Inequality, Electric Vehicles, and Energy Transition

Qianqian (Ivy) Tang
Emory University

Abstract: Who gains from electric vehicle (EV) adoption? High-income households disproportionally benefit more. I show that high-income counties adopt more EVs and their adoption accelerates faster over time. To capture affordability and potential market catering behaviors, I estimate a discrete choice model with non-homothetic preferences and an oligopoly EV market using California county-model-year registration data from 2015–2023. Model-specific technology improvement and high-quality model entry account for 67.1% of the rise in income gradients, while federal subsidies explain only 0.6%. The results suggest that unequal clean technology adoption is driven less by subsidy than by the direction of product innovation and availability.

1 Introduction

Who buys electric vehicles (EVs) in the United States? EV adoption is skewed toward high-income households, in large part because EVs remain substantially more expensive than comparable gasoline vehicles: as of 2024, popular EV models average over $40,000, compared with roughly $33,000 for leading gasoline models. A large literature studies EV adoption and policy (Holland et al. (2016); Gillingham et al. (2022); Holland et al. (2019); Davis et al. (2023)), but it often abstracts from income effects, largely because quasi-linear utility is a common benchmark. Related work allows for heterogeneity in price sensitivity across the income distribution (Sabal (2025), Grieco et al. (2024)), but typically treats income as shifting price elasticities rather than imposing an affordability constraint on discrete vehicle choices. As a result, existing approaches may miss how income shapes both adoption patterns and the distribution of welfare.

This paper asks which factors contribute to differences in EV adoption across income groups when income effects are explicitly modeled. I build a static model with non-homothetic preferences and an oligopolistic EV market to study which factors explain differences in EV adoption across income groups. Non-homotheticity allows income to directly affect con-

sumers' vehicle choices. In the model, vehicle quality is complementary to the consumption of normal goods within a budget constraint, so affordability limits which options are feasible. Vehicle quality reflects technology, reliability, driving range, charging availability, and tastes. Consumers receive zero utility from options they cannot afford, while higher-income households are willing and able to pay more for higher-quality EVs. Individual income is drawn from a log-normal distribution. Compared with Berry et al. (1995) (BLP), which incorporates income primarily through price sensitivity, my framework emphasizes affordability and its interaction with available product offerings.

I estimate the model using 2015–2023 California new vehicle registration data at the county level, which allows me to estimate demand and recover marginal costs. California is the largest EV market in the United States, accounting for about 32% of U.S. EV sales, and its EV share among new light-duty vehicle purchases rose from roughly 2% in 2009 to about 25% in 2024. Therefore, estimating the model in California provides a useful lens for understanding adoption patterns and the drivers of income gradients.

Two empirical patterns motivate the quantitative analysis. First, high-income counties not only have higher EV adoption levels but also experience faster adoption growth; that is, the Engel-curve slope increases over time. I define the slopes of EV adoption share with respect to log income per capita at county level as income gradients. Income gradients do not only capture pure income effect but also qualities of affordable vehicle options.

Second, product entry has shifted toward more expensive EV models relative to cheaper options. Income differences may also have supply-side implications. Greater willingness to pay among high-income households can encourage manufacturers to introduce higher-priced EV models. Such demand-driven product entry may generate a feedback loop that amplifies differences in adoption across income groups and places, with direct implications for the distributional effects of EV subsidies. This mechanism may help explain the entry of high-quality EVs. However, because of the complexity of modeling endogenous product choices, this paper abstracts from endogenous firm entry and exit and instead treats the observed product set as given.

Counterfactual simulations quantify the mechanisms behind these adoption differences. Model-specific technology improvement and the entry of higher-quality new models account for about 67.1% of the change in income gradients. Federal tax credits explain only 0.6% of the change, although they explain about 13% of the differences in adoption levels across income groups. The counterfactuals show that holding model quality fixed or restricting entry significantly reduces the EV income gradient.

My contribution to is to distinguish between two related but conceptually different outcomes. The first is the difference in EV adoption across income groups at a given point

in time. The second is the change in that difference as the EV market develops. Federal credits meaningfully affect adoption levels, but the model attributes little of the increase in the income gradient to those credits. Instead, the widening gradient is primarily associated with model-specific quality improvement and the entry of higher-quality vehicles. This distinction helps reconcile the finding that subsidies are distributionally unequal in levels with the finding that they are not the principal source of the growing income gradient.

The results contribute to three literatures. First, the paper contributes to research on EV policy and environmental outcomes. Holland et al. (2016) emphasize that the environmental benefits of EVs depend on emissions from local electricity generation, while recent work studies the welfare effects and policy design of EV subsidies. This paper adds an adoption-inequality channel: even when subsidies are motivated by environmental objectives, their distributional consequences depend on which households adopt the technology. Second, the paper contributes to the literature on non-homothetic demand and product quality. In the spirit of Fajgelbaum et al. (2011); Couture et al. (2024) and related work, income differences shape the demand for differentiated products. Third, the paper contributes to empirical methods for differentiated-products demand with sparse market shares (Gandhi et al. (2023)). A BLP-style inversion combined with Poisson pseudo-maximum likelihood (PPML) provides a practical way to estimate quality when many product-market cells have zero observed shares.

2 Literature Review

This paper relates to four strands of literature. First, a large literature studies factors of EV adoption, including electricity and gasoline prices, political ideology, and technology improvements (Bushnell et al. (2025); Archsmith et al. (2021); Davis et al. (2023); Forsythe et al. (2023)). However, the literature has not focused on differences in adoption cross the income distributions. My focus is on how income endogenously shapes adoption patterns through affordability.

Second, related work documents the income distribution of EV buyers and studies when low-income households benefit from EV adoption (Gillingham et al. (2023); Muehlegger and Rapson (2018, 2022); Bauer et al. (2021)). Statistical evidence shows that EV buyers have higher incomes, but without a structural framework it misses responses from demand and supply in counterfactual scenarios. Methodologically, I introduce non-homotheticity preference following Faber and Fally (2022), Couture et al. (2024), and Fajgelbaum et al. (2011).

Third, research examines EV adoption and charging infrastructure, including charging-

station availability and charger standards (Zhou and Li (2018); Arkolakis et al. (2025); Li (2023); Ou et al. (2020)). Standardization and expanded charging access can reduce range anxiety and lower adoption barriers. In my model, charging availability enters as a county-specific amenity, and I discipline income effects while controlling for these county-level factors.

Finally, a growing literature studies policy implications and optimal EV policy design, including work on subsidies and the Chinese EV market (Kwon (2023); Barwick et al. (2025); Allcott et al. (2024); Holland et al. (2016); Head et al. (2026); Kwon et al. (2025)). Using data from the California EV market, I study how subsidies affect adoption inequality across income groups through both demand and endogenous supply responses. The policy exercise in this paper is narrower than a complete long-run analysis of subsidy incidence. In particular, I isolate the direct demand effect of federal tax credits while holding product quality and availability fixed.

3 Data

This section describes the data and presents the empirical patterns that motivate the analysis. I first define markets and products and explain how to construct model-specific market shares and the outside option. I summarize the federal consumer subsidies for EV buyers over the sample period.

3.1 Data Sources

The analysis combines vehicle registration records with information on vehicle prices, federal purchase incentives, county demographics, and local charging infrastructure. The primary data come from the California Department of Motor Vehicles and report new vehicle registrations by county, make, model, and year. The main estimation sample covers 2015–2023. I classify battery electric vehicles and plug-in hybrid electric vehicles as EVs.

I compile manufacturer suggested retail prices for EV models from publicly available vehicle-specification sources. To exclude ultra-luxury vehicles that are unlikely to compete in the same market as most EVs, I drop models with an initial MSRP above $200,000.[1] I convert nominal prices to 2022 dollars and annualize them over a ten-year ownership horizon. All price variables are expressed in thousands of dollars.

Information on federal tax-credit eligibility is obtained from the Internal Revenue Service

[1]The excluded observations are associated with Lamborghini, McLaren, Bentley, Ferrari, Aston Martin, Maserati, Bugatti, and Rolls–Royce.

and matched to vehicles by model and year. The subsidy enters the demand model through the effective purchase price: the annualized value of the federal tax credit is subtracted from the annualized vehicle price. This construction allows changes in eligibility across models and years to affect the relative prices faced by consumers.

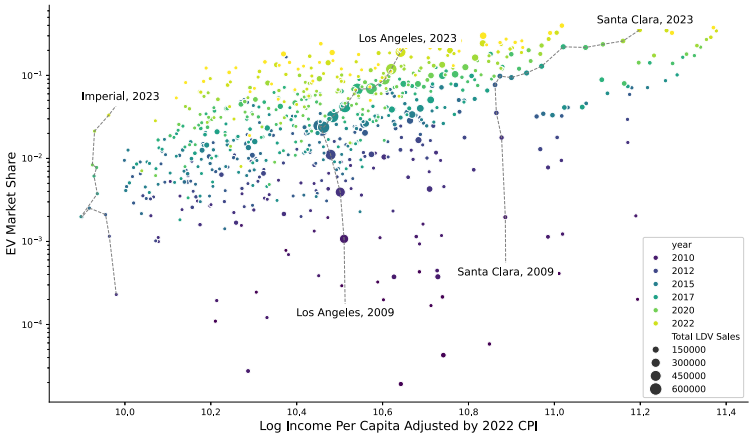
County-level demographic information comes from the American Community Survey and the Bureau of Economic Analysis. I use income per capita and the Gini coefficient to approximate the income distribution in each county and year as log-normal.

Data on charging infrastructure come from the Alternative Fuels Data Center. The station-level records contain information on location, charger type, opening date, and charging network. I use these records to construct county-year measures of charging availability, including the number of charging stations per capita and counts by charger type. I also construct county-model-year measures of access to charging stations affiliated with a vehicle model's parent company. The market definition is designed to exploit cross-county variation in adoption for the same model-year. Prices and federal subsidies vary by model-year, while income distributions, charging infrastructure vary across counties.

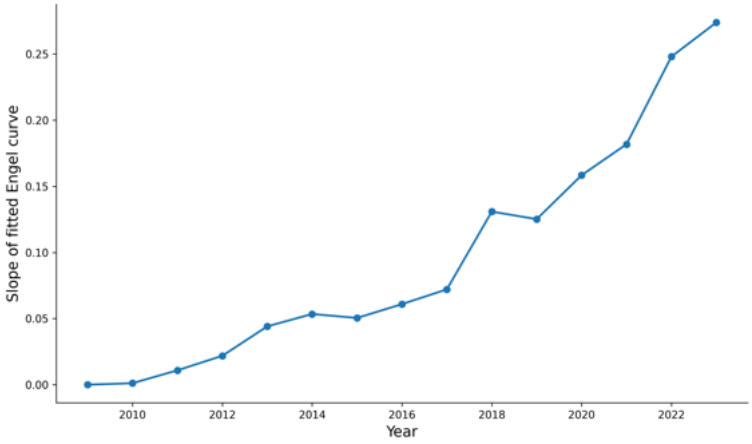
3.2 Federal Consumer Subsidies

Before the Inflation Reduction Act, federal support for electric vehicle purchases operated primarily through the Qualified Plug-in Electric Drive Motor Vehicle Credit under IRC §30D. The credit provided up to $7,500 for the purchase of a new qualifying plug-in electric vehicle, with the exact amount determined by the vehicle's battery capacity. In addition, eligibility was subject to a manufacturer-level phaseout: once an automaker sold 200,000 qualifying vehicles in the United States, the credit for that manufacturer's vehicles was gradually reduced and eventually eliminated.

The Inflation Reduction Act (IRA) of 2022 aims to accelerate the clean energy transition. Section 30D of the IRA provides tax credits for new, noncommercial EV purchases, subject to eligibility requirements. Under Section 30D, eligible vehicles must weigh less than 14,000 pounds (including sedans, sport utility vehicles, and pickup trucks) and must be assembled in North America. MSRP caps apply: $80,000 for SUVs, vans, and pickup trucks, and $55,000 for other vehicles. The credit is up to $7,500 and is split into two components tied to battery requirements: $3,750 based on critical mineral sourcing (from the United States or a country with a U.S. free trade agreement) and $3,750 based on battery component manufacturing/assembly (with a minimum 50% threshold required to be manufactured or assembled in North America in 2023). In addition, consumers' adjusted gross income must not exceed $300,000 for married couples filing jointly or $150,000 for single filers.



(a) Electric Vehicle Adoption over Income



(b) Slopes of Engel Curves over 2009-2023

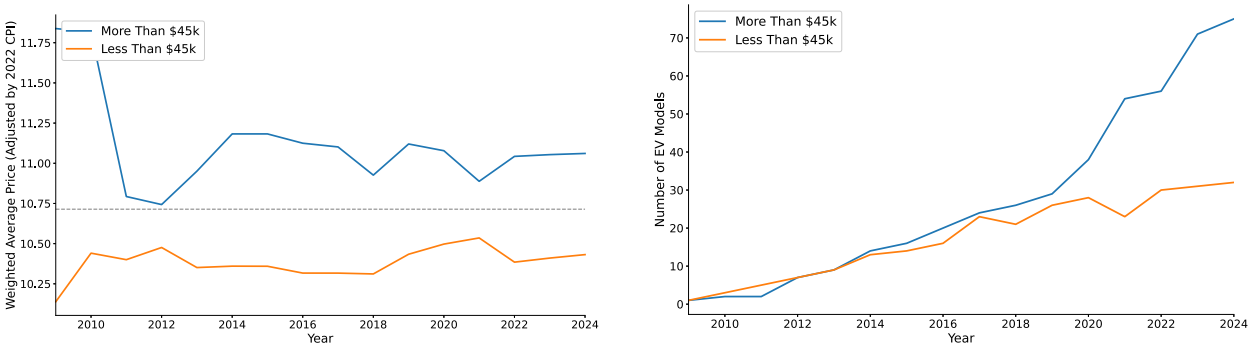
Figure 1: California EV Adoption Patterns

Notes: Panel (a) displays EV market shares across California counties from 2009 to 2023. Market size is measured by the number of new light-duty vehicle registrations at the county level. Each dot represents a county-year observation, and the size of the dot reflects the county's population. The color gradient from darker to lighter indicates the passage of time from 2009 to 2023. Income is adjusted to 2022 dollars. Panel (b) displays the slopes of the Engel curves, defined as $\frac{\text{Cov}(S_{EV}, \log \bar{y})}{\text{Var}(\log \bar{y})}$, calculated using observed EV market shares and county-level income per capita.

3.3 Stylized Facts

Figure 1a shows the relationship between new EV market shares and income per capita. Green transportation transition is uneven: high-income counties adopt more EVs than low-income counties in California. In Figure 1b, the slopes of income and EV adoption shares, income gradients, increase over time. EV adoption increases faster in high-income counties than in low-income counties, which indicates a growing income effect.

Demand from high-income groups also shapes the supply side. New EV models also tend to be more upscale. I classify vehicles into two groups, above $45,000 or lower. $45,000 is



(a) Weighted Average Prices

(b) Number of Unique Models

Figure 2: Market Entry

roughly the average price of an entry-level luxury vehicle in 2022. Market share weighted average prices for both high- and low-price vehicles have remained relatively stable since 2010 in Figure 2a[2]. However, the number of expensive EV models entering the market has grown faster than the number of cheaper models (Figure 2b). This pattern suggests that manufacturers increasingly target high-income consumers, which may widen EV adoption difference across income groups.

4 Model

This section presents a model of EV demand and vehicle supply. The main purpose of the model is to decompose the factors behind unequal EV adoption. The demand side incorporates non-homothetic preferences over vehicle purchases. The supply side is used to recover marginal costs and markups. I assume no endogenous entry or exit in this static model.

4.1 Demand System

Counties are indexed by $j \in \mathcal{J}$, EV models by $m \in \mathcal{M}$ and time by t. The outside option, denoted 0, is purchasing a non-EV. Household i has income y_{ijt}. Market size N_{jt}, the number of new light-duty vehicle registrations, is exogenous.

Each household purchases one new vehicle. Consumers take vehicle prices, subsidies, and available models as given. If household i in county j and year t chooses EV model m, direct

[2]Tesla Roadster was the only available pure battery electric vehicle in 2009.

utility is

$$u_{ijmt} = \log(q_{jmt}) + \alpha \log(c_{ijmt}) + \epsilon_{ijmt}$$
$$s.t.: \; c_{ijmt} = y_{ijt} - (p_{mt} - \tau_{mt})$$

where $p_{mt} - \tau_{mt}$ is effective EV purchase price is , p_{mt} is the vehicle price and τ_{mt} is the subsidy. Idiosyncratic taste shocks ϵ_{ijmt} follow a Type I extreme value distribution with scale parameter 1. α governs the importance of residual consumption relative to idiosyncratic vehicle tastes and vehicle quality. A larger α makes consumers more responsive to reductions in post-purchase consumption. This effect is particularly strong for consumers whose income is close to the vehicle's net price, since disposable income $y - p_{mt} + \tau_{mt}$ is small. q_{jmt} captures the quality, including vehicle physical attributes such as range, and vehicle type, as well as charging stations and political ideology.

The utility specification captures a simple form of non-homothetic demand. For a high-income household, the same EV price represents a smaller reduction in remaining consumption. For a low-income household, the same price can be a large burden. High-income households are willing to pay more on vehicles to obtain higher quality.

To evaluate non-income factors' contribution to EV adoption, I decompose qualities in the following way:

$$\tilde{q}_{jmt} = \bar{q} \times q_{mt} \times q_{jt} \times q^{int}_{jmt} \times q_t, \tag{1}$$

where $\bar{q}$ is average of $\tilde{q}_{jmt}$, q_{mt} is model-year quality, q_{jt} is county-year taste and infrastructure, q^{int}_{jmt} captures county-model interactions such as local charging networks tied to vehicle brands and charger types, and q_t is a general time effect.

Indirect utility of outside option is $U_{ij0t} = \alpha \log(y_{ijt}) + \epsilon_{ij0t}$. This normalization sets outside-option quality equal to one. All households can afford outside option. Households who cannot afford EV model m receive zero utility. Market shares integrate choice probabilities over the county income distribution:

$$s_{jmt}(G, p, \tau, q; \alpha) = \int \underbrace{\frac{q_{jmt} \left(y - p_{mt} + \tau_{mt}\right)^{\alpha}}{y^{\alpha} + \sum_{\ell \in \mathcal{M}} q_{j\ell} \left(y - p_{\ell t} + \tau_{\ell t}\right)^{\alpha}}}_{\text{Prob}_{mt}(y)} \, \mathrm{d}G_{jt}(y) \tag{2}$$

A distribution G_{jt} has location parameter μ_{jt} and scale parameter σ_{jt} which are estimated by average income and Gini Index.

To connect the structural model to the stylized facts, I define the county-level income gradient as the local slope of the aggregate EV Engel curve. Specifically, holding the dispersion of the county income distribution fixed, income gradient is the derivative of the total

EV share with respect to its location parameter, μ_{jt}:

$$\frac{\partial S_{jt}}{\partial \mu_{jt}} = \alpha \, \mathbb{E}_{G_j} \left[\text{Prob}_{j0t}(y, p, \tau, q; \alpha) \sum_{\ell \in \mathcal{M}} \text{Prob}_{j\ell t}(y, p, \tau, q; \alpha) \frac{p_\ell - \tau_\ell}{y - p_\ell + \tau_\ell} \right] > 0 \tag{3}$$

where $S_{jt} = \sum_{\ell \in \mathcal{M}} \int \text{Prob}_{j\ell t}(y, p, \tau, q; \alpha) dG_{jt}$ is the aggregate EV share, and Prob_{j0t} is the probability of choosing the outside option. The income gradient can change over time even when α is fixed because prices, subsidies, available models, and vehicle quality alter the choice probabilities in equation (3). Improvements in model quality can attract consumers at different points of the income distribution and thereby change the slope of the aggregate Engel curve. I use the quality decomposition in equation (1) to construct counterfactual income gradients and evaluate how model quality, county taste, model-county taste, and common year effects contribute to the evolution of the income gradient.

4.2 Production

I assume that each EV model is produced by a separate single-product firm. Let aggregate demand be $D_{mt}(\mathbf{p}_t, \boldsymbol{\tau}_t, \mathbf{q}_t; \alpha) = \sum_{j \in \mathcal{J}} N_{jt} s_{jmt}(G_{jt}, \mathbf{p}_t, \boldsymbol{\tau}_t, \mathbf{q}_t; \alpha)$. EV maker chooses price p_{mt} given product attributes and prices and attributes of competing goods

$$\begin{aligned} \max_{p_{mt}} \quad & p_{mt} x_{mt} - mc_{mt} x_{mt} \\ \text{s.t.} \quad & x_{mt} = D_{mt}(\mathbf{p}_t, \boldsymbol{\tau}_t, \mathbf{q}_t; \alpha) \end{aligned}$$

where mc_{mt} is the marginal production cost and x_{mt} is the total output. The first-order conditions for price is

$$p_{mt} - mc_{mt} = -\frac{x_{mt}}{\partial D_{mt} / \partial p_{mt}} \tag{4}$$

Aggregating this derivative across counties and income distributions gives

$$\Delta_{mt} \equiv \frac{\partial D_{mt}}{\partial p_{mt}} = -\sum_{j \in J} N_{jt} \int \frac{\alpha}{y - p_m + \tau_{mt}} \text{Prob}_{mt}(y)(1 - \text{Prob}_{mt}(y)) \, \mathrm{d}G_{jt}(y)$$

Increase in α raises semi-price elasticities Δ_{mt} when $\text{Prob}_{mt}(y) < \frac{1}{2}$ because consumers are more responsive to price change with large α.

4.3 Equilibrium

Given market sizes $\mathbf{N}$, qualities $\{q_{jmt}\}$ and marginal costs $\{c_{mt}\}$, a vector of EV prices $\mathbf{p}$ satisfies:

$$x_{mt}(\mathbf{p}) = \sum_j N_{jt} s_{jmt}(\mathbf{p})$$

Caplin and Nalebuff (1991) proves the existence and uniqueness of equilibrium under logit preferences and quasi-concave profit functions with single-product firms. Berry et al. (1995) and Berry (1994) show that mean-utility estimation does not rely on the uniqueness of the Nash–Bertrand equilibrium. In this model, consumer heterogeneity enters nonlinearly because income is not separable from product prices. Although I assume that each firm produces a single EV model, this assumption does not guarantee a unique equilibrium. As in Allcott et al. (2024), the numerical computations do not reveal evidence of multiple equilibria in this paper.

5 Estimation

In this section, I estimate demand system including α and qualities $\mathbf{q}$, recovering marginal cost $\{\text{mc}_{mt}\}$ by following Berry et al. (1995). I report the estimated parameters and the patterns of inverted fundamentals.

5.1 Demand Estimation

For a given value of α, I recover a unique vector of county-model-year qualities $q_{jmt}(\alpha)$ following Berry et al. (1995):

$$q_{jmt}(\alpha)^{n+1} \leftarrow q_{jmt}(\alpha)^n \times \frac{s_{jmt}^{\text{obs}}}{\hat{s}(G, \mathbf{q}, \mathbf{p}, \alpha)} \tag{5}$$

where $\hat{s}(G, \mathbf{q}, \mathbf{p})$ is model-implied market shares. The update continues until the predicted shares match the observed shares. Because the inversion is performed in levels rather than logs, observations with zero market shares can remain in the estimation sample.

The inversion sets $q_{jmt} = 0$ when the observed market share of model m in county j is zero. Using these inverted quality terms $\tilde{q}_{jmt}(\alpha)$ directly in counterfactual exercises would imply that the model is never chosen in that market, even when prices, subsidies, or other market conditions change. In addition, because the inverted qualities reproduce the observed shares, they contain sampling variation and other market-specific factors that are not directly related to product quality.

Therefore, I use PPML regression to decompose qualities into observables for two purposes: 1) avoiding omitting all zeros from taking logs; 2) finding the main contribution to EV adoption difference over time despite of income.

$$\frac{\tilde{q}_{jmt}(\alpha)}{\bar{q}(\alpha)} = \exp\left(X'_{jt}\beta_1 + X'_{mt}\beta_2 + X'_{jmt}\beta_3 + \kappa_t\right) + \vartheta_{jmt} \tag{6}$$

where X_{jt} contains county characteristics, X_{mt} contains model characteristics, and X_{jmt} contains variables that vary across counties and models. The year fixed effects κ_t capture changes in EV demand that are common across markets. The recovered quality terms are divided by their mean, $\bar{q}(\alpha)$, so that the fit depends on relative quality differences rather than the overall level of quality.

This second step serves two purposes. First, PPML allows zero recovered quality terms to remain in the sample. Second, the fitted values smooth market-specific noise and provide a measure of quality that can respond to observed product and county characteristics in counterfactual exercises.

I estimate the model over a range of values of α. For each candidate, I recover quality, estimate Equation (6), and calculate the Poisson deviance. The preferred estimate of α is the value that minimizes this deviance. I report several specifications with different sets of fixed effects and observed characteristics to assess the stability of the estimate and to examine how much of the variation in recovered quality can be explained by observable factors.

5.2 Estimation Results

The first panel of Table 1 reports the estimates of α, which governs the weight that consumers place on disposable consumption after purchasing a vehicle. A larger value of α implies a stronger response of vehicle demand to income and effective purchase prices. The estimates are positive in all specifications and range from 14.8 to 25.3. I use Column (5) as the preferred specification. It yields an estimate of 19.5, with a bootstrap standard error of 2.24.

The lower panel reports the PPML coefficients from the decomposition of recovered quality. Column (1) provides a saturated fixed-effects benchmark that includes county-by-year and model-by-year fixed effects. These fixed effects absorb time-varying demand conditions common to all EV models within a county and all characteristics common to a model within a year. The point estimate of α in this specification is 14.8 with a standard error 27.5, indicating that the parameter is estimated imprecisely.

Columns (2)–(4) replace the unrestricted model-by-year fixed effects with observed vehicle characteristics while retaining county-by-year fixed effects. Column (2) includes an

indicator for pure battery EVs and total driving range. Both variables are positively associated with recovered quality, and replacing the model-by-year fixed effects with these observable characteristics improves the precision of $\widehat{\alpha}$: its bootstrap standard error falls from 27.5 to 3.6.

Column (3) allows interaction between driving range and the indicator for pure battery electric. The coefficient on range is 0.137 for plug-in hybrids, while the corresponding coefficient for battery electric vehicles is the sum of the main and interaction effects, 2.168. Columns (4) and (5) yield a similar pattern, with implied range coefficients of 2.693 and 2.491 for battery EVs, respectively. Thus, longer driving range is associated with higher recovered quality for both types of vehicles, but the relationship is stronger for battery electric vehicles. This result is consistent with driving range being more important for pure battery EVs. Because range enters through an interaction and is not centered, the coefficient on the battery electric indicator in Columns (3)–(5) should not be interpreted separately from the range terms.

Column (4) further introduces indicators for vehicle body type. SUVs have a positive quality coefficient, whereas the coefficients for sedans and wagons, pickup trucks, and vans are negative. Among the reported categories, SUVs have the largest estimated association with recovered quality. The estimate of α also rises as driving-range interactions and body-type controls are introduced, from 14.8 in Column (2) to 25.3 in Column (4).

Column (5) replaces the county-by-year fixed effects with observed county characteristics and a common set of year fixed effects. This specification allows local demographics and charging infrastructure to account directly for differences in EV demand across counties. Counties with larger Democratic voter shares have higher recovered EV quality, consistent with the geographic sorting of EV demand documented by Davis et al. (2023). Population is also positively associated with recovered quality. The coefficient on parent-company charging stations per capita is positive and significant, suggesting that access to a manufacturer's charging network is an important correlate of demand for its vehicles. The coefficients on overall charging stations per capita and the number of charger types are positive but imprecisely estimated. The coefficient on the college share is also not statistically distinguishable from zero. These insignificant estimates are from colinearity of other controls. Overall, Column (5) provides a relatively precise estimate of α while accounting for model attributes and observable sources of geographic variation in EV demand.

Table 1: Demand System Parameter Estimates

	(1)	(2)	(3)	(4)	(5)
α	14.845	14.823	18.563	25.299	19.500
	(27.487)	(3.551)	(2.822)	(2.453)	(2.238)
	Quality Coefficients from PPML				
Pure Electric$_m$		1.500	-9.668	-12.473	-11.410
		(0.043)	(0.205)	(0.327)	(0.293)
ln(Total Range)$_{mt}$		0.601	0.137	0.123	0.120
		(0.027)	(0.013)	(0.014)	(0.014)
Pure Electric$_m$ $\times$ ln(Total Range)$_{mt}$			2.031	2.570	2.371
			(0.038)	(0.062)	(0.056)
Vehicle Type: Sedan/Wagon				-0.552	-0.491
				(0.060)	(0.061)
Vehicle Type: SUV				0.152	0.129
				(0.036)	(0.040)
Vehicle Type: Pickup				-1.118	-1.028
				(0.115)	(0.113)
Vehicle Type: Van				-0.136	-0.077
				(0.064)	(0.063)
ln(College Share)$_{jt}$					0.110
					(0.168)
ln(Democrat Voter Share)$_{jt}$					0.676
					(0.101)
ln(Charging Per Capita)$_{jt}$					24.364
					(20.821)
ln(Population)$_{jt}$					0.235
					(0.017)
ln(Charger Type Count + 1)$_{jmt}$					23.670
					(26.627)
ln(Parent Company Station Per Capita + 1)$_{jmt}$					182.768
					(33.682)
Year-Model FE	Yes				
Year-County FE	Yes	Yes	Yes	Yes	
Year FE					Yes
N	27,656	27,656	27,656	27,656	27,656
Deviance	62382.300	71640.650	64206.009	61363.951	63858.311
SER	1.547	2.513	2.203	2.022	2.148

Note: Standard errors are reported in parentheses. Standard errors of α are obtained by bootstrap.

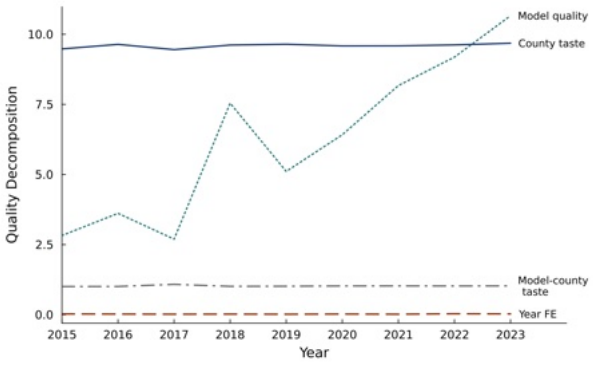
5.3 Quality Decomposition

The quality decomposition shows that the model component is the main source of the increase in recovered EV quality over time. Figure 3a reports the annual averages of model quality, county taste, model-county taste, and year effects. Average county taste remains close to 9.5 throughout the sample, while the model-county component changes little and the year effects remain near zero. In contrast, average model quality increases from approximately 3 in 2015 to more than 10 in 2023. County infrastructure and demographic information capture consistent tastes difference across locations, while upgrading of driving range and attributes in portfolios contributes to a growing trend of EV qualities.

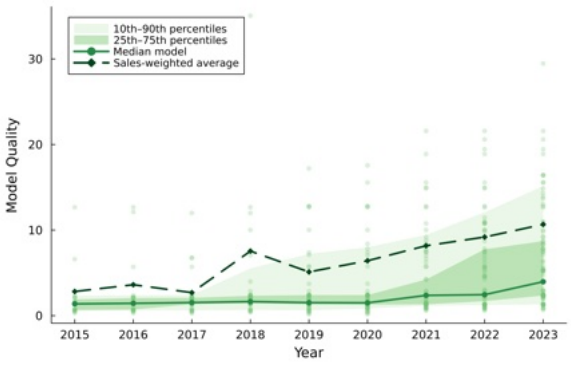
In my preferred specification of model quality, driving range is the only characteristic varying over time. Change in model quality mainly reflects on range improvement which benefits from battery cost deduction. The U.S. Department of Energy estimates that the high-volume manufacturing cost of a light-duty EV lithium-ion battery pack fell from $1,415 per kWh in 2008 to $139 per kWh in 2023, measured in 2023 dollars.[3] Lower battery costs made it less costly for manufacturers to install larger battery packs and extend driving range without a proportional increase in vehicle prices. It encourages more high-quality EVs to enter the market.

Figure 3b provides additional information about changes in the distribution of model quality.The 10th–90th percentile range of model quality widens substantially between 2017 and 2018, when the Tesla Model 3 was introduced, and continues to expand thereafter. The 25th–75th percentile range also widens after 2020, indicating an increase in the quality of models in the middle of the distribution. Together, these patterns provide a more complete picture of model entry and the expansion of the range of qualities available in the market. High-quality EV models increasingly enter the market, and the upper tail of the quality distribution expands sharply. This pattern motivates the counterfactual exercises that hold model quality fixed or remove model entry. These patterns could potentially indicate both an upward shift in model quality of existing models and an expansion of the range of qualities available to consumers. However, the distribution alone cannot distinguish the entry of new models from improvements in continuing models. This distinction motivates two separate counterfactual exercises: one that holds model quality fixed and another that holds the available product set fixed.

[3]The estimates refer to usable battery capacity and production of at least 100,000 units per year.

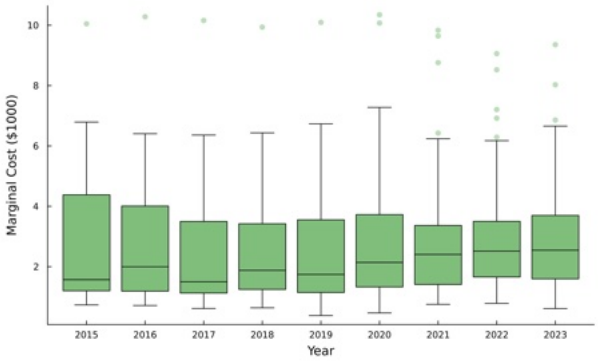


(a) Quality Decomposition

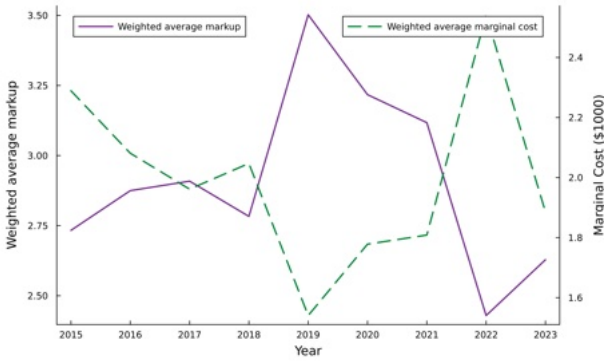


(b) Model Specific Quality Distribution

Figure 3: Quality Decomposition



(a) Marginal Cost Distribution



(b) Weighted Average Marginal Cost

Figure 4: Fundamentals of Firms

5.4 Marginal Cost and Markups

I recover model-level marginal costs and markups from the supply-side first-order conditions using observed vehicle prices. Figure 4a shows that the marginal cost distribution remains broadly similar over time, although the median increases modestly after 2019 and several high-cost outliers persist. Figure 4b shows that the weighted average marginal cost fluctuates across years without a sustained downward trend. Weighted average markups peak in 2019, decline through 2022, and partially recover in 2023. Overall, the estimates do not indicate a systematic decline in EV production costs over this period, and marginal costs are not the main force of EV adoption difference.

6 Counterfactual Analysis

Using the estimated demand system, I solve new equilibrium prices with recovered marginal cost. This section evaluates three counterfactual analyses: constrain quality change over time, removing new models entry and removing federal subsidies.

6.1 Common Quality Benchmark

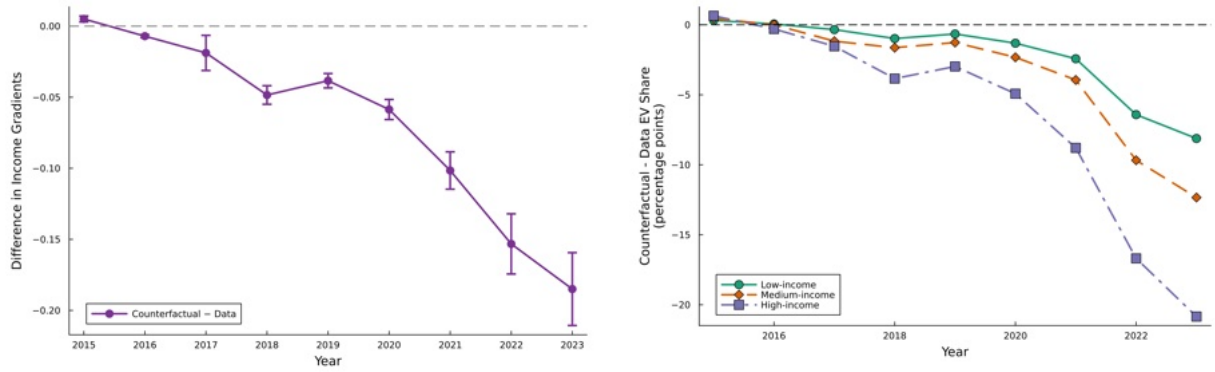
To measure how changes in model quality contributed to income gradient in EV adoption, I compare the baseline equilibrium estimated from data with a counterfactual in which every available EV model shares a common model quality equal to the median level in 2015, median($q_{m,2015}$). The observed set of available models, county income distributions, subsidies, marginal costs, and firm ownership remain unchanged. This exercise asks what EV adoption pattern would have looked like if EV model quality had not improved over years and all models shared the same quality.

Figure 5a shows the difference of income gradients under the common 2015 model quality from the baseline gradients. It stays negative after 2015, which indicates that removing quality upgrading and variations significantly reduces the speed of widening EV adoption inequality. The effect is especially large after 2020, when high-quality EV entry accelerates.

To disentangle the aggregated effect in the counterfactual scenario, Figure 5b displays average EV shares by counties classified into three income groups defined using income per capita in 2022 dollars. Lower model quality decreases adoption for all groups, but it decreases high-income adoption more. This is consistent with high-income households being more responsive to the arrival of high-quality EV models. Model-specific technology improvement is a major driver of unequal adoption, although the coefficients in the PPML regression are common across consumers. The heterogeneous response operates through vehicle affordability and substitution across models.

6.2 Fixed Product Set: No Entry

The second counterfactual remove model entry by holding the portfolio of EV models fixed over time as in 2015 (28 models). Only models available in 2015 are kept in consumers' choice set with their estimated qualities. It includes the effect of quality evolution of existing models and isolate the contribution of new entry on the equilibrium. To capture firms' price responses from product availability, I conduct two versions of this counterfactual, with and without price adjustment. In the second, prices are allowed to adjust endogenously to the fixed product set through the equilibrium.



(a) County EV-Income Gradient (b) Quality Contribution by Income Groups

Figure 5: Common Model Quality

Notes: Panel (a) reports the coefficient on log observed county per-capita income from a cross-county regression of model-implied EV shares. Each county share is obtained by integrating individual purchase probabilities over the county-specific income distribution. Panel (b) reports the change of counterfactual from baseline in EV share for county income groups defined using income per capita. Counterfactual prices are re-solved holding marginal costs, model availability, subsidies, and ownership fixed.

Figure 6 plots the predicted EV share for each county-year against log income per capita. The horizontal axis is a county-level characteristic rather than the income of an individual simulated consumer. Allowing prices to adjust leads to a modest additional reduction in markets shares due to decreasing competition, but the two counterfactual series remain close.

Holding the product set fixed substantially reduces EV adoption throughout the county income distribution. The reduction is particularly large among high-income counties. Under the baseline prediction, both the level and dispersion of EV shares increase sharply at the upper end of the county income distribution. Under the counterfactual without entry or exit, predicted shares remain concentrated at much lower levels. Thus, changes in product availability account for a substantial portion of both the increase in overall EV adoption and the steepening of the county-level income gradient.

The large difference from the baseline indicates that the expansion and changing composition of the EV choice set increased adoption on net, with larger gains in higher-income counties. This does not imply that product entry should be discouraged. Instead, it shows that the composition of entry matters. When entering models offer attributes that are disproportionately valued by higher-income consumers, product expansion can raise total EV adoption while also increasing the county-level income gradient. Entry in more affordable market segments may distribute the adoption gains more broadly.

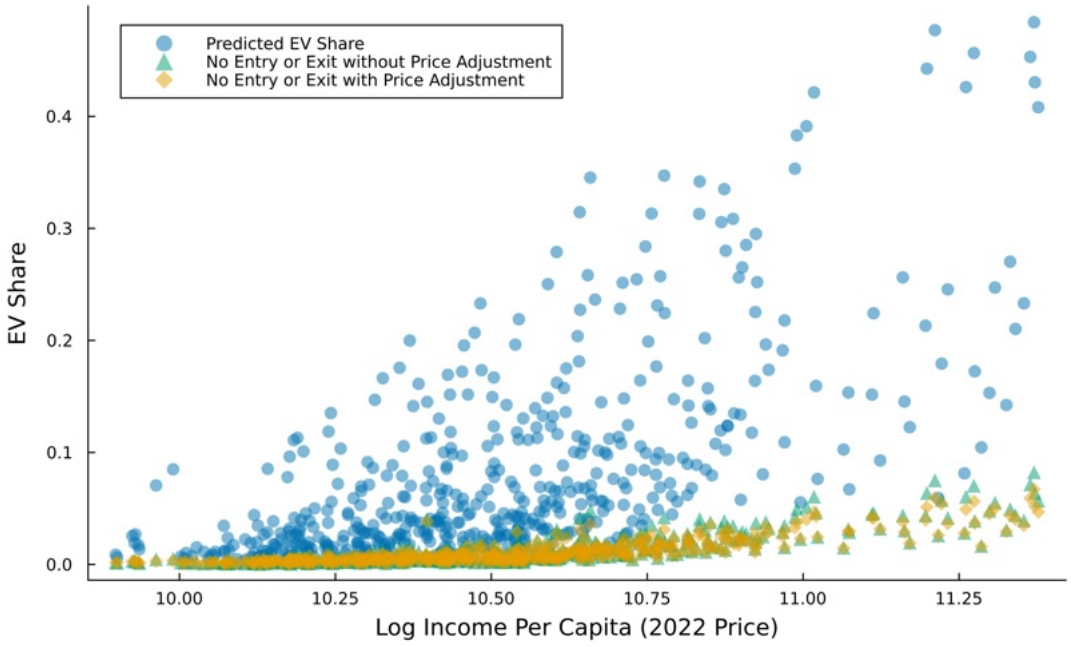


Figure 6: Predicted EV Shares Under a Fixed Product Set

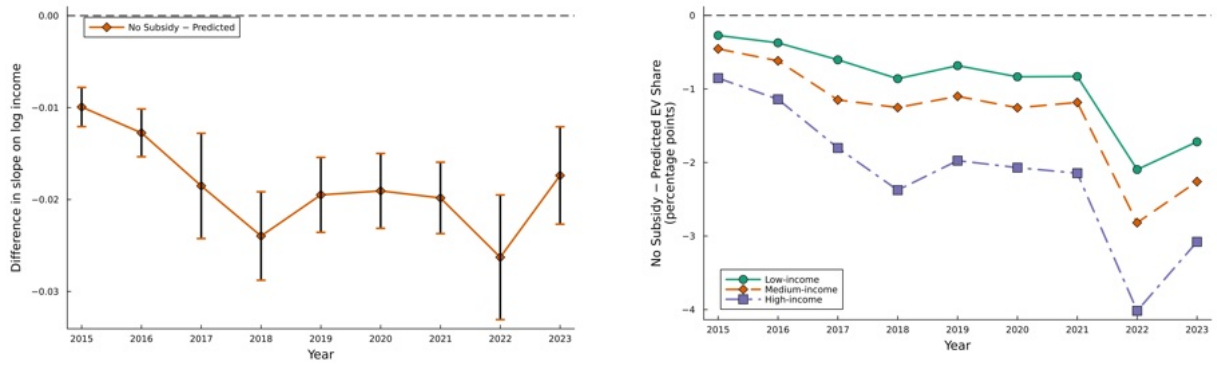
Notes: Each point represents a county-year observation. The horizontal axis reports observed log county income per capita in 2022 dollars. The vertical axis reports the model-predicted EV share after integrating household purchase probabilities over the county's estimated income distribution. The two counterfactuals hold the set of available EV models fixed, either maintaining baseline prices or allowing prices to adjust endogenously.

6.3 No federal tax credit

This counterfactual sets the federal tax credit to zero while leaving the set of available models and model-specific qualities at their baseline values in each year. Removing the credit increases consumers' effective purchase prices. The exercise isolates the direct demand effect of the federal tax credit. Panel (a) of Figure 7 reports the difference between income gradients in the no credit counterfactual and the baseline prediction. The differences are negative in every year, indicating that removing the federal tax credit makes the relationship between EV adoption and county income flatter. Federal tax credits therefore increase adoption more in absolute terms in higher-income counties and modestly steepen the income gradient. But the difference stays relatively stable compared to the first counterfactual scenario.

Panel (b) of Figure 7 shows the corresponding changes in predicted EV shares for counties grouped by income per capita. Removing the credit lowers EV adoption in all three county income groups. The decline is the largest in high-income counties and smallest in low-income counties. High-income households collect more subsidy benefits than the low income.

Quantitatively, federal tax credits account for approximately 13% of the adoption gap



(a) Change in the county income gradient (b) Change in EV shares by county income group

Figure 7: Effect of Removing the Federal Tax Credit

Notes: Panel (a) reports the income gradient coefficient in the no credit counterfactual minus the baseline-predicted coefficient. Panel (b) reports the no credit EV share minus the baseline-predicted EV share for counties grouped by income per capita. Negative values indicate that removing the credit reduces the corresponding outcome.

across county income groups, but only 0.6% of the increase in the income gradient over time. Thus, the credits contribute to cross-sectional differences in adoption levels but explain little of the increasing income gradient.

7 Conclusion

This paper examines how the income gradient in electric-vehicle adoption has changed and what explains its rise. Using California registration data from 2015 to 2023, I document that high-income counties adopt more EVs and their adoption speed is faster as well. Income gradients increased significantly over this period, suggesting an intensified income effects. I interpret these patterns using a discrete-choice demand with non-homothetic preferences and an oligopoly EV market. The model aggregates simulated individual purchase probabilities over each county's estimated income distribution. I recover model quality through a BLP-style inversion and use PPML to decompose the recovered quality terms.

Improvements in model-specific quality and the entry of new, higher-quality models together account for approximately 67.1% of the increase in income gradients. Federal tax credits account for only 0.6% of the change over time, although they explain approximately 13% of the difference in adoption levels across income groups. In the counterfactuals, holding model quality fixed or limiting the choice set to incumbent models substantially reduces the income gradient. Holding county-specific tastes or common time trends fixed has almost no effect (see plots in appendix).

These results show that quality upgrading can expand EV adoption unequally across in-

come groups due to affordability and consumers at high-income counties are more responsive to the arrival of high-quality and often higher-priced vehicles. Thus, Technology improvement contributed both to the growth of the EV market and to the widening inequality of adoption.

Distributional effects of clean energy transition depend not only on how quickly EV technology improves, but also on which products firms develop and bring to market. Purchase subsidies alone may do little to change adoption inequality when the available models remain concentrated in relatively luxury segments. Broader adoption may require complementary policies that support the development of affordable EV models, and reduce other barriers faced by low- and middle-income households. More broadly, assessments of the energy transition should account for how innovation, product entry, and policy jointly determine who adopts new technologies and when.

References

Allcott, H., Kane, R., Maydanchik, M. S., Shapiro, J. S., and Tintelnot, F. (2024). The Effects of "Buy American": Electric Vehicles and the Inflation Reduction Act.

Archsmith, J. E., Muehlegger, E., and Rapson, D. S. (2021). Future Paths of Electric Vehicle Adoption in the United States: Predictable Determinants, Obstacles and Opportunities.

Arkolakis, C., Gillingham, K. T., and Yoo, S.-Y. (2025). Optimizing Electric Vehicle Infrastructure. *AEA Papers and Proceedings*, 115:563–567.

Barwick, P. J., Kwon, H.-S., Li, S., and Zahur, N. B. (2025). Drive Down the Cost: Learning by Doing and Government Policies in the Global EV Battery Industry.

Bauer, G., Hsu, C.-W., and Lutsey, N. (2021). When might lower-income drivers benefit from electric vehicles? Quantifying the economic equity implications of electric vehicle adoption.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile Prices in Market Equilibrium. *Econometrica*, 63(4):841.

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

Bushnell, J., Muehlegger, E., and Rapson, D. S. (2025). Energy Prices and Electric Vehicle Adoption.

Caplin, A. and Nalebuff, B. (1991). Aggregation and Imperfect Competition: On the Existence of Equilibrium. *Econometrica*, 59(1):25.

Couture, V., Gaubert, C., Handbury, J., and Hurst, E. (2024). Income Growth and the Distributional Effects of Urban Spatial Sorting. *Review of Economic Studies*, 91(2):858–898.

Davis, L. W., Li, J., and Springel, K. (2023). Political Ideology and U.S. Electric Vehicle Adoption.

Faber, B. and Fally, T. (2022). Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data. *The Review of Economic Studies*, 89(3):1420–1459.

Fajgelbaum, P., Grossman, G. M., and Helpman, E. (2011). Income Distribution, Product Quality, and International Trade. *journal of political economy*.

Forsythe, C. R., Gillingham, K. T., Michalek, J. J., and Whitefoot, K. S. (2023). Technology advancement is driving electric vehicle adoption. *Proceedings of the National Academy of Sciences*, 120(23):e2219396120.

Gandhi, A., Lu, Z., and Shi, X. (2023). Estimating demand for differentiated products with zeroes in market share data. *Quantitative Economics*, 14(2):381–418.

Gillingham, K., Iskhakov, F., Munk-Nielsen, A., Rust, J., and Schjerning, B. (2022). Equilibrium Trade in Automobiles. *Journal of Political Economy*, 130(10):2534–2593.

Gillingham, K. T., Van Benthem, A. A., Weber, S., Saafi, M. A., and He, X. (2023). Has Consumer Acceptance of Electric Vehicles Been Increasing? Evidence from Microdata on Every New Vehicle Sale in the United States. *AEA Papers and Proceedings*, 113:329–335.

Grieco, P. L. E., Murry, C., and Yurukoglu, A. (2024). The Evolution of Market Power in the U.S. Automobile Industry. *The Quarterly Journal of Economics*, 139(2):1201–1253.

Head, K., Mayer, T., Melitz, M., and Yang, C. (2026). Industrial policies for multi-stage production: The battle for battery-powered vehicles.

Holland, S. P., Mansur, E. T., Muller, N. Z., and Yates, A. J. (2016). Are There Environmen-

tal Benefits from Driving Electric Vehicles? The Importance of Local Factors. *American Economic Review*, 106(12):3700–3729.

Holland, S. P., Mansur, E. T., Muller, N. Z., and Yates, A. J. (2019). Distributional Effects of Air Pollution from Electric Vehicle Adoption. *Journal of the Association of Environmental and Resource Economists*, 6(S1):S65–S94.

Kwon, H.-s. (2023). Subsidies versus Tradable Credits for Electric Vehicles: The Role of Market Power in the Credit Market.

Kwon, H.-s., Allcott, H., and Snyder, T. (2025). Equitable Energy Transitions? The Efficiency and Distributional Effects of Subsidies for Used Electric Vehicles. *SSRN Electronic Journal*.

Li, J. (2023). Compatibility and Investment in the U.S. Electric Vehicle Market.

Muehlegger, E. and Rapson, D. (2018). Understanding the Distributional Impacts of Vehicle Policy: Who Buys New and Used Alternative Vehicles?

Muehlegger, E. and Rapson, D. S. (2022). Subsidizing low- and middle-income adoption of electric vehicles: Quasi-experimental evidence from California. *Journal of Public Economics*, 216:104752.

Ou, S., Lin, Z., He, X., Przesmitzki, S., and Bouchard, J. (2020). Modeling charging infrastructure impact on the electric vehicle market in China. *Transportation Research Part D: Transport and Environment*, 81:102248.

Sabal, A. (2025). Product Entry in the Global Automobile Industry.

Zhou, Y. and Li, S. (2018). Technology Adoption and Critical Mass: The Case of the U.S. Electric Vehicle Market. *The Journal of Industrial Economics*, 66(2):423–480.

Appendix

A Data Sources

Vehicle registrations The main data source is new-vehicle registration data by county, make, model, and year from the California Department of Motor Vehicles. The sample covers 2015–2023 for the main estimation. I focus on electric vehicles, including battery electric vehicles and plug-in hybrid electric vehicles, and define the outside option as new non-EV vehicle purchases.

Vehicle prices I collect electric vehicle manufacturers' suggested retail prices from public vehicle-specification sources. Vehicles with an initial MSRP above $200,000 are excluded to avoid including ultra-luxury models whose adoption may reflect a different market segment. These makes are Lamborghini, McLaren, Bentley, Ferrari, Aston Martin, Maserati, Bugatti, and Rolls–Royce. Prices are adjusted using the 2022 CPI and transformed into annualized user costs based on 10 years of ownership. All prices are measured in units of $1,000.

Subsidies Federal tax-credit eligibility for each vehicle is obtained from the Internal Revenue Service. Annualized subsidies are incorporated by reducing the effective purchase price.

County demographics County demographic variables come from the American Community Survey and the Bureau of Economic Analysis. Income per capita and the Gini index are used to construct the county-level income distribution. Under the log-normal approximation, each county-year income distribution is summarized by mean income and a measure of inequality.

Charging infrastructure Charging-station data are obtained from the Alternative Fuels Data Center. These data include station location, charger type, opening date, and network information. I aggregate charging infrastructure to the county-year and county-model-year levels, including measures such as charging stations per capita, counts by charger type, and parent-company-owned charging stations.

B Estimation Algorithm

Following Sabal (2025), the log-normal distribution G_j is parameterized by the scale parameter

$$\sigma_j = \sqrt{2}\Phi^{-1}\left(\frac{\text{Gini}_j + 1}{2}\right),$$

and the location parameter

$$\mu_j = \log\left(\text{Income per capita}_j\right) - \frac{\sigma_j^2}{2}.$$

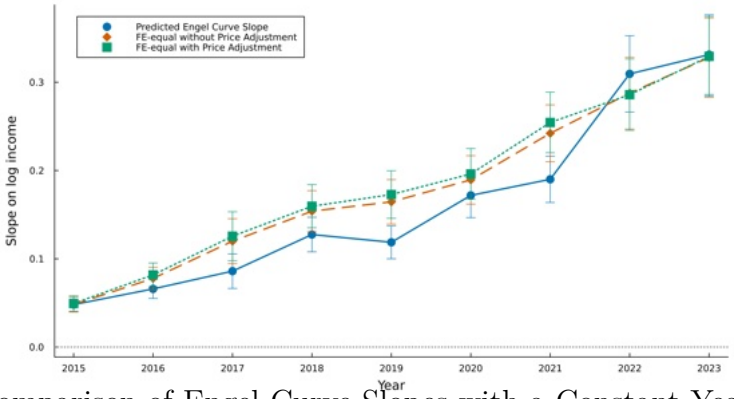


Figure C.1: Comparison of Engel Curve Slopes with a Constant Year Fixed Effect

For a given value of α, I recover a unique vector of county-model-year qualities $q_{jmt}(\alpha)$ following Berry et al. (1995):

$$q_{jmt}(\alpha)^{n+1} \leftarrow q_{jmt}(\alpha)^{n} \times \frac{s_{jmt}^{\text{obs}}}{\hat{s}(G, \mathbf{q}, \mathbf{p}, \alpha)}$$

The value of α is chosen to optimize the fit of the quality decomposition. For each candidate value of α, I recover $\widetilde{q}_{jmt}(\alpha)$ and estimate the PPML regression. The preferred α minimizes a deviance-based criterion:

$$\hat{\alpha} = \arg\min_{\alpha} D\left(\widetilde{q}(\alpha), \hat{q}(\alpha)\right),$$

where $D(\cdot)$ is the Poisson pseudo-maximum-likelihood (PPML) deviance, and

$$\hat{q}(\alpha) = \exp\left(X_{jt}'\hat{\beta}_1 + X_{mt}'\hat{\beta}_2 + X_{jmt}'\hat{\beta}_3 + \hat{\kappa}_t\right).$$

Standard errors for α are computed using bootstrap resampling.

Identification of α comes from the interaction among income distributions, prices, subsidies, and model-level adoption across counties. When high-income counties disproportionately adopt expensive, high-quality EVs, the model attributes part of this pattern to non-homothetic demand and part to recovered quality. The PPML decomposition disciplines the recovered quality estimates by linking them to observed model and county characteristics.

C Additional Counterfactuals

C.1 Constant Time Fixed Effect

The second counterfactual holds the general EV taste component fixed, i.e., $q_t = q_{2015}$. This exercise removes broad time-varying changes in EV tastes that affect all counties and models similarly. Examples include national trends in EV acceptance and general technological improvement.

Figure C.1 shows that removing the general time trend does not mitigate the upward

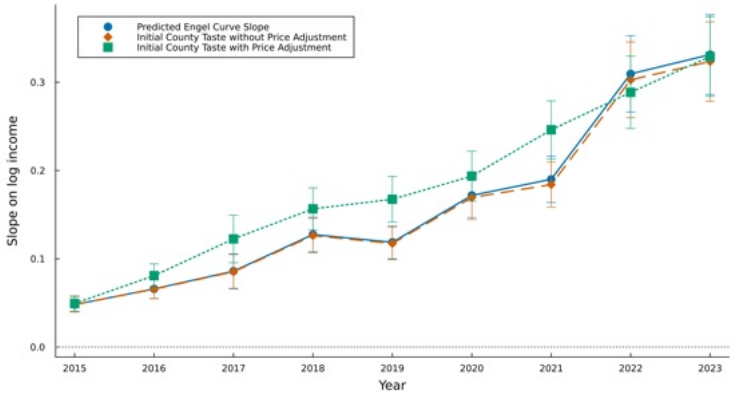


Figure C.2: Comparison of Engel Curve Slopes with Initial County-Specific Taste

trend in the income gradient. The counterfactual Engel curve slope closely follows the baseline slope. This result implies that broad growth in EV tastes or technology is not the main reason adoption grows faster in high-income counties. Instead, the inequality pattern is more closely tied to model-specific quality and entry.

C.2 Constant County-Specific Taste and Infrastructure

The third counterfactual holds county-specific taste and infrastructure fixed at their initial levels, $q_{jt} = q_{j,2015}$. This exercise asks whether improvements in county-level infrastructure and local EV tastes explain changes in adoption inequality.

Figure C.2 shows that holding county-specific taste and infrastructure fixed does not reduce the income gradient. The counterfactual slope remains close to the baseline slope. One interpretation is that county-level infrastructure improvements are correlated with adoption but are not the main source of the increasing income gradient. Meanwhile, more than 80% of EV users charge at home, according to the U.S. Department of Energy, and income distributions remain stable over the sample period.