Each realization averages 50 draws of the demand and marginal-cost shocks. "Entry only" varies the entry configuration across the 100 fixed-cost draws of the main specification, without prediction error, with its mean marked by the dashed line. The other two rows add an out-of-sample prediction error η to each Chinese EV's U.S. mean utility, drawn from the procedure described in the main text. "Prediction error and entry" re-solves the entry game for each of 100 draws of η, with five fixed-cost draws each (500 realizations). "Prediction error only" holds entry fixed at the configuration whose total-surplus effect without prediction error is the median across configurations (22 Chinese entrants) and varies η across 100 draws.

6.2.2 Accounting for Prediction Uncertainty

The results discussed so far use our model's best prediction of Chinese EVs' appeal in the United States. However, the validation exercises we conducted and the results in Table 2 suggest that, even though less biased and more precise than commonly used procedures, our predictions are still uncertain.

In Figure 12, we report the variability in predicted U.S. Chinese-EV share and change in total surplus once we take into account the out-of-sample error. The dashed lines in both figures plot our baseline results, which use mean effect across 100 fixed-cost realizations, setting the mean-utility prediction errors to 0, i.e., $\eta_{j,US} = 0$ for each Chinese EV in the United States. We draw a realization of η for the potential Chinese-EV entrants as follows. First, we randomly select a source country from our leave-one-country-out holdout exercise in Table 2, among the countries with at least five held-out Chinese EVs. Then, within that country and for each of the Chinese-EV entrants, we randomly pick the out-of-sample error from one of the three products in the source country that is closest to its predicted appeal.[21] Under each such draw, we solve the entry game with Chinese product j in the United States having mean utility shifted by $\eta_{j,US}$.[22] The first row

[21] This procedure allows us to account for the fact that prediction errors are correlated with appeal.

[22] Here, we still draw demand and marginal cost shocks from the empirical in-sample distribution of (ζ, ω), which we assume firms do not know at entry. Importantly, we assume that firms know η at entry, so that we capture uncertainty from the researcher's perspective.

in Figure 12 solves the entry game under 500 joint realizations of fixed-cost shocks and prediction errors. The second row sets prediction errors to zero and accounts only for uncertainty stemming from unobserved fixed-cost shocks. The third row holds entry fixed at the configuration whose total-surplus effect without prediction error is the median across configurations and only varies the prediction error.

As expected, out-of-sample errors are substantial and, interestingly, they vastly dominate the uncertainty that arises from fixed-cost shock draws. Accounting for both out-of-sample and fixed-cost uncertainty, the median Chinese EV share and the median change in total surplus are predicted to be 5.9 percent and $13.3 billion, respectively. These values align with the mean effect that we obtain absent prediction errors. However, the lower and upper quartiles of the predicted Chinese EV share accounting for both sources of uncertainty are 4.1 and 8.5, respectively. Setting prediction errors to zero, we would have instead obtained a predicted Chinese-EV lower and upper quartiles of 5.4 and 5.9 percent, respectively. The 95% interval accounting for both sources of uncertainty is [2.0, 11.8] percent. Thus, the Chinese-EV share varies significantly across scenarios. In terms of total surplus, we find that even after accounting for both errors, Chinese EVs in the United States would be expected to generate a total surplus of $3.6 billion in the most pessimistic scenario (2.5th percentile) and up to $34.9 billion in the most optimistic scenario (97.5th percentile).

These results demonstrate that out-of-sample uncertainty is quantitatively significant, and should be taken into account when using economic models to evaluate the impact of goods that have never entered a market. We find that this source of uncertainty is more significant in the context of Chinese-EV entry than the researcher's uncertainty about the true fixed entry costs.

7 Conclusion

We develop a framework for predicting demand for existing goods in new markets that weights the markets where a product is already sold by how closely their preferences correlate with those of the target market. Out of sample, it predicts demand more precisely and with less bias than standard approaches based on observable characteristics. Applied to Chinese EVs, it predicts that allowing them into the United States market would generate large gains for consumers and the environment, raising total surplus by $13.41 billion in 2030, at the cost of concentrated losses for American automakers and their workers. These predictions are uncertain, and most of the uncertainty comes from how much Americans would value Chinese EVs rather than from which Chinese EV products would enter. Still, the gains remain positive throughout the 95% interval. Our results show that forecasts about entry effects should account for how preferences correlate across markets, and out-of-sample uncertainty is quantitatively important in counterfactual predictions.

References

Allcott, Hunt, Reigner Kane, Maximilian Maydanchik, Joseph Shapiro, and Felix Tintelnot. 2024. "The Effects of "Buy American": Electric Vehicles and the Inflation Reduction Act." NBER Working Paper 33032, URL: https://www.nber.org/papers/w33032.

Athey, Susan, Mohsen Bayati, Nikolay Doudchenko, Guido Imbens, and Khashayar Khosravi. 2021. "Matrix Completion Methods for Causal Panel Data Models." *Journal of the American Statistical Association* 116 (536): pp. 1716–1730.

Bai, Jushan. 2009. "Panel Data Models with Interactive Fixed Effects." *Econometrica* 77 (4): pp. 1229–1279.

Barwick, Panle Jia, Jack Collison, Pinelopi Goldberg, Shanjun Li, and Yucheng Wang. 2026. "From Trade War to Green Transition: Optimal Electric Vehicle Tariffs with Revenue-Funded Subsidies." NBER Working Paper 35334, URL: https://www.nber.org/papers/w35334.

Barwick, Panle Jia, Hyuk-Soo Kwon, Shanjun Li, Yucheng Wang, and Nahim Bin Zahur. 2026. "Industrial Policies and Innovation: Evidence from the Global Automobile Industry." *International Journal of Industrial Organization* 104: pp. 103230.

Barwick, Panle Jia, Hyuk-Soo Kwon, Shanjun Li, and Nahim Bin Zahur. 2025. "Drive Down the Cost: Learning by Doing and Government Policies in the Global EV Battery Industry." NBER Working Paper 33378, URL: https://www.nber.org/papers/w33378.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (4): pp. 841–890.

Berry, Steven, James Levinsohn, and Ariel Pakes. 2004. "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market." *Journal of Political Economy* 112 (1): pp. 68–105.

Berry, Steven T.. 1994. "Estimating Discrete-Choice Models of Product Differentiation." *The RAND Journal of Economics* 25 (2): pp. 242–262.

BloombergNEF. 2026. "Electric Vehicle Outlook 2026." URL: https://about.bnef.com/insights/clean-transport/bloombergnefs-electric-vehicle-outlook-2026-global-ev-sales-set-for-another-record-breaking-year-but-growth-in-some-major-markets-slows/.

Broda, Christian and David Weinstein. 2006. "Globalization and the Gains from Variety." *Quarterly Journal of Economics* 121 (2): pp. 541–585.

Castro-Vincenzi, Juan Manuel, Eugenia Menaguale, Eduardo Morales, and Alejandro Sabal. 2025. "Market Entry and Plant Location in Multiproduct Firms." Working paper.

Castro-Vincenzi, Juanma, Jin Liu, Eduardo Morales, and Alejandro Sabal. 2026. "Tariffs and Production Relocation in a Multiproduct Oligopoly." Working paper, version dated June 26, 2026, URL: https://www.sabalalejandro.com/.

Coşar, Kerem, Paul Grieco, Shengyu Li, and Felix Tintelnot. 2017. "What Drives Home Market Advantage?." *Journal of International Economics* 110.

Eizenberg, Alon. 2014. "Upstream Innovation and Product Variety in the U.S. Home PC Market." *Review of Economic Studies* 81 (3): pp. 1003–1045.

Fan, Ying and Chenyu Yang. 2025. "Estimating Discrete Games with Many Firms and Many Decisions: An Application to Merger and Product Variety." *Journal of Political Economy* 133 (6): pp. 1886–1931.

Feenstra, Robert. 1994. "New Product Varieties and the Measurement of International Prices." *American Economic Review* 84 (1): pp. 157–177.

Grieco, Paul, Charles Murry, and Ali Yurukoglu. 2024. "The Evolution of Market Power in the U.S. Automobile Industry." *Quarterly Journal of Economics* 139 (2): pp. 1201–1253.

Hausman, Jerry. 1996. "Valuation of New Goods under Perfect and Imperfect Competition." In *The Economics of New Goods*, edited by Timothy Bresnahan and Robert Gordon, pp. 207–248. Chicago: University of Chicago Press.

Head, Keith, Thierry Mayer, Marc Melitz, and Chenying Yang. 2026. "Industrial Policies for Multi-Stage Production: The Battle for Battery-Powered Vehicles." NBER Working Paper 34884, URL: https://www.nber.org/papers/w34884.

Jia, Panle. 2008. "What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry." *Econometrica* 76 (6): pp. 1263–1316.

Mazumder, Rahul, Trevor Hastie, and Robert Tibshirani. 2010. "Spectral Regularization Algorithms for Learning Large Incomplete Matrices." *Journal of Machine Learning Research* 11 (80): pp. 2287–2322.

Moon, Hyungsik Roger, Matthew Shum, and Martin Weidner. 2018. "Estimation of Random Coefficients Logit Demand Models with Interactive Fixed Effects." *Journal of Econometrics* 206 (2): pp. 613–644.

Nevo, Aviv. 2003. "New Products, Quality Changes, and Welfare Measures Computed from Estimated Demand Systems." *Review of Economics and Statistics* 85 (2): pp. 266–275.

Our World in Data. 2026. "Carbon Intensity of Electricity." Data from Ember, URL: https://ourworldindata.org/grapher/carbon-intensity-electricity.

Petrin, Amil. 2002. "Quantifying the Benefits of New Products: The Case of the Minivan." *Journal of Political Economy* 110 (4): pp. 705–729.

Plötz, Patrick, Cornelius Moll, Georg Bieker, Peter Mock, and Yaoming Li. 2020. "Real-World Usage of Plug-in Hybrid Electric Vehicles: Fuel Consumption, Electric Driving, and CO2 Emissions." ICCT

White Paper, URL: https://theicct.org/publication/real-world-usage-of-plug-in-hybrid-electric-vehicles-fuel-consumption-electric-driving-and-co2-emissions/.

Sabal, Alejandro. 2026. "Product Entry in the Global Automobile Industry." Working paper, version dated June 26, 2026, URL: https://www.sabalalejandro.com/.

The Economist. 2024. "America's 100% Tariffs on Chinese EVs: Bad Policy, Worse Leadership." *The Economist*, May 15, URL: https://www.economist.com/leaders/2024/05/15/americas-100-tariffs-on-chinese-evs-bad-policy-worse-leadership.

U.S. Environmental Protection Agency. 2023. "Report on the Social Cost of Greenhouse Gases: Estimates Incorporating Recent Scientific Advances." U.S. Environmental Protection Agency Report, URL: https://www.epa.gov/system/files/documents/2023-12/epa_scghg_2023_report_final.pdf.

Varadhan, Ravi and Christophe Roland. 2008. "Simple and Globally Convergent Methods for Accelerating the Convergence of Any EM Algorithm." *Scandinavian Journal of Statistics* 35 (2): pp. 335–353.

Wagner, Gernot and Conor Walsh. 2024. "Don't Slam the Door on Inexpensive Chinese Electric Vehicles." *The New York Times*, May 15, URL: https://www.nytimes.com/2024/05/15/opinion/biden-china-trade-electric-vehicles-solar-panels-tariffs.html.

Wagner, Gernot and Shang-Jin Wei. 2024. "The Right Response to China's Electric-Vehicle Subsidies." *Project Syndicate*, April 5, URL: https://www.project-syndicate.org/commentary/western-tariffs-wrong-response-to-chinese-ev-subsidies-by-gernot-wagner-and-shang-jin-wei-2024-04.

Washington Post Editorial Board. 2026. "Ban Chinese Cars? They're Already Here." *The Washington Post*, September 25, URL: https://www.washingtonpost.com/opinions/2026/09/25/china-already-sells-cars-us-ban-would-hurt-american-buyers/.

Wollmann, Thomas. 2018. "Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles." *American Economic Review* 108 (6): pp. 1364–1406.

World Bank Group. 2026. "State and Trends of Carbon Pricing Dashboard." May 2026 data release, URL: https://carbonpricingdashboard.worldbank.org/about-us.

Online Appendix

Predicting the Effects of New Goods, with an Application to Chinese Electric Vehicles in the United States

Hunt Allcott, Luke Heeney, and Alejandro Sabal

Table of Contents

A Data Appendix **2**

A.1 Data Sources and Sample Coverage . . . 3
A.2 Vehicle Definitions and Sample Construction . . . 3
A.3 Prices, Market Size, and Income . . . 3
A.4 Survey Design and Second Choices . . . 5
A.5 Summary Statistics and Survey Weights . . . 5
A.6 Externality Inputs . . . 6

B Descriptive Facts Appendix **7**

C Model Appendix **9**

C.1 Demand Specification . . . 9
C.2 Estimator and Moment Conditions . . . 9
C.3 Additional Estimation Results . . . 10
C.4 Computation and Demand Validation . . . 12
C.5 Marginal Costs and Cost Prediction . . . 14
C.6 IFE and Validation . . . 15

D Counterfactual Simulations Appendix **19**

A Data Appendix

Table A1: Market-Level Data Summary Statistics

Statistic	(1) Mean	(2) SD	(3) Min	(4) Max	(5) N
Price ($000s)	52.7	27.9	4.1	149.9	70,937
Quantity (000s)	7.6	29.3	0.0	909.3	70,937
Market share (%)	0.109	0.220	0.001	9.356	70,937
EV market share (%)	9.0	16.2	0.0	98.3	391

Notes: Price, quantity, and product market share are summarized across product–market–year observations in the expanded-relaxed estimation panel. EV market share is summarized across country–year markets. Prices are in thousands of 2025 U.S. dollars and quantities are in thousands of vehicles.

Table A2: Vehicle Owner Survey Microdata Summary Statistics

Country	(1) Responses	(2) Valid 2nd	(3) Valid 2nd %	(4) 2nd chars. %	(5) First EV %	(6) First Chinese %
USA	4,196	3,588	85.5	93.6	19.1	0.0
UK	4,073	3,354	82.3	88.6	22.9	3.0
Germany	3,660	2,621	71.6	86.2	26.9	1.2
China	3,264	2,526	77.4	47.8	41.8	34.6
Japan	5,439	4,158	76.4	46.5	13.8	0.3

Notes: "Valid 2nd choice" reports the share of responses with a recorded second choice. "2nd chars." reports the share of valid second choices with matched or decoded vehicle characteristics.

A.1 Data Sources and Sample Coverage

Our market-level data combine annual vehicle sales from MarkLines with prices and characteristics from Teoalida and supplementary online sources. The sales data cover country, year, parent firm, brand, model, body type, and powertrain. The characteristics data report prices and physical attributes by trim.

The descriptive sales sample covers the largest 50 national markets by 2025 sales. We estimate the nonlinear demand parameters using 1,020 product–country observations from the United States, United Kingdom, and Germany in 2025. We then recover mean utilities for 70,937 product–country–year observations in 38 countries over 2015–2025. The sample for marginal-cost recovery contains 24 countries and 51,967 observations before the restrictions described in Appendix C.5. Table A3 reports coverage by country.

A.2 Vehicle Definitions and Sample Construction

We define a product as a brand–model–body-type–powertrain combination. Body types are cars, SUVs, and pickups. Powertrains are internal combustion, conventional hybrid, plug-in hybrid, and battery electric. EVs comprise plug-in hybrids and battery electric vehicles; GVs comprise internal combustion vehicles and conventional hybrids. We exclude vans, buses, and commercial trucks to limit commercial vehicles in our sample. Fuel-cell vehicles are dropped.

We standardize brand and model names across sources and take medians across trims for continuous characteristics and prices within each product and year. We match sales to characteristics by brand, model, body type, and powertrain. Where MarkLines combines powertrains, we allocate sales equally among the reported categories.

For marginal-cost recovery, we use prices from the sales market within 1 year of the sales year. For demand estimation in the remaining countries, we also use prices from other markets within 3 years. Supplementary online prices fill missing prices. We retain country-years in which products with prices account for at least 55 percent of sales. We drop products with market shares below 10^{-5} and prices above \$150,000 in 2025 U.S. dollars.

A.3 Prices, Market Size, and Income

We adjust prices for inflation using the source country's consumer price index and convert them into 2025 U.S. dollars at the base-year exchange rate.

Market size is

$$M_{mt} = \frac{\text{households}_{mt} \times \text{vehicles per household}_{mt}}{5}, \tag{11}$$

where both terms are measured in country m and year t. The denominator is an assumed average ownership duration of 5 years, as in Grieco, Murry, and Yurukoglu (2024). Where household counts are unavailable,

Table A3: Country and Year Coverage

Country	(1) Years	(2) Observations	(3) Cost sample
Argentina	2015–2025	1,173	No
Australia	2015–2025	2,163	No
Austria	2015–2025	2,890	Yes
Belgium	2015–2025	2,618	Yes
Brazil	2018–2021, 2023–2025	530	No
Canada	2015–2025	2,533	No
China	2015–2024	2,031	No
Colombia	2015–2024	984	No
Czech Republic	2015–2025	2,319	Yes
Denmark	2015–2025	2,183	Yes
Egypt	2015–2025	804	No
France	2015–2025	2,801	Yes
Germany	2015–2025	2,973	Yes
Greece	2015–2025	1,856	Yes
India	2015–2022	429	No
Ireland	2015–2025	2,194	Yes
Israel	2015–2024	1,958	No
Italy	2015–2025	2,686	Yes
Japan	2015–2022	749	Yes
Malaysia	2015–2023	354	Yes
Mexico	2015–2025	2,152	No
Netherlands	2015–2025	2,678	Yes
Norway	2016–2025	1,731	Yes
Philippines	2015–2025	1,354	No
Poland	2015–2025	1,786	Yes
Portugal	2015–2025	2,663	Yes
Romania	2015–2025	2,169	Yes
Saudi Arabia	2015–2025	1,799	No
South Africa	2022, 2025	357	Yes
Spain	2015–2025	2,663	Yes
Sweden	2015–2025	2,605	Yes
Switzerland	2015–2025	2,473	Yes
Thailand	2017–2025	580	No
Turkey	2015–2025	1,490	Yes
United Arab Emirates	2015–2025	1,849	Yes
United Kingdom	2015–2025	2,866	Yes
United States	2015–2025	3,014	Yes
Vietnam	2015–2025	480	No

Notes: Observations are product–country–years. The cost sample uses prices from the sales market and is subject to the additional restrictions in Appendix C.5. Nonlinear demand estimation uses the United States, United Kingdom, and Germany in 2025.

we infer them from population and average household size. Where vehicles per household are unavailable, we construct the vehicle stock from population and vehicles per capita.

Following Coşar et al. (2017) and Sabal (2026), we use PPP GDP per capita as a proxy for mean income and fit a lognormal distribution using that mean and the country's Gini coefficient. GDP per capita comes from the IMF World Economic Outlook and is rebased with the U.S. GDP deflator. Gini coefficients come from the World Bank. Let $\bar{y}_{mt}$ denote the mean-income proxy and Gini_{mt} the Gini coefficient. The lognormal parameters satisfy

$$\tau_{mt} = \sqrt{2}\,\Phi^{-1}\left(\frac{1+\text{Gini}_{mt}}{2}\right), \qquad \mu_{mt} = \log \bar{y}_{mt} - \frac{\tau_{mt}^2}{2}, \qquad \log y_{imt} \sim \mathcal{N}(\mu_{mt}, \tau_{mt}^2). \tag{12}$$

Income is measured in 2025 GDP-PPP international dollars.

A.4 Survey Design and Second Choices

The vehicle-owner survey covers China, Germany, Japan, the United Kingdom, and the United States and oversamples EV owners. It records household vehicles, purchase years, whether vehicles were bought new or used, powertrains, demographics, and second choices. We ask what respondents would have bought if their chosen brand and model had been unavailable. Chinese-brand owners whose second choice is also Chinese receive a follow-up question removing all Chinese brands.

The first-choice vehicle is generally the respondent's most recent purchase. For the Chinese-brand-owner questions outside the United States, we use the Chinese vehicle specified in the question. We match reported brands, models, and powertrains to the market data to assign body types. A response enters a second-choice moment when the required vehicle characteristics are observed and the respondent would have purchased another vehicle.

The estimation sample includes both new- and used-vehicle purchases in 2025, except that the Chinese-origin moment uses all purchase years. These choices increase the number of observations given the limited survey sample. We assume that the resulting conditional substitution patterns apply to the modeled new-vehicle market.

A.5 Summary Statistics and Survey Weights

Tables A1 and A2 summarize the market and survey data. Table A4 reports eligible survey respondents by estimation moment and country. To adjust for EV oversampling, we assign weight $\frac{1}{3}$ to first-choice EV respondents and 1 to other respondents. We use these weights for the empirical moments and their country weights, as described in Appendix C.2.

Table A4: Eligible Survey Respondents by Estimation Moment

Moment	(1) United States	(2) United Kingdom	(3) Germany
Car second-choice share	216	283	221
Chinese-origin second-choice share	0	106	31
European-origin second-choice share	112	372	368
EV second-choice share	197	257	218
Pickup second-choice share	66	1	0
US-origin second-choice share	284	85	41

Notes: Unweighted counts of respondents with a valid first vehicle, an inside second choice, and the characteristics required for each moment. The Chinese-origin moment uses all purchase years; the remaining moments use purchases in 2025. Both new and used purchases are included.

A.6 Externality Inputs

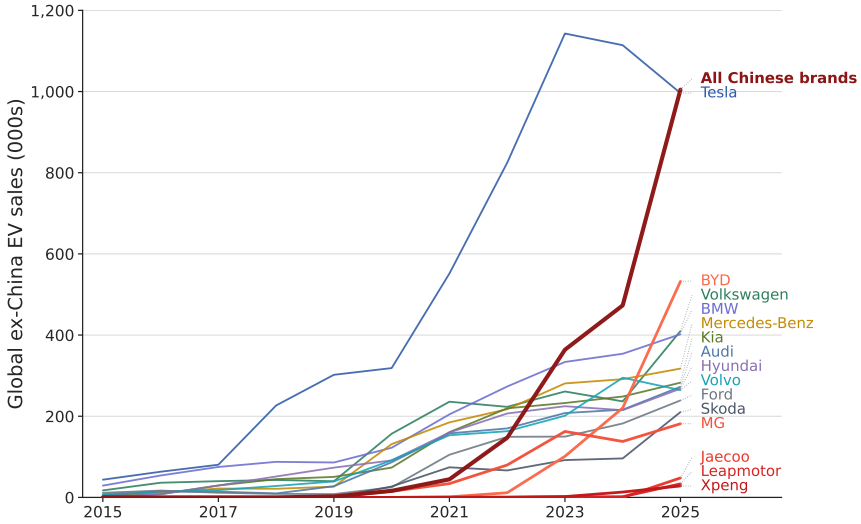
We calculate climate damages from vehicle use and manufacturing. Lifetime mileage is 150,000 miles. For plug-in hybrids, 63 percent of mileage uses gasoline and 37 percent uses electricity, following Plötz et al. (2020). We apply gasoline fuel economy and electric energy use to the corresponding mileage.

Electricity carbon intensity comes from the Our World in Data series sourced from Ember (Our World in Data 2026). We use country-average grid emissions as a proxy for marginal generation emissions. We impute missing vehicle efficiency using medians for similar vehicles, beginning with vehicles of the same powertrain, body type, and brand in the same country-year.

The social cost of carbon is $253 per metric ton in 2025 dollars. We obtain this value by adjusting the $241 per metric ton used by Allcott et al. (2024) from July–August 2023 dollars to 2025 dollars using the U.S. Consumer Price Index (CPI-U). We subtract country-year carbon prices for electricity and road fuel using the World Bank's implemented carbon taxes and emissions trading systems, weighted by their emissions coverage (World Bank 2026). We add powertrain-specific manufacturing damages, adjusted for inflation. The outside option represents continued use of an existing vehicle, with use-phase emissions based on the observed mix of gasoline vehicles and conventional hybrids.

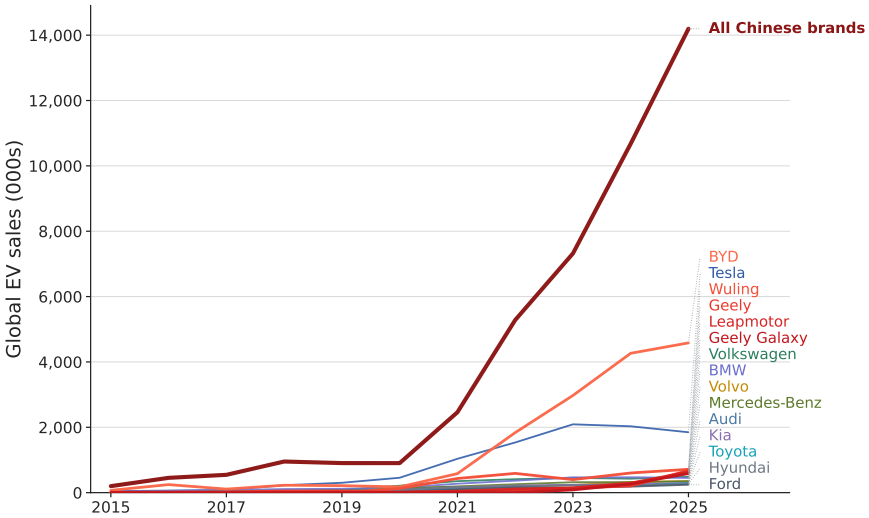
B Descriptive Facts Appendix

Figure A1: EV Sales by Brand, Worldwide Excluding China



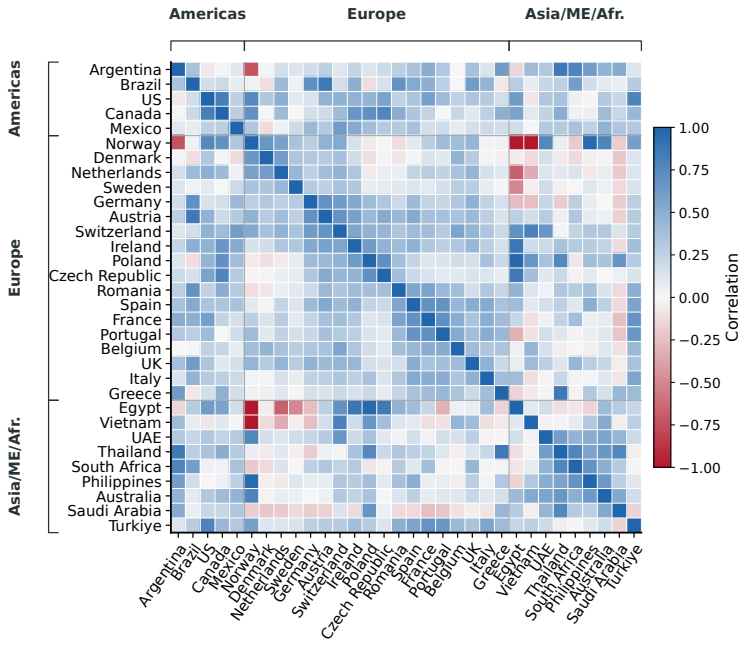
Notes: This figure presents the top 5 Chinese brands and the top 10 non-Chinese brands by 2025 EV sales in the top 50 global automotive markets, excluding China.

Figure A2: EV Sales by Brand, Worldwide Including China



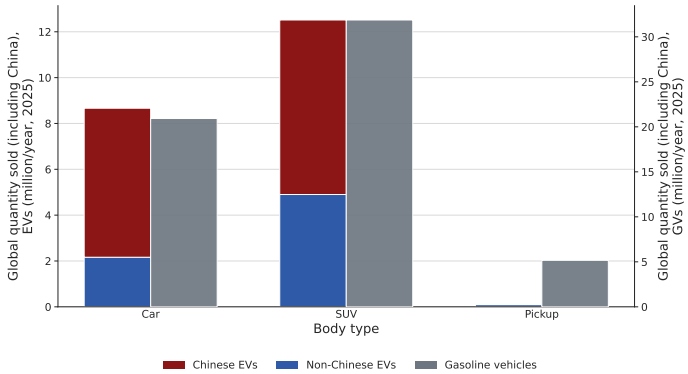
Notes: This figure presents the top 5 Chinese brands and the top 10 non-Chinese brands by 2025 EV sales in the top 50 global automotive markets, including China.

Figure A3: Market Share Correlations Across Countries (Overlapping Products Only)



Notes: This figure shows the correlation in 2025 product market shares between each country pair in our sample. Products that are not available in a country are coded as missing.

Figure A4: Sales by Body Type



Notes: This figure reports 2025 vehicle sales across the top 50 automotive markets, including China, by body type. Stacked bars separate Chinese and non-Chinese EVs, using the left axis; gray bars show gasoline vehicles, using the right axis. Both axes report millions of vehicles. EVs include battery electric and plug-in hybrid vehicles; gasoline vehicles include conventional hybrids.

C Model Appendix

C.1 Demand Specification

We estimate a random-coefficients demand system following Berry, Levinsohn, and Pakes (1995).

We allow consumers to differ in their willingness to buy a vehicle and their preferences for powertrain, body style, and brand origin. The vector Z_{jmt} contains an inside-good indicator and indicators for EVs, cars, pickups, and Chinese, U.S., and European brand origin. Their random coefficients have standard deviations σ. Income affects price sensitivity through the marginal utility of money α/y_{imt} and shifts the utility of all inside goods by $\beta_y \log y_{imt}$.

Utility net of the logit shock is $\bar{u}_{ijmt} = \delta_{jmt} + \beta_y \log y_{imt} - (\alpha/y_{imt})p_{jmt} + \sigma \cdot \nu_i \cdot Z_{jmt}$. Taste shocks ν_i are independent standard normals, independent of income. Normalizing the outside good's mean utility to zero, choice probabilities and market shares are

$$s_{ijmt} = \frac{\exp(\bar{u}_{ijmt})}{1 + \sum_{\ell \in \mathcal{J}_{mt}} \exp(\bar{u}_{i\ell mt})}, \qquad s_{jmt} = \int s_{ijmt}\, dG_{mt}(y, \nu). \tag{13}$$

We take market sizes and income distributions as given and estimate common preference parameters using the 2025 U.S., U.K., and German markets. At these estimates, we recover mean utilities for the full set of countries and years for which we have data.

C.2 Estimator and Moment Conditions

For each candidate parameter vector $\gamma = (\log \alpha, \beta_y, \sigma')'$, we recover mean utilities that match observed market shares and compute the model's elasticity and second-choice moments. We choose parameters to minimize the squared differences between these moments $q(\gamma)$ and their empirical targets $\widehat{q}$:

$$\widehat{\gamma} = \arg\min_{\gamma} (q(\gamma) - \widehat{q})' (q(\gamma) - \widehat{q}), \qquad s_{jm}(\delta(\gamma), \gamma) = s_{jm}^{\text{obs}}. \tag{14}$$

We omit the year subscript in the estimation equations because these markets share the same year.

The own-price and market elasticities jointly inform price sensitivity and heterogeneity in willingness to purchase a vehicle. Second-choice shares inform heterogeneity in preferences for powertrains, body types, and brand origins: consumers with strong preferences for a category tend to choose another vehicle in that category when their first choice is unavailable.

Price Elasticities. We target the sales-weighted mean own-price elasticity and the elasticity of total new-vehicle demand to a common proportional increase in prices. Let $1 - s_{0m} = \sum_j s_{jm}$ be the total inside share, where s_{0m} is the outside share. Holding mean utilities fixed, these elasticities are

$$E_m^{\text{own}} = \frac{\sum_j p_{jm}\, \partial s_{jm}/\partial p_{jm}}{1 - s_{0m}}, \qquad E_m^{\text{mkt}} = \frac{\sum_j \sum_\ell p_{\ell m}\, \partial s_{jm}/\partial p_{\ell m}}{1 - s_{0m}}. \tag{15}$$

For the U.S. market, we use targets of -5.36 following Grieco, Murry, and Yurukoglu (2024) and -1 following Berry, Levinsohn, and Pakes (2004), respectively.

Second Choices. The remaining 6 moments measure the share of buyers whose second choice is in the same category as their first choice, conditional on buying another vehicle. The categories are EVs, cars, pickups, and Chinese, U.S., and European brand origin. Let A_{nk} indicate that respondent n's first choice is in category k, the second choice is another vehicle, and the required characteristics are observed. Let Y_{nk} indicate that the second choice is also in category k, and let a_n be the survey weight. The weighted eligible respondent count N_{mk} in country m and the empirical moment are

$$N_{mk} = \sum_{n:m_n=m} a_n A_{nk}, \qquad \widehat{q}_k = \frac{\sum_n a_n A_{nk} Y_{nk}}{\sum_m N_{mk}}. \tag{16}$$

We construct model second choices by removing product j and holding the remaining prices and mean utilities fixed. If $s_{\ell m}^{(-j)}$ is product ℓ's share after removing j, diversion to that product and total diversion to other vehicles are

$$d_{j\to\ell,m} = \frac{s_{\ell m}^{(-j)} - s_{\ell m}}{s_{jm}}, \qquad d_{jm}^{\text{inside}} = \sum_{\ell\in\mathcal{J}_m\setminus\{j\}} d_{j\to\ell,m}. \tag{17}$$

We compute the same conditional share in the model and average across countries using the survey's weighted eligible counts. For the set of products $\mathcal{J}_{km}$ in category k, the country and pooled model moments are

$$q_{mk}(\gamma) = \frac{\sum_{j\in\mathcal{J}_{km}} s_{jm} \sum_{\ell\in\mathcal{J}_{km}\setminus\{j\}} d_{j\to\ell,m}}{\sum_{j\in\mathcal{J}_{km}} s_{jm} d_{jm}^{\text{inside}}}, \qquad q_k(\gamma) = \sum_m \frac{N_{mk}}{\sum_{m'} N_{m'k}} q_{mk}(\gamma). \tag{18}$$

C.3 Additional Estimation Results

Table A5: Empirical and Model-Implied Moments

Moment	(1) Target	(2) Target SE
Mean own-price elasticity (US, 2025)	-5.360	–
Market elasticity (US, 2025)	-1.000	–
Car second choice	0.643	0.020
Chinese-origin second choice	0.156	0.040
European-origin second choice	0.758	0.017
EV second choice	0.659	0.032
Pickup second choice	0.471	0.069
US-origin second choice	0.572	0.030

Notes: Second choice moments are conditional on choosing another vehicle. Survey moments pool the United States, United Kingdom, and Germany. Target standard errors are reported for the survey moments.

Standard Errors. We compute parameter standard errors by the delta method, treating the elasticity targets as fixed. Let $J = \partial q(\gamma)/\partial \gamma'$ at $\widehat{\gamma}$, allowing mean utilities to adjust to match observed shares. The sensitivity matrix and parameter covariance are

$$S = (J'J)^{-1}J', \qquad \widehat{\mathrm{Var}}(\widehat{\gamma}) = S\widehat{\Omega}S'. \tag{19}$$

Here $S = J^{-1}$ when J has full rank. We estimate $\widehat{\Omega}$ from survey sampling variation, accounting for respondents who contribute to multiple moments and for variation in the eligible country weights.

Let $N_k = \sum_m N_{mk}$. The respondent contribution is

$$\chi_{nk} = \frac{a_n A_{nk}}{N_k} \left\{ Y_{nk} - q_{m_n k}(\widehat{\gamma}) - (\widehat{q}_k - q_k(\widehat{\gamma})) \right\}, \tag{20}$$

where $\widehat{q}_k - q_k(\widehat{\gamma})$ is the weighted mean survey residual for moment k. We center these contributions within country and sum their outer products. Standard errors condition on country respondent counts, market data, income distributions, and integration draws. For the empirical targets in Table A5, we instead use contributions $a_n A_{nk}(Y_{nk} - \widehat{q}_k)/N_k$, centered within country.

Sensitivity. Table A6 shows how parameter estimates respond to proportional changes in empirical targets. Price sensitivity and the inside-good dispersion respond primarily to the elasticity targets. EV and body-type dispersions are particularly sensitive to their corresponding second-choice shares. The Chinese-origin dispersion also responds to the EV and European-origin moments, reflecting the overlap between powertrain and brand-origin substitution.

Table A6: Sensitivity of Demand Estimates to Empirical Moments

Parameter	(1) Mean own-price elasticity	(2) New-vehicle market elasticity	(3) Car second-choice share	(4) Chinese-origin second-choice share	(5) European-origin second-choice share	(6) EV second-choice share	(7) Pickup second-choice share	(8) US-origin second-choice share
$\log \alpha$	0.406	0.203	4.43e-4	1.04e-5	1.43e-4	0.004	0.006	-0.012
σ_{inside}	1.133	-1.611	0.033	3.97e-5	0.029	0.007	-0.004	0.071
σ_{EV}	0.089	-0.007	0.057	0.004	0.034	2.352	-0.013	-0.002
σ_{car}	0.118	-0.107	2.342	0.004	-0.013	0.025	-0.015	0.002
σ_{pickup}	-0.012	0.022	-0.074	-1.88e-4	-0.060	-0.066	1.958	-0.249
σ_{China}	0.029	-0.026	0.143	0.800	0.020	-0.284	0.003	-0.127
σ_{US}	0.092	-0.133	0.029	-5.77e-4	-0.512	0.016	-0.026	2.523
σ_{Europe}	0.082	-0.137	-0.035	-0.001	6.466	0.007	-0.018	-0.293

Notes: Entries are the local elasticities $\partial \log \widehat{\gamma}/\partial \log \widehat{q}_k$ of each parameter estimate with respect to each empirical target, allowing mean utilities to adjust to match observed shares. The price row reports proportional changes in $\log \alpha$.

C.4 Computation and Demand Validation

We recover mean utilities using the Berry (1994) contraction, accelerated by SQUAREM (Varadhan and Roland 2008). We integrate over income and tastes to compute shares, moments, elasticities, and welfare. To improve accuracy for Chinese-EV demand and welfare, we oversample high-income households and consumers with strong preferences for Chinese EVs.

We partition income into intervals of width 10,000 in 2025 international dollars, with a final interval $[400,000, \infty)$. We also stratify the standard-normal taste index

$$\bar{\nu}_i = 0.8\nu_{i,\text{EV}} + 0.6\nu_{i,\text{China}}, \tag{21}$$

using intervals with boundaries at percentiles 50, 90, 99. Each pair of an income interval and a taste interval is a stratum h. In each stratum, we draw a set $\mathcal{I}_h$ of 32 consumers using scrambled Sobol sequences and the conditional distributions of income and $\bar{\nu}_i$. The orthogonal taste component and remaining shocks retain their standard-normal distributions. We use 5,248 draws per market.

We approximate each integral over income and tastes by weighting each stratum by its probability $\Pr_{mt}(h)$. For example, market shares are

$$s_{jmt} = \int s_{ijmt}\, dG_{mt}(y,\nu) \approx \sum_h \Pr_{mt}(h)\, \frac{1}{|\mathcal{I}_h|} \sum_{i \in \mathcal{I}_h} s_{ijmt}. \tag{22}$$

Weighting by stratum probabilities adjusts for oversampling. We hold the draws fixed during estimation.

Demand Validation. We assess substitution from Chinese EVs using the survey question that removes all Chinese brands. This question provides additional evidence on substitution toward other EVs, gasoline vehicles, and competing brand origins. We compare these responses with model predictions at fixed prices, conditional on buying another vehicle and pooled using the survey's weighted eligible counts.

Appendix Table A7 summarizes price responses and counterfactual second choices. When we remove all Chinese brands, as in the survey question, the model's EV second-choice share differs from the survey by less than seven percentage points in both China and the pooled German and U.K. sample. In China, the model predicts lower diversion to U.S. and European brands than the survey reports.

Table A7: Key Elasticities and Diversion Ratios

	(1) US (2025)	(2) Germany and UK (2025)	(3) China (2024)	(4) Brazil (2025)	(5) World sample
Panel A. Elasticities					
Mean own-price elasticity	-5.36	-7.76	-9.85	-5.69	-7.35
Aggregate new-vehicle demand elasticity	-1.00	-1.50	-1.60	-1.51	-1.41
Aggregate EV demand elasticity	-3.08	-3.43	-3.44	-3.85	-3.38
Panel B. Chinese EV Owners' Second Choices if All Chinese Brands Removed (%)					
EVs					
Model	–	64.4	48.0	5.6	47.7
Survey	–	70.8	50.3	–	–
(SE)	–	(5.4)	(2.5)	–	–
U.S. brands					
Model	–	10.8	27.1	16.0	25.9
Survey	–	11.7	24.3	–	–
(SE)	–	(3.7)	(1.8)	–	–
European brands					
Model	–	63.9	55.4	34.8	57.2
Survey	–	64.9	42.1	–	–
(SE)	–	(5.4)	(2.1)	–	–
Survey observations, powertrain	–	72	396	–	–
Survey observations, brand origin	–	77	560	–	–

Notes: Column headers give each country's latest observed model year; the world sample includes all model countries at their latest years. Panel A reports the model's sales-weighted mean product own-price elasticity and the elasticities of aggregate new-vehicle and EV demand to a common proportional price increase for the named set of vehicles, holding mean utilities and other prices fixed. EVs include electric and plug-in hybrid vehicles. Panel B restricts to buyers whose first choice is a Chinese EV and conditions on choosing another vehicle. The model rows remove all Chinese brands jointly at fixed prices. The survey asks about all Chinese brands and uses all valid purchase years. Its Germany–U.K. figures pool respondents without survey weights; the corresponding model shares use the same country sample mix separately for powertrain and brand origin. European brand headquarters follow the survey category. Parentheses give survey standard errors in percentage points; the final rows give survey sample sizes. Japan's latest model year contains no Chinese EV, so it cannot enter the modeled pooled comparison. Dashes indicate no baseline Chinese EV or no corresponding survey sample; the world column reports model results for the sample countries, not a world survey estimate.

C.5 Marginal Costs and Cost Prediction

We recover marginal costs from multiproduct Bertrand first-order conditions using prices from the sales market. Let $f(j)$ denote the parent firm of product j, and define $\mathcal{H}_{j\ell,mt} = -\mathbb{1}\{f(j) = f(\ell)\}\,\partial s_{\ell mt}/\partial p_{jmt}$. Marginal costs satisfy $c_{mt} = p_{mt} - \mathcal{H}_{mt}^{-1} s_{mt}$.

To predict costs for products offered in additional countries, we regress log marginal costs on country, product, year, body-type, and powertrain fixed effects and vehicle characteristics. The characteristics are log footprint, log horsepower, log gasoline fuel economy for GVs, log electricity use for battery electric vehicles, and log electric range for plug-in hybrids. We also include a home-country indicator and Chinese-EV calendar-year indicators. With these covariates in $\boldsymbol{W}_{jmt}$, we estimate $\log c_{jmt} = \boldsymbol{W}_{jmt}\phi + \omega_{jmt}$. Table A8 reports the estimates of ϕ.

For potential offerings in countries with cost data, we average product characteristics across markets using sales weights for the same year, or across years when necessary. We exponentiate fitted log costs to predict costs for these offerings and use recovered costs for products already offered.

Table A8: Marginal Cost Coefficients

Coefficient	(1) Estimate	(2) Standard error
SUV	-0.838	0.037
Pickup	-1.015	0.044
Internal combustion	0.782	0.416
Hybrid	1.380	0.400
Plug-in hybrid	0.882	0.368
$\ln(\text{size})$	1.015	0.207
$\ln(\text{horsepower})$	0.391	0.040
$\ln(\text{fuel efficiency})$, non-electric	-0.086	0.026
$\ln(\text{kWh}/100\text{ km})$, electric	0.340	0.107
$\ln(\text{electric range})$, PHEV	-0.029	0.037
Domestic brand	-0.033	0.007
Chinese EV $\times$ 2020	0.036	0.101
Chinese EV $\times$ 2021	0.030	0.099
Chinese EV $\times$ 2022	-0.010	0.099
Chinese EV $\times$ 2023	-0.007	0.110
Chinese EV $\times$ 2024	0.019	0.116
Chinese EV $\times$ 2025	0.035	0.122

Notes: The entries are the coefficients ϕ of the log marginal-cost projection in equation (3), estimated by OLS on products with recovered marginal costs, with standard errors clustered by brand. Car and battery-electric vehicles are the omitted body-type and powertrain categories, and Chinese-EV year effects are relative to all other products. The intercept and country, product, and year fixed effects are included but not displayed.

C.6 IFE and Validation

Specification. We use matrix completion to predict demand for products in countries where they are unavailable, combining their characteristics with demand in other markets. We arrange mean utilities in a matrix with products in rows and country-years in columns, using recovered mean utilities for offered products with the required characteristics. Products with at least 5 observations enter the low-rank fit; we predict mean utilities for other products using their observable characteristics.

We allow the coefficients on log footprint, body type, and powertrain to vary by region. The coefficients on log horsepower, log fuel economy for GVs, log electricity use for battery electric vehicles, and log electric range for plug-in hybrids are common across regions. We also include brand and country-year fixed effects, a home-country indicator, and Chinese-EV calendar-year indicators.

We estimate equation (8), selecting the nuclear-norm penalty by 4-fold cross-validation to minimize the average root mean squared prediction error across folds. We then refit using all observed mean utilities.

Estimation Algorithm. The solution to the problem is well studied in statistics. We import the tools from Mazumder, Hastie, and Tibshirani (2010) and Athey et al. (2021) and find the Γ and θ that solve this program using a soft-impute algorithm that iterates the following mapping:

$$\theta \leftarrow (\boldsymbol{X}'\boldsymbol{X})^{-1}\boldsymbol{X}'\text{vec}(\tilde{\Delta} - \Gamma), \tag{23}$$

$$\Gamma \leftarrow \mathcal{T}_{\psi}(\tilde{\Delta} - \boldsymbol{X}\theta), \tag{24}$$

$$\tilde{\Delta} \leftarrow P_{\mathcal{O}}(\Delta) + P_{\mathcal{O}}^{\perp}(\boldsymbol{X}\theta + \Gamma), \tag{25}$$

where $\tilde{\Delta}$ is the completed matrix of mean utilities and, for an argument with singular-value decomposition $U\Sigma V'$, $\mathcal{T}_{\psi}$ returns $U \max\{\Sigma - \psi I, 0\}V'$. The objective is jointly convex; the iteration decreases it monotonically and its fixed points are exactly the first-order conditions of objective (8), so the algorithm converges to a global optimum regardless of initialization. Steps (23)-(25) are intuitive. Step (23) is standard least-squares. Given a matrix Γ and completed matrix $\tilde{\Delta}$, the solution for $\hat{\theta}$ is the OLS estimator. Step (24) is the imputation step. Given a rank-penalty ψ, $\tilde{\Delta}$ and θ, it selects the low-rank component of the residual via singular-value decomposition. The solution is only to keep directions that have singular values larger than ψ, shrink the retained directions by ψ, and to treat the rest as noise. Step (25) updates the completed matrix. For observed product-market entries, it selects the observed δ; for unobserved entries, it sums the observable and low-rank component.

Relation to Interactive Fixed Effects. The estimator (8) is exactly the penalized version of the interactive fixed effects model of Section 4. By the variational characterization of the nuclear norm (e.g., Mazumder,

Hastie, and Tibshirani 2010),

$$\|\Gamma\|_* = \min_{\Lambda,F:\Lambda F'=\Gamma} \frac{1}{2}\left(\|\Lambda\|_F^2 + \|F\|_F^2\right),$$

so (8) can be rewritten as

$$\min_{\theta,\Lambda,F} \frac{1}{2}\|P_{\mathcal{O}}(\Delta - \boldsymbol{X}\theta - \Lambda F')\|_F^2 + \frac{\psi}{2}\left(\|\Lambda\|_F^2 + \|F\|_F^2\right).$$

That is, our estimator is equivalent to interactive fixed effects with ridge-penalized loadings and factors, in which the number of factors is not imposed but estimated. The penalty sets to zero all factors whose explanatory power does not exceed ψ. In particular, the solution takes the factor form of the model in Section 4: writing the singular value decomposition $\hat{\Gamma} = U\Sigma V'$ with $\hat{R} = \text{rank}(\hat{\Gamma})$, we obtain $\hat{\Gamma}_{jmt} = \hat{\lambda}_j'\hat{F}_{mt}$ with $\hat{\lambda}_j = \Sigma^{1/2}U_{j\cdot}'$ and $\hat{F}_{mt} = \Sigma^{1/2}V_{mt\cdot}'$.[23] Relative to least-squares interactive fixed effects (Bai 2009), the nuclear-norm formulation (i) is convex, so the algorithm converges to a global optimum from any initialization; (ii) selects the number of factors within the estimation rather than through a preliminary information criterion; and (iii) remains well-defined under the incomplete observation pattern of our panel. On a complete panel the two approaches coincide up to shrinkage, i.e., the solution equals the principal-components estimator with each retained singular value reduced by exactly ψ.

Identification. The identification arguments of the (relative) mean utilities Δ follow from Berry, Levinsohn, and Pakes (1995). The identification arguments regarding the separability of θ and Γ are discussed in Moon, Shum, and Weidner (2018). The key formal argument for the identification of θ is that, under the true number of factors R, $\boldsymbol{X}$ survives projecting out the factor-loading directions. This intuition is similar to the identification of two-way fixed effects, where coefficients on time-invariant regressors are typically not identified. Here, the effects of regressors that move in the same directions as the true factor loadings cannot be identified either.

Relative to Moon, Shum, and Weidner (2018), we do not impose a fixed number of factors, but rather use the nuclear norm to select the rank. While this is a benefit, it also means that our θ estimates may contain a penalty normalization bias. Intuitively, the estimator prices high rank via the penalty term, so it may bias coefficients on observables to reduce the rank of the residual. While this means that the interpretation of θ is not without caveats, our main goal is to understand the counterfactual effects of Chinese EV entry, for which only the identified joint value of $\boldsymbol{X}\theta + \Gamma$ is required.

[23] As in any factor model, (Λ, F) are identified only up to an invertible rotation, $\Lambda F' = (\Lambda Q)(FQ^{-1\prime})'$; the SVD provides a normalization, and all objects we use — predictions $\hat{\lambda}_j'\hat{F}_{mt}$ and similarity measures — are rotation-invariant.

Table A9: Linear Demand Coefficients

Characteristic	(1) N. America	(2) Latin America	(3) Central Europe	(4) Eastern Europe	(5) Southern Europe	(6) Nordics	(7) Developed APAC	(8) Emerging
ln(size)	6.549	3.892	5.712	11.756	5.327	7.107	0.451	14.803
SUV	2.260	1.726	1.243	1.336	1.481	1.186	2.086	1.394
Pickup	0.611	1.072	-1.013	-2.349	-1.119	-1.489	1.855	-1.893
Hybrid	-0.650	0.340	-0.253	-0.623	-0.256	0.408	0.282	-0.051
Electric	-3.820	-3.982	-2.371	-4.662	-3.116	0.556	-3.890	-1.629
Plug-in hybrid	-5.092	-3.303	-2.667	-4.721	-3.001	-0.285	-4.751	-0.980
ln(horsepower)	3.022							
ln(fuel efficiency), non-electric	0.416							
ln(kWh/100 km), electric	0.670							
ln(electric range), PHEV	0.525							
Domestic brand	1.391							
Chinese EV × 2016	4.996							
Chinese EV × 2017	4.961							
Chinese EV × 2018	5.130							
Chinese EV × 2019	3.706							
Chinese EV × 2020	2.768							
Chinese EV × 2021	2.771							
Chinese EV × 2022	3.349							
Chinese EV × 2023	3.456							
Chinese EV × 2024	3.687							
Chinese EV × 2025	3.499							

Notes: The entries are the coefficients θ on observed characteristics in the IFE projection of mean utility. Car and internal-combustion vehicles are the omitted body-type and powertrain categories. The region columns apply to size, body type, and powertrain; the other slopes are common across regions. Brand and country-year fixed effects are included but not displayed. Chinese-EV year effects are relative to the first sample year.

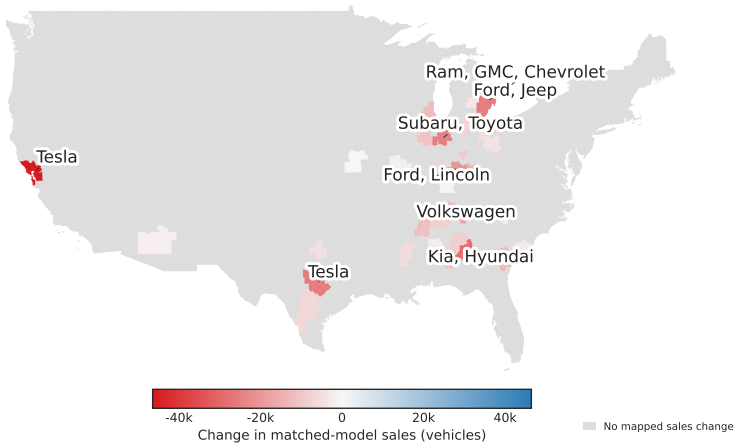
Table A10: How Well Appeal in Other Markets Predicts Appeal in a Held-Out Market

			Mean absolute error	
	Obs.	Correlation	OLS	Method
Panel A: Each method's adjustment to OLS in the out-of-sample holdouts				
US models, random half				
IFE	28,669	0.66	1.14	0.86
OLS + product FE	28,669	0.30	1.14	1.25
Chinese EVs, leave one country out (excluding China)				
IFE	784	0.64	1.95	1.47
OLS + product FE	784	0.34	1.95	1.82
Panel B: Predicting a model's US residual from its residuals in other markets				
Mean over all other markets	409	0.30		
Canada only	333	0.68		
Mean over European markets	235	0.17		
Mean weighted by similarity to the US	409	0.48		
Panel C: US random-half holdout, by number of countries where the model is sold				
Countries	Obs.		OLS	Product FE	IFE
1 (United States only)	530		0.82	0.99	0.76
2–3	3,560		1.16	1.16	0.91
4–8	8,340		0.99	1.11	0.79
9–15	3,300		1.39	1.50	0.99
16 or more	12,939		1.19	1.30	0.86

Notes: Panel A uses the held-out observations of Figure 7; for Chinese EVs, these are the countries in Panel (a) other than China, each with at least five held-out observations. For each method, x is its prediction minus the OLS prediction and r is the OLS residual; Correlation is between x and r. Panel B reports the correlation between each product's U.S. residual from the full-sample OLS fit and a summary of its residuals in other countries, averaging over years; the equal-weighted mean approximates a product fixed effect, similarity weights are countries' latent-demand correlations with the United States (Figure 6, Panel (b)), and the residuals are in-sample. Panel C reports U.S. holdout MAEs for OLS, OLS + product FE, and IFE by the number of countries where the model is sold.

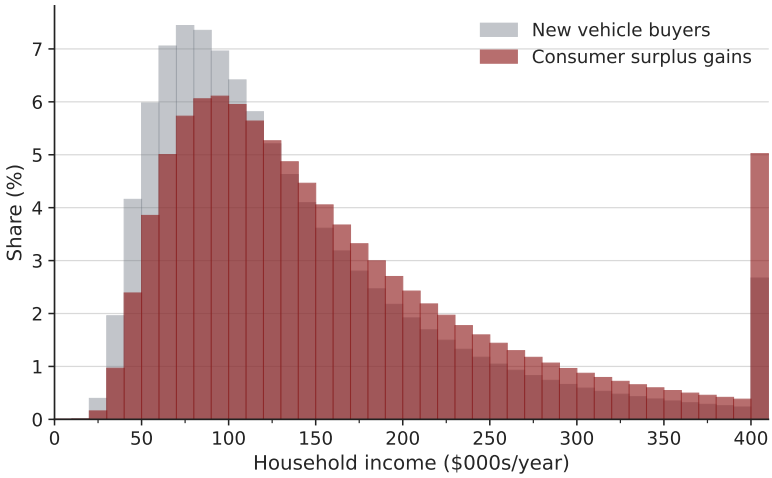
D Counterfactual Simulations Appendix

Figure A5: Matched-Model Sales Changes under Chinese EV Entry in 2030



Notes: Predicted changes in annual U.S. sales under Chinese EV entry in 2030 relative to the scenario without Chinese entry. Sales changes are matched to U.S. assembly plants, allocated using 2025 production weights, and aggregated by commuting zone. Labels identify the brands with the largest losses in the seven commuting zones with the largest losses. Gray areas have no mapped sales change.

Figure A6: Consumer Surplus Gains by Income



Notes: Consumer-surplus effects are the difference in world means across 5,000 simulated 2030 equilibrium configurations per world. Demand integration uses the saved cross-stratified Sobol design with $10,000 income strata. Incomes at or above $400,000 are grouped in the right-most bin.