NBER WORKING PAPER SERIES

TARIFFS, GLOBAL VALUE CHAINS, AND THE INCIDENCE OF PROTECTION: EVIDENCE FROM US AUTOMOBILES

Luke Heeney
Christopher R. Knittel
Jasdeep Mandia

Working Paper 35023
http://www.nber.org/papers/w35023

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2026

We are grateful for helpful comments and suggestions from the participants at the MIT CEEPR Seminar 2025, Association for Public Policy Analysis and Management (APPAM) Conference 2025. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Tariffs, Global Value Chains, and the Incidence of Protection: Evidence from US Automobiles
Luke Heeney, Christopher R. Knittel, and Jasdeep Mandia
NBER Working Paper No. 35023
April 2026
JEL No. F13, L13, L62

ABSTRACT

In many modern industries, firms compete in differentiated-product markets while relying on complex global value chains for intermediate inputs. In such settings, trade policies such as tariffs on vehicles and parts operate not only through consumer substitution and firm pricing, but also through firms' cost structures and sourcing decisions. We develop a structural model of the U.S. automobile market that integrates random-coefficients demand, multiproduct firm pricing, and a flexible supply-side framework in which shocks to the cost of imported parts transmit imperfectly into manufacturers' marginal costs. The model is disciplined by novel model-level data on imported-parts exposure and exploits exchange-rate variation to identify cost pass-through. Our counterfactual analysis quantifies the effects of alternative tariff policies on prices, profits, and welfare. First, tariffs on imported vehicles alone reallocate demand toward domestically assembled products and increase U.S. producer surplus, generating a gain of approximately$1 billion for U.S.-headquartered firms, while reducing consumer surplus by about $14 billion. Second, extending tariffs to imported intermediate inputs fundamentally alters these effects: consumer surplus losses roughly double, and producer surplus for U.S.-headquartered firms declines by about $2.6 billion. These aggregate effects mask substantial heterogeneity: firms with greater exposure to imported parts experience losses, whereas those relying more on domestic inputs are better able to increase profits. Overall, the results show that tariff incidence depends critically on firms' exposure to global value chains and cannot be inferred from final assembly locations alone.

Luke Heeney
Massachusetts Institute of Technology
heeney@mit.edu

Jasdeep Mandia
Massachusetts Institute of Technology
jmandia@mit.edu

Christopher R. Knittel
Massachusetts Institute of Technology
MIT Sloan School of Management
Department of Economics
and NBER
knittel@mit.edu

1 Introduction

A central question in the analysis of trade and industrial policy is how protection affects prices, profits, and welfare in markets characterized by market power and globalized production. In many modern industries, firms compete in differentiated-product markets while simultaneously relying on complex global value chains for intermediate inputs. In such settings, policy instruments that target trade—including tariffs (vehicles and/or parts), subsidies, and domestic content requirements—operate not only through consumer substitution and firm pricing, but also through firms' cost structures and sourcing decisions. As a result, the incidence of protection is no longer determined solely by the location of final assembly or the elasticity of demand, but depends critically on the exposure of domestic producers to foreign intermediate inputs.

This paper studies how intermediate-input exposure alters the incidence and welfare consequences of tariffs in oligopolistic differentiated-goods markets. We show that when domestically assembled products rely heavily on imported inputs, policies seemingly intended to protect domestic producers can instead raise their marginal costs, weaken their competitive position, and reduce producer surplus. More broadly, our analysis highlights that the distinction between "domestic" and "foreign" goods becomes blurred once production is fragmented across borders, and that ignoring global value chains can lead to misleading conclusions about the effects of trade policy.

We study these mechanisms in the context of the U.S. automobile market, a setting that is particularly well-suited to analyzing the importance of global production networks. Almost half of the vehicles sold in the United States are assembled abroad, and of those assembled domestically, a substantial fraction of the parts used are imported. Furthermore, the percentage of parts imported varies across models, ranging from 25% to 90% in 2024. Consequently, how tariffs on vehicles and automotive parts affect outcomes depends in part on heterogeneity in exposure to foreign parts imports across domestic firms, along with assembly location decisions and firms' pricing power.

To quantify these effects, we develop a structural model of the U.S. automobile market that integrates differentiated-products demand with firm pricing and global input sourcing. On the demand side, we estimate a random-coefficients discrete-choice model that captures realistic substitution patterns across hundreds of vehicle models and allows for rich heterogeneity in consumer price sensitivity and preferences for vehicle attributes, following Nevo (2001), Berry, Levinsohn, and Pakes (1995), and Grieco, Murry, and Yurukoglu (2024).

On the supply side, we depart from the common assumption that the marginal costs of domestically assembled products are insulated from shocks to foreign input prices. Instead, we allow changes in the cost of imported automotive parts—driven by tariffs or exchange-rate movements—to transmit imperfectly into manufacturers' marginal costs. Rather than modeling the high-dimensional sourcing problem directly, we discipline the relevant cost channel by estimating the reduced-form pass-through of changes in the cost of foreign parts to changes in marginal cost. This pass-through (cost shocks to marginal cost) parameter serves as a sufficient statistic for the combined effect of upstream cost absorption by suppliers, medium-run adjustments in firms' sourcing decisions, and the share of marginal cost owing to parts. Exchange-rate variation provides the identifying variation for this channel, allowing us to capture how global cost shocks propagate to domestically assembled vehicles, while accounting for endogenous changes in sourcing decisions without requiring a fully specified model of parts sourcing.

Our analysis focuses on the medium run: we hold product offerings and assembly locations fixed while explicitly allowing firms to adjust prices and quantities in response to policy-induced cost changes, and accounting for potential endogenous changes in parts sourcing via our reduced-form sufficient statistic. This horizon is particularly relevant in industries such as automobiles, where model lineups and production locations are determined years in advance, but pricing and input sourcing exhibit greater medium-run flexibility[1]. Our counterfactuals are well-suited to evaluating the effects of trade policies before potential changes in production location occur.

We estimate the model using a novel model-year level dataset for the U.S. light-duty automobile market covering 2015–2024. The dataset combines sales, vehicle characteristics, and manufacturing data (assembly location and vehicle foreign content) from sources such as S&P Polk, DataOne, and the National Highway Traffic Safety Administration (NHTSA), and is supplemented with data on EV subsidies, consumer demographics, and exchange rates.

Our counterfactual analysis yields three core results. First, tariffs on imported final goods alone tend to reallocate demand toward domestically assembled products. While the price increase for domestically assembled vehicles is minimal, a 5.7 percentage point increase in

---

[1] News reporting indicates medium-term flexibility in where automakers source certain parts from. See, for example, this 2022 article from Reuters.

market share raises the producer surplus of these firms by $1 billion.[2] In this case, domestic producers benefit from protection despite higher overall prices. The consumer surplus losses of around $14 billion are offset by an equivalent increase in tariff revenue, resulting in a net US gain of $1.8 billion. Second, extending tariffs to imported intermediate inputs reverses these US gains. By raising marginal costs for domestic assemblers that rely on foreign parts, parts tariffs can erode or reverse these competitive gains, leading to higher prices and reductions in domestic producer surplus. This results in a doubling of consumer welfare losses and a reversal of US producer surplus (from a gain to a loss of $2.6 billion). However, these aggregate effects mask substantial heterogeneity among firms: while many U.S.-headquartered firms (e.g., Chevrolet, GMC, and Ford) experience losses due to their exposure to imported parts, firms with lower reliance on foreign inputs (e.g., Tesla, Jeep, Honda, and Toyota) are able to increase profits as their products benefit from lower marginal costs relative to competitors. In both scenarios, higher-income households account for a substantial share of consumer welfare losses, with the top quintile alone representing about 40%. Third, reinstating EV subsidies generates sizable welfare gains. Consumer surplus increases by approximately $4.8 billion, while U.S.-headquartered firms' producer surplus rises by about $1.8 billion. These gains are accompanied by an expansion in EV market share, total vehicle sales, and the number of vehicles assembled in the United States.

The paper contributes to several literatures in industrial organization and trade. First, we apply differentiated-products IO tools to evaluate trade policy in the automobile market while explicitly incorporating global value chains, extending analyses of the auto industry (Berry, Levinsohn, and Pakes 1995; Grieco, Murry, and Yurukoglu 2024) by allowing policy to operate through imported intermediate inputs as well as through final-goods trade. A key ingredient is new, model-level data that links every vehicle model sold in the United States to the value share of its parts imported, including the primary foreign-sourcing country. This enables us to measure exposure to trade policy and foreign cost shocks at the level at which pricing and substitution occur.

Second, we introduce a flexible supply-side pass-through component that connects foreign cost shocks to vehicle marginal costs. We estimate how marginal costs respond to shocks as a

---

[2] The market share calculation includes the outside good share in the denominator.

function of a model's imported-parts exposure.

Third, we quantify the medium-run incidence of alternative tariff designs, distinguishing between tariffs applied only to imported vehicles and tariffs applied to both imported vehicles and imported parts. In doing so, we complement recent quantitative work on automobile tariffs and subsidies (Grieco, Murry, and Yurukoglu 2024; Allcott, Kane, Maydanchik, Shapiro, and Tintelnot 2024) by emphasizing transmission through firms' exposure to input sourcing rather than through changes in conduct or long-run capacity adjustments.

Our analysis is most closely related to the independent and concurrent work of Duarte, Magnolfi, Quint, Sullivan, and Sølvsten (2025). That paper uses the demand system from Grieco, Murry, and Yurukoglu (2024), with preferences calibrated to 2018, to study firm conduct and scale economies. Our approach differs along three key dimensions. First, we capture upstream absorption and sourcing adjustments through an estimated pass-through sufficient statistic. Second, we study the current market environment using data from 2015 to 2024, estimating contemporary consumer preferences and the prevailing vehicle choice set. This is important in a rapidly evolving market: between 2018 and 2024, the SUV share increased from 40% to 50%, while EV adoption rose from negligible levels to 7.8%. Accounting for these shifts in demand and product availability is quantitatively important. Third, we incorporate a more precise representation of the current tariff regime by allowing tariffs to vary across countries rather than imposing a uniform rate.

Our findings speak to the incidence of trade policy in the presence of firm heterogeneity and fragmented production. The downstream effects of tariffs on consumers and domestic producers depend on market structure and input exposure. By embedding global sourcing within a differentiated-products equilibrium framework, we show that policies that appear protective when focusing on final goods trade can impose substantial costs when intermediate inputs are taken into account. While our empirical analysis focuses on automobiles, the mechanisms we document are likely to operate in other industries with global value chains, such as electronics, machinery, and clean energy technologies, where recent trade policies have targeted both final goods and key intermediate inputs.

The remainder of the paper proceeds as follows. Section 2 describes the data. Section 3 presents the demand model and supply-side framework. Section 4 describes the estimation procedure, and section 5 discusses the results. Section 6 reports counterfactual tariff simulations and welfare implications. Section 7 concludes.

2 Data and descriptive statistics

2.1 Data

We construct a novel model-year level dataset for the U.S. light-duty automobile market covering 2015–2024 by combining sales, vehicle characteristics, and manufacturing data. Our primary sources are sales data from S&P Polk, vehicle characteristics from DataOne, and manufacturing data—including assembly location and foreign content—from the National Highway Traffic Safety Administration (NHTSA). We also incorporate data on EV subsidies, consumer demographics, and exchange rates. This section describes each data source and outlines the procedures used to construct and merge the dataset.

Vehicle Sales Data S&P Polk provides annual data on automotive registrations for 2015–2024. The dataset reports, for each year, the total number of registered vehicles by trim (and model year) at the ZIP-code level. We construct a proxy for new sales by computing the year-over-year change in registrations within each ZIP code, restricting attention to vehicles whose model year matches the calendar year. This restriction helps limit spurious "sales" driven by households re-registering an existing vehicle after moving to a different ZIP code. However, it likely understates true new-vehicle sales because new purchases may include vehicles from prior model years.[3]

We define products at the model-year level and aggregate trim-level sales to the model level by summing sales across all trims associated with a given model. We aggregate zip-code-level data across the US. To convert sales into market shares, we define the market size as the number of US households in a given year divided by 6, reflecting an average vehicle replacement cycle of six years, following Coşar, Grieco, Li, and Tintelnot (2018) and Allcott et al. (2024).[4]

---

[3] For example, we observe 11.27 million vehicle sales for 2024, while reported new-vehicle sales were 15.85 million (WardsAuto (2025)).

[4] We also test using a denominator of 2.5, as suggested by Grieco, Murry, and Yurukoglu (2024) and Sabal (2025), and allowing for a random coefficient on mean utility for the inside good. The results are qualitatively similar.

Vehicle Characteristics Data Vehicle characteristic data are obtained from DataOne, which provides trim-level information on a wide range of attributes, including price, horsepower, weight, length, engine type, and assembly location. We convert these trim-level data to the model level by taking the median of each characteristic across all trims within a given model.

Manufacturing: Assembly Location and Vehicle Foreign Content Data The American Automotive Labeling Act (AALA), implemented under Part 583 of the US Code of Federal Regulations, requires manufacturers of consumer vehicles sold in the United States to report, among other information, the vehicle's assembly location and the share of parts content (by value) originating from the United States and Canada.[5] We obtain these data for 2015–2024 from the National Highway Traffic Safety Administration (NHTSA), which compiles manufacturers' reports on assembly locations and parts origins. These data are available for free on the NHTSA website.

Electric Vehicle Subsidy Data Federal EV purchase incentives in the United States changed substantially over the study period. The long-standing Section 30D federal tax credit, which provided up to $7,500 for eligible new plug-in vehicles, featured manufacturer-specific phase-outs once cumulative US sales exceeded 200,000 units. The Inflation Reduction Act (IRA), enacted in August 2022, restructured the program by conditioning eligibility on final assembly in North America and, beginning in 2023, on additional sourcing requirements. Vehicles acquired after September 30, 2025, are no longer eligible for the credit (Internal Revenue Service (2025)).

Additionally, beginning January 1, 2023, Section 45W credits of up to $7,500 for light-duty electric vehicles were made available to businesses. Consumers were able to access this credit via the 'leasing loophole' that is explicitly modeled in Allcott et al. (2024).

We construct a vehicle-level EV subsidy series using data from the US Department of Energy's Tax Incentive Data Services (Department of Energy (2025)). For each product, we assign the pre-2023 statutory credit for model years up to 2022, a date-weighted average of pre- and post-2023 credits for model year 2023 (reflecting the mid-April implementation of the new rules), and the post-2023 credit for model year 2024. We then apply the statutory manufacturer

---

[5] The data do not distinguish between US and Canadian parts.

phase-out schedules during the 2019–2022 period and implement the IRA assembly eligibility rules. From 2023 onwards, all vehicles that were previously eligible for a Section 30D credit are assumed to receive the Section 45W credit in full. Appendix Table B2 reports the average EV subsidy by manufacturer in each year.

Consumer Demographic Data To incorporate geographic heterogeneity while keeping the model computationally tractable, we aggregate the nine US Census divisions into six broader regions. Appendix Figure A1 maps our modified divisions and the states included in each.

We obtain household demographic data from Ruggles, Flood, Sobek, Backman, Cooper, Rivera Drew, Richards, Rogers, Schroeder, and Williams (2025), including the number of households (nationally and by division-year), household income, and household location. We use these data to construct the demographic moments described later.

Vehicle Data Matching We merge the three datasets (sales, vehicle characteristics, and manufacturing) by model name and model year. To match, we use a combination of formulaic and manual matching, achieving coverage of at least 97% of annual sales in every year of the sample. Sales prices are deflated to constant 2015 US dollars. This gives us year-model level data on sales, median vehicle characteristics, and manufacturing.

2.2 Descriptive Statistics

Sales by Assembly Location and Vehicle Type Our merged vehicle data indicate that roughly half of vehicles purchased in the United States are assembled domestically. Among imported vehicles, Figure A2 shows that in 2024, the largest source countries are Japan and Mexico. Over the sample period, Canada's share declines substantially.

Sales by Vehicle Type and Assembly Location Figure 1a plots the evolution of sales shares by vehicle type from 2015 to 2024. Two notable shifts occur over the sample period. First, electric vehicles (EVs) expand rapidly: in 2015, they account for a negligible share—less than 0.5% of purchases—rising to 7.5% by 2024. Second, SUVs become substantially more prevalent, increasing from 31% of sales in 2015 to 49% in 2024. Truck sales also gain share over this period, though the increase is more modest.

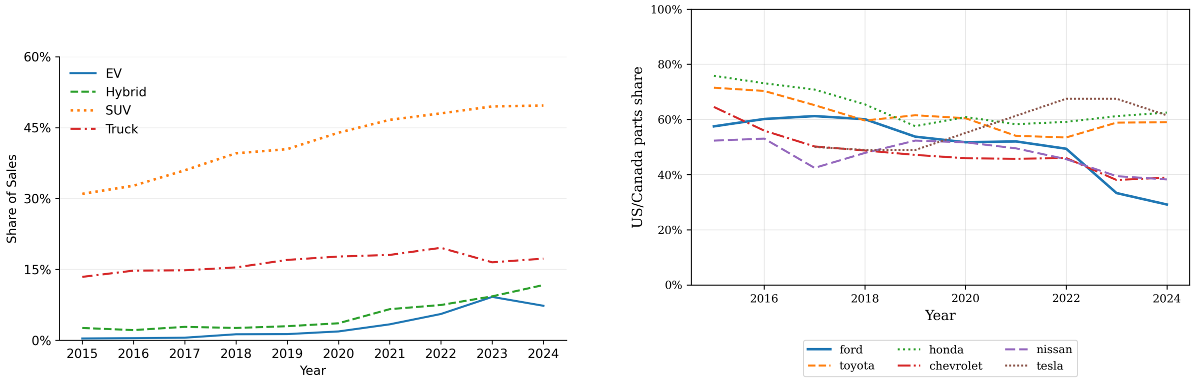

(a) Trends in sales by vehicle type

(b) Share of domestic content in domestically assembled vehicles (by value)

Figure 1: Trends: sales by vehicle type and share of domestic content

Notes: Panel (a) shows the evolution of vehicle sales by type over time. Panel (b) reports the average share of domestic (U.S. and Canadian) content in vehicles assembled in the United States, measured by value. The figures highlight shifts in demand composition and changes in the sourcing of intermediate inputs over the sample period.

Appendix Table B1 describes the matched data after the exclusion of vehicles priced above \$100,000 USD (in 2015 dollars) and vehicles with market share below 0.001% [6]. Over the sample period, 75% of cars are imported, while only 8% of trucks are foreign-assembled.

Foreign Content For vehicles assembled in the United States, the share of parts value sourced from the United States and Canada declines over the sample period, though the pace of decline varies substantially across firms. Figure 1b plots trends in the domestic (U.S. and Canadian) parts share for the five largest US automakers—defined by sales of domestically assembled vehicles—along with Tesla, which we include to represent the EV segment.

By 2024, the firms shown in Figure 1b display a pronounced bimodal distribution in domestic parts sourcing. This pattern reflects persistent differences in firms' production architectures and supply-chain configurations, which generate substantial heterogeneity in exposure to foreign input cost shocks. We exploit this cross-firm and time-series variation in input exposure to identify how changes in the cost of imported parts are transmitted into product-level marginal costs. This variation provides the key empirical leverage for isolating the cost channel through

---

[6] This trimming amounts to a loss of less than 1% of the total market share of inside goods in any given year.

which global value chains mediate the incidence of trade policy.

In addition to reporting the share of parts value originating in the United States and Canada (henceforth, "domestic"), manufacturers are required to report the source country and associated share for any foreign country whose parts account for more than 15% of total parts value. Among the reported foreign sources, Japan and Mexico account for the largest shares. Appendix figure A3 shows how the composition of reported foreign content evolves over our sample period. Because reporting is subject to the 15% threshold, we do not observe the full set of foreign-sourced parts: depending on the year, the source country is unreported for between 17% (2016) and 27% (2024) of parts value. The AALA data also have incomplete coverage across models, requiring us to impute missing values for domestic and foreign parts shares.[7] When a firm or model is missing observations for a subset of years, we impute in three steps. First, if the model is observed in adjacent year(s), we assign the closest available value (or the average of the nearest observations on either side). Second, if the model is never observed, we use the make-level average domestic parts share, using the closest year(s) in which that make-level average is available. Finally, if a make is entirely unobserved, we assign the overall average domestic parts share among U.S.-assembled vehicles (this occurs only for Rivian).

3 Model

3.1 Demand

We model US consumer demand for vehicles using a random-coefficients discrete choice framework following Berry, Levinsohn, and Pakes (1995) (BLP). A key advantage of this approach is that it allows consumers to have heterogeneous preferences over vehicle attributes—such as price, size, performance, and powertrain type (e.g., EV vs. ICE)—which generates flexible substitution patterns across products and heterogeneous willingness-to-pay. We estimate the demand-side parameters using observed market shares (and additional demographic moments described below) and match them to the shares implied by the model.

---

[7] One firm, Rivian, does not appear in the data at all.

3.1.1 Consumer Choices and Market Shares

In each year t, consumer $i \in \mathcal{I}$ chooses among products $j \in \mathcal{J}_t$ and purchases the product that delivers the highest utility. Consumers may also choose the *outside good*—that is, make no vehicle purchase in year t—which we normalize to have utility zero and include in $\mathcal{J}_t$ with index $j = 0$.

We model consumer-specific utility as follows. In each market (year) t, consumer i derives utility from vehicle j given by:

$$u_{ijt} = \phi_{it} + \beta_i' \mathbf{X}_{jt} + \alpha_{it}\big(p_{jt} - \tau_{jt}\big) + \xi_{jt} + \varepsilon_{ijt}, \tag{1}$$

where ϕ_{it}, $\boldsymbol{\beta}_i$, and α_{it} denote consumer-specific preferences for the inside good, observed product characteristics $\mathbf{X}_{jt}$, and price, respectively. The price faced by the consumer is the manufacturer's suggested retail price (MSRP) p_{jt}, set by the firm, net of any applicable subsidy for electric or hybrid vehicles τ_{jt}.[8] The term ξ_{jt} captures an unobserved product-year utility shock common across consumers, and ε_{ijt} is an idiosyncratic consumer–product-specific utility shock assumed to follow a Type I extreme value distribution.

The vector of observed characteristics $\mathbf{X}_{jt}$ includes horsepower, vehicle footprint, miles per gallon, indicators for vehicle type (car, truck, SUV, van), and powertrain type (ICE, EV, hybrid), as well as brand fixed effects and indicators for luxury and European brands. We also include interactions between year and EV, SUV, and hybrid indicators to capture evolving consumer preferences for these vehicle types, as well as time-varying quality changes—particularly for EVs and hybrids—arising from factors not fully captured by observed characteristics, such as improvements in charging infrastructure.

To capture heterogeneity in consumer tastes, we allow both observed demographics and

---

[8] Following Allcott et al. (2024), we assume that each EV transaction reflects the full \$7,500 federal tax credit. This abstracts from income-based eligibility constraints as well as domestic content requirements that may bind for direct purchases but not for leases (the so-called "leasing loophole"). This assumption is consistent with modeling the price faced by a marginal consumer who is leasing a vehicle. We also note that our identification strategy instruments for price using exchange rate movements, so any discrepancy between statutory and realized subsidies does not affect our estimates, provided it is orthogonal to the instrument.

unobserved preference heterogeneity to affect preferences for vehicle characteristics. Specifically:

$$\phi_{it} = \bar{\phi}_t + \phi \cdot \log(\text{income}_i), \tag{2}$$

$$\beta_{ik} = \bar{\beta}_k + \sigma_k \nu_{ik} + \sum_d \pi_{kd} \mathbf{1}\{D_i = d\}, \tag{3}$$

$$\alpha_{it} = \bar{\alpha} + \alpha \cdot \log(\text{income}_i), \tag{4}$$

where ν_{ik} captures unobserved taste heterogeneity, D_i denotes consumer demographic group membership, and k indexes vehicle characteristics. The terms ν_{ik} are i.i.d. consumer-specific taste shocks drawn from a standard normal distribution. The variable D_i denotes the U.S. regional division in which consumer i resides, and $\mathcal{T}$ denotes the subset of characteristics for which we allow region-specific preference shifters (SUV, truck, and EV indicators). Not all characteristics are permitted to exhibit unobserved heterogeneity or demographic interactions; for characteristics without such variation, the corresponding parameters σ_k and/or π_{kd} are set to zero.[9]

Berry, Levinsohn, and Pakes (1995) shows that, under the Type I extreme value assumption on ε_{ijt}, market shares are obtained by integrating the individual-level logit choice probabilities over the distribution of consumer heterogeneity. Accordingly, the market share of product j in market t is given by:

$$s_{jt} = \int \frac{\exp(\phi_{it} + \beta_i' \mathbf{X}_{jt} + \alpha_{it}(p_{jt} - \tau_{jt}) + \xi_{jt})}{1 + \sum_{k \in \mathcal{J}_t \setminus \{0\}} \exp(\phi_{it} + \beta_i' \mathbf{X}_{kt} + \alpha_{it}(p_{kt} - \tau_{kt}) + \xi_{kt})} \, dF(i), \tag{5}$$

where $F(i)$ denotes the joint distribution of consumer-specific attributes.

3.1.2 Firm Pricing and Marginal Costs

Under the assumption of static Nash-Bertrand pricing, the estimated demand system allows us to recover product-level marginal costs and markups, following Berry, Levinsohn, and Pakes (1995).

Equation (5) expresses market shares $s_t(\mathbf{p}_t)$ as a function of the vector of prices set by firms

---

[9] We divide the United States into six regional divisions: Northeast, North Central, South Pacific, South Central, Mountain, and Pacific.

in market t. Let $s_t(\mathbf{p}_t)$ denote the $J_t \times 1$ vector of product market shares, and let $\frac{\partial s_t}{\partial \mathbf{p}_t'}$ denote the $J_t \times J_t$ Jacobian matrix of shares with respect to prices. This Jacobian is computed from the random-coefficients demand system and integrated over consumer heterogeneity. Firms are assumed to jointly maximize profits across all products in their portfolios.

We define the ownership matrix Ω_t as:

$$(\Omega_t)_{jk} = \begin{cases} 1 & \text{if products } j \text{ and } k \text{ are owned by the same firm in market } t, \\ 0 & \text{otherwise.} \end{cases}$$

For each market t, the first-order condition for profit maximization is

$$s_t(\mathbf{p}_t) \; + \; \left[\Omega_t \odot \frac{\partial s_t(\mathbf{p}_t)}{\partial \mathbf{p}_t'} \right] (\mathbf{p}_t - \mathbf{c}_t) \; = \; \mathbf{0}, \tag{6}$$

where $\odot$ denotes the element-wise (Hadamard) product, and s_t, $\mathbf{p}_t$, and $\mathbf{c}_t$ are J_t-dimensional vectors of market shares, prices, and marginal costs, respectively.

In practice, the derivatives $\partial s_t / \partial \mathbf{p}_t'$ are computed numerically using `PyBLP`, and marginal costs are recovered by inverting equation (6).

3.2 Supply

Having recovered product-level marginal costs from firms' pricing behavior, we study how exposure to foreign intermediate inputs shapes the cost channel through which trade policy affects equilibrium outcomes. In markets with fragmented production, tariffs on intermediate inputs can raise the marginal costs of domestically assembled goods, altering firms' competitive positions and the incidence of protection. Our objective in this section is to characterize how shocks to the cost of imported parts transmit into marginal costs for vehicles assembled in the United States, and how this transmission varies with firms' reliance on foreign inputs.

Because we are interested in estimating the effect of US tariffs on vehicle costs, we design the supply-side analysis to identify how changes in the prices of foreign parts affect the marginal costs of vehicles assembled in the United States, allowing the effect to vary with each vehicle's exposure to imported inputs. We model marginal costs only for domestically assembled vehicles. For these vehicles, we observe (i) the share of parts value sourced from the United States or

Canada and (ii) the share of parts value imported from specific foreign countries when the value imported from a given country exceeds 15% of total parts value.

In this section, we define a product j as a make–model pair observed over multiple years (e.g., *Ford Bronco*), while allowing product characteristics to vary over time. This definition is required for our empirical specification and implies that we drop products that do not appear—or cannot be reliably matched—across multiple years.

Let $\rho_{d,jt}$ denote the share of parts value sourced domestically and $\rho_{f,jt}$ denote the share sourced from foreign suppliers. Since $\rho_{d,jt} + \rho_{f,jt} = 1$, it is sufficient to characterize sourcing using $\rho_{f,jt}$. We assume that firms choose their sourcing strategies to minimize the cost of producing a vehicle with a given set of characteristics. While we do not observe the inputs or prices that directly enter this decision, we do observe the resulting domestic and foreign sourcing shares. This sourcing choice is a function of many unobserved factors, including the relative prices of domestic and foreign parts. We denote the firm's optimal sourcing decision as:

$$\rho^*_{jt} = \left(\rho^*_{d,jt}, \rho^*_{f,jt}\right).$$

Estimating the supply side raises several challenges, including unobserved input prices and endogenous sourcing decisions. We discuss these challenges and our estimation strategy in the next section.

A natural extension of our framework would be to fully endogenize firms' input sourcing decisions by explicitly modeling the choice between domestic and foreign parts. We do not pursue this approach for two related reasons. First, full structural modeling of parts sourcing would require specifying firms' optimization problems over a large number of heterogeneous intermediate inputs, each sourced from multiple potential locations under distinct contractual, technological, and logistical constraints. In the automobile industry, a single vehicle incorporates thousands of components, often governed by long-term supplier relationships, platform-specific compatibility, and regulatory requirements. Modeling these decisions explicitly would require detailed data on part-level prices, substitution elasticities, and contractual rigidities that are not observed, and would necessitate strong functional-form assumptions that are difficult to validate empirically. Second, and most importantly, fully modeling sourcing choices is not necessary to answer the core incidence question that motivates this paper. For the purpose of evaluating how tariffs affect equilibrium prices, profits, and welfare, what matters is not the detailed mechanics of sourcing decisions per se, but the extent to which shocks to foreign input costs are transmitted into firms' marginal costs. We therefore summarize firms' sourcing responses using a reduced-form pass-through parameter that maps foreign cost shocks into marginal costs as a function of observed input exposure. This parameter serves as a sufficient statistic for the combined effects of upstream cost absorption and medium-run sourcing adjustment within existing production structures.

By estimating this pass-through using exchange-rate variation and lagged input exposure, our approach allows sourcing responses to influence equilibrium outcomes while avoiding the need to fully specify the high-dimensional sourcing problem. In this sense, we discipline the relevant cost channel directly, rather than imposing additional structure on firms' sourcing decisions that would be weakly identified. This modeling choice allows us to remain agnostic about the precise mechanisms underlying sourcing adjustment while still capturing their equilibrium implications for tariff incidence.

4 Estimation

We estimate the demand parameters following Berry, Levinsohn, and Pakes (2004) and Petrin (2002) using the generalized method of moments (GMM). Identification comes from matching simulated outcomes to their observed counterparts along three dimensions: (i) product-level market shares, (ii) demographic–characteristic micro moments, and (iii) second-choice micro moments.

To implement this approach, we partition the vector of demand parameters, denoted by θ, into linear parameters that enter mean utility:

$$\theta_1 = (\bar{\phi}_t, \bar{\beta}_k, \bar{\alpha}),$$

and nonlinear parameters that govern preference heterogeneity:

$$\theta_2 = (\phi, \alpha, \sigma_k, \pi_{kd}).$$

We can rewrite equation (1) as:

$$u_{ijt} = \delta_{jt} + \mu_{ijt}(\theta_2) + \varepsilon_{ijt}, \tag{7}$$

where:

$$\delta_{jt} = \bar{\phi}_t + \bar{\beta}'\mathbf{X}_{jt} + \bar{\alpha}(p_{jt} - \tau_{jt}) + \xi_{jt},$$

and $\mu_{ijt}(\theta_2)$ captures the remaining nonlinear components of utility driven by consumer heterogeneity.

Our estimation procedure proceeds as follows and is implemented using the `PyBLP` package developed by Conlon and Gortmaker (2020), with micro-moment functionality described in Conlon and Gortmaker (2025). For any candidate value of the nonlinear parameter vector θ_2, we first invert observed market shares to recover mean utilities $\delta(\theta_2)$ via contraction mapping. Conditional on $\delta(\theta_2)$, the linear mean-utility parameters θ_1 are then concentrated out using an instrumental variables regression of $\delta(\theta_2)$ on X_1, with instruments Z used to address price endogeneity. The resulting implied demand shocks $\xi(\theta_2)$ enter the GMM moment conditions, which are stacked together with the micro moments to form the full GMM objective function.

4.1 Simulated Shares and Mean Utility Inversion

For any candidate value of θ_2, the model implies a mapping from mean utilities to predicted market shares. We approximate predicted market shares using simulation. Specifically, we draw a set of agents $\{(\nu_i, \text{income}_i, D_i)\}_{i=1}^{I}$ from the empirical distribution of consumer demographics and the assumed distribution of unobserved heterogeneity. Draws of ν_i are generated using Halton sequences (Halton 1960), following Nevo (2001).

Simulated market shares are computed as:

$$\tilde{s}_{jt}(\delta, \theta_2) \;=\; \sum_{i=1}^{I} w_{it}\, P_{ijt}(\theta_2), \tag{8}$$

where:

$$P_{ijt}(\theta_2) = \frac{\exp(\delta_{jt} + \mu_{ijt}(\theta_2))}{\sum_{k\in\mathcal{J}_t} \exp(\delta_{kt} + \mu_{ikt}(\theta_2))} \tag{9}$$

is the probability that consumer i in year t chooses product j. In each year, we draw 400 simulated consumers per regional division and weight each draw using w_{it} to match the empirical demographic distribution.

Following Berry, Levinsohn, and Pakes (1995), mean utilities $\delta(\theta_2)$ are recovered via contraction mapping[10]:

$$\delta_{jt} \leftarrow \delta_{jt} + \log s_{jt} - \log \tilde{s}_{jt}(\delta, \theta_2). \tag{10}$$

4.2 Mean utility projection, unobserved quality, and the price instrument

Mean utility. With δ in hand, we can recover the mean utility parameters and the unobserved demand shock by projecting mean utility on observed product characteristics:

$$\delta_{jt}(\theta_2) \ = \ \hat{\bar{\phi}}_t + \hat{\bar{\beta}}\mathbf{X}_{jt} + \hat{\bar{\alpha}}(p_{jt} - \tau_{jt}) + \xi_{jt}(\theta_2). \tag{11}$$

However, we must account for price endogeneity. A product's price is likely correlated with its unobserved quality, as firms are likely to charge higher prices for better products. All other characteristics are assumed to be exogenous.

Price instrument. Our price instrument is the lagged real exchange rate (RER) between the United States and the country in which the vehicle was assembled. This choice follows the example of Grieco, Murry, and Yurukoglu (2024), who argue that the real exchange rate, lagged to account for planning timelines, captures shifts in manufacturing costs and firm pricing decisions, making it a valid instrument for prices. To construct the series of real exchange rates between the US and each source country, we use data from the International Monetary Fund's International Financial Statistics database, collated by the World Bank (International Monetary Fund (2025)).

We demonstrate the strength of the price instrument by estimating a homogeneous (representative-consumer) discrete-choice demand model at the model–market level. For product j in market t. Define the mean utility as the difference between the log of product j's share and the log of the outside option's share:

$$\delta_{jt}^{hom} = \log s_{jt} - \log s_{0,t}.$$

---

[10] The iteration proceeds until the norm of the update falls below 10^{-13}.

We project δ_{jt}^{hom} onto the non-consumer specific coefficients of our demand model:

$$\delta_{jt}^{hom} = \phi_t + \beta' \mathbf{X}_{jt} + \alpha p_{jt} + \gamma_f + \epsilon_{jt}, \tag{12}$$

where γ_f are firm fixed effects and ϕ_t are year fixed effects, and epsilon is the error term.

Table 1 supports the choice of instrument by comparing OLS and IV estimation of equation (12). We see that the coefficient on price becomes notably more negative, as expected for downward-sloping demand, and the F-statistic is indicative of a strong instrument.

Table 1: Demand Estimates: OLS vs. IV (2SLS)

	OLS	**IV (2SLS)**
Price	-0.333 $(0.038)^{***}$	-1.275 $(0.258)^{***}$
HP/Weight	-0.619 $(0.237)^{*}$	1.365 $(0.564)^{**}$
Size	2.799 $(0.404)^{***}$	7.864 $(1.417)^{***}$
MPG	0.456 $(0.204)^{*}$	-0.398 (0.297)
SUV	0.925 $(0.062)^{***}$	0.979 $(0.067)^{***}$
Truck	0.385 $(0.084)^{**}$	-0.902 $(0.379)^{**}$
Van	-0.699 $(0.116)^{***}$	-2.109 $(0.421)^{***}$
Hybrid	0.427 $(0.198)^{*}$	-1.236 $(0.508)^{**}$
ICE	1.394 $(0.321)^{**}$	-0.760 (0.647)
Fixed effects	Year and Firm	
Instrument	—	Lag of RER
F-Statistic	—	47.78
Implied own-price elasticity	—	-5.12

Notes: Entries are coefficients with standard errors in parentheses. Standard errors are clustered by market. Fixed effects are omitted. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Macro Moments The real exchange rate, together with exogenous product characteristics $\mathbf{X}_{jt}$, is stacked to form the instrument vector Z_{jt}, which allows equation (11) to be estimated using instrumental variables. Conditional on θ_2, the estimated coefficients then permit recovery of the unobserved demand shock:

$$\xi_{jt}(\theta_2) = \delta_{jt}(\theta_2) - \hat{\hat{\phi}}_t - \hat{\hat{\beta}}(\theta_2)' \mathbf{X}_{jt} - \hat{\hat{\alpha}}(\theta_2)(p_{jt} - \tau_{jt}).$$

We construct the first set of moments entering the GMM objective by imposing the standard BLP orthogonality condition that unobserved demand shocks are mean independent of the instruments:

$$g^{\text{quality}}(\theta) \; = \; \frac{1}{N} \sum_{t} \sum_{j} Z_{jt} \, \xi_{jt}(\theta_2), \tag{13}$$

where N denotes the number of product–market observations.

4.3 Demographic Micro Moments

To help discipline consumer taste heterogeneity and the implied substitution patterns, we incorporate two types of microdata: (i) moments linking consumer demographics to product choices, and (ii) moments capturing the correlation between the characteristics of consumers' first and second choice vehicles.

The demographic micro moments match simulated analogues to empirical expectations of (i) the probability that a consumer purchases any vehicle, (ii) the probability that a consumer purchases an EV, SUV, or truck, and (iii) the expected expenditure on a purchased vehicle, each conditional on consumer demographics.

For each value of θ, we compute the simulated analogue of the observed moments, denoted $\tilde{m}(\theta)$, and compare it to its empirical counterpart $\hat{m}$. Intermediate calculations required to construct these moments are denoted by $v(\theta)$.

As an illustration, we describe the construction of the expected difference in purchase prices between consumers in each income quintile relative to the first quintile.[11] Let Q_i denote consumer i's income quintile, with $q \in \mathcal{Q} = \{1, 2, 3, 4, 5\}$. We simulate the expected purchase price for consumers in income quintile q, conditional on purchasing an inside good, as:

$$v_{p,t}(\theta, q) = \frac{\sum_i w_{it} \mathbf{1}\{Q_i = q\} \sum_{j \neq 0} p_{jt} \, P_{ijt}(\theta)}{\sum_i w_{it} \mathbf{1}\{Q_i = q\} \sum_{j \neq 0} P_{ijt}(\theta)}, \tag{14}$$

where the restriction $j \neq 0$ excludes the outside good. The simulated moment for the difference in prices paid by consumers in the third and first income quintiles is then:

$$\tilde{m}_{p,t,3}(\theta) = v_{p,t}(\theta, 3) - v_{p,t}(\theta, 1). \tag{15}$$

---

[11] Here, "price paid" corresponds to the pre-subsidy MSRP p_{jt}, since the data used to construct the observed price–demographic moments do not include EV subsidies, which consumers receive as tax credits.

The corresponding empirical micro moment is defined as:

$$m_{p,t,3} = \mathbf{E}_t[\text{price} \mid Q_i = 3, j \neq 0] - \mathbf{E}_t[\text{price} \mid Q_i = 1, j \neq 0]. \tag{16}$$

This demographic micro moment enters the vector $g^{\text{demo}}(\theta)$, together with analogous moments capturing differences across income quintiles in vehicle purchase probabilities and the probability of purchasing an EV, SUV, or truck conditional on region of residence. Formally:

$$g^{\text{demo}}(\theta) = \hat{m}^{\text{demo}} - \tilde{m}^{\text{demo}}(\theta). \tag{17}$$

4.4 Second-Choice Moments

Second-choice micro moments use survey evidence on consumers' first- and second-choice vehicles to provide additional information on substitution patterns across products. We incorporate these moments by matching the simulated correlation between the characteristics of consumers' first and second choices to their empirical counterparts.

Because we do not have direct access to the underlying survey microdata, we instead use published second-choice moments for the U.S. automotive market from Grieco, Murry, and Yurukoglu (2024) for 2015 and Allcott et al. (2024) for 2022.[12] Since the second-choice data record only inside-good alternatives, all second-choice moments are conditional on an inside-good purchase.

For a candidate parameter vector θ, let $P^{in}_{ijt}(\theta_2)$ denote the model-implied probability that consumer i chooses inside good j in market t, and let $P^{in}_{ih(-j)t}(\theta_2)$ denote the probability that the same consumer would choose product h when product j is removed from the inside-good choice set. These probabilities are given by:

$$P^{in}_{ijt}(\theta_2) = \frac{\exp(\delta_{jt} + \mu_{ijt}(\theta_2))}{\sum_{\ell \in \mathcal{J}_t \setminus \{0\}} \exp(\delta_{\ell t} + \mu_{i\ell t}(\theta_2))}, \tag{18}$$

and,

$$P^{in}_{ih(-j)t}(\theta_2) = \frac{\exp(\delta_{ht} + \mu_{iht}(\theta_2))}{\sum_{\ell \in \mathcal{J}_t \setminus \{j,0\}} \exp(\delta_{\ell t} + \mu_{i\ell t}(\theta_2))}. \tag{19}$$

The implied joint probability that consumer i has first choice j and second choice h is

---

[12] For computational simplicity, we apply the Allcott et al. (2024) moments only to the 2022 market.

therefore:

$$P^{in}_{i,(j,h)t}(\theta) = P^{in}_{ijt}(\theta)\, P^{in}_{ih(-j)t}(\theta), \qquad j \in \mathcal{J}_t \setminus \{0\},\ h \in \mathcal{J}_t \setminus \{j, 0\}, \tag{20}$$

where $\mathcal{J}_t$ denotes the set of products available in market t, and the outside good $j = 0$ is excluded.

For each vehicle characteristic x_{jt} for which second-choice moments are available from Grieco, Murry, and Yurukoglu (2024) in the 2015 market (horsepower, miles per gallon, SUV indicator, truck indicator, luxury brand indicator, and European brand indicator), we compute the model-implied correlation between first- and second-choice characteristics, denoted $\widetilde{\mathrm{Corr}}_\theta(x_{\text{first}}, x_{\text{second}})$.

Given the joint distribution over first–second choice pairs, we compute these correlations by treating (x_{jt}, x_{ht}) as a bivariate random variable with weights proportional to $P^{in}_{i,(j,h)t}(\theta)$. Define normalized weights:

$$\omega_{i,jh,t}(\theta) = w_{it}\, P^{in}_{i,(j,h)t}(\theta), \tag{21}$$

where normalization follows from $\sum_i w_{it} = 1$ and $\sum_{j \in \mathcal{J}_t \setminus \{0\}} \sum_{h \in \mathcal{J}_t \setminus \{j,0\}} P^{in}_{i,(j,h)t}(\theta) = 1$.

Let $x^{(1)}_{i,jh,t}$ denote the value of characteristic x for consumer i's first-choice vehicle and $x^{(2)}_{i,jh,t}$ the corresponding value for the second-choice vehicle. The model-implied correlation is then:

$$\widetilde{\mathrm{Corr}}_\theta(x_{\text{first}}, x_{\text{second}}) = \frac{\sum_{i,j,h} \omega_{i,jh,t}(\theta)\big(x^{(1)}_{i,jh,t} - \bar{x}^{(1)}(\theta)\big)\big(x^{(2)}_{i,jh,t} - \bar{x}^{(2)}(\theta)\big)}{\sqrt{\sum_{i,j,h} \omega_{i,jh,t}(\theta)\big(x^{(1)}_{i,jh,t} - \bar{x}^{(1)}(\theta)\big)^2}\sqrt{\sum_{i,j,h} \omega_{i,jh,t}(\theta)\big(x^{(2)}_{i,jh,t} - \bar{x}^{(2)}(\theta)\big)^2}}, \tag{22}$$

where the weighted means are given by

$$\bar{x}^{(1)}(\theta) = \sum_{i,j,h} \omega_{i,jh,t}(\theta)\, x^{(1)}_{i,jh,t}, \qquad \bar{x}^{(2)}(\theta) = \sum_{i,j,h} \omega_{i,jh,t}(\theta)\, x^{(2)}_{i,jh,t}. \tag{23}$$

All sums are taken over consumers i, first-choice products $j \in \mathcal{J}_t \setminus \{0\}$, and second-choice products $h \in \mathcal{J}_t \setminus \{j, 0\}$.

The corresponding second-choice micro moment is defined as the difference between the observed target correlation and its simulated analogue:

$$g^{sc}_{2015}(\theta) = \widehat{\mathrm{Corr}}(x_{\text{first}}, x_{\text{second}}) - \widetilde{\mathrm{Corr}}_\theta(x_{\text{first}}, x_{\text{second}}). \tag{24}$$

For the 2022 moments borrowed from Allcott et al. (2024), specifically, the share of EV buyers whose second choice is also an EV and the share whose second choice is an EV of the same vehicle type—we instead construct match probabilities.[13] These moments take the form:

$$g^{sc}_{2022} = \widehat{\text{Pr}}[\text{second choice EV} \mid \text{first choice EV}] - \widetilde{\text{Pr}}_{\theta}[\text{second choice EV} \mid \text{first choice EV}]. \quad (25)$$

4.4.1 Construction of Moments

Before defining the GMM objective function, we describe how the empirical moments used for estimation are constructed.

We collect US household demographic data from Ruggles et al. (2025), including the number of total households in the US (and in each division) in each year, household incomes, and household location. The number of households in the US defines the numerator of our market size in each year. [14] To generate our representative agents, we randomly sample 400 household incomes from each division in each year and define the agent's weight as $w_i = \frac{1}{400}\frac{\text{num. households in division}_t}{\text{num. households in US}_t}$.

We construct several sets of micro-moments to discipline heterogeneity in purchase behavior and EV adoption. For income-dependent micro-moments, we target (i) the average transaction price paid conditional on income quintile and (ii) the probability of purchasing a vehicle conditional on income quintile. To discipline geographic heterogeneity in EV demand, we additionally target a *division-dependent* micro-moment: the EV adoption rate (EV share of sales) by Census division. Details on the construction of all moments are provided in Appendix C.0.1.

4.5 GMM Objective and Optimization

The GMM objective function is constructed by stacking the macro moments and the two sets of micro moments into a single vector:

$$g(\theta) = \begin{bmatrix} g^{\text{quality}}(\theta) \\ g^{\text{demo}}(\theta) \\ g^{\text{sc}}(\theta) \end{bmatrix}. \quad (26)$$

---

[13] This choice aligns the simulated moments with the empirical definitions reported in the source.

[14] Recall we define market size as total households / 6.

We estimate the demand parameters by minimizing the associated quadratic form:

$$\hat{\theta} \;=\; \arg\min_{\theta} \; g(\theta)^{\top} W g(\theta). \tag{27}$$

The `PyBLP` package implements a two-step GMM procedure in which the weighting matrix W is updated after an initial minimization. In the first step, we use an initial positive definite weighting matrix. In the second step, the weighting matrix is set equal to the inverse of a consistent estimate of the moment covariance matrix:

$$W \;=\; \widehat{\mathrm{Var}}\big(g(\hat{\theta}^{(1)})\big)^{-1}, \tag{28}$$

where $\widehat{\mathrm{Var}}(\cdot)$ is computed using a cluster-robust estimator at the model level.[15]

The complete set of micro moments, along with their empirical targets and simulated analogues, is reported in Table 3.

4.6 Supply-side regression

This section outlines how we estimate how shocks to the cost of foreign intermediate inputs are transmitted into the marginal costs of vehicles assembled in the United States. We do not model firms' input sourcing decisions explicitly. Instead, we discipline the relevant cost channel by estimating the pass-through of foreign input cost shocks into marginal costs as a function of each vehicle's exposure to imported parts. This approach allows medium-run sourcing adjustment and upstream cost absorption to influence equilibrium outcomes without imposing additional structure on firms' high-dimensional sourcing problems.

Our empirical specification faces three primary challenges. First, contemporaneous sourcing shares are endogenous to relative input prices: firms may adjust sourcing in response to cost shocks within the year. Second, prices of individual automotive parts are unobserved, requiring a proxy for foreign input cost shocks. Third, available data report foreign sourcing locations only when imports from a given country exceed a reporting threshold, limiting the dimensionality of observed exposure. We address these challenges.

First, contemporaneous sourcing shares are endogenous to relative input prices: a domestic

---

[15] We solve the GMM objective with a convergence tolerance of 10^{-6}.

or foreign cost shock can mechanically change the value of imported parts, and firms may also re-optimize sourcing within the year in response. To mitigate this concern, we measure exposure using lagged sourcing shares. Specifically, for each domestically assembled vehicle j, we use the share of parts value sourced from its primary foreign supplier in the previous year, $\rho_{j,t-1}$. This lagged measure captures predetermined cross-model differences in reliance on foreign inputs while reducing endogeneity from contemporaneous sourcing adjustments.

Second, we do not observe prices for individual automotive parts, so we require a proxy for foreign input cost shocks. We use the bilateral real exchange rate (RER) between the United States and the vehicle's sourcing country as a summary measure of source-country cost competitiveness. Exchange-rate movements shift the dollar cost of imported inputs and, conditional on fixed effects and observed characteristics, are plausibly orthogonal to unobserved shocks to vehicle demand and domestic production costs. Under this interpretation, exchange-rate variation provides quasi-exogenous cost shocks whose transmission into marginal costs depends on a vehicle's imported-parts exposure.

Third, the data report foreign sourcing locations only when imports from a country exceed a reporting threshold (15% of total parts value), which limits the dimensionality of observed exposure. In practice, for nearly all vehicles, only one foreign sourcing country exceeds this threshold. We therefore focus on each vehicle's primary foreign sourcing location. To preserve a stable mapping between exposure and exchange-rate variation, we exclude the 14 U.S.-assembled vehicles for which the primary sourcing country changes over time. We denote the share of parts value sourced from this primary foreign location as $\rho^{(1)}_{f,j,t}$ and the corresponding bilateral real exchange rate as $RER^{(1)}_{jt}$.

With these considerations in mind, we estimate the following specification for domestically assembled vehicles:

$$\log(mc_{jt}) = \gamma' \log(X_{jt}) + \eta \left(\rho^{(1)}_{f,j,t-1} \cdot \log(RER^{(1)}_{jt})\right) + \lambda_t + \lambda_j + \epsilon_{jt}, \tag{29}$$

where mc_{jt} denotes marginal cost recovered from the pricing first-order conditions, X_{jt} is a vector of observed vehicle characteristics, $\rho^{(1)}_{f,j,t-1}$ is the lagged share of parts value sourced from the vehicle's primary foreign supplier, and $RER^{(1)}_{jt}$ is the corresponding bilateral real exchange rate. Year fixed effects λ_t absorb common cost shocks affecting all vehicles, while product fixed effects λ_j control for time-invariant differences in production costs across models.

The coefficient of interest, η, governs the pass-through of foreign input cost shocks into marginal costs, scaled by exposure to imported inputs. This parameter summarizes the extent to which a given percentage change in the cost of foreign parts translates into higher production costs for domestic assemblers. Importantly, η reflects the net effect of two mechanisms: direct cost transmission from foreign suppliers and medium-run adjustments in sourcing or supplier pricing that attenuate these shocks. As such, η serves as a sufficient statistic for the responsiveness of marginal costs to foreign input price movements within existing production structures.

Our identification relies on three assumptions. First, exchange-rate movements affect marginal costs primarily through their impact on the dollar price of imported inputs. Second, conditional on fixed effects and observed characteristics, exchange-rate fluctuations are orthogonal to unobserved shocks to vehicle-level marginal costs. Third, lagged input exposure is predetermined with respect to contemporaneous exchange-rate movements, ruling out anticipatory sourcing responses. Under these assumptions, the estimated pass-through captures the economically relevant cost channel through which tariffs on intermediate inputs affect equilibrium prices and welfare.

In Section 6, we use this estimated pass-through to evaluate how alternative tariff regimes propagate through prices, quantities, and welfare in equilibrium.

5 Results

5.1 Demand Results

5.1.1 Parameter Estimates

We present our parameter estimates in Table 2. The estimated β coefficients display sensible signs for the continuous attributes of horsepower and vehicle size, as expected, although the coefficients on miles per gallon are estimated imprecisely.[16] The mean taste coefficients for

---

[16] We split the miles-per-gallon characteristic between EVs and ICE/hybrid vehicles, as the second-choice moments from 2015 are based on a market in which EVs were largely absent.

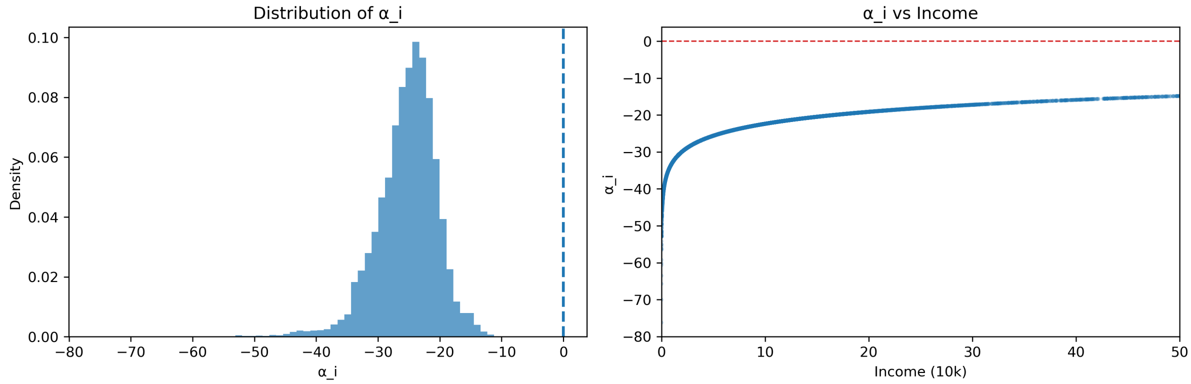

Figure 2: Distribution of consumer distaste for prices

trucks, vans, and EVs are all strongly negative; however, the large estimated σ parameters for these attributes—along with SUVs—indicate substantial heterogeneity in preferences for these vehicle types. The year fixed effects for EVs indicate that the average utility of electric vehicles has increased markedly over the sample period.

The demographic interactions reported in Panel B indicate that higher-income consumers have a stronger taste for vehicles and are less price sensitive, as expected. These coefficients, combined with the observed income distribution, imply a distribution of price disutility shown in Figure 2. The regional interaction terms further show that residents of the Pacific region exhibit a relatively strong preference for electric vehicles, while residents of the North Central, South Central, and Mountain regions display stronger preferences for trucks.

Our simulated micro moments closely match their empirical counterparts. The moments with the weakest fit are those capturing differences in purchase probabilities and expected purchase prices across income quintiles. This is not surprising, given that the model includes only a single coefficient on log income to govern heterogeneity in preferences for the inside good and price sensitivity. The observed and simulated micro moments are reported in Table 3.

5.1.2 Elasticities and Markups

Our demand system implies a share-weighted mean own-price elasticity of -6.764, meaning that a 1% increase in prices leads to a 6.76% reduction in sales, and a market elasticity of -0.420 in 2024. Our estimate of the own-price elasticity is in line with the literature, lying between the estimate of -5.06 reported by Grieco, Murry, and Yurukoglu (2024) and the range reported for

Table 2: BLP demand estimates with income and regional heterogeneity

Panel A. Std. devs of random coefficients σ		**Panel B. Demographic heterogeneity π**	
	Coef (s.e.)		Coef (s.e.)
$\log(\text{mpg}_{\text{ICE/Hyb}})$	5.753 (1.049)	Intercept $\times \log(\text{income}_{10k})$	3.935 (0.931)
log(hp)	1.957 (0.184)	Price$-$subsidy $\times \log(\text{income}_{10k})$	4.679 (1.563)
Van	6.581 (1.672)	Truck × North Central	3.512 (2.051)
Truck	12.456 (3.192)	Truck × South Central	3.229 (2.103)
SUV	2.640 (0.323)	Truck × Mountain	3.595 (2.806)
EV	1.996 (0.987)	SUV × North East	0.625 (0.531)
Euro brand	2.074 (0.392)	SUV × North Central	0.553 (0.505)
Luxury brand	2.707 (0.468)	EV × North East	0.893 (1.056)
		EV × South Atlantic	0.693 (1.113)
		EV × Mountain	1.290 (1.383)
		EV × Pacific	3.204 (1.059)

Panel C. Mean tastes β and EV×year interactions

Mean tastes β		EV×year interactions	
Variable	Coef (s.e.)	Term	Coef (s.e.)
Price$-$subsidy	-33.147 (5.205)	EV ×2016	1.201 (0.972)
log(size)	1.428 (0.139)	EV ×2017	-0.206 (0.951)
log(hp)	1.240 (0.347)	EV ×2018	1.678 (1.024)
$\log(\text{mpg}_{\text{ICE/hyb}})$	-1.147 (0.930)	EV ×2019	1.348 (0.795)
$\log(\text{mpg}_{\text{EV}})$	1.457 (1.508)	EV ×2020	2.841 (0.918)
Hybrid	-3.875 (0.553)	EV ×2021	3.193 (1.029)
EV	-11.989 (3.705)	EV ×2022	4.522 (1.098)
Van	-11.099 (2.984)	EV ×2023	4.169 (1.167)
Truck	-17.668 (4.679)	EV ×2024	3.266 (1.162)
SUV	-0.594 (0.328)		

Notes: Standard errors are clustered by model (2,982 clusters). $\log(\text{income}_{10k})$ is household income in \$10,000 units. Prices and subsidies are in \$100,000 units. Year and firm fixed effects and additional year-by-SUV, year-by-Hybrid interactions are included but omitted from the table for brevity.

Table 3: Empirical and estimated micro-moments

Moment	Observed	Estimated	Difference
Panel A. Income–price moments			
Mean price (Q2) − Mean price (Q1)	0.011	0.015	−0.004
Mean price (Q3) − Mean price (Q1)	0.017	0.031	−0.014
Mean price (Q4) − Mean price (Q1)	0.011	0.046	−0.035
Mean price (Q5) − Mean price (Q1)	0.063	0.071	−0.008
$P(\text{purchase} \mid Q_2)/P(\text{purchase} \mid Q_1)$	2.022	2.426	−0.404
$P(\text{purchase} \mid Q_3)/P(\text{purchase} \mid Q_1)$	2.804	3.510	−0.706
$P(\text{purchase} \mid Q_4)/P(\text{purchase} \mid Q_1)$	3.470	4.510	−1.040
$P(\text{purchase} \mid Q_5)/P(\text{purchase} \mid Q_1)$	5.841	5.448	+0.393
Panel B. Second-choice and match-on-characteristics moments			
2015 second-choice moments (from Grieco, Murry, and Yurukoglu (2024))			
$P(\text{second is van} \mid \text{first is van})$	0.720	0.721	−0.001
$P(\text{second is truck} \mid \text{first is truck})$	0.872	0.889	−0.017
$P(\text{second is SUV} \mid \text{first is SUV})$	0.690	0.692	−0.002
$P(\text{second is luxury} \mid \text{first is luxury})$	0.550	0.533	+0.017
$\text{corr}(\log \text{mpg}_{\text{first}}, \log \text{mpg}_{\text{second}})$	0.611	0.638	−0.027
$\text{corr}(\log \text{hp}_{\text{first}}, \log \text{hp}_{\text{second}})$	0.674	0.669	+0.005
$P(\text{second is Euro-brand} \mid \text{first is Euro-brand})$	0.413	0.406	+0.007
2022 EV second-choice moments (from Allcott et al. (2024))			
$P(\text{second is EV} \mid \text{first is EV})$	0.520	0.500	+0.020
$P(\text{second is EV, same class} \mid \text{first is EV})$	0.330	0.256	+0.074
Panel C. EV and body-type shares by division, 2021			
EV share among purchasers			
$P(\text{EV} \mid \text{purchase}, \text{Mountain}, 2021)$	0.032	0.033	−0.001
$P(\text{EV} \mid \text{purchase}, \text{North Central}, 2021)$	0.015	0.016	−0.001
$P(\text{EV} \mid \text{purchase}, \text{North East}, 2021)$	0.025	0.023	+0.002
$P(\text{EV} \mid \text{purchase}, \text{Pacific}, 2021)$	0.083	0.102	−0.019
$P(\text{EV} \mid \text{purchase}, \text{South Atlantic}, 2021)$	0.024	0.030	−0.006
$P(\text{EV} \mid \text{purchase}, \text{South Central}, 2021)$	0.014	0.008	+0.006
Truck share among purchasers			
$P(\text{Truck} \mid \text{purchase}, \text{Mountain}, 2021)$	0.212	0.210	+0.002
$P(\text{Truck} \mid \text{purchase}, \text{North Central}, 2021)$	0.191	0.196	−0.005
$P(\text{Truck} \mid \text{purchase}, \text{North East}, 2021)$	0.122	0.146	−0.024
$P(\text{Truck} \mid \text{purchase}, \text{Pacific}, 2021)$	0.134	0.144	−0.010
$P(\text{Truck} \mid \text{purchase}, \text{South Atlantic}, 2021)$	0.148	0.163	−0.015
$P(\text{Truck} \mid \text{purchase}, \text{South Central}, 2021)$	0.212	0.236	−0.024
SUV share among purchasers			
$P(\text{SUV} \mid \text{purchase}, \text{Mountain}, 2021)$	0.439	0.452	−0.013
$P(\text{SUV} \mid \text{purchase}, \text{North Central}, 2021)$	0.505	0.492	+0.013
$P(\text{SUV} \mid \text{purchase}, \text{North East}, 2021)$	0.513	0.488	+0.025
$P(\text{SUV} \mid \text{purchase}, \text{Pacific}, 2021)$	0.421	0.470	−0.049
$P(\text{SUV} \mid \text{purchase}, \text{South Atlantic}, 2021)$	0.456	0.436	+0.020
$P(\text{SUV} \mid \text{purchase}, \text{South Central}, 2021)$	0.425	0.445	−0.020

Notes: Panel B: the first seven moments (van, truck, SUV, luxury, and the two correlation moments plus Euro-brand) are for 2015 and are taken from Grieco, Murry, and Yurukoglu (2024). The last two Panel B moments (EV and EV in the same class) are for 2022 and are taken from Allcott et al. (2024). Panel C: entries show a subset of division-level moments for 2021; analogous EV, truck, and SUV moments are included for 2022–2024 but omitted here for brevity.

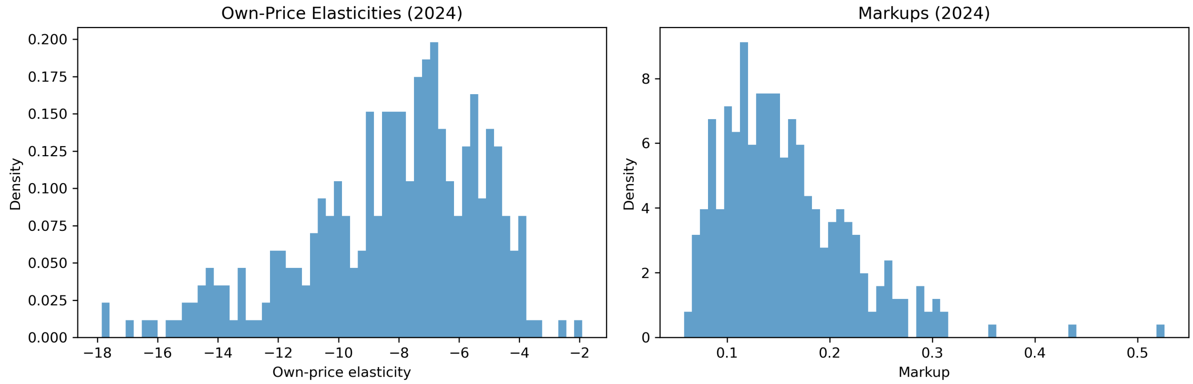

Figure 3: Distribution of own-price elasticities and markups, 2024

individual products by Coşar et al. (2018), which spans from -14.09 to -15.73. Our estimate of the market elasticity—the percentage change in total vehicle sales following a 1% increase in the price of all vehicles—is smaller in magnitude than those typically found in the literature, which are generally closer to -1 (Grieco, Murry, and Yurukoglu (2024), Allcott et al. (2024), Berry, Levinsohn, and Pakes (1995)).

These elasticities imply markups and marginal costs, which we recover by inverting equation (6) to obtain marginal costs and then computing markups as $(p_j - mc_j)/p_j$. The share-weighted average markup across the full sample is 17.2%. Figure 3 presents the distribution of own-price elasticities and markups in the 2024 market.

5.1.3 Diversion Ratios

A key strength of random-coefficients discrete choice models is their ability to generate realistic substitution patterns across products. We rely on these implied substitution patterns to compute changes in consumer and firm welfare in the counterfactual analysis. Here, we assess the plausibility of the model by examining the diversion ratios implied for a set of representative products.

The diversion ratio d_{jk} is defined as the fraction of consumers who, when switching away from product j in response to a price change, substitute toward product k (Conlon and Mortimer

2021). Formally,

$$d_{jk} = -\frac{\partial s_{kt}/\partial p_{jt}}{\partial s_{jt}/\partial p_{jt}}. \tag{30}$$

Table 4 reports the diversion ratios for six representative products in the 2024 market. The top alternatives are qualitatively similar to the base product in each panel. In most cases, substitution occurs within the same vehicle category (e.g., electric vehicles, trucks), and the leading substitutes correspond to vehicles that would intuitively be considered close competitors.

Tesla Model 3 EV		**Toyota RAV4 Hybrid**		**Ford Bronco**	
Alternative	**Share**	**Alternative**	**Share**	**Alternative**	**Share**
Outside Good	38.2%	Outside Good	30.1%	Outside Good	18.0%
Tesla Model Y EV	14.6%	Honda CR-V Hybrid	9.3%	Jeep Grand Cherokee	3.9%
BMW i4 EV	2.4%	Hyundai Tucson Hybrid	4.1%	Chevrolet Tahoe	3.0%
Tesla Model S EV	1.5%	Kia Sportage Hybrid	2.2%	Toyota 4Runner	2.7%
BMW i5 EV	1.5%	Toyota Grand Highlander Hybrid	1.8%	Kia Telluride	2.6%
Cadillac Lyriq EV	1.3%	Lexus RX Hybrid	1.8%	Ford Expedition	2.3%

Audi A5		**Honda Civic**		**Ford F-150**	
Alternative	**Share**	**Alternative**	**Share**	**Alternative**	**Share**
Outside Good	12.1%	Outside Good	10.3%	Outside Good	6.8%
Mercedes-Benz C-Class	4.9%	Toyota Corolla	5.8%	Chevrolet Silverado	29.2%
BMW 3 Series	2.7%	Nissan Rogue	4.9%	GMC Sierra	15.6%
BMW 5 Series	2.3%	Subaru Crosstrek	4.8%	Toyota Tundra	6.0%
BMW X3	2.2%	Nissan Sentra	4.2%	Toyota Tacoma	5.5%
Volvo S60	1.9%	Toyota Camry	3.4%	Ram 1500	5.5%

Table 4: Outside-good diversion and top five diversion destinations (2024)

5.2 Supply-Side Results

We present the results for the baseline specification described above, along with an alternative first-differences specification to assess robustness, in Table 5.

The preferred specification is reported in column (1), which includes vehicle characteristics as controls. In both the levels and first-differences specifications, we obtain sensible coefficients on the characteristics and a statistically significant coefficient on the interaction between the

Table 5: Cost-Side Regressions Results

	Levels		**First-differences**	
	(1)	**(2)**	**(3)**	**(4)**
ln(size)	0.518^{***}		0.575^{***}	
	(0.190)		(0.191)	
ln(weight)	0.459^{*}		0.705^{**}	
	(0.272)		(0.333)	
ln(hp)	0.223^{**}		0.265^{**}	
	(0.087)		(0.101)	
ln(mpg)	-0.088		-0.025	
	(0.074)		(0.063)	
$\rho^{(1)}_{f,j,t-1} \cdot \log\big(RER^{(1)}_{jt}\big)$	0.715^{***}	0.764^{***}	0.662^{**}	0.860^{***}
	(0.258)	(0.250)	(0.257)	(0.297)
Make-model FE	Yes	Yes	No	No
Year FE	Yes	Yes	Yes	Yes
Observations	323	323	207	207
R^2	0.986	0.981	0.362	0.159
Within R^2	0.306	0.087	0.272	0.040

Notes: The dependent variable is ln(costs). The real exchange rate is the bilateral real exchange rate between the US and the supplier country, normalized to 1 for each country in 2015; an increase in the RER variable indicates an appreciation of the foreign currency. Standard errors (clustered by make-model) are in parentheses below coefficients. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

lagged import share and the real exchange rate. The estimated coefficient of interest implies that, for a domestically assembled vehicle that sources 100% of its parts from abroad, a 1% increase in the cost of those parts leads to a 0.715% increase in the vehicle's marginal cost.

This estimate is very close to external assessments of the share of vehicle costs attributable to parts. In a case study of General Motors, Helper and Henderson (2014) reports that approximately 70% of a vehicle's cost is attributable to parts. In contrast, Menk, Chen, and Cregger (2012) estimates that roughly 71% of the value of finished vehicles reflects parts costs. In this context, we interpret our estimate as consistent with limited medium-run re-optimization of sourcing decisions by firms in response to price-level shocks.

6 Counterfactuals

This section quantifies the equilibrium implications of trade policy by combining the three central components of our framework: consumer substitution across differentiated products, firm pricing under oligopolistic competition, and the transmission of foreign input cost shocks into marginal costs through global value chains. The counterfactuals are designed to isolate how these margins interact to shape prices, quantities, and welfare under alternative tariff and subsidy regimes. By varying whether tariffs apply to final goods, intermediate inputs, or both, we clarify how protection reallocates demand, alters firms' cost structures, and ultimately determines the incidence of trade policy across consumers, producers, and regions.

The Trump Administration has prioritized tariffs as a trade policy instrument, implementing new measures affecting imports of automotive parts and finished vehicles. In March 2025, the U.S. government announced additional tariffs of 25% on imported vehicles and automotive components under Section 232 of the Trade Expansion Act of 1962, with preferential rates subsequently applied to imports from the United Kingdom (10%), Japan (15%), and South Korea (15%) (United States Congress (2025)).

At the same time, the Trump Administration repealed EV subsidies enacted under the Biden Administration's Inflation Reduction Act (IRA), effective September 2025.

We use our model to evaluate the counterfactual effects of these policy changes by applying them to the simulated 2024 U.S. automotive market. A central objective is to disentangle the impacts of the policy package on heterogeneous U.S. firms and consumers.

In addition to a baseline scenario featuring no tariffs and the repeal of IRA EV subsidies, we consider five main counterfactual scenarios. These consist of "subsidy reinstated" and "subsidy repealed" versions of three tariff regimes: (i) no tariffs, (ii) tariffs applied to all imported vehicles and automotive parts, and (iii) tariffs applied to vehicles only, with automotive parts exempted.

The counterfactual results should be interpreted as medium-run outcomes. In our setup, firms can adjust marginal costs through re-sourcing decisions within the same calendar year as the tariff implementation, as described in our cost-side analysis, and can re-optimize product prices. However, the model does not capture potentially important medium-run adjustments in supplier networks or long-run firm responses such as entry and exit decisions in vehicle assembly or parts manufacturing. The scenario labeled "vehicles and parts tariffs, subsidy removed" corresponds to the full automotive policy package implemented in 2025 under the Trump Administration.

Finally, data from the National Highway Traffic Safety Administration do not distinguish between parts sourced from the United States and Canada. As a result, our simulations effectively treat Canadian parts as tariff-exempt in all scenarios. Vehicles assembled in Canada, however, are subject to tariffs on the model.

6.1 Solution method

6.1.1 Counterfactual costs, prices, and shares

Leveraging the estimation of equation (29), we model the pass-through of a percentage tariff on intermediate goods to the marginal cost of domestically assembled vehicles as:

$$mc^{(d)}_{j,CF} = mc^{(d)}_{j,2024}(1 + \eta \cdot \rho_{f,j,2024} \cdot tariff_{parts}) \tag{31}$$

where $\rho_{j,2024}$ is the percentage of value of parts imported, $\eta = 0.715$ is as estimated in (29). Since we do not fully observe the source locations of foreign parts, we do not allow for country-specific tariff rates for intermediate parts; the full 25% tariff is applied to all imported parts.

For foreign assembled vehicles, we treat the tariff as increasing the marginal cost of vehicle

j and calculate the counterfactual marginal costs as:

$$mc_{j,CF}^{(f)} = mc_{j,2024}^{(f)}(1 + \tau \cdot tariff_{j,vehicles}) \tag{32}$$

where in this case the subscript j on tariffs represents the country-specific tariff rate. Here τ is a parameter that governs the pass-through of full vehicle tariffs to marginal costs. We borrow the estimate from Coşar et al. (2018) and set $\tau = 0.682$. As explained by the authors of this paper, τ being less than one indicates that, even for imported vehicles, a material portion of marginal costs is in addition to the value on which the tariff is imposed, such as marketing, intra-US distribution, and customer relations.

Given the perturbed marginal cost vector $mc_{j,CF} = (mc_{j,CF}^{(d)}, mc_{j,CF}^{(f)})$ in the counterfactual, we re-solve for the equilibrium prices and shares using the contraction mapping algorithm from Morrow and Skerlos (2011).

6.1.2 Firm and Consumer Surplus Computation

Firm surplus: We compute firm-level profits using simulated equilibrium prices and shares under each counterfactual scenario. For each product j in market t, let p_{jt} denote the equilibrium price, c_{jt} marginal cost, and s_{jt} the simulated market share. Product-level per-capita profit is:

$$\pi_{jt} = (p_{jt} - c_{jt}) \cdot s_{jt}. \tag{33}$$

Firm-level profits aggregate per-capita profit across products owned by firm f, which are then scaled by market size M_t:

$$\Pi_{ft} = M_t \cdot \sum_{j \in \mathcal{J}_f} \pi_{jt}. \tag{34}$$

Firm-level changes are computed as differences between counterfactual and baseline values. US producer surplus is the sum of profit changes for US-headquartered firms.

Consumer surplus: Consumer surplus is computed using the random-coefficients logit inclusive value, evaluated at simulated equilibrium prices in each market. For consumer i in market t, we simulate utility for product j:

$$u_{ijt} = \delta_{jt} + \mu_{ijt} + \varepsilon_{ijt}, \tag{35}$$

as in equation 7. The inclusive value for agent i is:

$$IV_{it} = \log\left(\sum_{j\in\mathcal{J}_t} \exp(\delta_{jt} + \mu_{ijt})\right), \tag{36}$$

The compensating-variation measure of consumer surplus for agent i is:

$$CS_{it} = \frac{1}{|\alpha_i|} IV_{it}, \tag{37}$$

with α_i the consumer-specific marginal utility of price implied by the estimated random coefficients. In practice, we compute IV_{it} using the simulated draws for μ_{ijt} and the equilibrium prices for each scenario, then aggregate over simulated agents using their weights w_{it} and scale by market size M_t:

$$CS_t = M_t \cdot \sum_i w_{it}\, CS_{it}. \tag{38}$$

We report changes in consumer surplus as:

$$\Delta CS_t = CS_t^{cf} - CS_t^0. \tag{39}$$

6.2 Counterfactual Results

We discuss two sets of counterfactuals—changes in tariffs and changes in the IRA EV subsidy—relative to a common baseline in which the IRA EV subsidy is repealed, and no additional tariffs are imposed. We begin by holding EV subsidies off and comparing two tariff regimes: (i) a 25% tariff on imported finished vehicles only, with some country-specific lower rates, and (ii) a 25% tariff on both imported vehicles and imported parts.[17] We then examine the welfare implications of re-enacting the IRA EV subsidy. Table 6 summarizes the headline outcomes across these counterfactuals. For each scenario, we report consumer surplus (overall and by

---

[17] We only apply country-specific tariff rates to final vehicles; all parts receive the 25% tariff.

income quintile), US producer surplus, tariff revenue, the fiscal cost of the EV subsidy (when applicable), and changes in prices, markups, and production.

We next unpack the mechanisms behind these aggregate results. We first decompose the tariff effects by assembly location, examining how marginal costs, prices, markups, and market shares adjust for US-assembled versus foreign-assembled vehicles. We then study the distribution of producer-side impacts across firms by reporting firm-level profit changes. To isolate the incidence of the tariff policy itself, we emphasize the counterfactual with no EV subsidy and a 25% tariff on imported vehicles. Finally, we repeat the tariff analysis under the re-enacted EV subsidy. Results for all five counterfactuals, together with the baseline, are reported in Appendix Table B3.

Table 6: Counterfactual Tariff and Subsidy Scenarios: 2024 Market Outcomes

	C1	C2	C3
	Without Subsidy		With Subsidy & No Tariff
Tariff Status	Vehicles-only	Parts & vehicle	
Δ Price (avg, %)	5.71	10.07	-0.05
Δ Markup (avg %)	-0.5	-0.8	0
Δ US Producer Surplus (b USD)	1.03	-2.58	1.80
CS Δ total (b USD)	-14.619 (-4.6%)	-32.831 (-10.4%)	4.817 (1.5%)
CS Δ Q1 (b USD)	-0.551 (-5.5%)	-1.224 (-12.3%)	0.010 (0.1%)
CS Δ Q2 (b USD)	-1.625 (-5.9%)	-3.478 (-12.6%)	0.334 (1.2%)
CS Δ Q3 (b USD)	-2.750 (-6.4%)	-6.195 (-14.5%)	0.694 (1.6%)
CS Δ Q4 (b USD)	-3.699 (-5.1%)	-8.352 (-11.5%)	1.333 (1.8%)
CS Δ Q5 (b USD)	-5.994 (-3.5%)	-13.583 (-7.9%)	2.446 (1.4%)
Δ vehicles sold (m)	-0.410	-0.898	0.201
EV share (% sales)	3.18	3.15	7.39
US-assembled share (% sales)	66.8	57.1	53.5
Δ US assembled (m)	1.248	-0.076	0.144
Tariff revenue (b USD)	15.38	41.20	0.000
EV subsidy cost (b USD)	0.000	0.000	6.68
Δ Net US outcomes (b USD)	1.80	5.79	-0.07

6.2.1 Trade policies

Consumer and firm welfare As shown in Table 6, expanding the policy from a 25% tariff on imported finished vehicles only (C1) to a 25% tariff that also applies to imported parts (C2) substantially amplifies the equilibrium cost shock. Average prices increase by 5.7% under the vehicles-only tariff (C1), but by 10.1% when parts are included (C2). Consistent with the larger price increase, total vehicle sales fall by 0.41 million in C1 versus 0.90 million in C2.

These larger price and quantity effects translate into much larger consumer welfare losses. Aggregate consumer surplus falls by \$14.6 billion in C1, compared to \$32.8 billion in C2. The losses are borne disproportionately by higher-income households in absolute dollar terms in both counterfactuals: the top income quintile (Q5) accounts for around 40% of the consumer welfare loss. This pattern is consistent with higher-income consumers accounting for a larger share of vehicle purchases (and therefore a larger share of total dollar surplus), even if lower-income households are more price-sensitive and may experience larger losses relative to their baseline surplus.

On the production side, the difference between C1 and C2 is also sharp. Under the vehicles-only tariff (C1), US producer surplus rises by \$1.0 billion and US-assembled output increases by 1.3 million vehicles, consistent with substitution away from imported finished vehicles toward domestically assembled alternatives. Once imported parts are also tariffed (C2), this substitution channel is muted: US-assembled output is essentially unchanged (-0.076 million), and US producer surplus falls by \$2.58 billion, indicating that higher input costs for domestic assembly offset (and more than offset) the demand-diversion gains from taxing imported final vehicles.

Finally, tariff revenues are large in the broader-tariff scenario: \$15.38 billion in C1 versus \$41.2 billion in C2. Mechanically, this raises the measured "net US outcomes" (CS + US PS + tariff revenue) from +\$1.8 billion in C1 to +\$5.79 billion in C2. Importantly, however, these net gains come from revenue collection and do not, by themselves, eliminate the large consumer surplus losses—whether households are ultimately better off depends on how tariff revenues are recycled (or not) back to consumers and/or domestic firms.

Change by assembly location Figure 4 decomposes the equilibrium effects of the two tariff counterfactuals (both evaluated without the IRA EV subsidy) by vehicle assembly location (US-assembled vs. foreign-assembled). Each bar reports the change in prices, marginal costs, markups (percentage points), and market shares (percentage points) relative to the no-tariffs, without the EV subsidy baseline. Panels (a) and (b) show the results for counterfactual C1 (25% tariff on foreign-assembled vehicles only) and C2 (25% tariff on vehicles and parts), respectively. The vehicle-only tariff primarily affects foreign-assembled models: their marginal costs rise by 14.87% and their prices by 11.80%. The fact that the price response is smaller than the cost shock implies incomplete pass-through, and correspondingly, foreign-assembled markups compress by 2.10 pp. A notable feature is that US-assembled prices change little (+0.28%), whereas foreign-assembled prices increase sharply. Under Nash–Bertrand pricing, this implies that domestic assemblers gain little additional pricing power: competition among US models and substitution to the outside option constrain markups. This is consistent with the outside good holding a large baseline share (48.4%) and meaningful diversion out of the market. The tariff induces a sharp reallocation of demand across assembly locations: foreign-assembled market share falls by 7.52 pp, while US-assembled market share rises by 5.66 percentage points (with the remaining 1.86 percentage points shifting to the outside option, since shares are defined over the full market, including non-purchase).

Panel (b) of figure 4 shows the counterfactual results when the tariff is extended to imported parts, resulting in the propagation to domestically assembled vehicles as well. US-assembled marginal costs increase by +9.67%, and prices rise by +8.35%, again implying incomplete pass-through and a 1.0 percentage point decline in US-assembled markups. Foreign-assembled vehicles continue to experience larger increases (+14.87% in costs and +11.98% in prices) and a larger markup decline (1.97 pp), indicating that the overall burden remains heavier on foreign-assembled models. Unlike C1, however, market shares fall for both types of assembly locations (US-assembled 0.34 pp, foreign-assembled 3.73 pp), implying a substantial shift toward the outside option. Intuitively, once imported parts are taxed, domestic assembly becomes more expensive, limiting the extent to which US-assembled vehicles can absorb demand diverted away from imports; as a result, the policy operates more like a broad cost shock that contracts the market rather than a pure reallocation from foreign to domestic assembly.

Overall, these assembly-location-level mechanisms align with the aggregate patterns: C1 mostly reallocates demand from foreign-assembled vehicles toward US assembly, while C2 raises costs for both, compresses markups across the board, and shifts a larger share of consumers toward non-purchase.

Profit by firms Table 6 reports that the change in aggregate US producer surplus increases under the vehicle-only tariff (C1) but turns negative when the tariff is extended to parts (C2). Figure 5 shows that these aggregate patterns conceal substantial firm-level heterogeneity. It

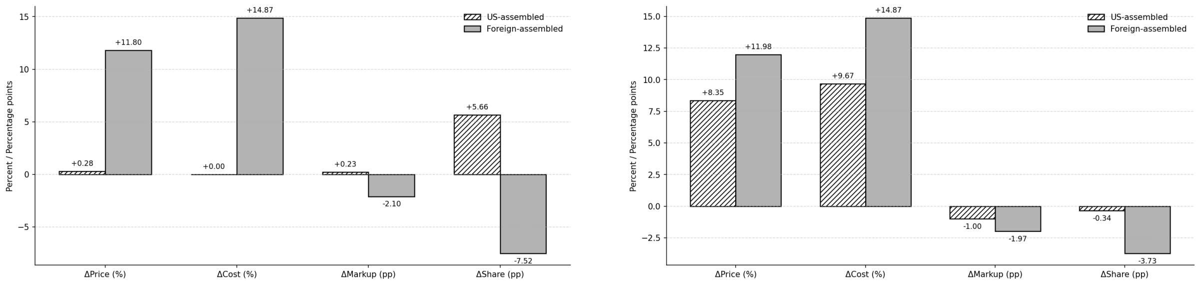

(a) **Counterfactual: Vehicle-only tariff** (b) **Counterfactual: Parts and vehicle tariffs**

Figure 4: Changes by assembly-location: Tariff counterfactuals

Notes: These figures compare changes in prices, costs, markups, and market shares by assembly location of the vehicle (US-assembled vs. foreign-assembled). Panel (a) reports counterfactual C1 (25% tariff on imported vehicles only) and Panel (b) reports counterfactual C2 (25% tariff on imported vehicles and imported parts). Each change is relative to the "no-tariffs, without EV subsidy" baseline. "Assembled" categorization is distinct from firm headquarters location. These assembly-location-level mechanisms align with aggregate patterns: C1 mostly reallocates demand from foreign-assembled vehicles toward US assembly, while C2 raises costs for both, compresses markups across the board, and shifts a larger share of consumers toward non-purchase.

plots percentage changes in firm-level profits relative to the baseline under two counterfactual tariff scenarios: C1 (a 25% tariff on imported finished vehicles only) on the x-axis and C2 (a 25% tariff applied to both imported vehicles and imported parts) on the y-axis.

Three patterns stand out. First, larger firms (larger bubbles) tend to lie closer to the origin, indicating smaller percentage changes in profits. Second, under C1, most non-U.S. firms (blue) experience profit declines—26 firms overall have negative changes—reflecting the direct impact of tariffs on imported vehicles. A few exceptions, such as Honda, Acura, Nissan, and BMW, benefit due to their substantial US assembly presence. Third, extending tariffs to parts (C2) reverses gains for many firms, particularly U.S.-headquartered manufacturers (e.g., Chevrolet, GMC, and Ford) that are more exposed to imported inputs. In contrast, firms with lower reliance on foreign parts (e.g., Tesla, Jeep, Honda, and Toyota) are more likely to benefit.

Finally, the quadrant structure highlights heterogeneity in tariff incidence. Many foreign firms lie in the bottom-left quadrant (losers under both policies), while several US firms lie in the bottom-right quadrant (gaining under C1 but losing under C2), underscoring the importance of input sourcing in shaping the incidence of protection. Appendix figure A5a shows the labels of the remaining firms (with a lesser market share). Appendix figure A5b plots the change in profit in USD instead of the percentage value.

Additional results are reported in the Appendix Figures. Appendix Figure A4 highlights het-

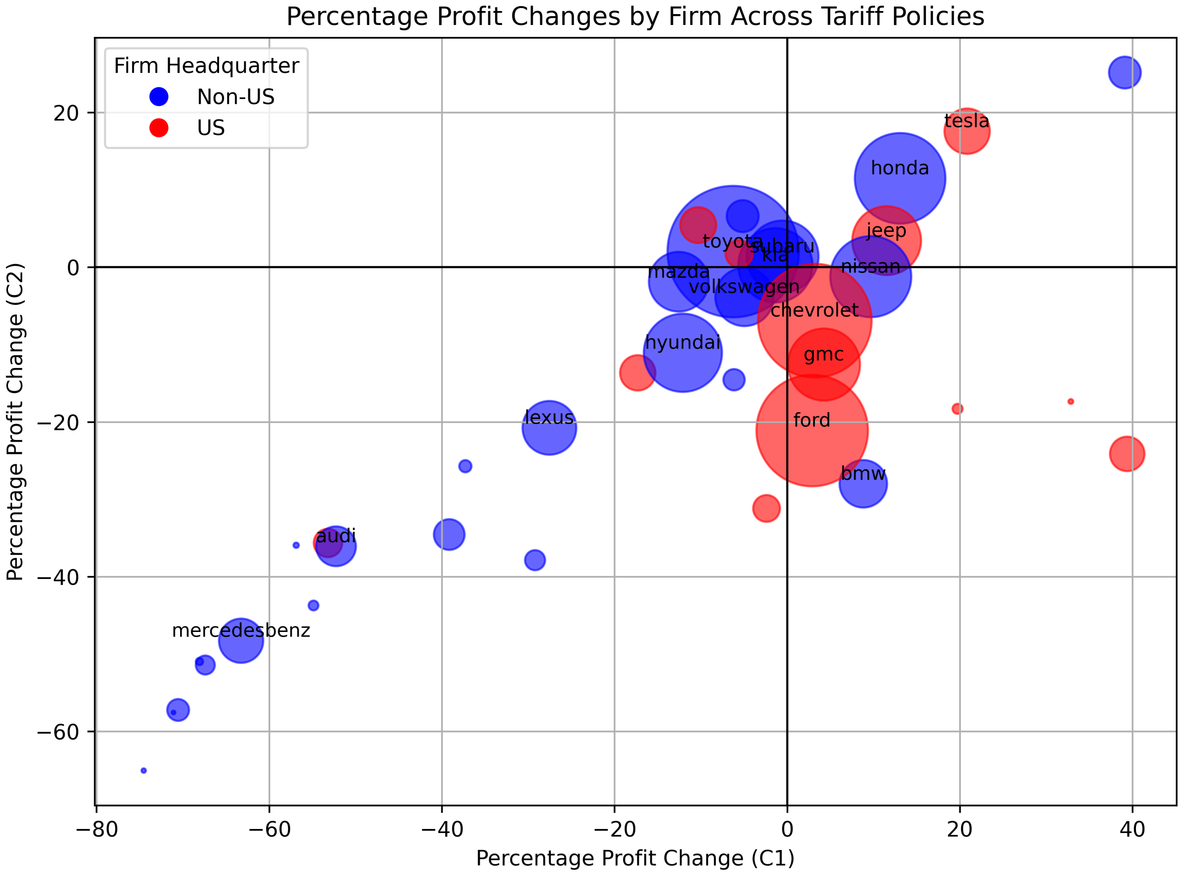

Figure 5: **Percentage profit changes across firms under tariff counterfactuals**

Notes: The figure plots percentage changes in firm-level profits relative to the baseline under two counterfactual tariff scenarios: C1 (25% tariff on imported finished vehicles only) on the x-axis and C2 (25% tariff applied to both imported vehicles and imported parts) on the y-axis. Each bubble represents one of the 38 firms in the sample, of which 13 are U.S.-headquartered. Bubble size is proportional to the firm's market share. Colors indicate firm headquarters (red: U.S.; blue: non-U.S.), while assembly locations may differ from headquarters, reflecting global production networks.

erogeneity in producer-side impacts by relating US-headquartered firms' profit changes (relative to the no-tariffs, without EV subsidy baseline) to their exposure to imported inputs, measured by the sales-weighted imported-parts share in their US-assembled vehicles. Appendix Figure A6 maps the state-level distribution of consumer-surplus changes under the parts-and-vehicles tariff (without EV subsidies maintained), showing that consumer impacts are geographically broad but uneven across states. Finally, Appendix Figure A7 visualizes how tariff counterfactuals reallocate vehicle production across assembly locations, reporting counterfactual production levels and the corresponding percentage change from baseline.

6.2.2 EV subsidy re-enactment

We now examine the welfare implications of re-enacting the IRA EV subsidy and report how the resulting welfare changes are distributed across income quantiles. Table 6 summarizes the headline outcomes across all counterfactuals. In C3 (no tariffs, EV subsidy in place), the policy operates primarily as a demand-side stimulus with almost no effect on average prices (+0.05%). Table The subsidy increases consumer surplus by \$4.8b and US producer surplus by \$1.8b, alongside a modest expansion in market activity (+0.2m vehicles sold) and +0.144m additional US-assembled vehicles. The consumer gains are skewed toward higher-income households in dollar terms—Q5 gains \$2.4b versus \$0.01b for Q1—reflecting that higher-income consumers account for a larger share of vehicle purchases. Accounting for the fiscal cost of the program (\$6.68b), the net US outcome is almost zero (-\$0.07b), implying that most of the gross surplus gains are offset by subsidy outlays.

On the producer side, re-enacting the EV subsidy generates pronounced heterogeneity in firm outcomes. Figure 6 plots the distribution of profit changes across firms and shows that Tesla is the largest beneficiary, with additional gains for US-based EV producers such as Jeep and Cadillac. This dispersion is consistent with the subsidy differentially raising demand for qualifying domestic EV models, shifting purchases away from unsubsidized ICE options, foreign (unsubsidized) EVs, and the outside good.

State-based impacts Re-enacting the EV subsidy generates a geographically broad, but relatively modest, increase in consumer welfare. Figure 7 shows that most states experience

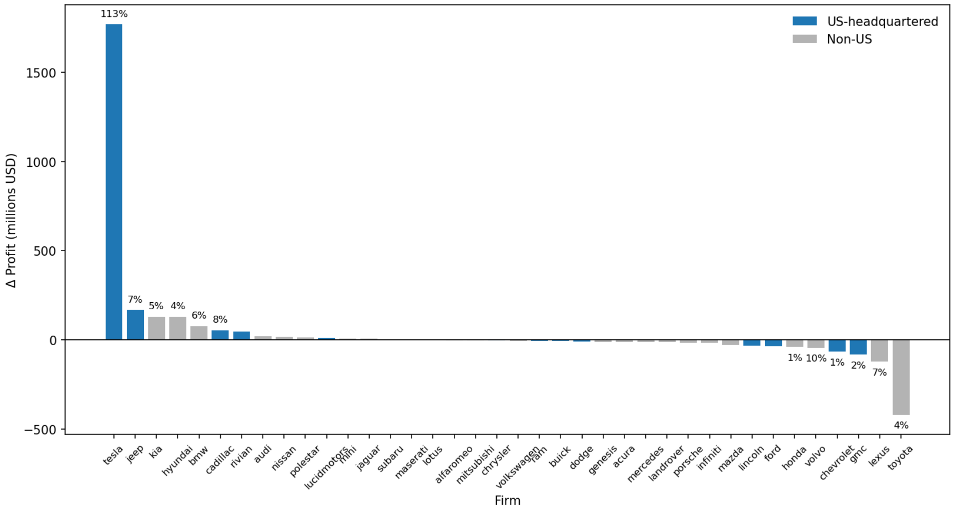

Figure 6: Firm profit changes: No tariffs, EV subsidy re-enacted (C3)

Notes: Bar height is profit change in millions of USD 2015 compared to the ’no-tariffs, without EV subsidy’ baseline. Bar labels are the corresponding % change in firm profits.

consumer surplus gains on the order of roughly 0.6–1.8% in percentage terms. Gains are somewhat larger in a small set of states and are concentrated in regions with a stronger preference for EVs—most notably the Pacific states of Oregon (+4.8%), Washington (+3.1%), California (+3.2%), and Alaska (+4%). Each of these Pacific states shares the same EV preference in our model; differences in welfare gains arise from heterogeneity in incomes and the corresponding distaste for price changes.

7 Conclusion

This paper studies the incidence of trade policy in a differentiated-products market with market power and global value chains. Using the U.S. automobile industry as a laboratory, we show that exposure to imported intermediate inputs fundamentally alters how tariffs affect prices, profits, and welfare. Our counterfactual analysis yields three key results. First, tariffs on imported vehicles increase U.S. producer surplus by about $1 billion but reduce consumer surplus by roughly $14 billion. Second, extending tariffs to imported parts reverses these gains: consumer losses roughly double and U.S. producer surplus declines by about $2.6 billion. These aggregate effects mask substantial heterogeneity: firms with greater exposure to imported parts

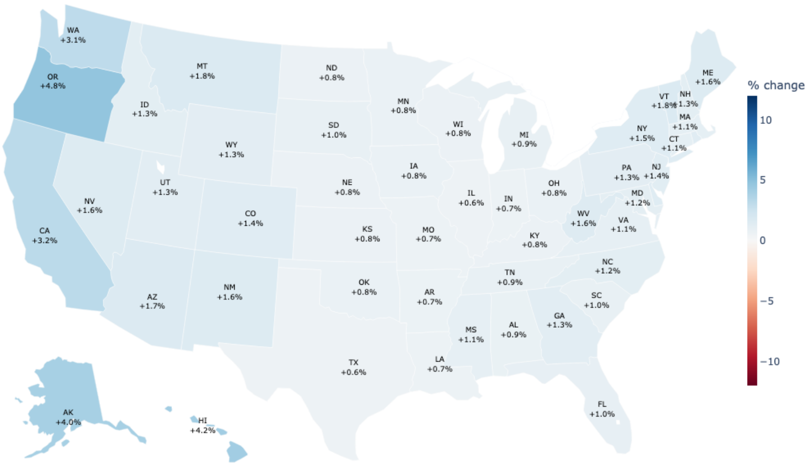

Figure 7: State consumer surplus impacts: No tariffs, subsidies reinstated

Notes: Consumer surplus change is relative to the 'no-tariffs, with EV subsidy' baseline. Labels for Rhode Island and Delaware have been removed.

experience losses, whereas those relying more on domestic inputs are better able to increase profits. Third, reinstating EV subsidies increases consumer surplus by about \$4.8 billion and raises U.S. producer surplus by about \$1.8 billion.

Our results highlight that the reversal in producer surplus for the domestic firms upon expanding the tariffs on the parts also arises from the interaction of three forces: substitution across differentiated products, oligopolistic pricing, and cost exposure through global value chains.

Although our empirical analysis focuses on automobiles, the mechanisms we document are not specific to that sector. Many industries affected by recent trade policy—including electronics, machinery, and clean energy technologies—feature differentiated products, market power, and heavy reliance on imported intermediate inputs. In such settings, policies that tax inputs as well as final goods are likely to reshape marginal costs and competitive positions in ways that differ from intuition based solely on final-goods trade. More broadly, our findings underscore that evaluating the incidence of trade policy requires explicit attention to global value chains and the interaction between cost exposure and market structure.

References

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

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. “Automobile Prices in Market Equilibrium.” *Econometrica* 63 (4):841–890. URL https://www.jstor.org/stable/2171802. Publisher: [Wiley, Econometric Society].

———. 2004. “Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market.” *Journal of Political Economy* 112 (1):68–105. URL https://www.jstor.org/stable/10.1086/379939. Publisher: The University of Chicago Press.

Bureau of Labor Statistics. 2025. “Public Use Microdata (PUMD).” URL https://www.bls.gov/cex/pumd.htm.

Conlon, Christopher and Jeff Gortmaker. 2020. “Best practices for differentiated products demand estimation with PyBLP.” *The RAND Journal of Economics* 51 (4):1108–1161. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/1756-2171.12352.

———. 2025. “Incorporating Micro Data into Differentiated Products Demand Estimation with PyBLP.” *Journal of Econometrics* :105926URL https://www.sciencedirect.com/science/article/pii/S030440762400277X.

Conlon, Christopher and Julie Holland Mortimer. 2021. “Empirical Properties of Diversion Ratios.” *The RAND Journal of Economics* .

Coşar, A. Kerem, Paul L. E. Grieco, Shengyu Li, and Felix Tintelnot. 2018. “What drives home market advantage?” *Journal of International Economics* 110:135–150. URL https://www.sciencedirect.com/science/article/pii/S0022199617301356.

Department of Energy. 2025. “Tax Incentive Data Services.” URL https://www.fueleconomy.gov/feg/ws/tax-data-services.shtml.

Duarte, Marco, Lorenzo Magnolfi, Daniel Quint, Christopher Sullivan, and Mikkel Sølvsten. 2025. "Conduct and Scale Economies: Evaluating Tariffs in the US Automobile Market."

Grieco, Paul L E, Charles Murry, and Ali Yurukoglu. 2024. "The Evolution of Market Power in the U.S. Automobile Industry." *The Quarterly Journal of Economics* 139 (2):1201–1253. URL `https://doi.org/10.1093/qje/qjad047`.

Halton, J. H. 1960. "On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals." *Numerische Mathematik* 2 (1):84–90. URL `http://link.springer.com/10.1007/BF01386213`.

Helper, Susan and Rebecca Henderson. 2014. "Management Practices, Relational Contracts, and the Decline of General Motors." *Journal of Economic Perspectives* 28 (1):49–72. URL `https://pubs.aeaweb.org/doi/10.1257/jep.28.1.49`.

Internal Revenue Service. 2025. "Clean vehicle tax credits | Internal Revenue Service." URL `https://www.irs.gov/clean-vehicle-tax-credits`.

International Monetary Fund. 2025. "International Financial Statistics database." URL `https://data.worldbank.org`.

Menk, Debbie, Yen Chen, and Joshua Cregger. 2012. "Methodology for Creating a Matrix to Assess the Domestic Content of a Vehicle by Make and Model." *Center for Automotive Research* .

Morrow, W. Ross and Steven J. Skerlos. 2011. "Fixed-Point Approaches to Computing Bertrand-Nash Equilibrium Prices Under Mixed-Logit Demand." *Operations Research* 59 (2):328–345. URL `https://pubsonline.informs.org/doi/10.1287/opre.1100.0894`. Publisher: INFORMS.

Nevo, Aviv. 2001. "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica* 69 (2):307–342. URL `https://onlinelibrary.wiley.com/doi/abs/10.1111/1468-0262.00194`. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1468-0262.00194.

Petrin, Amil. 2002. "Quantifying the Benefits of New Products: The Case of the Minivan." *Journal of Political Economy* 110 (4):705–729. URL `https://www.jstor.org/stable/10.1086/340779`. Publisher: The University of Chicago Press.

Ruggles, Steven, Sarah Flood, Matthew Sobek, Daniel Backman, Grace Cooper, Julia A. Rivera Drew, Stephanie Richards, Renae Rogers, Jonathan Schroeder, and Kari C.W. Williams. 2025. "IPUMS USA."

Sabal, Alejandro. 2025. *Product Entry in the Global Automobile Industry*. Ph.D. thesis, Princeton University.

United States Congress. 2025. "Section 232 Automotive Tariffs: Issues for Congress." URL `https://www.congress.gov/crs-product/IN12545`.

WardsAuto. 2025. "Most Automakers Anticipate Bright 2025 | WardsAuto." URL `https://www.wardsauto.com/news/most-automakers-anticipate-bright-2025/798797/`.

A Appendix Figures

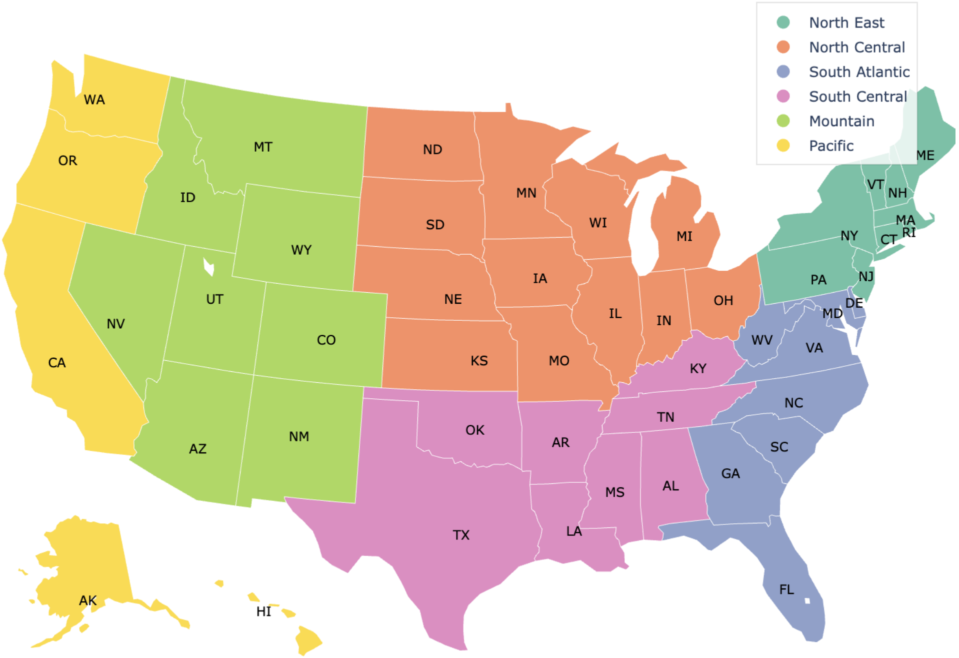

Figure A1: US regional divisions

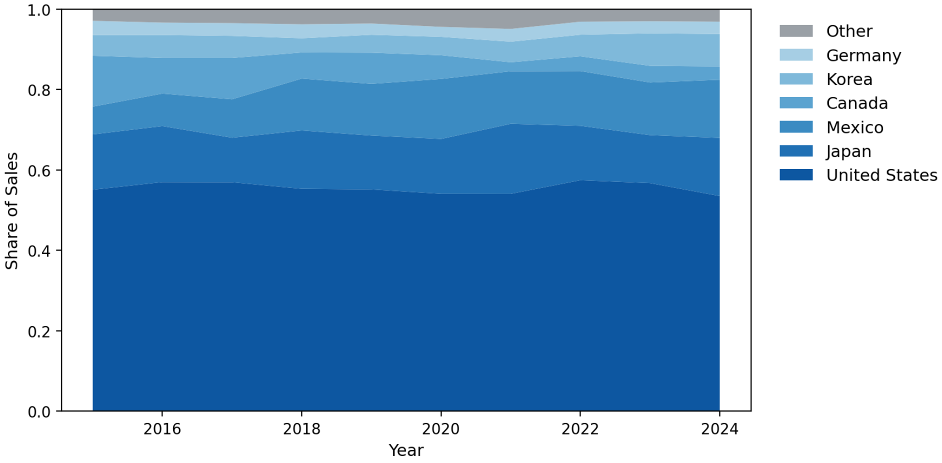

Figure A2: Assembly location (country) of vehicles sold in the United States

Notes: The figure shows the distribution of assembly locations for vehicles sold in the United States, weighted by sales, over time. It illustrates the importance of foreign production in supplying the U.S. market and the evolution of global production patterns.

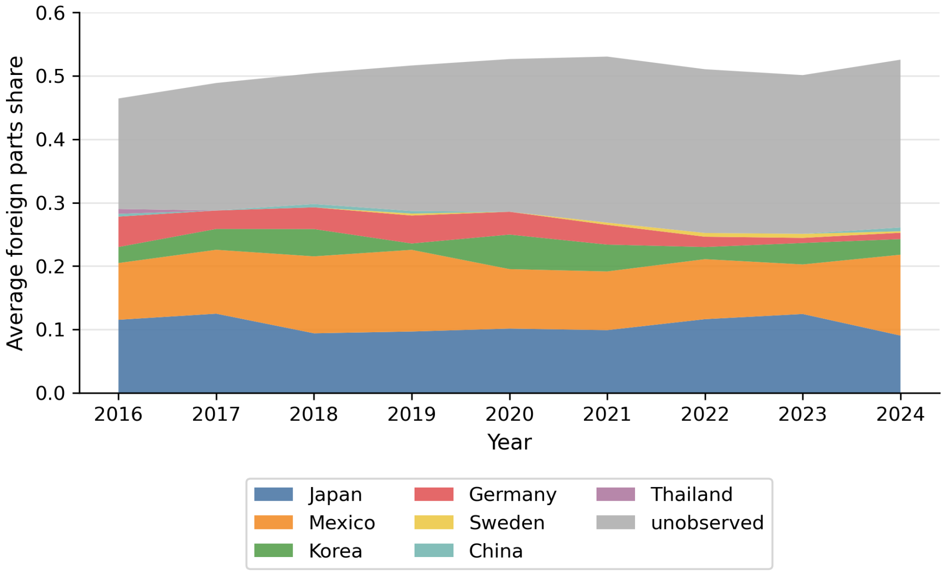

Figure A3: Observed foreign sources of vehicle parts

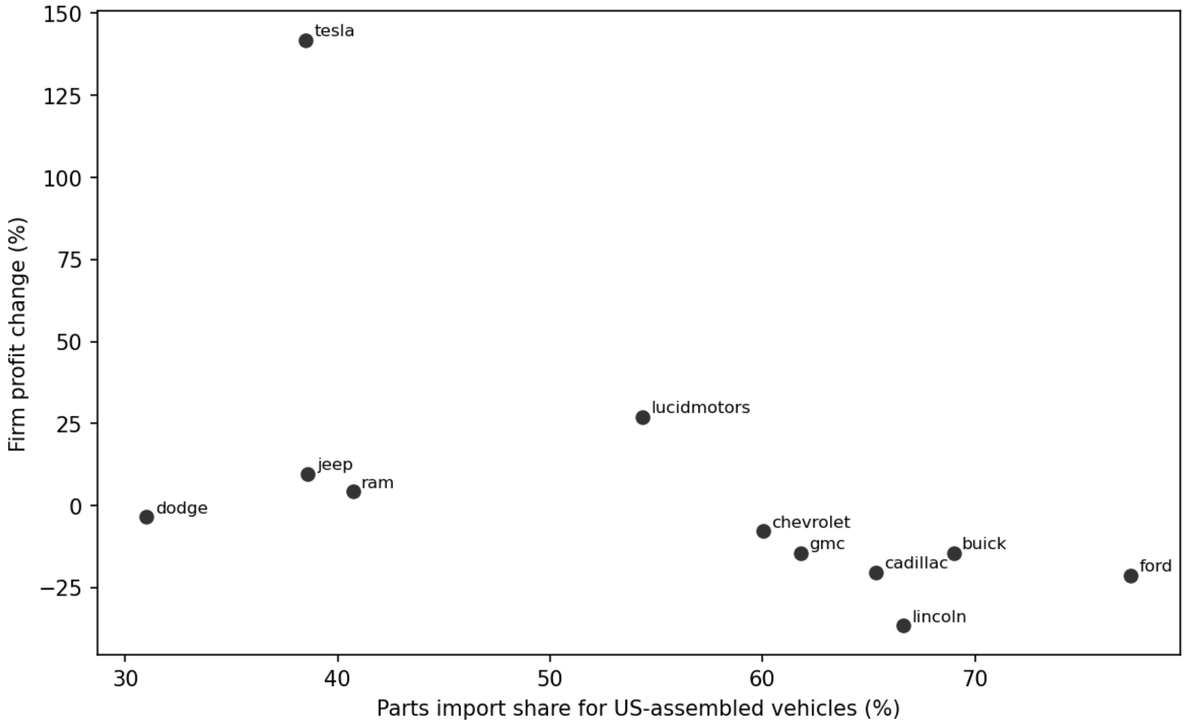

Figure A4: US firms' profit change and parts import share: Vehicles only tariff, with EV subsidies

Notes: Only US-headquartered firms are show. Profit change in % is compared to the 'no-tariffs, with EV subsidy' baseline. The x-axis is the sales-weighted share of parts imported for the US-assembled vehicles of each US-headquartered firm. The y-axis is the total profit change for the firm (including from any vehicles assembled abroad). Rivian has been removed from this chart as its import share % was imputed as the average of all vehicles.

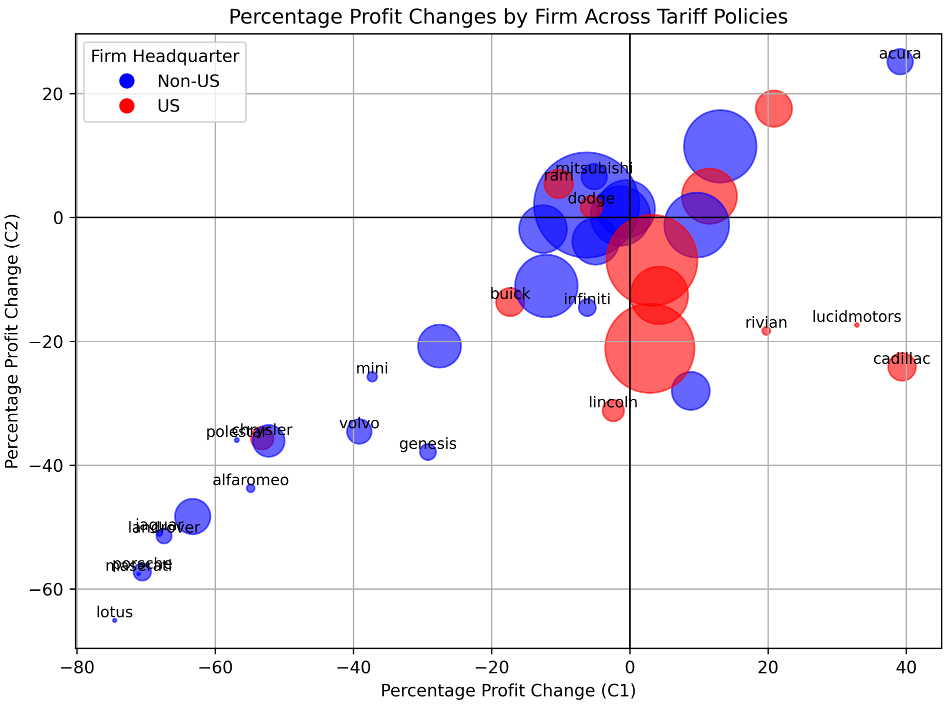

(a) **Percentage profit changes across firms under tariff counterfactuals (remaining labels)**

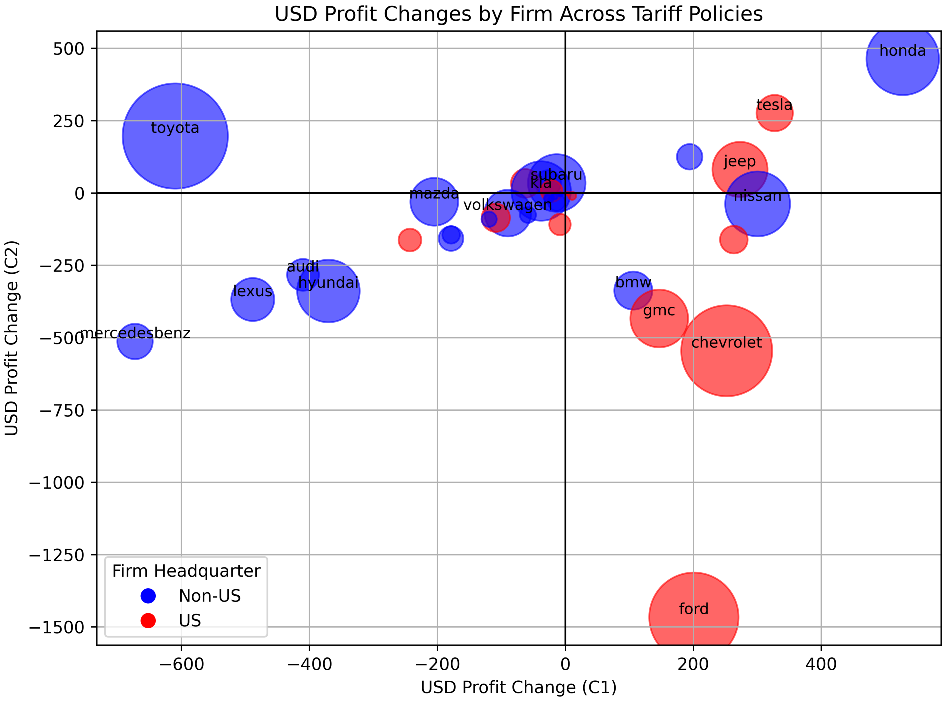

(b) **Profit (USD) changes across firms under tariff counterfactuals**

Figure A5: Changes by assembly-location: Tariff counterfactuals

Notes: These figures plot changes in firm-level profits relative to the baseline under two counterfactual tariff scenarios: C1 (25% tariff on imported finished vehicles only) on the x-axis and C2 (25% tariff applied to both imported vehicles and imported parts) on the y-axis. Panel (a) reports the firm-level percentage profits changes with the label of the firms with less market share (labels missing in the Figure 5). Panel (b) reports the change in profits in USD.

Table B1: Summary Statistics by Vehicle Type

		Cars, $N = 4178$	Trucks, $N = 455$	SUVs, $N = 3015$	Vans, $N = 403$
Sales	Mean	15,686	57,234	22,869	19,918
	Std. dev.	33,280	92,476	38,741	25,606
	Min	15	38	18	18
	Max	334,818	529,238	374,263	130,780
Price (2015 USD, $100,000)	Mean	0.36	0.35	0.43	0.31
	Std. dev.	0.17	0.07	0.16	0.05
	Min	0.13	0.19	0.18	0.21
	Max	1.00	0.74	1.00	0.47
Horsepower (100s)	Mean	2.46	3.41	2.91	2.65
	Std. dev.	1.04	0.97	0.89	0.64
	Min	0.70	1.91	1.38	1.31
	Max	8.45	8.35	8.35	4.01
Footprint (square ft, 100s)	Mean	0.97	1.31	1.06	1.27
	Std. dev.	0.12	0.19	0.12	0.20
	Min	0.45	1.02	0.81	0.86
	Max	1.27	1.84	1.40	1.92
Curb weight (pounds, 1000s)	Mean	3.54	5.06	4.43	4.57
	Std. dev.	0.57	0.87	0.76	0.64
	Min	1.81	3.52	3.02	3.25
	Max	5.92	9.10	6.92	5.99
Imported (%)	Mean	0.75	0.08	0.55	0.59
Electric (%)	Mean	0.06	0.04	0.05	0.00

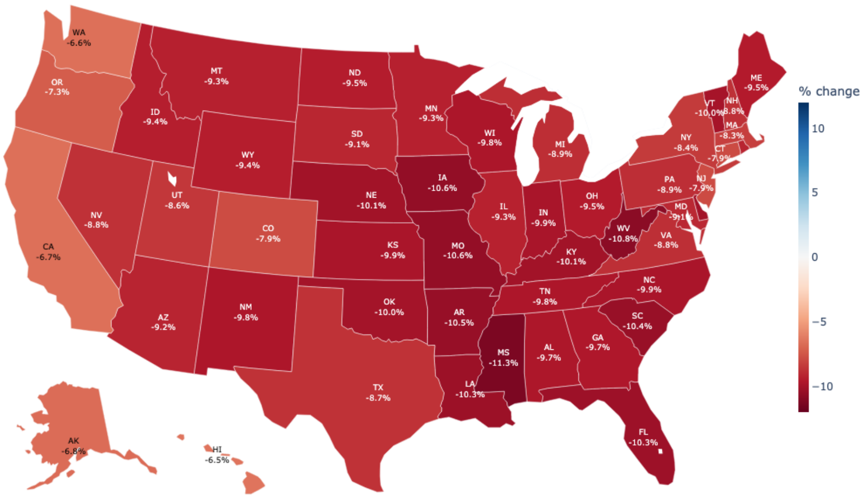

Figure A6: State consumer surplus impacts: Parts and vehicle tariffs, with EV subsidies

Notes: Consumer surplus change is relative to the 'no-tariffs, with EV subsidy' baseline. Labels for Rhode Island and Delaware have been removed.

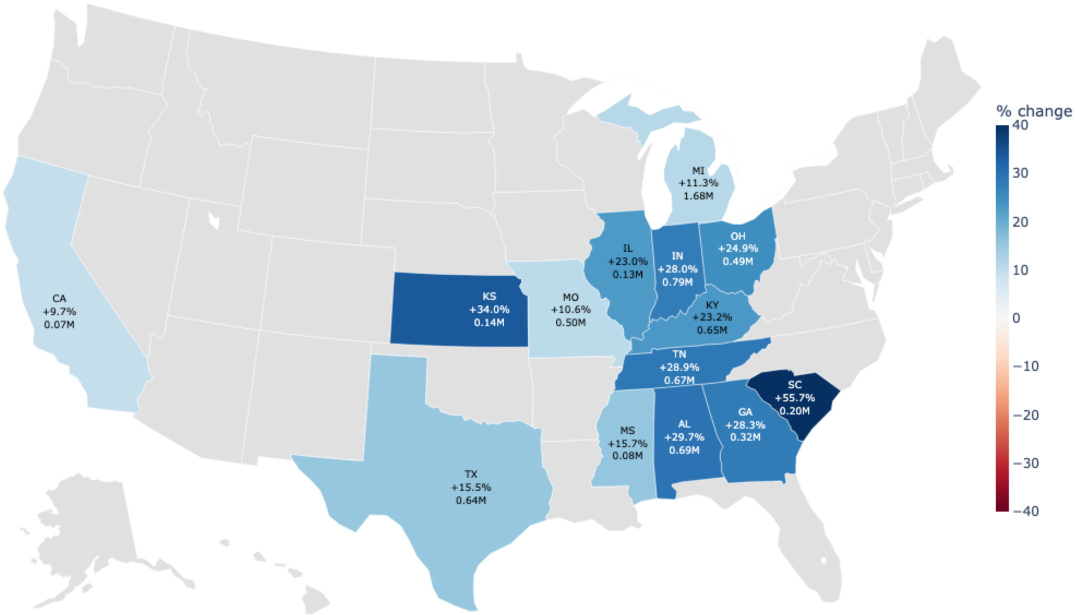

(a) State assembly impacts: Vehicle-only tariff (C1)

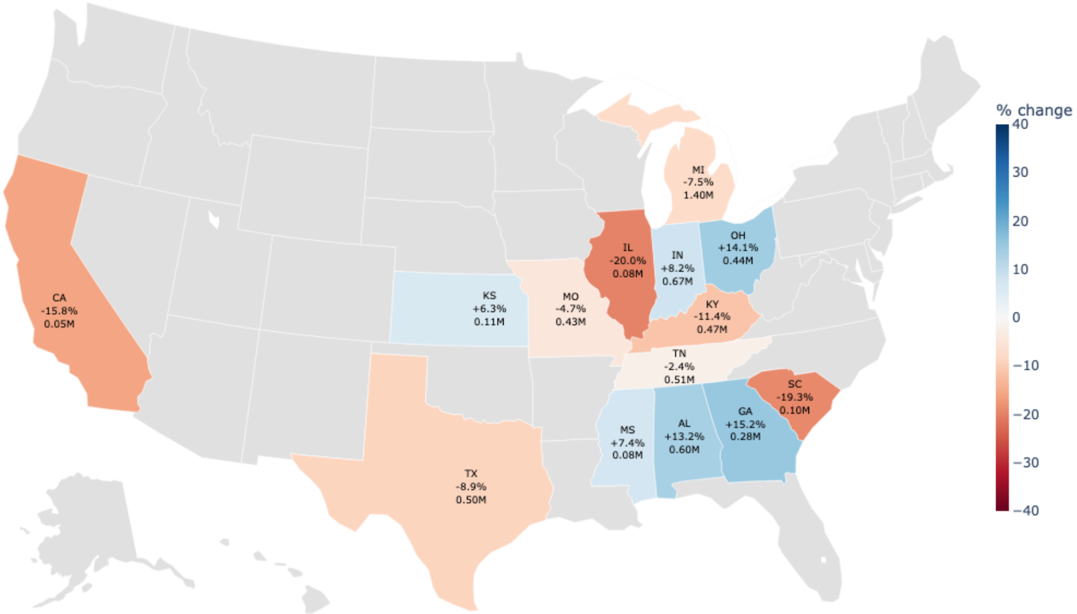

(b) State assembly impacts: Vehicle and parts tariffs (C2)

Figure A7: State assembly impacts: Tariff counterfactuals

Notes: Values shown are total counterfactual vehicle production and % change from 'no-tariffs, without EV subsidy' baseline. Arizona production has been removed as only the Lucid Air, with a negligible share, is produced there.

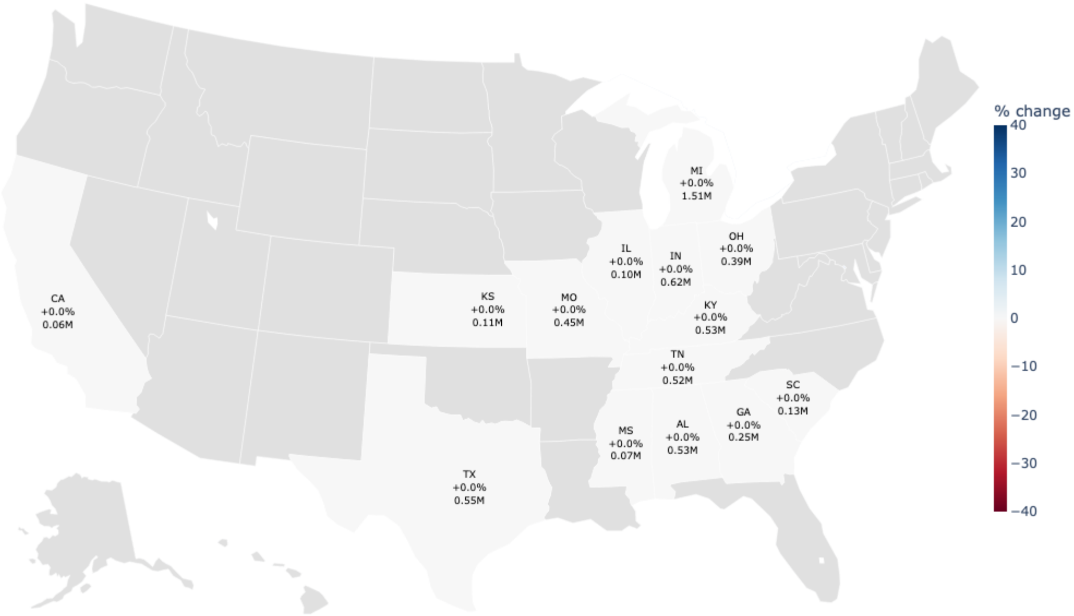

Figure A8: State assembly impacts: No tariffs, subsidies removed.

Notes: Values shown are total counterfactual vehicle production and % change from baseline. Arizona production has been removed as only the Lucid Air, with a negligible share, is produced there.

B Appendix Tables

Table B2: Average EV Subsidy Available by Producer (45W-adjusted in 2023–2024)

	Market year									
Producer	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024
Cadillac	0	0	0	0	0	0	0	0	7500	7500
Chevrolet	7500	7500	7500	7500	4200	465	0	0	7500	7500
Ford	7500	0	7500	0	0	0	7500	4664	7500	7500
Hyundai	0	0	7500	7500	7500	7500	7500	4664	7500	7500
Kia	7500	7500	7500	0	7500	7500	7500	4664	7500	7500
Nissan	7500	7500	7500	7500	7500	7500	7500	4664	7500	7500
Tesla	7500	7500	7500	7500	2805	0	0	0	7500	7500
Volkswagen	7500	7500	7500	0	7500	0	7500	4664	7500	7500

Notes: Entries are average EV subsidy amounts (USD) by producer and market year. Values are rounded to the nearest dollar.

C Appendix Text

C.0.1 Construction of Moments

Income-dependent micro-moment construction We use the CEX Interview Survey micro-data to construct income-dependent automobile purchase moments (Bureau of Labor Statistics (2025)). The CEX micro-data reports monthly expenditure on new vehicles linked to household income, which we aggregate to the annual level to construct our observed income micro parts:

$$\mathbf{E}_t[\text{price paid}_{it}|\text{income quintile}_{it}],$$

where price paid is the sum of net new-vehicle expenditures in that calendar year before the receipt of any appropriate electric vehicle subsidy.

To build purchase probability micro-moments we define an annual purchase indicator for each consumer unit as $\mathbf{1}[\text{purchase}] = \mathbf{1}[\text{new vehicle spend ¿ 0}]$ and define the micro parts,

Table B3: Counterfactual Tariff and Subsidy Scenarios: 2024 Market Outcomes

	B0	C1	C2	C3	C4	C5
		Without Subsidy			With Subsidy	
Tariff Status	No	Vehicles-only	Parts & vehicle	No	Vehicles-only	Parts & vehicle
Δ Price (avg, %)	0.00	5.71	10.07	-0.05	5.64	10.00
Markup (avg %)	18.9	18.4	18.1	18.9	18.4	18.1
Δ Markup (avg %)		-0.5	-0.8	0	-0.5	-0.8
Δ US Producer Surplus (b USD)	0.00	1.03	-2.58	1.80	3.11	-0.62
CS Δ total (b USD)	0.000 (0.0%)	-14.619 (-4.6%)	-32.831 (-10.4%)	4.817 (1.5%)	-10.219 (-3.2%)	-28.751 (-9.1%)
CS Δ Q1 (b USD)	0.000 (0.0%)	-0.551 (-5.5%)	-1.224 (-12.3%)	0.010 (0.1%)	-0.543 (-5.4%)	-1.219 (-12.2%)
CS Δ Q2 (b USD)	0.000 (0.0%)	-1.625 (-5.9%)	-3.478 (-12.6%)	0.334 (1.2%)	-1.311	-3.218 (-11.7%) (-4.8%)
CS Δ Q3 (b USD)	0.000 (0.0%)	-2.750 (-6.4%)	-6.195 (-14.5%)	0.694 (1.6%)	-2.129 (-5.0%)	-5.683 (-13.3%)
CS Δ Q4 (b USD)	0.000 (0.0%)	-3.699 (-5.1%)	-8.352 (-11.5%)	1.333 (1.8%)	-2.526 (-3.5%)	-7.317 (-10.1%)
CS Δ Q5 (b USD)	0.000 (0.0%)	-5.994 (-3.5%)	-13.583 (-7.9%)	2.446 (1.4%)	-3.711 (-2.2%)	-11.313 (-6.6%)
Δ vehicles sold (m)	0.000	-0.410	-0.898	0.201	-0.225	-0.724
EV share (% sales)	3.36	3.18	3.15	7.39	6.75	6.65
US-assembled share (% sales)	53.2	66.8	57.1	53.5	67.2	57.6
Δ US assembled (m)	0.000	1.248	-0.076	0.144	1.412	0.077
Tariff revenue (b USD)	0.000	15.38	41.20	0.000	15.67	41.86
EV subsidy cost (b USD)	0.000	0.000	0.000	6.68	6.00	5.56
Δ Net US outcomes (b USD)	0.00	1.80	5.79	-0.07	2.56	6.93

Notes: Δ entries report counterfactual outcomes relative to the 'no-tarff and no EV subsidies' baseline. Dollars are USD 2015. Δ Net US outcomes is the change in US producer and consumer surplus, plus tariff revenue, minus (plus) additional EV subsidy expenditure (savings) compared to baseline. US Producer Surplus counts profit changes for US-Headquartered firms. Consumer surplus (CS) changes are in billion USD; parentheses report percentage changes. "With subsidy" and "Without subsidy" refer to whether the EV subsidy policy is in place in the counterfactual.

namely:

$$\mathbf{E}_t[\mathbf{1}[\text{purchase}_{it}]|\text{income quintile}_{it}].$$

Income quintiles are constructed from the FMLI annual income, using survey weights to form weighted quintile cutoffs within each year.

Division dependent micro-moments We construct division micro-moments by aggregating state-based sales to the division level. We define the probability of a household purchasing an electric vehicle, conditional on their associated division, as:

$$P(EV|purchase, d, t) = \frac{\text{total EV sales}_{d,t}}{\text{total vehicle sales}_{d,t}},$$

where d denotes the division. Observed moments for Truck and SUV sales are computed analogously.