

# Heterogeneous Diffusion of Electric Vehicles in China: Demand, Learning, Product Entry, and the Incidence of Industrial Policy

Yu (Jasmine) Hao\*

Jinge Li †

June 2026

## Abstract

China's electric vehicle (EV) sales share rose from about 1% in 2015 to approximately 45% in 2024. We evaluate this transition using an equilibrium differentiated-products model and two complementary counterfactual exercises. A Shapley decomposition attributes the 2015–2024 rise in EV sales share primarily to improved product attributes (45.5%), expanded EV product variety (14.8%), and battery cost reductions (8.2%), with offsetting effects from the phaseout of direct subsidies (–13.6%), residual demand and cost shifts (–8.0%), and changes in income, market size, and demographics (–3.3%). To distinguish the historical contribution of subsidy phaseout from the value of the subsidy program, we then study a dynamic counterfactual without subsidies. Eliminating subsidies throughout the transition from 2015 to 2024 lowers the 2024 EV share by between 23% and 33% and implies a private surplus return of RMB 3.38 per RMB of subsidy spending. These gains are unevenly distributed: per capita surplus losses from subsidy removal are about five times larger in Tier 1 cities—the largest and highest-income metropolitan areas—than elsewhere. About half of the aggregate subsidy effect arises through indirect Wright's law learning rather than the subsidy's direct price effect. Under subsidy removal, EV-focused automakers retain between 16% and 27% of their 2024 EV sales, compared with only 11% for traditional state-owned automakers, highlighting that learning effects are concentrated among EV specialists.

**Keywords:** electric vehicles, technology diffusion, Shapley decomposition, industrial policy, differentiated products, BLP demand

---

\*Hao: Faculty of Business and Economics, The University of Hong Kong, haoyu@hku.hk;

†Li: Harris School of Public Policy, The University of Chicago, jingeli@uchicago.edu.

We thank seminar participants at the HKU Brown Bag Seminar and the SUFE Empirical IO Conference, as well as Shanjun Li, Mark Jacobsen, Mark Roberts, Xiaodong Zhu, and Hyuk-soo Kwon, for helpful comments.

**JEL Classification:** D12, H23, L62, O33, Q48, L13

# 1 Introduction

Technological diffusion is straightforward to observe but its underlying drivers are difficult to disentangle. Diffusion depends not only on the arrival of a new technology, but also on whether firms incorporate it into attractive products, whether production costs fall, whether consumers substitute toward the new products, and whether policy support arrives at the right stage of market development. The empirical challenge is therefore not simply to document diffusion, but to distinguish among these forces when they move together. This distinction is especially difficult in markets with endogenous prices and differentiated products, where technology, entry, costs, and demand interact in equilibrium. A new technology may have limited effect when product variety is thin or costs remain high, but the same technology may diffuse rapidly once quality improves, firms enter, and consumers face a broader and more competitive choice set.

This paper studies these mechanisms in the context of China’s electric vehicle transition. China provides an important setting because EVs grew from a niche technology into a central part of the passenger vehicle market within a single decade. The national EV share rose from 1.03% in 2015 to 44.66% in 2024. This rapid diffusion occurred alongside major technological improvement, substantial product entry, declining battery costs, changing subsidy policy, and significant reallocation of demand across cities and firms. Given a transition of this magnitude, two questions follow. First, how much of the EV share rise from 2015 to 2024 can be attributed to each policy and market force? Second, what are the welfare and reallocation consequences when the policy stack that supported this transition is removed? We answer both questions using an equilibrium structural model of the Chinese auto market.

The attribution question is difficult because these forces are complementary rather than additive. Declining battery costs matter more when firms offer attractive EV products. Product entry matters more when EV quality has improved and consumers are willing to consider EVs. Subsidies also have different effects in an early market with few EV options than in a mature market with many high quality products. For this reason, a decomposition that introduces channels in one arbitrary sequence can assign too much importance to some forces and too little to others.

To address this complementarity and order dependence, we estimate an equilibrium differentiated products model of the automobile market. The model links product data and market data to equilibrium prices, quantities, and substitution across gasoline vehicles and EVs. We use the model in two ways. First, we decompose the EV transition from 2015 to 2024 using a Shapley decomposition, based on the cooperative game theory Shapley value [Shapley, 1953, Shorrocks, 2013]. The decomposition switches groups of economic primitives from their 2015 values to their endpoint year values and separates six channels: Quality, Variety, Battery, Subsidy, Residual, and Market. The Quality and Battery blocks

capture product side productivity gains, namely longer range and lower marginal cost per unit of vehicle service. The other blocks determine how these gains translate into adoption across products, firms, and cities. Second, we use the same equilibrium model to conduct no subsidy counterfactuals that remove policy support and trace the consequences for adoption, welfare, and reallocation.

The decomposition shows that product improvement was the central force behind China's EV transition. The Quality block contributes 45.49% to the EV share change from 2015 to 2024. This block captures the within product attribute trajectory, including driving range, engine power, vehicle size, and their interactions, for products present in both endpoint sets. The implication is that the 2024 EV was not simply a cheaper version of the 2015 EV. It was a substantially different product, with longer range, more competitive power and size, and a different body and size segment composition. Product market expansion and cost decline were important complements to this quality improvement. The Variety block contributes 14.81%, reflecting the expansion of the EV choice set, while the Battery block contributes 8.20%, reflecting battery related cost reductions. These channels are not independent: lower battery costs matter most when firms can embody them in appealing EV products, and entry matters most when EV quality is already high.

The policy results show why the historical decomposition and the policy counterfactuals must be interpreted separately. In the Shapley decomposition, the Subsidy block compares the 2024 subsidy schedule with the 2015 schedule. Because direct purchase subsidies were phased down between these dates, this block contributes  $-13.63\%$ . The negative sign does not imply that subsidies reduced EV demand. The no subsidy counterfactual asks a different question: what would have happened if subsidy support had been removed? In that exercise, subsidies accelerate adoption by increasing early demand, supporting cumulative Wright's law learning, and sustaining product availability. The counterfactual also shows that the incidence of subsidy policy depends on the stage of diffusion. During the active subsidy years, EV adoption was concentrated in high income cities, so historical subsidy benefits accrued disproportionately to richer markets. By 2024, EV adoption had diffused much further down the city income distribution, but direct purchase subsidies had been phased out. Applying the 2015 subsidy level to the 2024 market produces a more balanced distribution of benefits across city income tiers. The same policy instrument can therefore have very different distributional effects depending on when it is applied.

In addition, the no subsidy counterfactual reveals how policy removal affects welfare and producer reallocation. The welfare loss from removing subsidies operates through both a direct cash transfer channel and an indirect Wright's law learning channel. These two channels contribute similar shares, 52% direct and 48% indirect. The indirect share is higher in Tier 1 cities than in the Rest tier because the Wright channel works through equilibrium EV prices, which generate larger consumer surplus effects where EV adoption

is concentrated. On the producer side, subsidy removal reallocates EV production across firm types. EV native firms, including BYD, Tesla, and the New Forces, retain from 16% to 27% of their 2024 EV business, while traditional state owned OEMs retain only 11%. This producer incidence asymmetry reflects a technology driven reallocation: the same equilibrium forces that deliver consumer surplus to high tier cities also concentrate producer surplus among firms whose product portfolios are most exposed to the EV margin.

The paper contributes to three strands of literature. The first is the empirical industrial organization literature on differentiated product markets, new products, and product design. We build on the demand and supply framework of [Berry et al. \[1995\]](#), [Nevo \[2001\]](#), and [Berry et al. \[2004\]](#), in which observed and unobserved product attributes shape substitution patterns and equilibrium markups are recovered from multiproduct pricing. [Berry and Haile \[2014\]](#) clarify the market level identification problem in this class of models. We also relate to work showing that new products and changes in product characteristics can have large welfare and market structure effects [[Petrin, 2002](#), [Fan, 2013](#), [Sweeting, 2013](#), [Wollmann, 2018](#)]. In automobile markets, [Reynaert \[2021\]](#) shows that firms respond to environmental regulation by reshaping product portfolios rather than only adjusting prices. Our contribution is not to propose a new demand estimator, but to use an equilibrium differentiated products framework to decompose a large technology transition into quality improvement, variety expansion, cost decline, policy, and demand reallocation.

The second strand studies environmental and energy technology policy. [Newell et al. \[1999\]](#), [Popp \[2002\]](#), and [Acemoglu et al. \[2012\]](#) provide the broader logic for directed technical change, while [Ryan \[2012\]](#) and [Fowlie et al. \[2016\]](#) show that environmental policy can reshape industry dynamics rather than only static prices. In the Chinese context, [Aghion et al. \[2015\]](#) show that industrial policy subsidies raise productivity more when they target more competitive sectors. In the EV context, [Li et al. \[2017\]](#) and [Springel \[2021\]](#) emphasize indirect effects through complementary infrastructure, installed base, and expectations. On subsidy incidence, [Sallee \[2011\]](#) shows that hybrid vehicle tax credits accrue disproportionately to high income early adopters, while [Muehlegger and Rapson \[2022\]](#) estimate how EV subsidies affect adoption among low and middle income households. We contribute by placing direct purchase subsidies inside a broader equilibrium transition that also includes quality improvement, product entry, battery cost decline, and demand reallocation across cities and firms.

The third strand studies dynamic adoption, learning, and decomposition methods. Durable goods demand is dynamic because consumers can wait, firms can adjust future offerings, and current sales can affect future market conditions [[Gowrisankaran and Rysman, 2012](#)]. Our model is not a fully dynamic durable goods model, but the interpretation of subsidies and range is dynamic in this sense: early EV sales can affect later battery costs,

product availability, and endogenous quality choices. We capture the learning channel using Wright’s law, following the learning curve tradition that begins with [Wright \[1936\]](#) and the broader evidence on learning and scale in industrial production [[Benkard, 2004](#), [Thompson, 2012](#), [Bollinger and Gillingham, 2019](#)]. For EV batteries, [Nykqvist and Nilsson \[2015\]](#) and [Ziegler and Trancik \[2021\]](#) estimate learning rates for lithium ion battery prices; the latter provides the benchmark against which we report our  $22.8\% \pm 2.3\%$  Wright fit to BNEF and IEA data. Finally, we use the Shapley framework [[Shapley, 1953](#), [Shorrocks, 2013](#)] because the mechanisms are complementary rather than separable. The Shapley decomposition provides a transparent average over possible sequences and makes order dependence visible. In this sense, the wide Shapley ranges are an empirical finding about complementarity across channels, not simply an estimation inconvenience.

The rest of the paper is organized as follows. Section 2 describes the policy background and data. Sections 3 and 4 present the demand and supply models. Section 5 uses the estimated model to decompose EV diffusion from 2015 to 2024 and reports diagnostics for the decomposition. Section 6 studies policy counterfactuals and incidence. Section 7 concludes.


## 2 Policy Background and Data

### 2.1 Policy Background

China’s EV policy combined direct consumer incentives with broader industrial policy, charging infrastructure, and local restrictions on internal combustion vehicles. For EV adoption, the most visible demand side policies were direct purchase subsidies and purchase tax benefits. Their generosity changed substantially over the sample period. In the early sample years, national and local cash subsidies substantially lowered EV purchase prices. After 2022, national cash subsidies ended, and the main direct consumer incentive observed in the data was the purchase tax exemption. Figure 1 summarizes the direct purchase subsidies and tax benefits used in the analysis.

Direct fiscal incentives were only one part of this policy environment. Local license plate rules created another important demand side incentive. Several large cities restrict the issuance of license plates for internal combustion vehicles through nonprice lotteries or paid auctions, including Beijing, Shanghai, Guangzhou, Shenzhen, Tianjin, and Hangzhou. These restrictions make plate access scarce. Because eligible EVs often receive plate access outside these systems, EV buyers receive an in kind benefit in addition to any fiscal subsidy. In Shanghai, for example, a license plate for an internal combustion vehicle trades for roughly RMB 90,000 in the auction, while eligible EVs are exempt from the auction. The implicit value of plate access can therefore be comparable to early cash subsidies in restricted cities. The direct subsidy measure in Figure 1 does not include this

Figure 1: Direct purchase subsidies and tax benefits for electric vehicles



Notes: The panels plot direct consumer incentives for electric vehicle purchases. Panel (a) plots the national cash subsidy per vehicle. Panel (b) plots the total direct purchase benefit per vehicle, separating the national cash subsidy, local cash subsidy, and purchase tax exemption. The local cash subsidy is the additional cash subsidy provided by local governments on top of the national subsidy in the early sample years. Cash subsidies declined in steps and ended after 2022. In 2023 and 2024, the measured direct consumer benefit consists mainly of the purchase tax exemption. All values are reported in ten thousand RMB per vehicle.

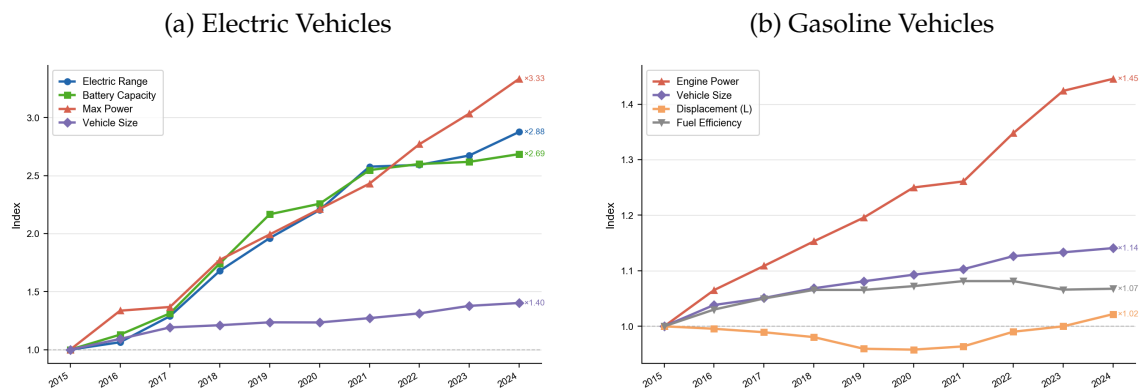
in kind plate benefit. Its value varies across cities and over time with auction prices, non-price lottery odds, and local eligibility rules, and reliable data are not available for all city and year combinations in the sample. Rather than assign an imputed plate value to every city and year, the empirical specification treats plate priority as part of the local EV policy environment, captured by demand fixed effects and interactions between the EV indicator and city indicators. This choice keeps the fiscal subsidy series directly tied to observed monetary policy instruments. It also means that the measured subsidy is conservative in restricted cities. As a scale check, applying a RMB 90,000 plate value to 0.045 million EV sales in Tier 1 cities in 2024 gives about RMB 4 billion in annual plate value. This calculation is not included in the subsidy measure, but it illustrates that omitted in kind benefits can be economically meaningful.

This distinction helps interpret the empirical results. At a point in time, subsidies and tax benefits lower consumer net prices and raise EV demand. Over the full sample period, however, direct fiscal support became less generous while the market also changed through quality improvements, product entry, battery cost reductions, and income and demographic shifts. As a result, the historical contribution of the direct subsidy schedule can be negative even though subsidies raise EV demand within a given policy environment.

## 2.2 Progress in Vehicle Product Attributes

The policy changes described above occurred alongside rapid improvement in EV product attributes. This product progress is important because EV adoption was not driven only by lower purchase prices. Over the sample period, consumers also faced EVs with longer driving ranges, larger battery packs, higher power, and broader product offerings. Figure 2 summarizes these changes by comparing attribute growth for EVs and gasoline vehicles. The figure shows that EV attributes improved much more rapidly than comparable gasoline vehicle attributes. Among EVs, electric driving range, battery capacity, and maximum power all rose sharply between 2015 and 2024, while vehicle size increased more gradually. These patterns indicate that EVs became more capable products, not simply cheaper products. The improvement in driving range and battery capacity is especially relevant because range anxiety and charging needs were major barriers to early EV adoption. Gasoline vehicles also improved, but along different dimensions and at a slower pace. Engine power increased steadily, while vehicle size, displacement, and fuel efficiency changed only modestly. This contrast suggests that the relative quality of EVs improved substantially over the sample period. The empirical model therefore allows observed product attributes to shift demand and allows EV technology attributes, such as driving range and battery capacity, to play a direct role in consumer choice.

Figure 2: Vehicle Attribute Improvements



*Notes:* The panels plot annual medians of product attributes, indexed to their 2015 values. Panel (a) reports electric vehicle attributes, including electric driving range, battery capacity, maximum power, and vehicle size. Panel (b) reports gasoline vehicle attributes, including engine power, vehicle size, displacement, and fuel efficiency. For gasoline vehicles, fuel efficiency is measured so that a higher index indicates lower fuel consumption. Text labels at the end of each series report the 2024 index value relative to 2015. All medians are computed across products within each year.

## 2.3 Data

The main estimation sample is a panel of passenger vehicles observed by product, city, and year from 2015 to 2024. A product is defined as a vehicle model and fuel type combination. Fuel types include battery electric vehicles, plug in hybrids, range extended electric vehicles, internal combustion vehicles, and hybrids. The demand sample contains 476,787 observations across 790 city and year markets.

Table 1 reports summary statistics for the three empirical units used in the analysis. Panel A describes the demand estimation sample, Panel B describes the supply estimation sample, and Panel C describes city level market covariates.

The demand system combines consumer net prices, observed vehicle attributes, local market characteristics, and policy variables. Consumer net price equals MSRP less any applicable national cash subsidy, local cash subsidy, and purchase tax exemption. Product attributes include power, vehicle size, body type, size segment, fuel type, and firm group, as well as technology specific measures such as EV driving range, battery capacity, electricity consumption, and gasoline fuel consumption. Local market variables include income, GDP per capita, population density, urbanization, population, EV license priority, pilot city status, EV policy strength, and oil prices. Market size and the outside option are constructed consistently at the city and year level and are used throughout the analysis.

For the supply side, the data are aggregated to the model and year level. Battery costs are measured using the BloombergNEF annual battery pack cost series. The battery pack cost falls from USD 373 per kWh in 2015 to USD 115 per kWh in 2024. This cost series enters the marginal cost regression directly as an observed cost shifter. Thus, references to battery learning should be read as a reduced form label for the observed battery cost path, which reflects industry learning, scale, and input cost dynamics. The analysis does not separately estimate a structural Wright's law learning elasticity.

Together, these data provide variation across products, technologies, cities, and years. The product variation identifies how consumers value vehicle attributes and prices, while the city and year variation captures differences in income, market size, policy exposure, and local market structure.

Table 1: Summary Statistics

	GVs			EVs		
	# Obs.	Mean	Std. Dev.	# Obs.	Mean	Std. Dev.
<b>Panel A: Product, city, and year observations for demand estimation</b>						
Sales (cars)	366,706	21.25	52.80	110,081	10.90	45.27
MSRP (10k RMB)	366,706	15.86	10.29	110,081	24.13	13.70
Net price (10k RMB)	366,706	15.74	10.18	110,081	20.79	12.39
Vehicle size (m <sup>3</sup> )	366,706	13.72	2.00	110,081	13.45	2.64
Horsepower	366,706	117.00	34.60	110,059	118.70	55.63
L/100km	366,706	6.92	1.25	—	—	—
kWh/100km	—	—	—	110,081	14.29	5.73
Driving range (km)	—	—	—	106,354	337.97	192.01
Battery capacity (kWh)	—	—	—	106,354	46.91	25.27
National cash subsidy (10k RMB)	—	—	—	110,081	0.92	1.22
Local cash subsidy (10k RMB)	—	—	—	110,081	0.07	0.41
Purchase tax exemption (10k RMB)	—	—	—	110,081	1.84	1.10
<b>Panel B: Model and year observations for supply estimation</b>						
Sales (cars, national)	4,762	1,636.32	2,721.50	1,782	673.60	1,861.90
MSRP (10k RMB)	4,762	15.79	10.43	1,782	24.08	13.21
Net price (10k RMB)	4,762	15.68	10.32	1,782	20.30	11.90
Vehicle size (m <sup>3</sup> )	4,762	13.75	2.02	1,782	13.17	2.71
Horsepower	4,762	116.84	34.86	1,781	112.14	55.83
L/100km	4,762	6.96	1.26	—	—	—
kWh/100km	—	—	—	1,782	14.49	6.98
Driving range (km)	—	—	—	1,733	325.73	184.27
Battery capacity (kWh)	—	—	—	1,733	45.51	24.55
<b>Panel C: City and year observations for market structure</b>						
	# Obs.	Mean	Std. Dev.	Min	Max	
Income per capita (10k RMB)	790	4.62	1.44	2.17	9.31	
GDP per capita (10k RMB)	790	8.93	4.31	2.22	26.92	
Population density (10k/km <sup>2</sup> )	790	0.09	0.12	0.00	0.90	
Urbanization rate	790	0.68	0.18	0.00	1.00	
EV license priority (dummy)	790	0.09	0.29	0.00	1.00	
Pilot city (dummy)	790	0.47	0.50	0.00	1.00	
EV policy strength	790	9.17	26.48	-23.01	93.91	
Oil price (year level, USD/bbl)	790	65.90	17.32	42.00	99.00	
# of firms	790	76.29	6.32	57.00	92.00	
# of GV models	790	418.94	53.96	228.00	533.00	
# of BEV models	790	94.85	74.60	0.00	261.00	
# of PHEV models	790	34.77	30.01	0.00	102.00	
# of REEV models	790	4.61	7.69	0.00	25.00	

*Notes:* The table reports summary statistics for the estimation samples. Panel A uses product, city, and year observations for demand estimation; Panel B uses model and year observations for supply estimation; and Panel C uses city and year observations for local market covariates. The sample covers 79 cities from 2015 to 2024. GV's include internal combustion vehicles and hybrids. EV's include battery electric vehicles, plug in hybrids, and range extended electric vehicles. Prices are reported in 10,000 RMB, and net price subtracts applicable cash subsidies and purchase tax exemptions. Vehicle size equals exterior length times width times height. Fuel consumption is reported as L/100km for GV's and kWh/100km for EV's. Dashes indicate variables that are not applicable for that vehicle type.  $\textcircled{g}$

### 3 Demand-Side Model

The demand side follows a differentiated products discrete choice model in the spirit of [Berry et al. \[1995\]](#). A market is defined as a geographic market by year pair, indexed by  $m = (c, t)$ . The geographic market  $c$  is one of the 79 markets listed in [Table Appendix D.1](#): Tier 1 cities, New Tier 1 cities, Tier 2 cities, and province level residual markets in the Rest tier. For simplicity, we refer to all geographic markets as cities, although the Rest tier consists of province level residual cells that aggregate prefectures not separately identified in the BLP panel. In each market  $m$ , consumer  $i$  chooses among the set of vehicle products  $J_m$  available in that market and an outside option.

The utility of consumer  $i$  for product  $j$  in market  $m$  is given as consists of three components: (i) a consumer-invariant mean utility component that captures the average attractiveness of a product, (ii) a consumer-specific component that generates heterogeneous preferences across consumers, and (iii) an idiosyncratic taste shock. Formally,

$$u_{ijm} = \delta_{jm} + \mu_{ijm} + \varepsilon_{ijm}, \quad (1)$$

$$\delta_{jm} = \beta_p \log p_{jm} + \beta_r \log r_{jm}^{\text{EV}} + \beta_{rd} \log r_{jm}^{\text{EV}} D_{cm} + x'_{jm} \beta + \eta_j^B + \eta_j^S + \eta_j^F + \eta_j^G + \eta_t + \xi_{jm}, \quad (2)$$

$$\mu_{ijm} = \pi_p \frac{1}{Y_{im}} \log p_{jm}. \quad (3)$$

Equation (1) decomposes utility into a mean utility component  $\delta_{jm}$ , a consumer-specific deviation  $\mu_{ijm}$ , and an idiosyncratic taste shock  $\varepsilon_{ijm}$ . The idiosyncratic shock is assumed to follow a Type I extreme-value distribution.

Equation (2) specifies the mean utility component. The variable  $p_{jm}$  denotes the consumer net purchase price after direct subsidies. The variable  $r_{jm}^{\text{EV}}$  denotes EV driving range (measured in kilometers), while  $D_{cm}$  is the log population density of city  $c$ . The coefficient  $\beta_r$  captures the average valuation of EV driving range, and  $\beta_{rd}$  allows the value of range to vary with local population density. The vector  $x_{jm}$  contains additional observed product and market characteristics including engine power for gasoline vehicles, interactions between EV status and local policy variables (license-plate restrictions, EV pilot-city designation, and policy strength), interactions between EV status and city demographics (log household income, GDP per capita, log population density, and urbanization rate), as well as oil prices, log population density, GDP per capita, urbanization rate, and log city income. The fixed effects  $\eta_j^B$ ,  $\eta_j^S$ ,  $\eta_j^F$ ,  $\eta_j^G$ , and  $\eta_t$  denote body-type, size-segment, fuel-type, firm-group, and year fixed effects, respectively. The term  $\xi_{jm}$  captures product-market specific demand shocks observed by consumers and firms but unobserved by the econometrician.

Equation (3) introduces consumer heterogeneity through income-dependent price sen-

sitivity. Household income is denoted by  $Y_{im}$ . The parameter  $\pi_p$  shows how the marginal utility of price varies across consumers. Because the price coefficient is scaled by  $1/Y_{im}$ , lower-income households are more sensitive to vehicle prices than higher-income households. A negative estimated value of  $\pi_p$  implies that a given increase in vehicle prices reduces utility more strongly for lower-income households than for higher-income households.

In summary, the model-specific component of utility is determined by vehicle prices, vehicle attributes, local market conditions, policy variables, and unobserved product characteristics. Consumers differ only in their sensitivity to prices. Specifically, the consumer-specific component  $\mu_{ijm}$  allows the marginal disutility of price to vary with household income  $Y_{im}$ . This specification is motivated by the idea that vehicle purchases represent a larger share of the budget for lower-income households, making them more sensitive to price changes than higher-income households.

In contrast, the model does not include an unobserved random coefficient on EV status. Instead, the vector  $x_{jm}$  includes interactions between the EV indicator and a vector of market demographics  $Z_m$ , consisting of household income, GDP per capita, population density, and urbanization. We also drop the random coefficient on log driving range. In preliminary estimates that allowed both an unobserved EV taste random coefficient  $\sigma_{EV}$  and a range random coefficient  $\sigma_r$ , the optimizer drove both toward zero whenever the parameter bounds were relaxed, indicating that the data identify mean-utility effects rather than additional unobserved heterogeneity once the rich observable interactions and fixed effects are in the specification. We therefore report a specification that drops both unobserved random coefficients and identifies EV-related heterogeneity through observable interactions and the random coefficient on the income-scaled price term in Equation (3). The estimated income-price coefficient  $\pi_p$  in the canonical specification is interior to the optimizer bounds, but earlier specifications that imposed tighter bounds delivered estimates at the boundary; the heterogeneity in price sensitivity across income groups should therefore be read as a quantitative property of the IV moment system at the chosen instrument set and bound configuration, not as a tightly identified data feature. This is the basis on which all subsequent cross-tier welfare results in Section 6.2 rest, and we report a calibrated demand-elasticity sensitivity exercise in Section 6.1 that brackets the no-subsidy effects under  $|\alpha| = 2.5$ ; readers who view the heterogeneity in price sensitivity as a calibration rather than an estimate should read the cross-tier welfare gradient in the same calibrated spirit. This specification allows EV demand to vary systematically across markets with different demographic and economic characteristics. For example, EV adoption may differ between richer and poorer cities because of differences in purchasing power, between denser and less dense cities because of differences in driving patterns and charging convenience, and between more and less urbanized areas because of differences in infrastructure and consumer preferences.

Given the Type I extreme-value assumption, the predicted market share of product  $j$  is

$$s_{jm}(\theta) = \int \frac{e^{\delta_{jm} + \mu_{ijm}}}{1 + \sum_{k \in J_m} e^{\delta_{km} + \mu_{ikm}}} dF_m(Y_i).$$

where  $\theta$  denotes the vector of demand parameters and  $F_m(Y_i)$  denotes the market-specific distribution of household income used to simulate consumer heterogeneity. The outside option is normalized to have mean utility equal to zero. For a candidate value of the nonlinear parameter  $\pi_p$ , the vector of mean utilities  $\delta_m$  is recovered by matching predicted and observed market shares using the BLP contraction mapping.

### 3.1 Identification and Estimation

Following [Whitefoot et al. \[2017\]](#), we view vehicle design as a multistage process in which firms choose some product attributes well before observing market-specific demand shocks. Characteristics such as body type, size segment, and engine power reflect fundamental design choices are costly to modify once a model enters production. We therefore assume that these attributes are chosen prior to the realization of market-specific demand shocks and are uncorrelated with  $\xi_{jm}$ . However, in contrast, EV manufacturers can adjust driving range more easily than fundamental vehicle characteristics by varying battery capacity or adopting improvements in battery technology. Consistent with this idea, the data show that manufacturers frequently update the driving range of an existing vehicle models over their lifecycle while leaving its fundamental design characteristics unchanged, as illustrated in [Figure 13](#). Firms may therefore choose driving range in response to demand factors that are observed by manufacturers but not by the econometrician. As a result, driving range may be correlated with the unobserved demand shock  $\xi_{jm}$ .

Vehicle prices are also likely to be endogenous because products with favorable unobserved characteristics command higher willingness to pay and therefore higher equilibrium prices. To address these concerns, we treat both vehicle prices and EV driving range as endogenous product characteristics and estimate the model using instrumental variables.

To address the endogeneity of vehicle prices and EV driving range, we follow the instrumental-variable approach developed in [Berry et al. \[1995\]](#) and subsequently applied in the automobile demand literature. The estimation relies on moment conditions of the form  $E[\xi_{jm} | Z_{jm}] = 0$ , where  $\xi_{jm}$  is the unobserved demand shock and  $Z_{jm}$  denotes a vector of exogenous product characteristics, market controls, fixed effects, and excluded instruments.

We utilize three sets of excluded instruments. First, following the BLP differentiation-instrument approach, we construct, for each product attribute  $x_{jm}^{(k)}$  in  $\{\log r, \log \text{power}^{GV}, \log r \times$

density}, the leave-one-out rivals' mean within the market:  $z_{jm}^{\text{diff},(k)} = x_{jm}^{(k)} - \bar{x}_{-j,m}^{(k)}$  where  $\bar{x}_{-j,m}^{(k)}$  averages over all products  $j' \neq j$  in market  $m$ . These instruments shift equilibrium markups and prices through competitive interactions while remaining orthogonal to product-specific demand shocks. Second, following the nested differentiation approach used in the automobile literature, we construct the same leave-one-out rivals' mean within the product's fuel-type segment (BEV, PHEV, REEV, HEV, ICE separately). These instruments capture variation in local competitive pressure arising from the composition of products within each fuel category. Third, we exploit variation in battery production costs <sup>1</sup>. Because batteries constitute a major component of EV production costs, changes in battery prices affect both vehicle prices and the cost of providing longer driving range. We therefore construct instruments based on battery costs, battery capacity, and EV technology type. Conditional on observed vehicle characteristics, market controls, and fixed effects, the identifying assumption is that battery-cost shocks influence demand only through their effects on prices and range.<sup>2</sup>

The market-specific income distribution  $F_m(Y_i)$  used to integrate Equation (3) is taken from the corresponding year's *China Statistical Yearbook* city-level household-income panel. We discretize  $F_m$  into  $N_S = 25$  quasi-Monte Carlo draws per market using a lognormal calibration matched to the mean and Gini coefficient of the city-year cell, and we use the same set of draws throughout the contraction mapping, the BLP inversion, and all counterfactual share computations. The random coefficient parameter governing income-dependent price sensitivity is identified from variation in household income distributions across markets. Because the consumer-specific utility component depends on household income, markets with similar product offerings but different income distributions imply different substitution patterns and price responses. This variation provides identifying information for the nonlinear parameter governing heterogeneous price sensitivity.

The model is estimated using the generalized method of moments (GMM). For a candidate value of the nonlinear parameter  $\pi_p$ , the BLP contraction mapping is used to recover the mean utility levels  $\delta_{jm}$  that match observed and predicted market shares. Conditional on the recovered mean utilities, the linear parameters are estimated using the moment conditions described above. The nonlinear parameter is then chosen to minimize the resulting GMM objective function. Standard errors are clustered at the market level. The estimation sample contains 476,787 product-market observations.

<sup>1</sup>The battery-cost instruments include battery costs interacted with BEV, PHEV, and REEV indicators, battery capacity, and interactions between battery costs and battery capacity.

<sup>2</sup>We explored discrete range-tier fixed effects and a Hausman-style leave-out instrument; neither delivered a well-conditioned moment system. The canonical IV set is bracketed with a calibrated demand-elasticity robustness exercise in Section 6.1.

**Instrument strength.** Table 2 reports the first-stage F-statistics for the three instrument groups separately and jointly. The within-market BLP differentiation instruments deliver a partial  $F$  of 1,655 on log price; the within-fuel-type nested instruments deliver 5,144; and the battery-cost instruments deliver 15,728. The joint  $F$  across all eleven excluded instruments is 18,479. All three groups individually and the joint test exceed the conventional Stock–Yogo weak-instrument threshold of 10 by orders of magnitude, so price endogeneity is strongly resolved by the instrument set. The identifying variation has three economic interpretations matched to the three IV groups: within-market substitution patterns ( $Z_{\text{diff}}$ ), within-fuel-type substitution patterns ( $Z_{\text{nest}}$ ), and time-series cost shocks transmitted through battery-input prices ( $Z_{\text{bat}}$ ). The joint power of the three groups reflects that price variation in the panel is driven by genuinely distinct sources, none of which is mechanically tied to the demand residual  $\xi_{jm}$ .

Table 2: First-stage F-statistics for the price endogeneity instruments

Instrument group	# excluded IVs	First-stage $F$
Within-market diff IVs (3 cols)	3	1,655.0
Within-fuel-type diff IVs (3 cols)	3	5,143.7
Battery-cost IVs (5 cols)	5	15,728.1
<b>All excluded IVs (joint)</b>	<b>11</b>	<b>18,479.3</b>

*Notes:* The dependent variable is log consumer net price. The three IV groups are the within-market BLP differentiation instruments  $Z_{\text{diff}}$ , the within-fuel-type nested differentiation instruments  $Z_{\text{nest}}$ , and the battery-cost shifters  $Z_{\text{bat}} = \text{BNEF}_t \times \{1[\text{BEV}], 1[\text{PHEV}], 1[\text{REEV}], \log \text{batcap}, \text{BNEF}_t \times \log \text{batcap}\}$ . The per-group  $F$  statistic tests joint significance of the group conditional on the exogenous X-controls and the other two groups. The overall first-stage  $F$  tests joint significance of all 11 excluded IVs. All groups individually and jointly exceed the conventional Stock–Yogo threshold of 10 by orders of magnitude, indicating that price is strongly identified.

We do not report a Hansen  $J$  overidentification test. In a panel with  $N \approx 477,000$  observations, the asymptotic  $\chi^2(10)$  distribution mechanically rejects almost any non-trivial moment condition because the test statistic scales linearly with  $N$ . Conventional applied-IO practice in large- $N$  panel BLP settings is to treat Hansen  $J$  as uninformative in this regime and to rely on first-stage strength plus the economic argument for instrument validity. The economic argument is straightforward in our setting:  $Z_{\text{diff}}$  and  $Z_{\text{nest}}$  are functions of competing products’ observable characteristics, which firms take as given in their own pricing problem;  $Z_{\text{bat}}$  is the BNEF battery-pack price, which is a national time-series shock that no individual product can manipulate. Conditional on the rich fixed-effect structure and the demographic interactions in Equation (3), the residual variation in  $Z$  is orthogonal to product-specific demand shocks by construction.

### 3.2 Demand Estimation Results

Table 3: Demand Estimation Results

	OLS Logit		IV Logit		Random Logit	
	Coef.	S.E.	Coef.	S.E.	Coef.	S.E.
<i>Vehicle Attributes</i>						
Log Price	-0.076***	(0.024)	-4.032***	(1.458)	-2.925***	(0.568)
Log Engine Power (GV)	0.410***	(0.033)	0.514	(1.306)	0.869**	(0.383)
Log EV Range (km)	0.716***	(0.038)	8.767***	(1.214)	8.191***	(0.338)
EV Range $\times$ Log Density	0.198*	(0.108)	-1.531***	(0.571)	-1.428***	(0.381)
<i>Policy and Market Variables</i>						
EV $\times$ License Restriction	0.014	(0.024)	0.058*	(0.035)	0.065**	(0.031)
EV $\times$ Pilot City	0.112***	(0.021)	0.058	(0.038)	0.040	(0.038)
EV $\times$ Policy Strength	0.024	(0.015)	0.035	(0.045)	0.022	(0.032)
EV $\times$ Log Income	7.522***	(0.861)	8.774	(12.601)	14.129***	(2.966)
EV $\times$ GDP Per Capita	-0.238***	(0.088)	-0.416	(0.544)	-0.559***	(0.185)
EV $\times$ Log Pop. Density	-0.043	(0.043)	0.515**	(0.219)	0.467***	(0.142)
EV $\times$ Urb. Rate	-0.008	(0.124)	0.491	(0.385)	0.469**	(0.207)
Oil Price	33.463***	(1.603)	81.865***	(20.788)	84.148***	(5.646)
Log Pop. Density	-0.059**	(0.028)	-0.056*	(0.034)	-0.037	(0.032)
GDP Per Capita	0.387***	(0.075)	0.435***	(0.131)	0.499***	(0.082)
Urbanization Rate	-0.026	(0.052)	-0.052	(0.067)	-0.078	(0.055)
Log City Income	0.013	(0.047)	-0.006	(0.195)	-1.165***	(0.415)
<i>Random Coefficients</i>						
$\pi(\text{Income} \times \text{Price})$					-1.980**	(0.859)
Fixed Effects	✓		✓		✓	
Observations	476,787		476,787		476,787	
$R^2$	0.198					
First-Stage $F$			3623.7			
GMM Objective					0.6903	

Notes: Standard errors clustered by market (OLS and IV) or 2-step optimal GMM market-clustered (RC BLP). \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ . All specifications include fuel type, body type, year, firm group, and EV  $\times$  year fixed effects. Policy variables are standardized to unit variance. IV and RC instruments: BLP differentiation IVs on product characteristics, within-fuel-type nested IVs, and BNEF battery cost shifters.

Table 3 reports demand estimates from the OLS logit, IV logit, and random coefficients BLP specifications. The OLS estimates do not account for the endogeneity of vehicle prices and EV driving range. Products with favorable unobserved demand shocks may command higher prices and may also have better unobserved quality. This creates a positive correlation between price and the unobserved demand component  $\xi_{jm}$ , biasing the OLS price coefficient toward zero. Consistent with this concern, the OLS log price coefficient is small in magnitude, at  $-0.076$ .

After instrumenting for prices and EV driving range, the estimated price coefficient becomes substantially larger in magnitude. The IV logit estimate of the log price coefficient is  $-4.032$ , compared with  $-0.076$  in the OLS specification. The coefficients on key vehicle attributes also increase. In particular, the coefficients on log EV range and log engine power are positive, indicating that consumers value longer EV range and greater gasoline vehicle engine power.

The preferred specification is the random coefficients BLP model reported in the final two columns. The mean log price coefficient remains negative and statistically significant, at  $-2.925$ . The income and price interaction is also negative and statistically significant. Interpreted according to the specification reported in the table, this means that the price coefficient becomes more negative for consumers with higher values of the income interaction term. This result should therefore be described as evidence of heterogeneous price sensitivity, rather than as direct evidence that lower income households are more price sensitive.

The estimates indicate that consumers place substantial value on EV driving range. The coefficient on log EV range is positive and large in the preferred specification. The interaction between EV range and population density is negative, implying that the marginal value of additional range is lower in denser cities. This pattern is consistent with shorter driving distances and greater charging availability reducing the benefit of additional range in dense urban markets.

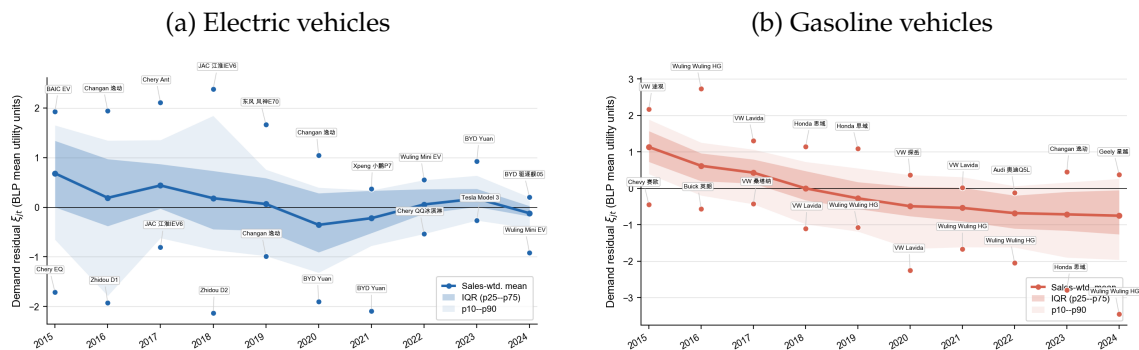
Observable market characteristics explain substantial heterogeneity in EV demand across cities. In the preferred specification, EV demand is higher in cities with higher household income, greater population density, and higher urbanization rates. The coefficient on EV interacted with GDP per capita is negative, even though the EV income interaction is positive. Because income, GDP per capita, density, and urbanization are highly correlated across cities, this coefficient should be interpreted as a conditional relationship rather than as evidence that richer cities have weaker EV demand overall.

The policy coefficients suggest that non financial EV incentives matter most clearly through license restrictions. In the preferred specification, EVs receive a positive and statistically significant utility premium in markets with license plate restrictions from which EVs are exempt. This finding is consistent with registration constraints increasing the relative attractiveness of EV ownership in large urban markets. By contrast, the pilot city and policy strength coefficients are not statistically significant in the random coefficients specification, so we interpret them mainly as controls for the broader local policy environment.



changing consumer perceptions in the EV market. The figure therefore provides reduced-form evidence that time-varying residual demand shifts play an important role in the transition from gasoline vehicles to EVs.

Figure 4: Within Product Deviations in Recovered Mean Utility by Vehicle Type



Notes: Panel (a) plots electric vehicle products, and Panel (b) plots gasoline vehicle products. For each product  $j$ , we compute  $\tilde{\delta}_{jt} = \delta_{jt} - \bar{\delta}_j$ , where  $\bar{\delta}_j$  is product  $j$ 's sales weighted mean recovered BLP mean utility across all years in which the product is observed. The solid line reports the sales weighted mean of  $\tilde{\delta}_{jt}$  by year. The dark band reports the interquartile range, and the light band reports the 10th to 90th percentile range across products within each year.

## 4 Supply-Side Model

Following Berry et al. [1995], we recover marginal costs from firms' static multi-product Bertrand pricing decisions. The demand model is specified in terms of the consumer net purchase price  $p_{jm}$ , which is the price paid by consumers after direct purchase subsidies. Let  $\tau_{jm}$  denote the per-unit purchase subsidy, and let  $\tilde{p}_{jm}$  denote the corresponding producer price, or the price received by the firm. These prices are related by  $\tilde{p}_{jm} = p_{jm} + \tau_{jm}$ . Let  $p_m = (p_{1m}, p_{2m}, \dots, p_{J_m m})'$  denote the vector of consumer net prices in market  $m$ . The demand system implies market shares  $s_{jm}(p_m; \hat{\theta})$ , where  $\hat{\theta}$  denotes the estimated demand parameters. In the supply side, product characteristics, market characteristics, fixed effects, and demand shocks are held fixed, so market shares are written as functions of the price vector chosen by firms.

Because the subsidy is taken as fixed by the firm within a market, a change in the producer price  $\tilde{p}_{jm}$  changes the consumer net price  $p_{jm}$  one-for-one. Firm  $f$  owns the set

of products  $\mathcal{J}_{fm}$  in market  $m$  and chooses prices to maximize variable profits:

$$\Pi_{fm} = \sum_{j \in \mathcal{J}_{fm}} (\tilde{p}_{jm} - \text{mc}_{jm}) M_m s_{jm}(p_m), \quad (4)$$

where  $\text{mc}_{jm}$  is marginal cost,  $M_m$  is market size, and  $s_{jm}(p_m)$  is the market share predicted by the estimated demand system.

For each product  $j \in \mathcal{J}_{fm}$ , firm  $f$  chooses the producer price  $\tilde{p}_{jm}$ . The first-order condition with respect to  $\tilde{p}_{jm}$  is

$$\frac{\partial \Pi_{fm}}{\partial \tilde{p}_{jm}} = M_m s_{jm}(p_m; \hat{\theta}) + M_m \sum_{k \in \mathcal{J}_{fm}} (\tilde{p}_{km} - \text{mc}_{km}) \frac{\partial s_{km}(p_m; \hat{\theta})}{\partial \tilde{p}_{jm}} = 0. \quad (5)$$

Dividing by market size  $M_m$  gives

$$s_{jm}(p_m; \hat{\theta}) + \sum_{k \in \mathcal{J}_{fm}} (\tilde{p}_{km} - \text{mc}_{km}) \frac{\partial s_{km}(p_m; \hat{\theta})}{\partial \tilde{p}_{jm}} = 0. \quad (6)$$

Since  $\tau_{jm}$  is fixed within market  $m$ , the derivative with respect to the producer price is equal to the derivative with respect to the consumer net price, and hence  $\frac{\partial s_{km}(p_m)}{\partial \tilde{p}_{jm}} = \frac{\partial s_{km}(p_m)}{\partial p_{jm}}$ . Also because the demand model is specified in log prices, this derivative is computed as  $\frac{\partial s_{km}(p_m)}{\partial p_{jm}} = \frac{\partial s_{km}(p_m)}{\partial \log p_{jm}} \frac{1}{p_{jm}}$ .

Let  $\sigma_{ijm}$  denote the individual choice probability of consumer  $i$  for product  $j$  in market  $m$ . The derivative with respect to log price is implied by the estimated demand system:

$$\frac{\partial s_{km}(p_m)}{\partial \log p_{jm}} = \int \sigma_{ikm} (\mathbf{1}\{j = k\} - \sigma_{ijm}) \left( \beta_p + \pi_p \frac{1}{Y_{im}} \right) dF_m(Y_i). \quad (7)$$

Stacking the first-order conditions across products within market  $m$  gives

$$s_m(p_m) + (\Omega_m \circ \Delta_m(p_m)) (\tilde{p}_m - \text{mc}_m) = 0, \quad (8)$$

where  $\Omega_m$  is the ownership matrix, with element  $\Omega_{jk,m} = 1$  if products  $j$  and  $k$  are owned by the same firm and zero otherwise. The symbol  $\circ$  denotes element-wise multiplication. The matrix  $\Delta_m(p_m)$  is the Jacobian matrix of market shares with respect to producer prices, with element  $\Delta_{jk,m} = \frac{\partial s_{km}(p_m)}{\partial \tilde{p}_{jm}}$ .

Marginal costs are therefore recovered by inverting the first-order conditions:

$$\text{mc}_m = \tilde{p}_m + (\Omega_m \circ \Delta_m(p_m))^{-1} s_m(p_m). \quad (9)$$

The recovered marginal costs are used only for products with valid first-order-condition solutions.

We then parameterize the recovered marginal costs in a single pooled regression that uses technology-specific interactions to let the cost-structure of EVs and gasoline vehicles differ:

$$\begin{aligned} \log mc_{jm} = & \gamma_0 + \gamma_s \log \text{size}_{jm} + \gamma_h \log \text{power}_{jm}^{\text{GV}} + \gamma_r \log r_{jm}^{\text{EV}} + \gamma_b \text{BatteryCost}_{jm}^{\text{EV}} \\ & + \lambda_{\text{fuel}(j)} + \lambda_{\text{body}(j)} + \lambda_{\text{firm}(j)} + \lambda_t + \omega_{jm}. \end{aligned} \quad (10)$$

The technology-specific covariates are constructed as masked interactions:  $\log \text{power}_{jm}^{\text{GV}} \equiv (1 - \mathbf{1}[j \in \text{EV}]) \cdot \log \text{power}_{jm}$ ,  $\log r_{jm}^{\text{EV}} \equiv \mathbf{1}[j \in \text{EV}] \cdot \log r_{jm}$ , and  $\text{BatteryCost}_{jm}^{\text{EV}} \equiv \mathbf{1}[j \in \text{EV}] \cdot \text{bnef}_t$ , where  $\text{bnef}_t$  is the BNEF annual battery-pack cost (/kWh). This pooled specification is mathematically equivalent to fitting separate regressions on the EV and GV subsamples for the technology-specific slopes, but it pools the fixed effects across fuel categories, which substantially improves the identification of the battery-cost coefficient by exploiting the EV-versus-GV contrast within each year rather than relying on residual variation in the time series of battery prices alone. The fixed effects control for fuel type, body type, firm group, and year. The coefficient  $\gamma_r$  on EV driving range captures the marginal cost of providing longer range;  $\gamma_h$  on gasoline-vehicle engine power captures the cost of providing higher engine power;  $\gamma_b$  on BatteryCost is the elasticity of EV marginal cost with respect to the BNEF battery-pack price and is the structural primitive that drives the Wright's-law channel in the counterfactual simulations of Section 6.1.

#### 4.1 Marginal Cost Estimation Results

Table 4 reports the pooled marginal cost estimates from Equation (10). The dependent variable is log marginal cost recovered from the multiproduct Bertrand first order conditions. The technology specific covariates are constructed as masked interactions, so the regression allows the cost slopes for gasoline vehicle engine power, EV driving range, and battery costs to differ by technology while pooling body type, fuel type, year, and firm group fixed effects across products. The estimates reported here are the cost primitives used in the decomposition and counterfactual simulations.

Table 4: Supply-Side Marginal Cost Estimates (Pooled OLS)

	Coef.	S.E.
<i>Product Characteristics</i>		
Log Size	1.5070***	(0.0055)
Log Engine Power ×1[GV]	0.9524***	(0.0024)
Log EV Range ×1[EV]	0.3967***	(0.0034)
BatteryCost ×1[EV] (USD/kWh)	0.0040***	(0.0000)
<i>Fuel Type Fixed Effects (base: BEV)</i>		
HEV	-1.7490***	(0.0228)
ICEV	-2.0170***	(0.0228)
PHEV	0.6893***	(0.0062)
REEV	2.5170***	(0.0221)
<i>Year Fixed Effects (base: 2015)</i>		
2016	-0.0603***	(0.0019)
2017	-0.0736***	(0.0019)
2018	-0.1037***	(0.0018)
2019	-0.1491***	(0.0018)
2020	-0.1637***	(0.0018)
2021	-0.1675***	(0.0019)
2022	-0.1912***	(0.0019)
2023	-0.2187***	(0.0019)
2024	-0.2376***	(0.0021)
Body Type FE		✓
Firm Group FE		✓
Observations		476,787
R <sup>2</sup>		0.802

*Notes:* The dependent variable is log marginal cost recovered from the multi-product Bertrand first-order conditions using observed consumer net prices, subsidies, and the estimated demand system. The regression is pooled across all 476,787 product–city–year observations (BEV, PHEV, REEV, HEV, and ICEV) and the technology-specific slopes are identified via masked interactions: Log Engine Power is zero for EV rows; Log EV Range and BatteryCost are zero for GV rows. Pooling the fixed effects across fuel categories identifies the BatteryCost coefficient from the within-year EV-versus-GV contrast, which carries non-degenerate variation, rather than from the perfectly collinear BNEF time series alone. Standard errors are HC1 heteroskedasticity-robust. \* $p < 0.10$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ .

The product characteristic coefficients have the expected signs. Larger vehicles have higher marginal costs: the coefficient on log size is 1.507, so a 10 percent increase in exterior vehicle volume is associated with roughly a 15 percent increase in marginal cost. For gasoline vehicles, the coefficient on log engine power is 0.952, implying that higher engine power is costly to provide. For example, doubling engine power raises log marginal cost by about 0.66 log points. For EVs, the coefficient on log driving range is 0.397, implying that longer range increases marginal cost. Doubling EV range raises log marginal cost by

about 0.28 log points. These estimates capture the cost side of the quality improvement documented in the demand estimates: consumers value longer EV range, but range is not a free demand shifter.

The BatteryCost coefficient is the key primitive for the Wright’s law learning channel in Section 6.1. The estimate is  $\gamma_b = 0.0040$  when BNEF battery pack cost is measured in USD/kWh. Because the dependent variable is log marginal cost and BatteryCost is measured in levels, this coefficient is a semi elasticity rather than an elasticity. Holding fixed product characteristics, fuel type, body type, firm group, and common year effects, a 100 USD/kWh increase in the BNEF battery price is associated with a  $0.0040 \times 100 = 0.40$  log point increase in EV marginal cost relative to gasoline vehicles. Conversely, the 258 USD/kWh decline in BNEF prices from 2015 to 2024 implies a reduction of about  $0.0040 \times 258 \approx 1.03$  log points in the EV specific cost component, or roughly a 64 percent decline in levels.

The identification of  $\gamma_b$  requires a precise interpretation. Because the BNEF battery price varies only by year, it would be collinear with year fixed effects in an EV only regression. The pooled EV and gasoline vehicle specification identifies the coefficient from changes over time in the EV versus gasoline vehicle marginal cost gap, after controlling for common year effects and fuel type fixed effects. Thus, the coefficient should be read as the EV specific cost response to the battery price series, not as a purely within year cross sectional effect.

The year fixed effects capture common cost changes across all products that remain after observed product characteristics and the EV specific battery cost channel are controlled for. Relative to 2015, the year effects decline from  $-0.060$  in 2016 to  $-0.238$  in 2024, corresponding to about a 21 percent decline in common marginal cost components over the decade. These common changes may reflect generic manufacturing improvements, input cost trends, exchange rate movements, or other cost shifters not separately modeled.

Overall, the supply estimates provide the cost side needed for the counterfactual analysis. The demand estimates show that consumers value EV range, while the supply estimates show that longer range is costly to provide. This combination creates a meaningful quality cost tradeoff and allows the decomposition to distinguish product quality improvement from battery cost decline.

## 4.2 Implied Elasticities and Price Cost Margins

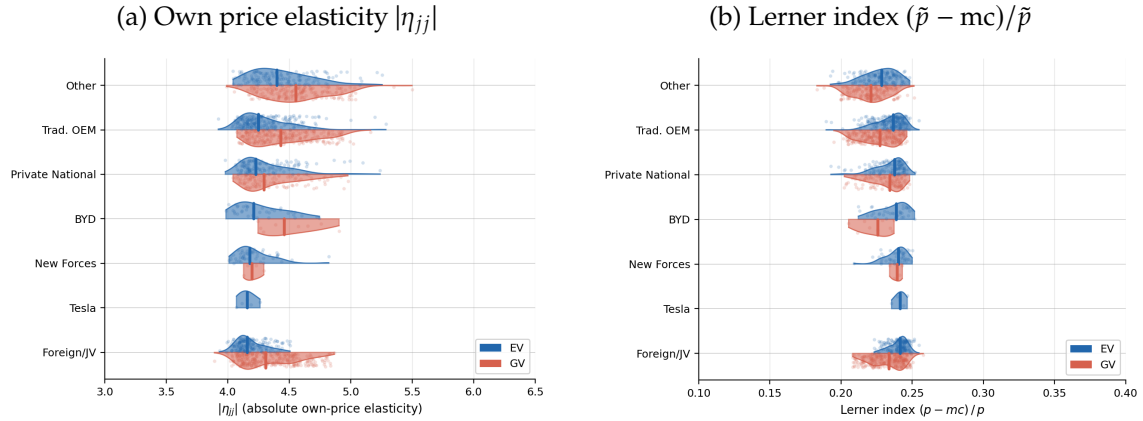
We combine the recovered marginal costs with the estimated demand system to evaluate the model’s pricing implications. Figure 5 reports two implied objects by firm group and fuel type. Panel a plots the distribution of absolute own price elasticities, defined as the percentage change in product  $j$ ’s sales induced by a one percent change in its own consumer net price. Larger absolute values indicate more price sensitive demand. These

elasticities summarize the substitution patterns implied by the demand model and enter the markup inversion through the price derivative matrix,  $\partial s_{km} / \partial \tilde{p}_{jm}$ .

Most own price elasticities in Panel a are concentrated between four and five in absolute value, indicating substantial but plausible price responsiveness. EV and GV elasticities overlap within most firm groups, suggesting that price sensitivity is not driven solely by fuel type. Instead, firm groups with broader product portfolios, such as traditional OEMs, private national firms, foreign joint ventures, and the residual category, display wider elasticity distributions. Tesla and New Forces have more concentrated EV elasticity distributions, reflecting their more focused EV product portfolios.

Panel b plots the corresponding Lerner indices,  $(\tilde{p}_{jm} - mc_{jm}) / \tilde{p}_{jm}$ , which measure the share of the producer price retained as a margin over marginal cost. The implied margins are mostly concentrated between 0.20 and 0.30. Differences between EV and GV margins are modest within many firm groups, but EV margins are more concentrated for EV native firms, including Tesla, BYD, and the New Forces. Legacy firm groups, including foreign joint ventures, traditional state owned OEMs, private national manufacturers, and the residual Other category, show wider margin dispersion. The seven class firm group taxonomy used throughout the paper is defined in Table Appendix B.4. Overall, Figure 5 suggests that the estimated model implies economically reasonable price responsiveness and price cost margins while preserving meaningful heterogeneity across firm groups.

Figure 5: Implied Own Price Elasticities and Lerner Indices



Notes: Panel (a) plots the distribution of absolute own price elasticities by firm group and fuel type. Panel (b) plots the distribution of Lerner indices by firm group and fuel type. Violin half plots show the distribution of product, city, and year observations; dots mark within group medians. Blue denotes EV products, and red orange denotes GV products.

## 5 Decomposition of EV Diffusion

Between 2015 and 2024, China’s EV market underwent one of the fastest large-scale technology transitions in recent history. Several forces moved at the same time: EV products improved rapidly, battery costs declined, firms expanded EV offerings, purchase subsidies were phased down, and consumer demand for EVs evolved across markets.

This section decomposes the rise in China’s EV market share into these forces. The exercise is a historical accounting decomposition, not a policy counterfactual. It asks how much of the 2015–2024 EV-share change can be attributed to product attributes, choice-set expansion, battery costs, subsidies, residual demand, and demographics. For each channel, we replace the corresponding primitives with their endpoint-year values and re-solve for equilibrium prices, quantities, and market shares.

The estimated demand and supply model is what makes this exercise possible. Each force operates through both sides of the market: product attributes and entry shift substitution; battery costs shift marginal costs and prices; subsidies shift consumer net prices and firm revenues; demographics shift market size and price sensitivity. The decomposition therefore captures the full equilibrium effect of each force, including price responses, markup adjustments, and reallocation between EVs and GVs.

### 5.1 Decomposition Design

The decomposition proceeds by switching groups of economic primitives from their 2015 values to their endpoint-year values. Let  $\mathcal{B}$  denote the set of decomposition blocks, each representing a potential driver of the EV transition. For any subset  $S \subseteq \mathcal{B}$ , let  $V_y(S)$  denote the model-implied national EV share in endpoint year  $y$  when the blocks in  $S$  are set to their year- $y$  values and the rest are held at 2015 values. So  $V_y(\emptyset)$  is the 2015 baseline share and  $V_y(\mathcal{B})$  is the full-endpoint share. The six blocks are Quality, Variety, Battery, Subsidy, Residual, and Market; Table [Appendix D.2](#) lists the primitives in each. We report the decomposition at the national aggregate (Section [5.3](#)) and separately for the four city tiers (Section [5.4](#)). The tier-specific decomposition is the central evidence for the heterogeneous-diffusion mechanism of the paper title. Each block’s contribution varies systematically across tiers, so “what drove EV adoption” has a different quantitative answer depending on where the city stands on the diffusion curve.

A simple sequential decomposition introduces blocks one at a time and assigns each its marginal contribution in a chosen order. But each block’s contribution depends on which other blocks have already been introduced (the channels are complementary, as discussed in the introduction). Any single ordering is therefore arbitrary.

To avoid this dependence on an arbitrary ordering, we use the Shapley value from cooperative game theory [[Shapley, 1953](#)]. Let  $\phi_{b,y}$  denote the contribution of block  $b$  to

the model-implied change in national EV market share between 2015 and endpoint year  $y$ . The Shapley value assigns this contribution as the average marginal effect of introducing block  $b$  across all possible sequences in which the decomposition blocks could be added:

$$\phi_{b,y} = \sum_{S \subseteq \mathcal{B} \setminus \{b\}} \frac{|S|! (|\mathcal{B}| - |S| - 1)!}{|\mathcal{B}|!} [V_y(S \cup \{b\}) - V_y(S)]. \quad (11)$$

The term  $V_y(S \cup \{b\}) - V_y(S)$  measures the marginal contribution of block  $b$  when the blocks in  $S$  have already been introduced. The weighting term  $\frac{|S|! (|\mathcal{B}| - |S| - 1)!}{|\mathcal{B}|!}$  is the probability that, in a random ordering of all blocks, exactly the blocks in  $S$  appear before block  $b$ . Thus, the Shapley value averages the marginal contribution of block  $b$  across all possible orderings. With six blocks, this corresponds to averaging over  $6! = 720$  possible orderings. The Shapley values satisfy the adding-up property

$$\sum_{b \in \mathcal{B}} \phi_{b,y} = V_y(\mathcal{B}) - V_y(\emptyset), \quad (12)$$

so the signed sum of the block contributions exactly equals the model-implied change in national EV market share between 2015 and endpoint year  $y$ .

For each coalition  $S$ , we solve for equilibrium prices and quantities. This equilibrium step matters because the channels interact through prices and substitution: a battery-cost decline lowers production costs but also changes equilibrium prices and the competitive position of EVs against GVs; entry changes the choice set and competitive pressure across fuel types.

Some coalitions generate very high EV shares, especially when attribute gains, low battery costs, and favorable residual demand are introduced together. To prevent the solver from entering unstable EV-saturation regions, we impose two caps. First, the product-level counterfactual quantity is capped at twice observed quantity. Second, the aggregate EV market share is capped at 1.5 times the observed 2024 EV share. The caps affect only extreme coalitions.

The saturation regions arise mechanically from the demand elasticity. Under log utility, the effective own-price elasticity at mean income is about 4.9, at the elastic end of the EV-demand literature. At this elasticity, coalitions that simultaneously improve attributes, lower marginal cost through battery learning, and turn on favorable demand residuals drive equilibrium prices toward marginal cost. Small numerical perturbations then push the fixed-point iteration into a high-share regime where a few cheap BEVs absorb a disproportionate share of national demand. The caps truncate these extremes.

The aggregate cap binds in 12 of the 64 coalitions at the 2024 endpoint. The binding coalitions sit in the high-V tail; for example, the full coalition minus Subsidy would push the EV share to roughly 80% without the cap. The per-product cap binds in a few coalitions

where a single cheap BEV would otherwise capture more than twice its observed national quantity. The Shapley closure  $\sum_b \phi_b = V(\text{full}) - V(\emptyset)$  remains exact under the capped  $V$ , so the caps shift individual  $\phi_b$  magnitudes but not the decomposition identity.

Before interpreting the block contributions, we verify that the model matches the aggregate EV-share transition. The block contributions sum to the model-implied endpoint change  $V(\text{full}) - V(\emptyset)$  by construction; they explain the observed change only to the extent that the model-implied change is close to the raw data. Under the 2015 environment, the model-implied EV share is 1.03%. Under 2024 values, it is 44.66%, an aggregate change of 43.64%. The observed aggregate change is 44.30%, so the endpoint gap is  $-0.75\%$ . Appendix Table [Appendix D.3](#) reports the same diagnostic for each intermediate endpoint year. The gaps are small, so the decomposition closely tracks the observed transition. The remaining  $-0.75\%$  gap reflects the BLP solver’s incomplete reproduction of observed 2024 shares, not a defect of the decomposition. The waterfall and cumulative figures use the model-implied change as the reference object so that the Shapley values close the endpoint exactly.


## 5.2 Order Sensitivity and Shapley Averaging

Figure 6 illustrates why averaging over orderings is important in this setting. The figure focuses on the main 2015–2024 transition. Panels [a](#) and [b](#) report two sequential decompositions that use the same six blocks but introduce them in different orders. Because the blocks interact, the contribution assigned to a given block depends on the order in which the counterfactual primitives are switched on.

This order dependence is especially clear for the Variety block, which captures changes in the set of EV products. In Panel [a](#), Variety is introduced first and receives little credit because new EV products are added to an otherwise 2015-like environment. In Panel [b](#), the same block is introduced last and receives much more credit because the market has already been transformed by improved EV quality, lower battery costs, changed subsidies, residual demand shifts, and market conditions. Thus, a sequential decomposition can assign very different importance to the same economic force simply because of the chosen ordering.

Panel [c](#) reports the Shapley decomposition. Rather than choosing one sequence, the Shapley value averages each block’s marginal contribution over all  $6! = 720$  possible orderings. This produces an order-invariant allocation of the model-implied 2015–2024 change across the six blocks. The figure shows that the EV transition is highly nonseparable: the effects of Quality, Variety, Battery, Subsidy, Residual, and Market conditions depend on which other primitives have already changed.

Figure 6: Order Sensitivity of Sequential Decompositions



*Notes:* The figure decomposes the model-implied change in national EV market share from 2015 to 2024. Panels (a) and (b) report sequential decompositions under two different orderings of the same six blocks. Panel (c) reports the Shapley value, which averages marginal contributions over all 720 possible orderings. The black line is the observed aggregate change in EV share, and the gray dashed line is the model-implied change,  $\sum_{b \in \mathcal{B}} \phi_{b,2024} = V_{2024}(\mathcal{B}) - V_{2024}(\emptyset)$ . Hatched areas indicate negative contributions. All changes are reported in percentage units of aggregate EV share.

### 5.3 National Decomposition Results

We now turn to the national decomposition for the endpoint year 2024. Table 5 and Figure 7 report the Shapley contributions to the model implied change in national EV market share from 2015 to 2024. The EV share rises from 1.03% in the 2015 baseline to 44.66% in the model implied 2024 endpoint, so the endpoint difference is 43.64%.

Table 5: Shapley Decomposition of the 2015–2024 EV-share Change

Block	Contribution	Interpretation
Quality	45.49%	Product-quality and attribute-related channel
Variety	14.81%	Larger and more EV-oriented product set
Battery	8.20%	Battery-related cost reductions
Subsidy	-13.63%	Direct subsidy phase-out relative to 2015
Residual	-7.96%	Residual demand shifts after observed controls
Market	-3.27%	Income, market size, and demographic changes
Signed sum	43.64%	Model-implied EV-share change

*Notes:* This table reports the Shapley decomposition of the change in national EV market share from 2015 to 2024. Contributions are reported in EV share units. The signed sum equals  $V_{2024}(\mathcal{B}) - V_{2024}(\emptyset)$  and is computed using unrounded Shapley values.

Table 5 shows that the largest contribution comes from the Quality block. This block switches the observable product attributes that vary between 2015 and 2024 for products present in both endpoint sets: log driving range, log engine power for GV products, log vehicle size, and the range  $\times$  density interaction. It excludes FuelType, BodyType, and SizeSegment fixed effects, which are category level intercepts and would otherwise be double counted with Variety. It also excludes prices, policy variables, residual demand, and demographic interactions, which are assigned to other blocks. The 45.49% contribution therefore measures the equilibrium effect of the within product attribute trajectory, not the effect of any single attribute and not the effect of the changing product set.

The second largest contribution is Variety, at 14.81%. This result reflects the expansion of the EV choice set between 2015 and 2024. Consumers in 2024 face many more EV options than in 2015, across models, body types, and firm groups. The interpretation is close to the new product channel in differentiated products demand [Petrin, 2002], where new products change substitution patterns and expand consumer choice. In this exercise, Variety is an accounting block for choice set expansion rather than a structural estimate of firms' entry values.


The Battery block contributes 8.20%. This block measures the equilibrium effect of switching the battery cost path from its 2015 value to its 2024 value, holding cumulative EV production fixed at the observed path. The larger dynamic effect of Wright's law learning, in which subsidy induced production feeds back into battery costs, is reported separately in Section 6.1. The Battery contribution is smaller than Quality and Variety, but it is still economically important because it complements product improvement: lower battery costs matter most when firms can embody them in attractive EV products.

The Subsidy block contributes  $-13.63\%$ . This negative contribution does not imply that subsidies reduce EV demand. The historical decomposition compares the 2024 subsidy environment with the 2015 subsidy environment, and direct purchase subsidies were phased down over this period. The Subsidy block therefore captures subsidy phaseout in the historical transition, not the value of subsidy support in the 2024 market. The value of subsidy support is measured in the no subsidy counterfactual in Section 6.1, and the relation between the historical decomposition and the policy counterfactual is discussed in Section 6.4.

Finally, the Residual and Market blocks are both negative, contributing  $-7.96\%$  and  $-3.27\%$ , respectively. These signs should be interpreted as conditional decomposition results. They do not imply that residual demand shifts or market demographics were irrelevant for EV adoption. Rather, they mean that, after averaging over all Shapley orderings, the remaining residual demand shifts and market level changes reduce the national EV share relative to the other 2024 primitives. In particular, the Market block combines changes in income, population, market size, and demographic variables that affect both demand and price sensitivity. Its negative sign is therefore a net effect in EV

share units, not evidence that higher income mechanically reduces EV demand.

Figure 7: Shapley Waterfall: 2015–2024 EV-share Change



Notes: The figure plots the Shapley value for each block. Bars report contributions to the change in national EV share relative to the 2015 baseline. Positive bars increase the cumulative total, and negative bars decrease it. The dashed line marks the model implied endpoint difference,  $44.66\% - 1.03\% = 43.64\%$ . The final gray bar reports the model implied 2024 endpoint. All values are reported in EV share units.

## 5.4 Heterogeneity Across City Tiers


The aggregate decomposition masks heterogeneity across city markets. China’s EV transition did not diffuse uniformly: large, high-income cities adopted EVs earlier; lower-tier cities became more important later. We therefore apply the same Shapley decomposition to city tiers.

Figure 8 reports the result; appendix Table Appendix D.4 reports the numerical values. The exercise uses the same six blocks and the same coalition equilibria as the national decomposition. For each coalition  $S$ , equilibrium prices and quantities are solved at the national product-market level. We then construct a tier-specific value function  $V_{t,2024}(S)$  by aggregating the resulting city-level EV shares within tier  $t$ . The adding-up condition

holds exactly within each tier:

$$\sum_{b \in \mathcal{B}} \phi_{b,t,2024} = V_{t,2024}(\mathcal{B}) - V_{t,2024}(\emptyset).$$

Figure 8: City Tier-specific Shapley Decomposition of the 2015–2024 EV-share Change



*Notes:* The figure reports Shapley values by block for each city tier and for the national aggregate. For each coalition, equilibrium prices and quantities are solved using the national product-market equilibrium. Tier-specific values are then computed by aggregating the resulting city-level EV shares within each tier. The black X marks the national Shapley value from Table 5. All values are reported in EV-share units.

Three patterns stand out. First, Quality is the largest positive block in every tier, but its contribution is larger in higher-tier city markets. The Quality contribution is especially large in Tier 1 and New Tier 1, consistent with EV product-quality improvements being especially valuable in larger and richer urban markets, where consumers were better positioned to substitute toward improved EV products.

Second, the Subsidy block is more negative in lower-tier cities. As in the national decomposition, this block measures the effect of moving from the 2015 subsidy environment to the 2024 subsidy environment. A more negative Subsidy contribution therefore means that the historical phase-down of direct purchase subsidies reduced model-implied EV share more strongly in that tier. In EV-share terms, the phase-down was more costly for lower-tier markets.

Third, the Residual block is much more negative in Tier 1 than in the Rest tier. After accounting for Quality, Variety, Battery, Subsidy, and Market, the remaining residual component does not add to Tier 1 EV growth. This does not mean Tier 1 EV demand weakened mechanically. The observable channels already account for most of the Tier 1 transition, leaving a negative average-order residual.

The tier results reinforce the heterogeneous-diffusion interpretation of the paper. The same national technology and policy transition generates different EV-share contributions across city tiers because cities differ in market size, income, product exposure, and stage of diffusion.

## 5.5 Dynamics of the Transition

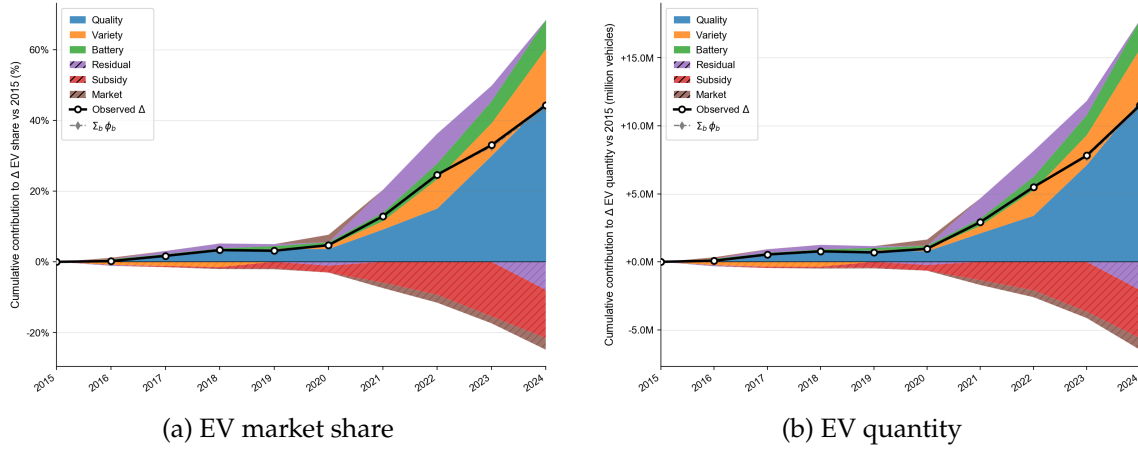
Figure 9 applies the Shapley decomposition year by year, using 2015 as the baseline. Panel a plots cumulative contributions to the change in EV market share, while Panel b plots the corresponding contributions to the change in EV quantities. In both panels, the black line reports the observed change relative to 2015, and the gray dashed line reports the model implied Shapley sum.

The figure shows that EV diffusion accelerates sharply after 2020. In the early years, all blocks have relatively small contributions because EVs remain a thin segment of the market. After 2020, the Quality contribution rises rapidly, reflecting improvements in observed EV attributes, especially driving range and vehicle size. Variety also becomes increasingly important as the EV choice set expands. Battery cost reductions add a meaningful contribution, especially in later years when lower costs can be embodied in a broader set of EV products.

The negative Subsidy block grows over time because the historical subsidy schedule becomes less generous relative to the 2015 baseline. The negative Residual and Market components partly offset the positive product and cost channels, but they are not large enough to reverse the overall transition. Panel b shows that the same mechanisms explain not only EV share growth but also the large increase in EV sales quantities.

The same order dependence provides an additional diagnostic. If the blocks were nearly separable, the marginal contribution of each block would be similar across coalitions. Instead, the marginal contributions vary widely depending on which other primitives have already changed. Appendix Table [Appendix A.1](#) reports the minimum and maximum marginal contribution of each block across coalitions. The wide ranges reinforce that EV diffusion reflects complementarities among product attributes, product variety, battery costs, policy, residual demand, and market conditions, rather than the sum of independent channels.

Figure 9: Cumulative Shapley decomposition from the 2015 baseline



Notes: Panel (a) plots cumulative Shapley contributions to the change in EV market share relative to 2015. Panel (b) plots cumulative Shapley contributions to the change in EV quantity relative to 2015. Stacked areas report block contributions; solid areas indicate positive contributions, and hatched areas indicate negative contributions. The black line reports the observed change relative to 2015, and the gray dashed line reports the model implied Shapley sum  $\sum_b \phi_b$ . The market share panel is reported in EV share units, while the quantity panel is reported in million vehicles.

## 6 Subsidy Support: Dynamic Effects and Incidence

The decomposition above explains the historical EV transition from 2015 to 2024. This section asks a different question: how much did direct subsidy support matter for EV adoption, and who benefited from it? The distinction is important. The Subsidy block in the Shapley decomposition measures the contribution of moving from the 2015 subsidy schedule to the less generous 2024 subsidy schedule. By contrast, the exercise below compares market outcomes along the observed subsidy path with outcomes under a counterfactual path without direct subsidy support.

### 6.1 Dynamic Counterfactual Without Direct Subsidy Support

Figure 10 evaluates the role of direct subsidy support in EV market development. The simulation compares the observed subsidy path with a counterfactual path in which direct EV subsidy support is removed in each year from 2015 to 2024. Direct subsidy support includes national cash subsidies, local cash subsidies, and the purchase tax exemption. The exercise incorporates two dynamic mechanisms: battery cost learning through Wright's

law, calibrated at an 18% cost reduction per doubling of cumulative production,<sup>3</sup> and endogenous product availability through entry and exit.

The exercise is dynamic in state variables, but not in agents' optimization. Within each year, consumers choose among available products given prices and attributes, and firms set prices in a static Bertrand Nash equilibrium. Cumulative production is then updated across years, which changes the subsequent battery cost path and product availability. The results should therefore be read as a sequence of static market equilibria linked by state transitions, rather than as a dynamic programming model with forward looking expectations.

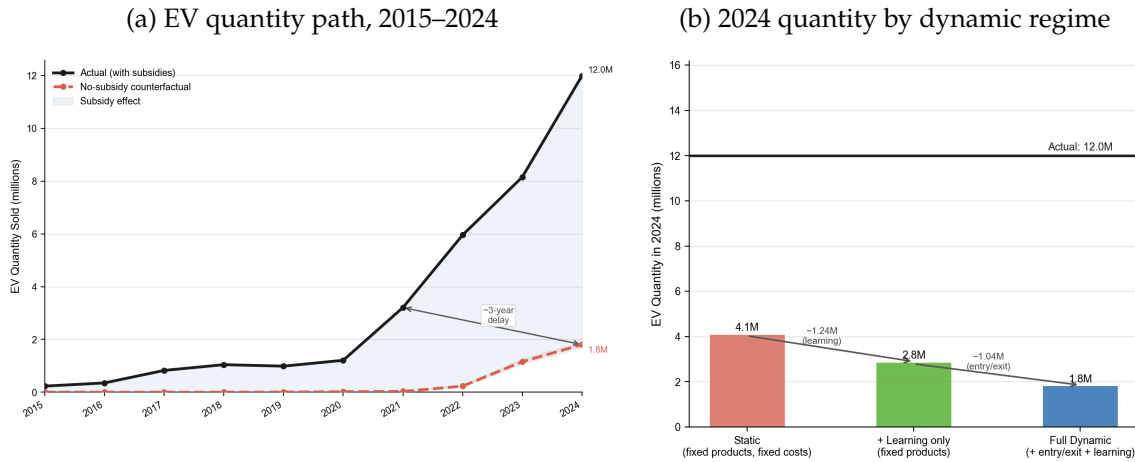
Panel [a](#) reports the EV quantity path under the full dynamic regime. Under observed subsidy support, the simulated path reaches 12.0 million EVs in 2024. The observed national NEV registration count is approximately 11.15 million, so the simulation overshoots observed quantity by about 7.6%. This gap is small relative to the counterfactual effect. Without direct subsidy support, the simulated 2024 EV quantity falls to 1.8 million. The counterfactual path remains far below the observed support path throughout the diffusion period and implies an adoption delay of about three years. Thus, subsidy support does not merely shift purchases within a static market. It accelerates the transition by raising early demand, supporting cumulative learning, and sustaining later product availability.

Panel [b](#) shows how dynamic mechanisms amplify the effect of subsidy removal. In the Static regime, with product availability and costs held fixed, the 2024 market would still sell 4.1 million EVs without direct subsidy support. Allowing battery costs to respond through cumulative learning reduces the counterfactual quantity to 2.8 million. Allowing both learning and entry and exit reduces it further to 1.8 million. The reduction from 4.1 million to 2.8 million captures the battery learning channel, while the additional reduction from 2.8 million to 1.8 million captures the entry and exit channel. These results show that weaker early demand slows battery cost declines and reduces the future variety of EV products. Appendix Table [Appendix B.1](#) reports the corresponding 2024 quantities by dynamic regime.

---

<sup>3</sup>Fitting  $\log c_t = \log A - b \log Q_t$  to the 2013 to 2024 BNEF battery cost series and global cumulative EV stocks gives an implied learning rate of 22.8% per doubling. We use the more conservative 18% rate in the main simulation. Using the fitted 22.8% rate changes the 2024 counterfactual EV share by less than 0.5% relative to the headline results.

Figure 10: EV Adoption with and without Direct Subsidy Support



Notes: The panels compare simulated EV quantities under observed direct subsidy support and under a counterfactual path without direct subsidy support. Panel (a) plots the quantity path from 2015 to 2024 under the full dynamic regime, which includes endogenous product availability and Wright’s law battery cost learning. Shading reports Monte Carlo uncertainty across 50 draws. The arrow marks the approximate adoption delay implied by subsidy removal. Panel (b) reports the 2024 counterfactual EV quantity under three regimes: Static, Learning only, and Full dynamic. The arrows show the additional reductions associated with battery cost learning and product entry and exit. Quantities are reported in million vehicles.

As a sensitivity check, we repeat the simulation under a less elastic demand calibration. The baseline estimate implies a price sensitivity of approximately  $|\alpha| = 4.9$  at mean income. We rescale this value to  $|\alpha| = 2.5$ , which is closer to the middle of estimates in the EV demand literature (Springel, 2021; Bollinger and Gillingham, 2019; Li et al., 2017). Under this less elastic calibration, the full dynamic counterfactual yields  $3.9 \pm 0.2$  million EVs in 2024, compared with  $1.8 \pm 0.07$  million under the baseline calibration. The corresponding EV share effect is  $-23.4 \pm 1.4\%$ . The magnitude is smaller, but the qualitative ranking is unchanged: entry and exit amplify battery learning, and battery learning amplifies the direct demand effect of subsidy removal.

## 6.2 Diffusion Timing and Subsidy Incidence Across Cities

The dynamic counterfactual shows that subsidy support affected aggregate EV adoption. We next ask how those benefits were distributed across cities. The key issue is timing. Cities entered the EV diffusion curve at different dates, while the national cash subsidy schedule was being phased down.

Figure 11 documents this uneven diffusion pattern. The national EV share hides large

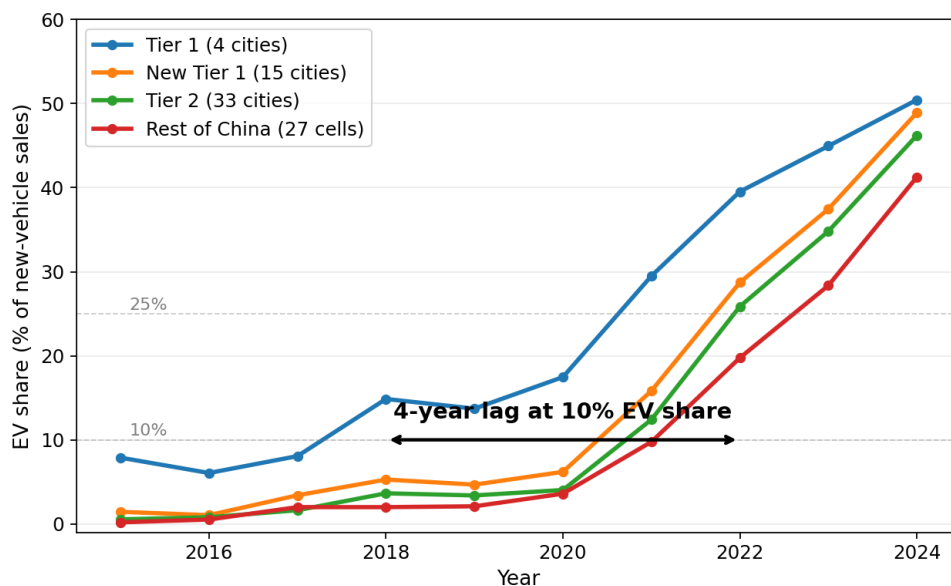
differences in adoption timing across city tiers. Tier 1 cities, Beijing, Shanghai, Guangzhou, and Shenzhen, crossed a 10% EV share in 2018. The Rest tier reached the same threshold in 2022, four years later. The lag narrows as the technology matures, falling to about two years at the 25% threshold and to about one year at the 40% threshold, as reported in Table 6. This timing pattern is central to the heterogeneous diffusion interpretation of the paper. The same national policy can have different effects across cities because cities are located at different points on the adoption curve when the policy is in place.

Table 6: Year Each Tier First crosses an EV-share Threshold

City tier	5% threshold	10% threshold	25% threshold	40% threshold
T1	2015	2018	2021	2023
NT1	2018	2021	2022	2024
T2	2021	2021	2022	2024
Rest of China	2021	2022	2023	2024
<i>Lag vs. Tier 1</i>	6 yr (Rest)	4 yr (Rest)	2 yr (Rest)	1 yr (Rest)

*Notes:* The table reports the first calendar year in which each city tier’s Q-weighted EV share (BEV + PHEV + REEV as a fraction of all new-vehicle sales in the tier) reaches each threshold. T1 = {Beijing, Shanghai, Guangzhou, Shenzhen}; NT1 = 15 new first-tier cities (Chengdu, Hangzhou, Wuhan, Chongqing, Nanjing, Tianjin, Suzhou, Xi’an, Changsha, Shenyang, Qingdao, Zhengzhou, Dalian, Dongguan, Ningbo); T2 = 33 second-tier cities in the panel; Rest of China = the 27 small-city residual cells. The bottom row reports the diffusion lag of the Rest tier relative to Tier 1.

Figure 11: EV Diffusion Timing across City Tiers



Notes: The figure plots EV shares in new vehicle sales by city tier. EVs include BEVs, PHEVs, and REEVs. Tier 1 includes Beijing, Shanghai, Guangzhou, and Shenzhen. New Tier 1 includes 15 large metropolitan markets below the four Tier 1 cities. Tier 2 includes 33 remaining second tier cities in the sample. Rest includes 27 smaller city cells. The arrow marks the difference between Tier 1 and Rest in the first year each tier exceeds a 10% EV share.

Having established that EV adoption arrived earlier in larger and richer markets, Figure 12 examines how this timing shaped the distribution of subsidy benefits. Cities are divided into low, middle, and high income groups using terciles of 2020 GDP per capita. The figure compares the historical allocation of EV sales and subsidy benefits with an accounting counterfactual that applies the 2015 subsidy level to the 2024 market.


Panel a shows the changing geography of EV adoption. In 2015, high income cities accounted for 73% of national EV sales, while low income cities accounted for only 13%. By 2024, EV adoption had spread much further down the income distribution: low income cities accounted for 35% of EV sales, while high income cities accounted for 37%. This diffusion occurred as national cash subsidies were being phased out. As a result, subsidy support was most generous when EV adoption was concentrated in richer markets, and much less generous by the time adoption had become more geographically balanced.

Panel b makes this timing channel explicit. Historically, high income cities received 54.5% of cumulative subsidy benefits, compared with 23.6% for low income cities and

21.8% for middle income cities. In the accounting counterfactual that applies the 2015 subsidy level to the 2024 market, the distribution is much more balanced: low income cities receive 34.5%, middle income cities receive 28.7%, and high income cities receive 36.8%. The same nominal subsidy is also larger relative to vehicle prices in lower income cities, so the effective subsidy rate is highest where affordability constraints are likely to matter most.

Taken together, the two panels show a timing mismatch in subsidy support. The historical schedule was most generous when early EV adoption was concentrated in high income cities. By the time EV demand had spread to lower income cities, the subsidy had largely been withdrawn. The realized incidence was therefore skewed toward richer cities, even though applying the same subsidy level to the later, more mature EV market would have produced a more balanced distribution of benefits. The welfare calculations below build on this timing pattern by separating the direct subsidy channel from the indirect learning channel.

Figure 12: EV Subsidy Incidence by City Income Group



*Notes:* The panels describe the distribution of EV adoption and subsidy benefits across city income groups. City income groups are defined by terciles of 2020 GDP per capita within the estimation sample. Panel (a) plots each group’s share of national EV sales from 2015 to 2024. The black line reports the national cash subsidy per vehicle on the right axis. Panel (b) reports each group’s share of cumulative subsidy benefits. The historical bars use realized EV sales and the cash subsidy schedule during the active subsidy years. The counterfactual bars apply the 2015 subsidy level to the 2024 market, holding quantities and product mix fixed. The annotations report the subsidy rate as a share of the average vehicle price in each income group.

### 6.3 Direct and Indirect Welfare Channels

The previous subsection showed that subsidy benefits were unevenly distributed across cities because EV adoption arrived earlier in larger and richer markets. We now separate the consumer surplus loss from subsidy removal into two welfare channels. The *direct* channel captures the contemporaneous price and demand effect of removing direct subsidy support, holding the battery cost path and product availability fixed. The *indirect* channel captures the Wright’s law feedback: lower counterfactual EV quantities in earlier years reduce cumulative production, raise the counterfactual battery cost path, increase EV marginal costs, and raise the equilibrium prices faced by consumers in 2024. The direct channel is therefore the effect of removing the fiscal transfer in a given market environment, while the indirect channel is the loss of compounded learning benefits created by weaker early demand.

Table 7 reports this decomposition by city tier. Two patterns stand out. First, the indirect Wright’s law channel accounts for 48.0% of the national consumer surplus loss. A welfare analysis that ignores cumulative learning would therefore miss roughly half of the consumer cost of subsidy removal. Second, the indirect share is slightly larger in Tier 1, New Tier 1, and Tier 2 cities (49% to 50%) than in the Rest of China (44%). The reason is that the learning channel works through EV prices, and EV demand remains more concentrated in higher tier cities. A given equilibrium price increase therefore moves more consumer surplus in those markets.

The level of consumer surplus should be interpreted with the demand model in mind. Consumer surplus is evaluated using the BLP inclusive value formula,

$$CS_i = \frac{1}{|\alpha_i|} \log \sum_j \exp(u_{ij}).$$

Because the price coefficient enters the denominator, the absolute monetary scale of consumer surplus depends on the elasticity specification. For this reason, the most robust objects in Table 7 are the direct versus indirect split and the relative differences across city tiers. These comparisons show that the welfare cost of subsidy removal is not only larger in aggregate when learning is included, but also much more concentrated in advanced urban markets. Per capita consumer surplus losses range from RMB 292 in Tier 1 cities to RMB 62 in the Rest of China.

The decomposition has a policy design implication. The direct channel is what a purchase subsidy buys in a fixed market environment. The indirect channel is the additional value of supporting scale and cumulative learning. The table does not rank alternative policy instruments, but it shows that roughly half of the consumer surplus benefit comes from the production scale margin rather than from the direct transfer alone. This sug-

Table 7: Direct and indirect channels of consumer surplus loss by city tier

City tier	Pop. (M)	CS loss (bn RMB)	Direct	Indirect	Per cap. (RMB)
Tier 1	83.6	24.4	50.7%	49.3%	292
New Tier 1	200.5	50.0	49.9%	50.1%	250
Tier 2	242.8	49.6	50.9%	49.1%	204
Rest of China	796.2	48.9	55.8%	44.2%	62
National	1,323.1	172.9	52.0%	48.0%	131

*Notes:* The table decomposes the 2024 consumer surplus loss from subsidy removal into direct and indirect channels. The direct channel removes direct subsidy support while holding the battery cost path and product availability fixed. Direct subsidy support includes the national cash subsidy, local cash subsidy, and purchase tax exemption. The indirect channel additionally incorporates the Wright’s law feedback from lower cumulative EV production to higher battery costs, higher EV marginal costs, and higher equilibrium prices. Direct and indirect columns report each channel’s share of the total consumer surplus loss within the city tier. Consumer surplus losses are reported in billion RMB, and per capita losses are reported in RMB per person. Values are computed under the log price BLP specification used in the main demand model.

gests that instruments that preserve early scale, including production side incentives or charging infrastructure, may capture part of the learning benefit without relying entirely on purchase subsidies.

#### 6.4 Reconciling the Two Subsidy Results

The negative Subsidy block in the Shapley decomposition and the positive effect of subsidy support in the no subsidy counterfactual answer different questions. The Shapley decomposition is a historical accounting exercise. It asks how much of the 2015 to 2024 EV share change is attributable to replacing the 2015 subsidy environment with the 2024 subsidy environment, averaged across coalitions of the other decomposition blocks. Because direct purchase support was phased down over this period, the Subsidy block is negative.

The no subsidy counterfactual asks a different question: how much did direct subsidy support matter relative to a path without that support? In this exercise, subsidies raise EV adoption by lowering consumer net prices, supporting cumulative learning, and sustaining product availability. There is therefore no contradiction between a negative historical Subsidy contribution in the decomposition and a positive role for subsidies in the dynamic counterfactual. The first result explains the historical transition, while the second evaluates the market consequences of removing direct subsidy support.

This distinction also clarifies the different subsidy magnitudes reported in the paper. The Subsidy Shapley value of  $-13.63\%$  measures the contribution of moving from the more generous 2015 direct subsidy schedule to the less generous 2024 schedule. The dynamic removal effects of  $-22.8\%$  to  $-32.8\%$  measure the effect of removing direct subsidy support

along the dynamic path, allowing battery learning and product availability to respond. The demand elasticity robustness estimate of  $-23.4\%$  measures the same dynamic removal exercise under a less elastic demand calibration, with  $|\alpha| = 2.5$ . These numbers differ because they correspond to different counterfactual questions.

The welfare accounting in this section is private surplus only. It includes consumer surplus, producer surplus, and the fiscal cost of cash subsidies and tax exemptions. It does not include environmental benefits from replacing internal combustion vehicles with EVs, such as reductions in tailpipe CO<sub>2</sub> or local air pollutants, nor does it include broader industrial policy spillovers in batteries or supply chains. Incorporating these external benefits would require assumptions about the counterfactual internal combustion fleet, vehicle usage, regional grid emissions, and local pollution damages. This distinction is especially important for EVs because the environmental value of electrification depends on the electricity generation mix and on where the vehicle is charged and driven. [Holland et al. \[2016\]](#) and [Holland et al. \[2019\]](#) make this point in the U.S. context; the same issue is relevant for China given spatial variation in the grid and local pollution exposure. We therefore leave a full social cost benefit analysis outside the scope of this paper.

The two subsidy results have a joint implication for policy design. The negative historical Shapley contribution reflects the timing of the policy schedule: direct support was more generous early in the transition and less generous by 2024. The positive no subsidy counterfactual shows that subsidy support still generated consumer side gains through both direct price effects and indirect learning effects. Together, the results imply that the timing and targeting of subsidy support are central design margins.

The 2015 to 2024 schedule delivered the largest benefits when EV adoption was concentrated in early adopting, high income cities, and it was largely phased down by the time EV demand had spread to lower tier markets. The realized incidence pattern therefore reflects the policy schedule as well as the underlying technology. A different schedule, for example one linked to city tier, buyer income, or vehicle price, would shift the distribution of benefits. A policy mix that combines purchase support with instruments that preserve early production scale could also capture part of the indirect learning channel. We do not solve that policy design problem here. The contribution of the counterfactual analysis is to quantify the margins that such a redesign would need to consider: timing, city incidence, direct price support, and indirect learning.

## 7 Conclusion

This paper provides a structural account of China's EV diffusion from 2015 to 2024. The evidence points to a transition driven not by any single force, but by the interaction of several complementary forces: improving EV product quality, falling battery costs,

entry of new products and firms, stronger urban EV demand, and direct policy support. China's experience therefore illustrates a broader feature of technology diffusion: adoption depends not only on the arrival of a new technology, but also on whether firms embody that technology in attractive products, whether costs fall with scale, whether consumers substitute toward the new products, and whether policy support arrives at the right stage of market development.

The main contribution of the paper is not a simple ranking of these mechanisms. Instead, the paper shows why such a ranking is intrinsically difficult when quality, cost, entry, demand, and policy are complements. The Shapley decomposition organizes a complex transition into transparent blocks and makes explicit how much each contribution depends on the order in which mechanisms are introduced. Product improvement is the largest aggregate driver of China's EV transition, but its effect is amplified by variety expansion and battery cost decline. Direct subsidies helped create and sustain the market in its early stages, but their historical contribution is negative because direct purchase subsidies were phased out over the sample period. This result highlights the importance of distinguishing a historical decomposition from a policy counterfactual: the first explains the observed transition, while the second evaluates what would have happened without policy support.

The policy counterfactuals show that subsidies mattered through both direct and indirect channels. Roughly half of the consumer surplus effect of subsidy support operates through the direct cash transfer, while the remaining effect operates through Wright's law learning and the equilibrium product responses that follow from higher early adoption. This finding links micro level adoption incentives to aggregate market performance. Policies that raise early scale can affect later costs, product availability, and welfare, even after the initial subsidy is removed. At the same time, the incidence of these gains is uneven. Welfare gains are concentrated in higher income cities, and producer surplus is concentrated among firms whose portfolios are more exposed to the EV margin, including BYD, Tesla, and the New Forces. China's EV transition therefore generated not only aggregate diffusion, but also meaningful reallocation across cities and firm groups.

These results have implications for technology policy. First, the regressive incidence of EV subsidies is not an inherent feature of the technology. It depends on the timing and design of the policy schedule. Early subsidies mainly benefit markets where the new technology is already commercially viable, while the same instrument applied in a more mature market can reach a broader set of consumers. Second, policies that support scale, learning, or product availability may capture part of the dynamic benefit of subsidies at lower fiscal cost than direct cash transfers alone. Third, because diffusion depends on complementarities among quality, cost, entry, and demand, evaluating technology policy requires an equilibrium framework that connects micro data to market wide outcomes.

Several limitations point to useful directions for future work. The model focuses

on the new vehicle market and does not incorporate the rapid development of China's used EV market. Used EV substitution could attenuate some of the no subsidy effects measured at the new vehicle margin. Charging infrastructure also enters only indirectly, through city level controls and the EV demand interactions. A richer dynamic model in which charging infrastructure, product entry, and EV adoption co evolve would close an additional feedback loop that is held fixed here. Finally, the analysis focuses on purchase subsidies and tax exemptions rather than comparing these instruments with mandates, emission standards, or research subsidies. These comparisons are natural next steps for evaluating the design of technology policy.

Taken together, the results suggest that China's EV transition was not simply the outcome of consumer subsidies or autonomous technological progress. It was a mutually reinforcing process in which policy supported scale, firm entry, quality upgrading, and learning reduced the effective cost of EV adoption over time. This dynamic complementarity helps explain both the speed of China's EV diffusion and the difficulty of assigning that diffusion to any single policy or market force. More broadly, the paper shows how technology diffusion can reshape economic performance through equilibrium changes in product quality, firm composition, consumer welfare, and market structure.

## References

- Daron Acemoglu, Philippe Aghion, Leonardo Bursztyn, and David Hemous. The environment and directed technical change. *American Economic Review*, 102(1):131–166, 2012.
- Philippe Aghion, Jing Cai, Mathias Dewatripont, Luosha Du, Ann Harrison, and Patrick Legros. Industrial policy and competition. *American Economic Journal: Macroeconomics*, 7(4):1–32, 2015.
- C. Lanier Benkard. A dynamic analysis of the market for wide-bodied commercial aircraft. *Review of Economic Studies*, 71(3):581–611, 2004.
- Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, 1995.
- Steven Berry, James Levinsohn, and Ariel Pakes. Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of Political Economy*, 112(1):68–105, 2004.
- Steven T. Berry and Philip A. Haile. Identification in differentiated products markets using market level data. *Econometrica*, 82(5):1749–1797, 2014.
- Bryan Bollinger and Kenneth Gillingham. Learning-by-doing in solar photovoltaic installations. *American Economic Review*, 109(6):2535–2567, 2019.
- Ying Fan. Ownership consolidation and product characteristics: A study of the u.s. daily newspaper market. *American Economic Review*, 103(5):1598–1628, 2013.
- Meredith Fowlie, Mar Reguant, and Stephen P. Ryan. Market-based emissions regulation and industry dynamics. *Journal of Political Economy*, 124(1):249–302, 2016.
- Gautam Gowrisankaran and Marc Rysman. Dynamics of consumer demand for new durable goods. *Journal of Political Economy*, 120(6):1173–1219, 2012.
- Stephen P. Holland, Erin T. Mansur, Nicholas Z. Muller, and Andrew J. Yates. Are there environmental benefits from driving electric vehicles? the importance of local factors. *American Economic Review*, 106(12):3700–3729, 2016.
- Stephen P. Holland, Erin T. Mansur, Nicholas Z. Muller, and Andrew J. Yates. Distributional effects of air pollution from electric vehicle adoption. *Journal of the Association of Environmental and Resource Economists*, 6(S1):S65–S94, 2019.

- Shanjun Li, Lang Tong, Jianwei Xing, and Yiyi Zhou. The market for electric vehicles: Indirect network effects and policy design. *Journal of the Association of Environmental and Resource Economists*, 4(1):89–133, 2017.
- Erich Muehlegger and David S. Rapson. Subsidizing low- and middle-income adoption of electric vehicles: Quasi-experimental evidence from California. *Journal of Public Economics*, 216:104752, 2022.
- Aviv Nevo. Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2):307–342, 2001.
- Richard G. Newell, Adam B. Jaffe, and Robert N. Stavins. The induced innovation hypothesis and energy-saving technological change. *Quarterly Journal of Economics*, 114(3): 941–975, 1999.
- Björn Nykvist and Måns Nilsson. Rapidly falling costs of battery packs for electric vehicles. *Nature Climate Change*, 5:329–332, 2015.
- Amil Petrin. Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy*, 110(4):705–729, 2002.
- David Popp. Induced innovation and energy prices. *American Economic Review*, 92(1): 160–180, 2002.
- Mathias Reynaert. Abatement strategies and the cost of environmental regulation: Emission standards on the European car market. *Review of Economic Studies*, 88(1):454–488, 2021.
- Stephen P. Ryan. The costs of environmental regulation in a concentrated industry. *Econometrica*, 80(3):1019–1061, 2012.
- James M. Sallee. The surprising incidence of tax credits for the toyota prius. *American Economic Journal: Economic Policy*, 3(2):189–219, 2011.
- Lloyd S. Shapley. A value for n-person games. In *Contributions to the Theory of Games II*, pages 307–317. Princeton University Press, 1953.
- Anthony F. Shorrocks. Decomposition procedures for distributional analysis: A unified framework based on the shapley value. In *Income Inequality: Economic Disparities and the Middle Class in Affluent Countries*, pages 59–90. Stanford University Press, 2013.
- Katalin Springel. Network externality and subsidy structure in two-sided markets: Evidence from electric vehicle incentives. *American Economic Journal: Economic Policy*, 13(4): 393–432, 2021.

- Andrew Sweeting. Dynamic product positioning in differentiated product markets: The effect of fees for musical performance rights on the commercial radio industry. *Econometrica*, 81(5):1763–1803, 2013.
- Peter Thompson. The relationship between unit cost and cumulative quantity and the evidence for organizational learning-by-doing. *Journal of Economic Perspectives*, 26(3): 203–224, 2012.
- Kate S Whitefoot, Meredith L Fowlie, and Steven J Skerlos. Compliance by design: Influence of acceleration trade-offs on co2 emissions and costs of fuel economy and greenhouse gas regulations. *Environmental science & technology*, 51(18):10307–10315, 2017.
- Thomas G. Wollmann. Trucks without bailouts: Equilibrium product characteristics for commercial vehicles. *American Economic Review*, 108(6):1364–1406, 2018.
- T. P. Wright. Factors affecting the cost of airplanes. *Journal of the Aeronautical Sciences*, 3(4): 122–128, 1936.
- Micah S. Ziegler and Jessika E. Trancik. Re-examining rates of lithium-ion battery technology improvement and cost decline. *Energy & Environmental Science*, 14(4):1635–1651, 2021.

## Appendix A Order Sensitivity Diagnostic

Table [Appendix A.1](#) reports the distribution of marginal contributions across coalitions for the 2024 Shapley decomposition. The Shapley value is the order-invariant average contribution, while the minimum and maximum columns show how much a block’s sequential contribution can vary depending on which other blocks have already been introduced.

Table Appendix A.1: Order sensitivity diagnostic for the 2024 Shapley values

Block	Shapley	Min	Max	Range	Equal weighted mean
Quality	45.49%	0.00%	67.89%	67.89%	40.85%
Variety	14.81%	-20.08%	66.85%	86.94%	12.14%
Battery	8.20%	0.00%	55.41%	55.41%	3.06%
Subsidy	-13.63%	-56.12%	0.00%	56.12%	-13.95%
Residual	-7.96%	-57.89%	19.06%	76.95%	-11.08%
Market	-3.27%	-23.32%	31.39%	54.70%	0.55%

*Notes:* The table reports the distribution of marginal contributions across coalitions in the 2024 Shapley decomposition. The Shapley column reports the weighted average contribution used in the main decomposition. The minimum and maximum columns report the smallest and largest sequential marginal contribution across coalitions. All values are reported in EV share units.

## Appendix B Additional Policy Counterfactual Results

### Appendix B.1 Additional Dynamic Counterfactual Results

Table [Appendix B.1](#) accompanies the discussion of dynamic amplification channels in Section 6.1 by reporting the no-subsidy 2024 EV quantity under each of the three regimes alongside the corresponding shortfall relative to actual. The reduction from 4.1 to 2.8 million between the Static and Learning regimes isolates the contribution of Wright’s-law battery-cost learning to subsidy efficacy; the further reduction from 2.8 to 1.8 million between the Learning and Full dynamic regimes isolates the contribution of endogenous entry and exit. The two channels deliver roughly comparable additional reductions of 1.2 and 1.0 million vehicles, indicating that neither channel dominates the dynamic-amplification mechanism.

Table Appendix B.1: No-subsidy EV quantity by dynamic regime

Regime	No-subsidy EV quantity	Shortfall from actual	Incremental reduction
Static	4.1 million	7.9 million	–
Learning only	2.8 million	9.2 million	1.2 million
Full dynamic	1.8 million	10.2 million	1.0 million

*Notes:* The actual 2024 EV quantity is 12.0 million. The Static regime holds product availability and battery costs fixed. The Learning only regime allows battery costs to respond through cumulative production while holding product availability fixed. The Full dynamic regime allows both battery-cost learning and entry/exit responses. The incremental reduction is measured relative to the previous row.

### Appendix B.2 Additional Welfare and Incidence Accounting

Tables [Appendix B.2](#), [Appendix B.3](#), and [Appendix B.4](#) report the welfare-accounting components, the city-tier subsidy pass-through ratios, and the firm-group EV-quantity-retained shares that underlie the discussion in Section 6.2. The three tables share a common scenario: the 2024 actual subsidy package is removed under the canonical endo\_endo regime of the dynamic simulation, and the resulting consumer-surplus, producer-surplus, and quantity changes are decomposed across consumer tiers and producer groups. Together they document that the same nominal subsidy distributes its benefit unevenly across cities (Table [Appendix B.3](#)), across firm types (Table [Appendix B.4](#)), and between consumers and producers (Table [Appendix B.2](#)).

Table Appendix B.2: Subsidy-removal welfare accounting and incidence

Object	Estimate
2024 EV subsidy spend	209.6 bn yuan/year
Producer surplus gain from subsidy	536.0 bn yuan/year
Consumer surplus gain from subsidy	172.9 bn yuan/year
Total private surplus gain (CS+PS)	708.9 bn yuan/year
Total surplus per yuan of subsidy	3.38
PS share of private surplus gain	76%
Mean subsidy per EV	17,813 yuan
Average per-capita CS loss under removal	131 yuan/year
<i>Memo: Implicit Tier 1 license-plate subsidy (excluded)</i>	
Tier-1 per-capita CS loss under removal	292 yuan/year
New Tier-1 per-capita CS loss under removal	250 yuan/year
Tier-2 per-capita CS loss under removal	204 yuan/year
Rest-of-China per-capita CS loss under removal	62 yuan/year

*Notes:* The table reports the 2024 subsidy-removal accounting, all values on a common national scale. Subsidy spend is the fiscal cost of the 2024 NEV cash schedule plus the 10% purchase-tax exemption. Consumer surplus is computed from the BLP inclusive-value formula on the 79-city panel and scaled to the national NEV registration count by multiplying by a market-size factor of 28.6 (the ratio of national 2024 NEV registrations  $\approx$  11.15 million to BLP-sample EV quantity  $\approx$  0.39 million). Producer surplus is computed at the national level from the recovered Bertrand markup and the model-implied national EV quantity. The total surplus per yuan of subsidy (3.38) and the within-PS ranking across firm groups in Table [Appendix B.4](#) are the headline distributional takeaways. The table does not include the implicit value of the EV license-plate priority in Tier 1 cities, which is reported separately in Table [Appendix B.3](#) and discussed in Section ??.

Table Appendix B.3: Subsidy pass-through to consumer surplus by city tier, 2024

City tier	Cities	EV qty (M)	Subsidy received (bn)	CS gain (bn)	Pass-through
Tier 1	4	1.29	29.1	24.4	83.8%
New Tier 1	15	3.15	63.6	50.0	78.6%
Tier 2	33	3.26	60.5	49.6	82.0%
Rest of China	27	4.43	62.5	48.9	78.3%
National	79	12.13	215.7	172.9	80.2%

*Notes:* Pass-through ratio is the fraction of nominal subsidy disbursement to EV buyers in each tier that translates into a consumer-surplus gain. EV qty is the BLP-sample EV quantity in the 2024 panel (a within-sample share-of-agent-pool measure that scales proportionally to the national NEV registration count). Subsidy received aggregates the per-product cash subsidy plus 10% purchase-tax exemption over the tier's 2024 EV products. CS gain is the actual-minus no-subsidy consumer-surplus change under the full-dynamic regime of Table 7. Pass-through is below 100% in all tiers because Bertrand–Nash firms capture part of the subsidy in higher markups, and pass-through is moderately higher in Tier 1 and Tier 2 (premium-product mix is less elastic to net price than the small-vehicle EV mix in NT1 and Rest tiers).

Table Appendix B.4: Firm-group EV quantity retained under subsidy removal

Firm group	Actual quantity	No-subsidy quantity	Retained share
BYD	4.26 million	0.69 ± 0.07 million	16.1%
Tesla	0.74 million	0.20 ± 0.05 million	26.7%
Foreign/JV	0.70 million	0.10 ± 0.02 million	14.7%
Trad. OEM	2.69 million	0.31 ± 0.05 million	11.4%
New Forces	1.53 million	0.27 ± 0.05 million	17.9%
Private National	1.98 million	0.25 ± 0.03 million	12.6%
Other	0.08 million	0.00 ± 0.00 million	4.7%
<b>Total</b>	<b>12.00 million</b>	<b>1.82 million</b>	<b>15.2%</b>

*Notes:* 2024 EV quantity by firm group under the actual-subsidy scenario and the endo\_endo no-subsidy counterfactual (50 Monte Carlo draws of the stochastic Stage D entry, reported as cross-draw mean ± one standard deviation). Retained share is the no-subsidy mean EV quantity divided by the actual quantity for each firm group. Firm groups follow the seven-class manufacturer classification described in the supply-side appendix: BYD is reported separately given its dominant market share; Tesla is reported separately for its premium-import exposure; Foreign/JV pools international joint ventures; Trad. OEM pools the state-owned legacy manufacturers; New Forces pools the post-2014 EV-only entrants (Nio, XPeng, Li Auto, etc.); Private National pools private-domestic firms; Other pools the remaining small-share manufacturers. The Tesla and New Forces groups retain the largest share of their 2024 EV business under subsidy removal, while small-share *Other* manufacturers and traditional OEMs are the most exposed.

### Appendix B.3 Supply-side and Robustness Diagnostics

Table [Appendix B.5](#) reports the quantity-weighted mean Lerner index recovered from the multi-product Bertrand-Nash first-order conditions, broken down by fuel type and calendar year. The Lerner index measures the price-cost margin as a share of the producer price. Two patterns stand out. First, all five fuel-type margins sit in the 0.20 to 0.25 band throughout the sample, indicating that the demand and supply specifications together imply economically reasonable pricing power that is similar in magnitude to U.S. and European automobile-BLP estimates. Second, the BEV and ICE Lerner ratios rise by about 3% between 2015 and 2024, consistent with consolidation and product-line specialization over the diffusion window. The PHEV margin declines slightly over the same window, reflecting the entry of small-battery PHEVs at lower markups. The REEV class appears in 2021 and stabilizes near 0.25.

Table Appendix B.5: Q-weighted Lerner index by fuel type and year

Year	BEV	HEV	ICE	PHEV	REEV
2015	0.213	0.219	0.205	0.234	—
2016	0.220	0.221	0.213	0.237	—
2017	0.217	0.225	0.215	0.235	—
2018	0.221	0.229	0.219	0.238	—
2019	0.232	0.236	0.227	0.241	—
2020	0.236	0.240	0.232	0.245	—
2021	0.238	0.242	0.235	0.244	0.249
2022	0.240	0.243	0.239	0.243	0.248
2023	0.243	0.244	0.240	0.241	0.247
2024	0.242	0.246	0.242	0.239	0.247

*Notes:* The Lerner index is computed at the product-year-city level from the recovered marginal costs and observed consumer net prices and then aggregated by quantity-weighted mean within each (fuel type, year) cell. ICE is ICEV.

Table [Appendix B.6](#) reports the canonical and lit-median robustness exercises side by side. The canonical column uses the estimated demand-elasticity  $|\alpha| \approx 4.9$  from the demand model of Section 3. The robustness column rescales  $|\alpha|$  to the EV-demand literature median of 2.5 before re-running the dynamic simulation. The structural ranking across the three dynamic regimes is preserved in both calibrations: the Static effect is smaller in magnitude than the Learning-only effect, which is in turn smaller than the Full-dynamic effect. The robustness column delivers subsidy effects roughly 30% smaller in magnitude than the canonical column, but the qualitative finding that subsidies remove a substantial share of 2024 EV adoption is preserved.

Table Appendix B.6: Sensitivity of subsidy-removal effects to the demand elasticity

Regime	$ \alpha  = 4.9$ (canonical)	$ \alpha  = 2.5$ (lit median)	Ratio
Static (exog_exog)	-22.75%	-13.60%	0.60
Learning (exog_endo)	-28.91%	-18.85%	0.65
Full dynamic (endo_endo)	$-32.76 \pm 0.65\%$	$-23.40 \pm 1.35\%$	0.71

*Notes:* 2024 subsidy-removal effect on EV share under the canonical demand-elasticity calibration ( $|\alpha| = 4.9$ ) and the literature-median calibration ( $|\alpha| = 2.5$ , obtained by scaling both  $\beta_{\log p}$  and  $\pi_p$  by the same factor). The robustness column is from a fresh dynamic-simulation run that uses identical data, BLP demand attributes, and Wright’s-law parameters as the canonical run, but with the rescaled price coefficients. Ratio is the robustness-column estimate divided by the canonical estimate. The endo\_endo standard deviations are across 50 stochastic-entry draws.

The  $|\alpha| = 2.5$  rescaling is applied only to the dynamic simulation. We do not re-estimate the static Shapley decomposition at the rescaled elasticity because the Shapley equilibrium-solver inputs include both the demand parameters and the recovered marginal costs, and a consistent re-derivation would require re-running the marginal-cost backout under the rescaled demand parameters before re-solving the coalition equilibria. Such a re-derivation is a substantive recomputation rather than a sensitivity exercise, and we leave it to future work that targets static Shapley robustness as a primary research question. The dynamic-sim robustness reported in Table [Appendix B.6](#) provides a calibrated bound on how the no-subsidy effects scale with  $|\alpha|$  and is the most policy-relevant sensitivity diagnostic that the present specification supports.

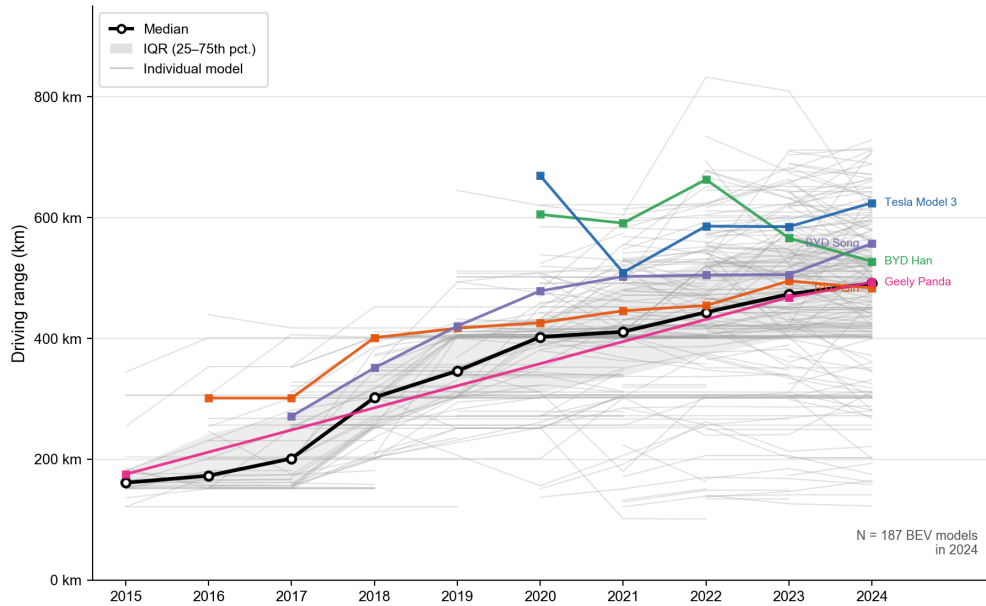
Table Appendix B.7: Firm-group definitions

Firm group	Brands	2024 EV %	Member brands
BYD	4	100%	BYD, Denza, Yangwang, Fang Cheng Bao.
Tesla	1	100%	Tesla (Shanghai Gigafactory).
Foreign / JV	28	7%	Foreign brands sold in China (often via joint ventures with Chinese SOEs): VW, Audi, Skoda, BMW, Mercedes-Benz, Mini, Smart; Toyota, Honda, Nissan, Mazda, Mitsubishi, Lexus, Acura, Infiniti; Buick, Chevrolet, Cadillac, Ford, Lincoln; Hyundai, Kia; Peugeot, Citroën, Fiat; MG, Jaguar, Land Rover, Lotus, Volvo, Polestar.
Trad. OEM (SOE)	28	50%	Brands under the seven traditional state-owned auto groups: SAIC (Roewe, MG, Maxus, IM, Rising), FAW (Hongqi, Bestune, Besturn), Dongfeng (Dongfeng, Voyah, Aeolus, M-Hero), Changan (Changan, Deepal, Avatr, Kaicene), GAC (Trumpchi, Aion, Hyper), BAIC (Beijing, ARCFOX, Stelato), Brilliance (Jinbei, Zhonghua).
New Forces	8	100%	EV-native startups founded after 2014: NIO, Li Auto, XPeng, Leapmotor, Nezha, Xiaomi, WM Motor, HiPhi.
Private National	19	39%	Brands under the five non-SOE Chinese auto groups: Geely (Geely, Geometry, Zeekr, Galaxy, Lynk & Co, Lotus, Polestar, Smart <sup>†</sup> ), Great Wall (Haval, WEY, Tank, ORA), Chery (Chery, Exeed, Jetour, iCAR), JAC, Seres.
Other	23	77%	Niche manufacturers and exits: Skywell, Leapmotor-spinoffs, Dorcen, Hanteng, Bestune-spinoffs, JMC, Haima, Yundu, Zotye, smaller LEV-segment producers.
Total	111	43%	

Notes: “Brands” counts distinct brand-level entries in the 2024 BLP panel. “2024 EV %” is the share of 2024 product-rows in the group that have FuelType ∈ {BEV, PHEV, REEV}, weighted equally across products. Group assignment follows code/\_brand\_groups.py. Foreign brands are classified by the consumer-perceived brand name, not by the manufacturing JV parent: e.g., “Volkswagen” is Foreign/JV even though it is produced by SAIC-Volkswagen and FAW-Volkswagen. <sup>†</sup>Smart, Polestar, and Lotus appear under Private National because Geely holds the majority stake; their consumer-perceived brand identity remains foreign and would alternatively place them in Foreign/JV. “Other” collects manufacturers that do not map to any of the five Chinese national parent groups and are not in Tesla, BYD, or the New Forces list; this group is dominated by exiting or niche players and is included to keep the row totals consistent with the BLP panel.

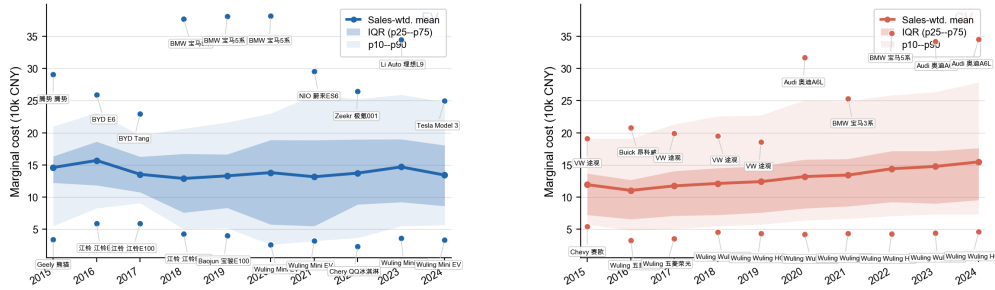
## Appendix C Figures

Figure 13: BEV driving-range trajectories by product, 2015–2024.



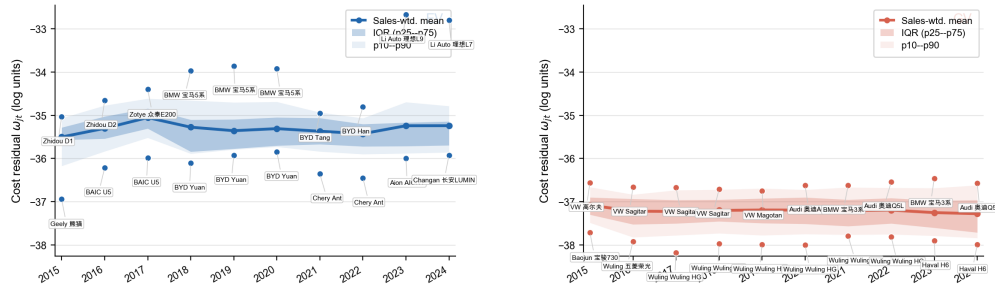
Notes: Each thin gray line traces one BEV model’s range (km) over the years it was sold. The black line is the cross-model median; the shaded band is the interquartile range (25th–75th percentile). Five representative models are highlighted. Range upgrades are product-specific in timing and magnitude, motivating the flexible per-product, per-year range treatment in demand estimation.

Figure 14: Distribution of recovered marginal costs, EV vs. GV



Notes: Distribution of the product-year-city marginal costs recovered from the multi-product Bertrand first-order conditions, plotted separately for EV (BEV + PHEV + REEV) and GV (ICE + HEV) products. The EV distribution is wider than the GV distribution, reflecting heterogeneity in battery size, electric drivetrain configuration, and the steeper trajectory of EV cost decline over 2015–2024 that is identified by the Year and EV × Year effects in Equation (10). The right tail of the EV distribution corresponds to large-battery long-range BEVs and premium PHEVs whose marginal cost levels remain comparable to mid-size GV products throughout the sample.

Figure 15: Distribution of unobserved marginal-cost residual  $\omega_{jm}$



Notes: Distribution of the unobserved supply-side residual  $\omega_{jm}$ , defined as the part of recovered log marginal cost not explained by the pooled OLS specification in Equation (10). The residual is mean-zero by construction and tightly concentrated for both EVs and GVs, indicating that the observable characteristics (vehicle size, engine power, EV driving range, BNEF battery cost, FuelType, BodyType, firm-group, and year fixed effects) absorb most of the cross-product variation in recovered marginal costs. The wider EV right tail reflects residual variation in battery chemistry and supply-chain procurement that is not captured by the BNEF aggregate-cost time series.

## Appendix D Tables

Table Appendix D.1: City tier definitions

Tier	Number of Cities	Member markets
Tier 1	4	Beijing, Shanghai, Guangzhou, Shenzhen.
New Tier 1	15	Chengdu, Hangzhou, Wuhan, Chongqing, Nanjing, Tianjin, Suzhou, Xi'an, Changsha, Shenyang, Qingdao, Zhengzhou, Dalian, Dongguan, Ningbo.
Tier 2	33	Foshan, Wuxi, Hefei, Fuzhou, Kunming, Xiamen, Jinan, Harbin, Changchun, Shijiazhuang, Quanzhou, Lanzhou, Urumqi, Guiyang, Nanning, Nanchang, Taiyuan, Haikou, Zhuhai, Wenzhou, Shaoxing, Jiaxing, Taizhou (Zhejiang), Huizhou, Zhongshan, Changzhou, Nantong, Tangshan, Xuzhou, Linyi, Weifang, Yantai, Jinhua.
Rest	27	Province-level residual cells, each labeled " <i>Province – Other</i> " and aggregating all prefectures in that province not separately identified in the BLP panel: Anhui, Fujian, Gansu, Guangdong, Guangxi, Guizhou, Hainan, Hebei, Heilongjiang, Henan, Hubei, Hunan, Inner Mongolia, Jiangsu, Jiangxi, Jilin, Liaoning, Ningxia, Qinghai, Shaanxi, Shandong, Shanxi, Sichuan, Tibet, Xinjiang, Yunnan, Zhejiang.
Total	79	

*Notes:* The Tier 1 / New Tier 1 / Tier 2 classification follows the China Business Network (CBN) annual city ranking, which is the standard tier definition in applied China-economics work. The Tier 2 row lists all prefecture-level cities that appear as their own market in the BLP panel but are not in the CBN Tier 1 or New-Tier 1 lists. The Rest row covers the 27 province-level residual cells used to absorb prefectures that the panel does not identify separately. Membership is fixed across all calendar years in the sample. The same 79 markets appear in every year from 2015 to 2024, yielding  $79 \times 10 = 790$  market-year observations.

Table Appendix D.2: Blocks in the Shapley Decomposition

Figure label	Block name	Economic content
Quality	Product attributes	Within-product values of log driving range (EV products), log engine power (GV products), log vehicle size, and the range $\times$ density interaction; toggled between 2015 and 2024 values for products present in both endpoint sets. Excludes fixed effects (which are year-invariant and handled within Variety), prices, policy variables, and demand residuals.
Variety	Entry and product variety	Changes in the available product set, including the expansion and reshaping of EV and GV product offerings.
Battery	Battery costs	Battery-related cost changes that affect EV marginal costs and the cost of providing longer driving range.
Subsidy	Subsidies and tax incentives	Changes in direct purchase subsidies, local subsidy components, and tax exemptions between 2015 and the endpoint year.
Residual	Residual demand and cost	Residual demand and cost components not explained by observed product attributes, prices, policies, demographics, or fixed effects.
Market	Income and market size	Changes in household income, population, market size, and demographic variables that affect demand and price sensitivity.

*Notes:* The first column reports the short labels used in the decomposition figures. Each block is switched from its 2015 value to its endpoint-year value while the remaining blocks are held fixed at their 2015 values. For each coalition of blocks, equilibrium prices, quantities, and market shares are recomputed before calculating the Shapley contribution.

Table Appendix D.3: Endpoint fit of the year-by-year decomposition

Endpoint year	$V(\emptyset)$	$V(\text{full})$	Observed $\Delta$	$\sum_b \phi_b$	Gap
2016	1.03%	1.10%	0.21%	0.08%	-0.13%
2017	1.03%	2.65%	1.69%	1.63%	-0.06%
2018	1.03%	4.22%	3.33%	3.20%	-0.13%
2019	1.03%	4.00%	3.14%	2.97%	-0.17%
2020	1.03%	5.71%	4.71%	4.69%	-0.02%
2021	1.03%	14.05%	12.88%	13.02%	0.15%
2022	1.03%	25.80%	24.63%	24.78%	0.15%
2023	1.03%	33.53%	33.02%	32.50%	-0.51%
2024	1.03%	44.66%	44.30%	43.64%	-0.75%

*Notes:* The baseline environment is 2015. In each row,  $V(\emptyset)$  is the model-implied EV share under the 2015 environment, while  $V(\text{full})$  is the model-implied EV share when all blocks are set to the listed endpoint-year values. The Shapley values satisfy  $\sum_b \phi_b = V(\text{full}) - V(\emptyset)$  by construction. The gap is  $\sum_b \phi_b$  minus the observed aggregate EV-share change. All values are reported in percentage units of aggregate EV share.

Table Appendix D.4: Tier-specific Shapley decomposition of the 2015–2024 EV-share change

Block	Tier 1	New Tier 1	Tier 2	Rest	National
Quality	55.59%	54.66%	44.12%	46.89%	45.49%
Variety	18.05%	21.41%	20.91%	16.13%	14.81%
Battery	9.41%	8.72%	8.34%	8.11%	8.20%
Subsidy	-10.38%	-13.05%	-14.81%	-18.90%	-13.63%
Residual	-31.33%	-20.63%	-6.93%	-5.43%	-7.96%
Market	0.69%	-4.33%	-6.78%	-6.39%	-3.27%
$\sum_b \phi_b$	42.02%	46.79%	44.85%	40.42%	43.64%
$V(\emptyset)$	8.71%	1.45%	0.55%	0.19%	1.03%
$V(\mathcal{B})$	50.73%	48.24%	45.40%	40.61%	44.66%

*Notes:* The table reports tier-specific Shapley decompositions of the model-implied 2015–2024 EV-share change at the 2024 endpoint. For each coalition, equilibrium prices and quantities are solved at the national product-market level. Tier-specific value functions are then computed by aggregating the resulting city-level EV shares within each tier. Each column closes within its own value function:  $\sum_b \phi_b = V(\mathcal{B}) - V(\emptyset)$ . Comparisons across tier columns show how the same channel contributes to EV-share changes within different city groups, rather than a mechanical allocation of the national Shapley value across tiers. Tier 1 includes Beijing, Shanghai, Guangzhou, and Shenzhen; New Tier 1 includes 15 new first-tier cities; Tier 2 includes 33 second-tier cities in the panel; Rest includes the remaining small-city cells. All values are reported in EV-share units.