Heterogeneous Diffusion of a New Technology: Subsidies, Reallocation, and the Rise of China's Electric-Vehicle Market, 2015–2024

Yu Hao*
The University of Hong Kong

April 29, 2026

Abstract

China's electric-vehicle share of new passenger-vehicle registrations rose from under one percent to 44% between 2015 and 2024, concentrating first in lower-income, lower-tier, budget segments before propagating up the income and quality distributions. We decompose this transition using a BLP demand system on a 79-market panel, a Shapley value over 2^8 counterfactual equilibria, and an Olley–Pakes split of the Lerner change. The mean Lerner rose by $+0.138$ points: within-firm growth accounts for 112%, reallocation -12%, while the Herfindahl index fell from $1{,}374$ to 887 — the inverse of the U.S. superstar pattern. Battery learning ($+27.3$ pp) and an EV-specific unobserved trend ($+15.1$ pp) drive most of the $+44.3$ pp adoption rise; the wealth-effect channel under 73% income growth pulls -22.9 pp. Direct subsidies contribute only -2.7 pp statically, but a never-existed-subsidy forward simulation gives 2024 EV share counterfactually -21.7 to -40.6 pp lower: subsidies were constitutive of the bootstrap, not a marginal price wedge.

Keywords: electric vehicles; technology diffusion; subsidy incidence; random-coefficient logit; equilibrium decomposition; China.
JEL codes: L62, O33, Q48, D12, L13.

*The University of Hong Kong, Hong Kong SAR. Email: haoyu@hku.hk.

1 Introduction

China's passenger-vehicle market underwent one of the fastest technology transitions documented in a major consumer-durable market: the combined battery-electric, plug-in hybrid, and range-extended share of new registrations rose from under one percent to nearly half in a single decade. Understanding which economic channels accounted for this transition within an equilibrium accounting framework—and which merely co-moved with it—is a step toward informing technology policy in other large markets contemplating similar shifts. Canonical models of technology diffusion (Griliches, 1957; Comin and Hobijn, 2010) emphasise profitability and information as correlates of adoption, but the Chinese EV case involves simultaneous movement along multiple margins—product entry, battery cost reduction, income growth, and direct purchase subsidies—whose relative contributions have not been decomposed within an equilibrium framework. We provide such a decomposition; we are explicit throughout that the resulting attributions are structural-accounting statements conditional on the maintained BLP functional form, not identified causal effects.

We characterise this transition as a process of *heterogeneous diffusion*. Unlike a uniform national rollout, the adoption of EV technology and the incidence of early subsidies were highly asymmetric across the cross-section of consumers and markets: diffusion was initially concentrated in lower-income demographics, lower-tier cities, and budget product segments, and only later propagated upward through the income distribution and the price/quality ladder. This bottom-up heterogeneity is central to understanding both the welfare implications of the transition and the political economy of the subsidy regime, because it determines where the equilibrium incidence of a given subsidy schedule lands and how that incidence interacts with subsequent supply-side product entry. The random-coefficient logit demand system below—which features a single income-heterogeneous price coefficient $\alpha_{ic} \propto y_{ic}^{-1}$—is designed precisely to capture this heterogeneous-diffusion pattern, by linking the cross-market income distribution to equilibrium adoption, markups, and welfare in a single demand structure.

This paper estimates a random-coefficient logit demand system on a panel of 79 urban markets over 2015–2024, recovers product-level marginal costs from the multi-market Bertrand–Nash first-order condition, and decomposes the equilibrium change into eight economic channels using the Shapley value over $2^8 = 256$ counterfactual coalitions. The demand model features a single income-heterogeneous price coefficient that generates both city-level and time-varying variation in price sensitivity (urban Chinese real disposable income grew 73% between 2015 and 2024, from ¥33.4k to ¥57.6k), linking the cross-market and over-time income distribution directly to equilibrium markups, EV adoption, and consumer welfare. The decomposition follows the logic of Grieco et al. (2024), extended from four blocks to eight to capture the rapid structural changes specific to a fast-moving technology transition in an emerging industry. We introduce a *Brand trajectory* block that isolates the systematic, per-brand annual evolution of unobserved product quality over time. This block absorbs the residual drift in brand appeal — capturing, for instance, the successful repositioning of domestic incumbents from budget to mainstream nameplates (BYD's $\Delta\xi = +2.69$ over the decade) and the foreign-incumbent retreat (Ford's $\Delta\xi = -1.56$) — and lets us separate genuine brand-level quality upgrading from aggregate EV-specific time trends. The technical justification for recovering this object via post-hoc projection of the BLP residual rather than as brand$\times$year fixed effects inside the BLP estimation is deferred to §D.2.1.

Within the decomposition — an accounting exercise conditional on the estimated BLP structure — a sharp separation emerges between the channels to which EV adoption is attributed and those to which

markup growth is attributed. On the adoption margin (Table 3), battery learning (A_{bat}, $+27.3$ pp, 62% of total) is the largest single contributor; the EV-specific unobserved trend (EV_trend, $+15.1$ pp, absorbing charging-infrastructure, smart-features, dealer-network buildout, and consumer-familiarity drift the observed covariates do not measure), the demand-environment bundle (D_{env}, $+12.0$ pp), Brand_trajectory ($+6.8$ pp), Entry_set ($+5.5$ pp), and A_{nonbat} ($+3.2$ pp) follow. Consumer_Composition — the per-city demand environment (population, income distribution, and GDP per capita; renamed from *Entry_demand* in earlier drafts because the block is unrelated to firm or product entry) — pulls in the *opposite* direction at -22.9 pp through the wealth-effect channel implied by the maintained $\alpha_{ic} = \hat{\pi}_p \cdot (y_{ic}/\bar{y})^{-1}$ form: when income rises, $|\alpha_{ic}|$ falls and the EV-vs-ICE price-gap response shrinks. Direct purchase subsidies contribute -2.7 pp on the static margin at the agent-integral specification.

On the markup margin, the industry mean Lerner rose from 0.254 to 0.423 ($+0.138$). Within-firm markup growth accounts for 112% of the firm-Q-proxy-weighted analogue ($+0.092$ Lerner; Table 4); across-firm reallocation contributes -12%. The within-firm component is broadly based and reflects three economic mechanisms the static framework cannot cleanly separate: (i) product-mix upgrading within incumbents (BYD's $\Delta\xi = +2.69$ from budget to premium nameplates is the clearest case); (ii) brand-premium growth as Chinese brands closed the perceived-quality gap with foreign JVs; (iii) lower model-implied price elasticity from the 73% income growth under $\alpha \propto y^{-1}$. The aggregate HHI fell from $1{,}374$ to 887: aggregate concentration declined while within-firm markups rose, the opposite of the U.S. superstar-firm pattern. Consumer_Composition absorbs most of the within-firm Lerner attribution under the maintained restriction; finite-mixture alternatives (Appendix B.1, R8) imply a substantially flatter income gradient, so we read the Consumer_Composition attribution as an upper bound and defer the full identified decomposition to a companion paper.

The closest methodological comparator is Grieco et al. (2024), who decompose the 1980–2018 U.S. automobile market with the same BLP toolkit. Their headline result is a re-characterisation of the implied marginal-cost path: when administrative and overhead costs are explicitly carved out from the FOC-implied marginal cost, technological progress *lowers* the per-vehicle attribute cost, but a rising administrative-cost share absorbs much of the price increase, leaving Lerner indices lower in 2018 than in 1980. They emphasise that the supply side is the dominant force in the U.S. four-decade sample, with income heterogeneity playing a small quantitative role because U.S. real income grew slowly. Our paper does *not* estimate the administrative-cost split GMY identify — we recover one composite marginal cost from the multi-market Bertrand–Nash FOC and decompose its movement into battery (A_{bat}) and non-battery (A_{nonbat}) attribute components. Both components contribute *negatively* to Lerner (-0.038 and -0.021), consistent with a cost-pass-through-dominated supply side over our window: battery cost fell from \$373/kWh in 2015 to \$115/kWh in 2024 (-69%), and the cost decline passes through to prices more than to markups. We focus on a different dimension from GMY: the *demand-side* wealth-effect channel that the maintained $\alpha \propto y^{-1}$ form generates under the 73% urban income growth observed over our decade — a mechanism that is muted in their U.S. window. In ten years, urban Chinese real disposable income grew 73% (¥33.4k to ¥57.6k), aggregate markups rose by 0.138 Lerner points, aggregate concentration *fell*, and the structural decomposition attributes the bulk of welfare growth to a decline in the model-implied market-wide price elasticity that the imposed $\alpha_{ic} \propto y_{ic}^{-1}$ demand system links to this income shift. Three patterns behind the reversal stand out relative to the U.S. benchmark, each of which we read as descriptive within the equilibrium accounting framework rather than as identified causal effects. *First*, the within-firm component of the Olley–Pakes markup decomposition is first-order: within-firm Lerner indices rise on incumbents'

existing products. The model-imposed $\alpha \propto y^{-1}$ form mechanically predicts that buyers' implied price-sensitivity falls under income growth, which can in turn rationalise higher within-firm markups; the static framework cannot separate this income channel from product-mix upgrading or brand-premium growth, and finite-mixture alternatives (Appendix B.1) deliver a much flatter income gradient than the y^{-1} form, so the within-firm Lerner attribution to the income channel should be read as an upper bound. *Second*, the reallocation component is *negative* (-12%): low-markup new entrants (Tesla, the New-Forces startups) gained share while high-markup incumbents lost share, the opposite sign from the superstar-firm reallocation documented for the U.S. (De Loecker et al., 2020; Autor et al., 2020). The combination of rising aggregate markups and falling concentration is unusual in the rising-markups literature, which more often documents concentration rising in parallel with markups in U.S. settings. *Third*, the static decomposition only captures the contemporaneous tail-truncation of the 2015–2024 phase-out — the change in the per-vehicle wedge between the schedule's 2015 and 2024 levels, evaluated on the otherwise-fixed 2024 economy. The dynamic counterfactual is a different question: what would 2024 have looked like if the regime had never existed at all, with products entering and exiting endogenously over the decade? (§5.5.) The static-incidence pattern concentrates in lower-income, lower-tier, budget segments — where the subsidy yuan directly landed. The dynamic incidence (the supply pipeline the regime called into existence) concentrates elsewhere along a *firm-origin* axis. Under a never-existed-subsidy forward simulation, the 2024 CF EV share lies between -21.7 pp (Scenario E in §5.5: empirical absolute-floor rule calibrated from 711 observed product exits, plus a 3-year grace period for post-2015 entrants) and -40.6 pp (full-BLP solver at $\theta = 0.50$, agent-integral). The aggregate dynamic-channel finding — subsidies were constitutive of the EV market's bootstrap, not a marginal price-wedge effect — is robust to the rule choice. The cross-firm distributional pattern, however, is rule-sensitive in a way that bears on the firm-origin narrative. Under the relative-threshold rule ($\theta = 0.50$), the dropout test penalises products with high observed profit, so high-margin nameplates (premium tiers and post-2015 New-Forces flagships) appear to "collapse" more than mass-market BYD products: New Forces retain 3–25% of baseline profit by tier, BYD retains 17–31%. Under the empirical rule with grace period, both BYD and New Forces retain essentially all of their lifecycle profit (BYD 98.1%, New Forces 99.8%, gap of -1.7 pp); the BYD-vs-New-Forces asymmetry collapses (and at corner cells reverses sign). The economically-grounded qualitative claim — BYD's competitive position predates the subsidy regime, and the post-2015 New Forces' market position emerged within it — is robust to both rules; the magnitude of any cross-firm asymmetry is not. The IV-based first stage in §D.7 ($F = 3.75$) cannot identify the causal direction of the static-to-dynamic propagation, and the cross-firm asymmetry magnitudes under the never-existed CF should be read as model-generated objects whose magnitude depends on the dropout rule.

Two layers of caveat bound interpretation. First, the within-firm Lerner rise is *attributed* to the income channel conditional on the maintained $\alpha \propto y^{-1}$ form, which Appendix B.1 shows overstates a data-driven alternative; the Consumer_Composition attribution is therefore an upper bound on the income channel's contribution to within-firm markup growth. Second, the firm-survival asymmetry above is identified within the structure of the never-existed forward simulation (§5.5); the cross-product profit ratios and their geographic implications are model-generated objects rather than directly observed counterfactuals, and they depend on the imposed Wright's-Law learning specification, the no-ICE-price-response simplification, and the post-hoc reweighting that maps national results onto each city's 2024 firm mix. A point-identified estimate of the within-firm Lerner mechanism — and a separately IV-identified estimate of the static-to-dynamic propagation channel — remain on the research agenda of a companion paper.

While our empirical setting is the Chinese passenger-vehicle market, the mechanisms we identify offer

broader lessons for technology-diffusion settings beyond China. The Chinese decade serves as a benchmark for *rapid, policy-supported transitions in high-growth economies*, and the architecture of our decomposition is portable to other emerging markets even when the quantitative magnitudes are not. Our finding that adoption is driven heavily by battery learning ($+27.3$ pp on ΔEV share) and an EV-specific unobserved trend ($+15.1$ pp), rather than by direct purchase subsidies (-2.7 pp statically), suggests that policymakers in other emerging markets contemplating similar transitions—India, Brazil, Indonesia, ASEAN—should prioritise supply-side product-pipeline development alongside consumer-side subsidies. Conversely, the contrast between our within-firm markup growth and the U.S. superstar-firm reallocation pattern (De Loecker et al., 2020; Autor et al., 2020) highlights how market structure mediates technology transitions differently across regulatory environments: in a high-growth setting with low entry barriers, an EV transition can raise markups while concentration falls. Section 6 discusses the boundary conditions that limit external validity.

The paper proceeds as follows. Section 2 introduces the registration panel, the BNEF battery cost series, and the LLM-extracted policy index. Section 3 sets out the demand and supply specifications and the identification argument. Sections 4–5 report first-stage estimates and the eight-block decomposition framework. Section 5 delivers the main findings: the within-firm-driven Lerner rise, the eight-block adoption decomposition with battery learning and EV-specific trend dominant, and the never-existed-subsidy forward simulation under both relative-threshold and empirically-identified dropout rules. Section 6 discusses the lifecycle framing and policy implications. All replication materials are documented in Appendix E.1.

1.1 Related literature

The closest paper methodologically is Grieco et al. (2024), who decompose the 1980–2018 evolution of the U.S. automobile market into supply-side and demand-side channels using a BLP demand–supply system. In their four-decade U.S. sample, prices rose but markups *fell*, and welfare gains came through product-quality improvement and falling marginal costs. In our ten-year Chinese sample, markups *rise* ($+0.138$ Lerner points at the agent-integral subsidy spec) and within the decomposition welfare gains are attributed to income growth—a divergence that the model's $\alpha_{ic} \propto y^{-1}$ mechanism links to the 73% growth in Chinese urban incomes, a first-order shift in the model-implied demand elasticity that has no counterpart in the slow-growth U.S. setting. *Why we pair BLP with an Olley–Pakes split, and GMY do not.* The Olley–Pakes within-firm/reallocation decomposition is the standard tool of the rising-markups literature (De Loecker et al., 2020; Autor et al., 2020; De Loecker and Warzynski, 2012) for distinguishing whether aggregate market power grows from within-firm margin expansion or from share reallocation toward dominant high-markup superstar firms. Because GMY's U.S. sample exhibits *falling* aggregate markups, the within-firm/reallocation question is not central to their setting and they do not perform an OP split. Our setting is the inverse: aggregate markups *rise* by $+0.138$ Lerner points (Shapley canonical, agent-integral subsidy), so an OP decomposition is necessary to determine whether the rise reflects within-firm growth or a superstar-firm reallocation. The decomposition (§5.3) reveals the opposite of the canonical U.S. pattern: aggregate concentration *falls* (HHI $1{,}374 \to 887$) while within-firm Lerner rises across the incumbent distribution. The pairing of a BLP Shapley with an OP split is, to our knowledge, novel, and is exactly the methodological complement that any rising-markups setting requires. Our eight-block structure extends GMY's four-block decomposition along three dimensions: (i) we separate the product-set and city-demand-environment channels into Entry_set and Consumer_Composition, which produce an opposite-sign attribution pattern across outcomes (battery learning A_{bat} and EV_trend jointly dominate the EV-share contribution at $+27.3$ and $+15.1$ pp, while Consumer_Composition absorbs the dominant positive contribution to the Lerner index, with

Entry_set contributing $+0.035$ Lerner and Consumer_Composition $+0.150$); (ii) we report a Brand_trajectory post-hoc *residual diagnostic* that projects ξ^{BLP} onto (Brand $\times$ Year) cell means, recovering descriptively where the within-brand annual drift lives (e.g., BYD (比亚迪) +2.69 net residual gain, FAW (一汽) -2.42, Ford (福特) -1.56). We do not claim this projection identifies a separate channel: a brand-level (Brand $\times$ Year) fixed-effect would be the proper specification but is numerically infeasible in our 10-year panel, so the projection is a description of residual structure rather than an identified primitive. Sub-firm-group brand-level trajectories survive in the EV_trend residual; (iii) we exhaustively project and quantile-remap what remains of ξ^{BLP} so that the residual EV_trend block contains only EV-specific unobserved time variation orthogonal to the brand-trajectory diagnostic and the (ev$\times$year) layer.

The paper also connects to the vehicle-choice literature that pioneered the BLP framework in automobile markets (Berry, 1994; Berry et al., 1995; Goldberg, 1995; Petrin, 2002; Train and Winston, 2007). Our estimated mean own-price elasticity of -6.23 sits at the upper end of the typical U.S. range of -3 to -6 (Berry et al., 1995; Petrin, 2002), consistent with the lower income levels of the Chinese urban population and the within-sample tier gradient (-3.64 in Tier 1 to -7.94 in the Rest tier). In the Chinese setting, Barwick et al. (2019) report elasticities of -2.40 to -4.96 for 2009–2014; our larger magnitudes reflect the extended income range and the EV segment's price-sensitive buyers.

The closest contemporary paper on China's EV subsidies is Hu et al. (2025), who estimate a dynamic demand–supply model of the BEV segment with forward-looking consumers and supply-side anticipation, and evaluate alternative subsidy schedules. The two papers answer fundamentally different questions about the same regime. Their counterfactual space is the *design* dimension — alternative schedules conditional on a subsidy regime existing — and they identify how schedule timing affects welfare under temporally-substituting consumer behaviour. Our counterfactual space is the *constitutive* dimension: a never-existed-subsidy counterfactual (§5.5) that asks whether the EV market would have reached self-sustaining scale at all without any subsidy regime over 2015–2022, and identifies the firm-survival exposure (BYD vs. post-2015 New Forces) inside the answer. In the lifecycle terms developed in §5.5, Hu et al. address the Phase 2 question (now that the market exists at scale, how should the residual schedule be designed?), while our paper diagnoses the Phase 1$\rightarrow$Phase 2 transition (was the historical regime constitutive of the market?). One paper is operational, the other diagnostic; together they cover the operational-vs-constitutive role of the same regime.

The paper builds on Hao and Hao (2026), which uses transaction-level micro-data from Xi'an to decompose BYD's market evolution using a seven-driver sequential structure. The present paper extends the toolkit from one firm in one city to the national market across 79 regions, replaces the single sequential ordering with the full 2^8 Shapley structure, and separates the product-set and city-demand channels that were confounded in the earlier composite Entry block.

Technology diffusion across income and geography. The paper sits within a long tradition on technology diffusion that emphasises heterogeneous adoption across consumer characteristics. The hybrid-corn S-curves of Griliches (1957) and the cross-country diffusion accounting of Comin and Hobijn (2010) document that diffusion is rarely uniform: profitability, information, and complementary inputs determine which adopters move first. The macro-development literature on agricultural technology in emerging economies—Foster and Rosenzweig (1995), Munshi (2004), Conley and Udry (2010), and the cross-country mobile-phone evidence of Suri (2011)—finds that adoption diffuses along learning networks defined by income, education, and geography. Our finding that the Chinese EV transition concentrates initially in lower-income,

lower-tier, budget segments—and propagates upward over time through the price/quality ladder—adds a contemporary, structural-decomposition complement to that literature: the heterogeneous diffusion of an emerging clean-technology durable, with the random-coefficient demand system identifying the income shape of incidence directly rather than inferring it from a reduced-form treatment effect.

Industrial policy, directed technical change, and network externalities. The paper also speaks to the industrial-policy and directed-technical-change literatures. Acemoglu (2002) and Acemoglu et al. (2012) formalise how policy can redirect innovation toward clean technologies, and Aghion et al. (2016) provide cross-country evidence that fuel-tax-equivalent policy effort accelerates clean-technology patenting. We do not estimate the dynamic-innovation channel directly—a task left to a companion paper—but document the equilibrium incidence of the contemporaneous Chinese subsidy regime, the empirical residual that any directed-innovation interpretation must reconcile with. The early concentration of the subsidy in budget segments aligns with the demand-coordination story of Murphy et al. (1989); the supply-side complementarities between charging infrastructure and EV adoption studied by Springel (2021) and Li et al. (2017) echo the network-externalities frameworks of Katz and Shapiro (1986) and Farrell and Saloner (1986). Our decomposition treats charging infrastructure as one channel within the broader EV_trend residual; §D.2.1 discusses why we cannot separate it cleanly from other unobserved EV-specific time variation under the available panel structure.

EV adoption, charging infrastructure, and clean-vehicle policy design. The empirical literature on EV adoption and the design of clean-vehicle policy is the closest substantive comparator. Beresteanu and Li (2011), Holland et al. (2016), Muehlegger and Rapson (2022), and Rapson and Muehlegger (2023) document the demand-side responsiveness of consumers to EV subsidies and the accompanying environmental externalities; Springel (2021), Li et al. (2017), and Bollinger and Gillingham (2019) estimate the supply–charging-network feedback that drives the dynamic propagation of subsidies. On subsidy pass-through specifically, Sallee (2011), Busse et al. (2013), Allcott and Wozny (2014), and Xing et al. (2021) provide cross-country evidence that consumer subsidies are partially absorbed by upstream firms via sticker-price adjustments. Our static decomposition uses the multi-product Bertrand–Nash FOC to allow exactly this strategic price response, with the Subsidy block measuring the residual price-wedge effect after firms re-optimise. Within environmental economics more broadly, Newell et al. (1999), Jaffe et al. (2002), and Popp (2002) establish that environmental policy stringency redirects innovation; our finding that Chinese EV markups rose through within-firm growth on incumbents—rather than through reallocation toward new high-markup entrants—adds an unusual data point on which firms capture the rents from a directed-innovation regime.

2 Data

We turn first to the construction of the estimation panel, describing the source data, sample restrictions, and the subsidy and policy variables that enter the demand system.

Indices used throughout. The unit of observation is a (model $\times$ fuel-type) product j in geographic market c in year t. Throughout the paper, j indexes products, c indexes geographic markets, t indexes years, and i

indexes simulated households inside the random-coefficient integral. Households are simulated rather than observed: they enter only via the within-market log-income distribution that drives the price coefficient.

Source data. The analysis combines four inputs:

1. A proprietary monthly registration panel covering new passenger-vehicle registrations in China's largest urban markets from 2014 through 2024. For each registration the panel reports model identifier, manufacturer, brand, fuel type, manufacturer's suggested retail price, and the technical characteristics that enter the demand system: vehicle size (length $\times$ width $\times$ height), engine power (kW), displacement (L), and electric range (km, where applicable).
2. The *China City Statistical Yearbook* 6.0, providing annual measures of GDP per capita, urbanisation rate, the share of the population with tertiary education, total population, and population density at the city-year level. The same source provides the household-count denominator that defines market size.
3. A hand-coded EV policy dataset built on the Chinese full-text law database PKULaw (北大法宝), containing $4{,}122$ central- and city-level documents related to electric-vehicle policy.
4. The official central New Energy Vehicle (NEV) subsidy schedule, combined with the local matching subsidy rules in force before the central government banned local matching at the start of 2019.

Sample, fuel-type categories, and aggregate scope. The estimation panel restricts to 2015–2024 and excludes imported vehicles. Every non-imported registration is assigned either to its own standalone market (52 prefecture-level cities) or to one of 27 *province-level residual* markets that aggregate all remaining prefectures within a province into a single composite. Because every province contributes exactly one residual market, the partition is exhaustive: the 79 geographic markets (52 standalone $+$ 27 residual) cover 100% of non-imported domestic new-vehicle registrations nationwide. The full panel contains $496{,}648$ product–market–year observations; the BLP estimation sample drops 57 singleton cells to construct the differentiation instruments, leaving $496{,}591$. Vehicles are classified into five fuel types: battery-electric (BEV), plug-in hybrid (PHEV), range-extended electric (REEV), conventional hybrid (HEV), and internal-combustion (ICEV). The first three are pooled into the EV indicator $\mathbb{1}[\text{EV}_j] = 1$ throughout the decomposition; HEV and ICEV are pooled into the ICE category.

Geographic-market tiers. A four-tier classification groups the 79 markets into Tier 1 (the four first-tier megacities: the two direct-administered municipalities Beijing and Shanghai plus the two Guangdong sub-provincial megacities Guangzhou and Shenzhen), New Tier 1 (15 large prefecture-level cities ranked as "new first-tier" by the YiCai Business Data city-tier index), Tier 2 (33 further standalone prefecture-level cities), and *Rest* (the 27 province-level residual markets). The complete city list and tier assignment are in Appendix A.4.1. Throughout the paper, "geographic market" or "market" refers generically to any of the 79 rows; "city" is reserved for the 52 standalone city-level markets.

Firm groups and product segments. Manufacturers are first mapped to a parent company via a hand-coded dictionary, then collapsed into six firm groups: BYD, Tesla, Traditional original-equipment manufacturer (Traditional OEM, the seven large state-owned enterprise [SOE] automotive groups including their foreign joint ventures), New Forces (post-2014 EV-pure-play startups), Independent Domestic (non-JV domestic OEMs predating the EV-subsidy era — note this includes both privately-held groups such as Geely and Great Wall and local/provincial SOEs such as Chery and JAC; grouped here by their independence from the

large foreign-JV state-conglomerate structure rather than by ownership form), and Other. Vehicles are also classified by body type (five categories: SUV, sedan, hatchback, multi-purpose vehicle [MPV], and crossover) and by price tier (three categories: Budget below ¥15 ten-thousand, Mid ¥15–30 ten-thousand, and Premium above ¥30 ten-thousand, computed at the (model, fuel-type) level using the quantity-weighted observed net price). Definitions and member lists for each grouping are in Appendices A.4.2 and A.4.3.

Sample-construction adjustments. Two deterministic cleaning steps are applied to the raw registration panel. First, 3,771 rows manufactured as range-extended electrics but filed as conventional ICE (because the national fuel-economy testing protocol allows extended-range manufacturers to report a blended fuel-consumption number) are reassigned to REEV. The reassignment rule is conservative: it triggers only on rows whose reported fuel economy lies below the mechanically-impossible threshold of 2 litres per 100 km. After this step, the REEV cohort matches the classification used in the Chinese auto industry press, capturing the Li Auto L-series, the AITO M-series, the Avatr 12, the Yangwang U8, and the Deepal S7. Second, 30 small Chinese ICE brands that each contribute fewer than 200 panel rows, each have zero presence in 2023–2024, and each ceased domestic operations during the sample (Qoros (观致), Zotye (众泰), Jinbei (金杯), Bisu (比速), Hawtai (华泰), Soueast (东南), and others) are pooled into a single `Other_exited` brand fixed effect. (Baojun (宝骏) is retained as a separate brand fixed effect throughout: though Baojun appears underrepresented in the registration panel for a portion of the sample due to partial absorption under the parent SGMW label, the brand remains commercially active in 2023–2024 and does not meet the pool-threshold criterion.) The pooled brands together account for 2,753 observations or 0.55% of the panel. Pooling reduces the lower-quintile dispersion of the estimated brand-effect distribution from 2.98 to 2.12 while leaving every active recent entrant — including small post-2022 brands such as Xiaomi, Yangwang, IM Motors, and Voyah — as separate fixed effects. After both cleaning steps, the panel contains 198 raw brand labels collapsed into 169 brand fixed effects.

Three patterns that motivate the decomposition. Three margins in the raw data motivate the structural decomposition. First, the EV price premium narrows sharply over the sample: the quantity-weighted average BEV+PHEV+REEV net price falls from ¥18.9 ten-thousand in 2015 to ¥17.8 ten-thousand in 2024, while the ICE average price rises from ¥15.1 to ¥20.5 ten-thousand. Second, electric range improves sharply: the quantity-weighted mean rises from 141 km in 2015 to 352 km in 2024, with a peak of 373 km in 2023. Third, the EV product set expands dramatically, from 52 nameplates in 2015 to 494 in 2024, while the ICE set contracts modestly. By the end of the sample, EV variety exceeds ICE variety. These three margins are precisely the channels that the structural decomposition later separates into battery-related improvement, non-battery attribute evolution, and product entry. Table 29 reports the year-by-year sample composition; Table 30 reports quantity-weighted product characteristics by year and fuel type.

2.1 Subsidy as a price pass-through wedge

The central NEV subsidy operates economically as a per-vehicle price reduction, and I treat it as such throughout the paper. For BEVs the central schedule is tiered by electric range and, in some years, by battery characteristics; PHEVs and REEVs receive a flat payment conditional on meeting the 50-km electric-range threshold. Each (model, fuel-type, year) cell is matched to its applicable central subsidy amount and, for 2015–2018, to the allowed local matching subsidy. The transaction price faced by the household is then the manufacturer's sticker price minus the per-vehicle subsidy that the household receives. Appendix A.2

reports the schedule in full and documents the mapping from administrative rules to per-vehicle subsidy amounts.

A long literature on commodity-tax incidence and on durable-goods subsidies has shown that an upstream subsidy is not generally absorbed one-for-one by the consumer: producers may capture part of the wedge by raising their sticker price, and the share that reaches the consumer (the pass-through rate) depends on the slope of the demand and supply curves at the relevant point of the market. Within the EV literature specifically, Sallee (2011), Busse et al. (2013), and Muehlegger and Rapson (2022) document partial subsidy pass-through in U.S. vehicle markets; Allcott and Wozny (2014), Xing et al. (2021), and Hu et al. (2025) extend this finding to China-specific and energy-specific settings. The present paper does not estimate the pass-through rate as a free parameter. Instead, the random-coefficient demand system identifies how much of an exogenous change in the deterministic transaction-price wedge $p_{jct}^{\text{net}}/p_{jct}^{\text{sticker}}$ would feed through to the equilibrium EV share, with the wedge treated as a deterministic input and the sticker price allowed to adjust endogenously through the multi-firm Bertrand–Nash FOC. The Subsidy block in the decomposition (Section 5) is exactly this counterfactual: it activates the 2015 wedge schedule on top of the 2024 product set and re-solves the equilibrium, so any within-market strategic price response of multi-product firms to the changed wedge is already embedded in the reported number. The block therefore measures the static pass-through margin of the subsidy regime, restricted to the deterministic price-wedge channel and conditional on the rest of the equilibrium structure being held fixed.

2.2 Policy text data and the policy-strength index

The license-plate-advantage indicator $\texttt{ev\_license}_{ct}$ equals one for any city-year in which an EV could be registered without entering the local restrictive license-plate lottery, and zero otherwise. Eight cities operate restrictive plate quotas during the sample period, and in each of them "green plates" for new energy vehicles are not subject to the lottery; the indicator captures this discontinuous policy advantage at a clean, observable margin.

The policy-strength index $\texttt{ev\_policy\_primary}_{ct}$ is a continuous text-derived measure of how aggressive each city's broader EV-policy environment is in a given year. Its construction involves text classification of the $4{,}122$ PKULaw documents, named-entity extraction to identify the issuing city, year, and policy instrument, and aggregation into a single index normalised by population-weighted standardisation across the city-year panel. Because the construction is non-trivial and is itself a methodological contribution of the paper, the full pipeline is documented in Appendix A.3. The index is then de-meaned at the population-weighted national mean before entering the demand system, so the reported coefficient on $\texttt{ev\_policy\_primary}$ is the effect of being one cross-sectional unit above the national average. The income-interacted term $\texttt{ev\_policy\_x\_lowinc}_{ct}$ multiplies the de-meaned index by the city-year share of urban households below the national median income, allowing the policy effect to differ for higher- and lower-income city populations. Together these two policy variables identify the income-dimensional heterogeneity of policy exposure that drives the main subsidy-incidence finding in Section 5.4.

3 Model and identification

The structural model combines a random-coefficient demand system with multi-market Bertrand–Nash pricing, following the tradition of Berry (1994), Berry et al. (1995), Goldberg (1995), and Nevo (2001). The

empirical framework is a random-coefficient logit demand system in the tradition of Berry et al. (1995) and Petrin (2002), paired with a static multi-market Bertrand–Nash pricing model on the supply side. The demand side is designed to recover heterogeneous price sensitivity across cities with different income distributions and to map the observed subsidy schedule into utility through net transaction prices. The supply side is designed to translate the estimated demand system into product-level marginal costs, equilibrium markups, and counterfactual prices for each coalition in the decomposition. The economic intuition that drives identification appears before the estimator mechanics in each subsection.

3.1 Demand

Utility. Indirect utility for household i choosing product j in market c in year t is

$$u_{ijct} \;=\; \delta_{jct} \;+\; \alpha_{ic}\,\log p^{\text{net}}_{jct} \;+\; \epsilon_{ijct}, \tag{1}$$

where $\delta_{jct} = x'_{jct}\boldsymbol{\beta} + \xi_{jct}$ is the mean utility, x_{jct} collects observed product characteristics, city-level demand shifters, policy variables, EV-year interactions, and brand / body / fuel fixed effects, ξ_{jct} is a residual product–market demand shock observed by firms and consumers but not by the econometrician, and ϵ_{ijct} is an i.i.d. Type-I extreme value taste shock. Outside-good utility is normalised to $u_{i0ct} = \epsilon_{i0ct}$. Price enters logarithmically rather than in levels because Chinese passenger-vehicle prices span more than an order of magnitude and a proportional price change has a more comparable utility effect across budget and premium products.

Heterogeneity through the price coefficient. Heterogeneity is restricted to a single random coefficient on $\log p^{\text{net}}_{jct}$:

$$\alpha_{ic} \;=\; \pi_p \left(\frac{y_{ic}}{\bar{y}}\right)^{-1}, \qquad \pi_p < 0, \tag{2}$$

where y_{ic} is the simulated income draw of household i in city c and $\bar{y}$ is the panel-average urban income. The functional form $\alpha_{ic} \propto y_{ic}^{-1}$ gives higher-income households smaller price sensitivity in absolute value, which is the empirical regularity in every BLP-style auto study with income data (Berry et al., 2004; Train and Winston, 2007; Barwick et al., 2019). The single-random-coefficient restriction is deliberate: in preliminary specifications, additional random coefficients on size, power, and EV range were weakly identified and produced little improvement in the GMM objective once the income–price interaction was included, because urban Chinese households' willingness to pay for both price and quality is already captured by income heterogeneity.

Income-draw distribution. Conditional on a city-year, y_{ic} is drawn from a log-normal distribution

$$\log y_{ic} \;\sim\; \mathcal{N}\big(\log y_{ct},\; \sigma^2_{\log y}\big), \qquad \sigma_{\log y} = 0.307, \tag{3}$$

where y_{ct} is the city-level urban per-capita disposable income (Urban Household Per-capita Disposable Income (城镇居民人均可支配收入)) from the *China City Statistical Yearbook*, and $\sigma_{\log y} = 0.307$ is the within-city log-income dispersion calibrated from Khor and Pencavel (2006)'s national urban estimate. The integral over $F(y \mid c, t)$ is approximated with 25 Halton draws per market. The mean y_{ct} ranges from ¥21,700 (lowest province-residual market in 2015) to ¥93,100 (highest standalone city in 2024), and the 73%

growth in the cross-market average (from ¥33,400 in 2015 to ¥57,600 in 2024) is the primitive to which the Consumer_Composition block's Lerner contribution is attributed under the model's $\alpha_{ic} \propto y^{-1}$ restriction.

Market shares. The implied market share of product j in (c,t) is the standard random-coefficient logit integral:

$$s_{jct} \;=\; \int \frac{\exp\!\big(\delta_{jct} + \alpha_{ic}\log p^{\text{net}}_{jct}\big)}{1 + \sum_{k\in\mathcal{J}_{ct}} \exp\!\big(\delta_{kct} + \alpha_{ic}\log p^{\text{net}}_{kct}\big)}\; \mathrm{d}F(y \mid c,t), \tag{4}$$

where $\mathcal{J}_{ct}$ is the set of products available in (c,t). This share equation, evaluated at the 25 Halton income draws, is the only nonlinearity in the model.

3.2 Supply

Multi-market Bertrand–Nash. A multi-product firm f owning products in $\mathcal{J}_f$ chooses each product's transaction price to maximise the sum of gross margins across markets, taking other firms' prices as given:

$$\max_{\{p_{jct}\}_{j\in\mathcal{J}_f}} \;\; \sum_{c} M_c \sum_{j\in\mathcal{J}_f} \big(p_{jct} - mc_{jt}\big)\, s_{jct}(\mathbf{p}_{ct}), \tag{5}$$

where M_c is the household-equivalent market size in city c, $\mathbf{p}_{ct}$ is the vector of all products' prices in (c,t), and mc_{jt} is the (national) marginal cost of producing product j in year t. Costs are taken to be common across cities for the same year because vehicles are produced at a small number of national plants and shipped country-wide, so transport costs are economically negligible relative to manufacturing costs.

First-order condition and recovered marginal costs. The first-order condition for the price vector at the multi-market equilibrium is

$$\mathbf{p}_t \;-\; \mathbf{mc}_t \;=\; -\Big(\sum_{c} M_c\, \boldsymbol{\Omega}_{ct}\Big)^{-1} \sum_{c} M_c\, \mathbf{s}_{ct}(\mathbf{p}_t), \tag{6}$$

where $\boldsymbol{\Omega}_{ct}$ is the city-c Jacobian of own- and cross-product shares with respect to price, masked to the firm-ownership pattern. Inverting the FOC at observed prices $\hat{\mathbf{p}}_t$ delivers product-level marginal costs $\hat{\mathbf{mc}}_t$ without imposing any further functional form on the cost side. The Lerner index for product j in (c,t) is $L_{jct} = (p_{jct} - mc_{jt})/p_{jct}$ (cf. De Loecker and Warzynski, 2012), and the industry mean Lerner reported throughout the paper is the quantity-weighted average of L_{jct} across all products and markets.

3.3 Identification

Three sources of variation. Identification of the demand parameters $(\boldsymbol{\beta}, \pi_p)$ relies on three sources of exogenous variation:

1. *Within-segment relative price variation*, identified from BLP-style differentiation instruments: a product's relative markup is identified by how isolated it is in characteristic space within its body–fuel–year segment. Two products that are characteristic-near substitutes compete head-to-head and earn small markups; an isolated product earns more.
2. *Cross-market interaction between price and the income distribution*, identified from the random-coefficient structure: cities with higher mean income display systematically smaller quantity responses to price

variation, and this heterogeneity traces out the dispersion of α_{ic} around its mean π_p. Concretely, the parameter π_p is identified from the interaction of (a) the within-market price wedge (shifted by the differentiation instruments) and (b) the cross-market income distribution (drawn from the Yearbook 6.0 panel). The exclusion restriction is that within-segment product differentiation shifts relative prices but does not directly affect the income–price interaction term conditional on the city fixed effects and macro controls already in δ.

3. *Cross-market policy strength interacted with the low-income share*, identified from independent variation in the city-year PKULaw policy index and the Yearbook 6.0 share of below-median-income urban households. The interaction term identifies the income-dimensional heterogeneity of policy exposure that drives the subsidy-incidence finding in Section 5.4.

Differentiation instruments. For each continuous characteristic $x \in \{\log(\text{size}),\ \text{Displacement},\ \text{electric efficiency},\ \text{fuel efficie}$ define the leave-one-out segment mean

$$\bar{x}_{-j,s(j)} \;=\; \frac{1}{|\mathcal{J}_{s(j)}| - 1} \sum_{k \in \mathcal{J}_{s(j)},\, k \neq j} x_{kct}, \tag{7}$$

where the segment $s(j)$ is body-type $\times$ fuel-type $\times$ year. The corresponding differentiation instrument for product j is

$$Z^{x,\text{diff}}_{jct} \;=\; \left| x_{jct} - \bar{x}_{-j,s(j)} \right|, \tag{8}$$

adapting Berry et al. (1995) to Gandhi and Houde (2019)'s within-segment formulation. The instrument shifts a product's relative price through the identity of its neighbours rather than through its own characteristics, satisfying the exclusion restriction conditional on the segment fixed effects already in δ_{jct}.

EV-vintage instrument. For EV products, the EV-vintage instrument

$$Z^{\text{ev-vintage}}_{jct} \;=\; \mathbb{1}[\text{EV}_j] \cdot \left(t - t^{\min}_j\right) \tag{9}$$

captures how long the EV nameplate has been in the market, with $t^{\min}_j$ the first sample year in which j is observed. Older EV nameplates have had more time to descend the production-side learning curve and to build brand recognition, generating systematic relative-price variation conditional on observed characteristics. For ICE products the instrument is identically zero.

Policy control-function residual. The policy-strength variable $\texttt{ev\_policy\_primary}_{ct}$ is plausibly correlated with unobserved city-level EV promotion that might also enter demand. Following Petrin (2002), a first-stage regression of `ev_policy_primary` on the differentiation instruments and the EV-vintage instrument is run, and the residual $\hat{\eta}^{\text{pol}}_{ct}$ is included as an additional regressor in the demand equation. Including the residual purges the estimated policy coefficient of any remaining endogeneity through the unobserved-promotion channel.

Estimation. Given the moment conditions implied by the instruments above, the demand parameters are estimated by two-step generalised method of moments (GMM) in `PyBLP`. The inner loop is the standard Berry et al. (1995) contraction mapping that inverts the share equation (4) for δ_{jct} at tolerance 10^{-14}; the outer loop minimises the GMM criterion over π_p using BFGS. The first-step weighting matrix is the identity;

the second-step weighting matrix is the optimal heteroskedasticity-robust matrix formed from the first-step moments. Section 4 reports the resulting parameter estimates and validates the implied own-price elasticities and Lerner indices against the Chinese and U.S. BLP literature; Appendix A.1 reproduces the full coefficient vector and the first-stage instrument diagnostics.

Income endogeneity and robustness. A potential concern is that city-level income may be correlated with unobserved demand shifters—for example, cities with faster GDP growth may also have stronger pro-EV sentiment or more aggressive local industrial policy, so that the income–price interaction π_p captures omitted demand heterogeneity rather than true price sensitivity. Three features of the specification mitigate this concern. First, the BLP estimation includes year-specific EV fixed effects ($\texttt{ev\_x\_yr}_{2016}$ through $\texttt{ev\_x\_yr}_{2024}$) that absorb any aggregate time-varying EV taste, so the income interaction is identified from *cross-city within-year* variation in the income distribution, not from the common trend. Second, the specification includes city-level controls for urbanisation, education, population density, and the PKULaw policy index, each of which would absorb a direct channel from local economic structure to EV demand. Third, as a robustness check, Appendix B.2.1 re-estimates the model under a homogeneous price coefficient (imposing $\pi_p = 0$, so that α_{ic} is the same for all agents within and across markets) and re-runs the full Shapley decomposition. The Consumer_Composition block's ΔLerner contribution falls sharply under the homogeneous specification, consistent with the markup-growth attribution to the wealth effect being driven by the estimated income interaction rather than by other features of the model.

4 Estimation and first-stage results

We estimate the model by two-step GMM in PyBLP (Conlon and Gortmaker, 2020). This section reports the joint demand and supply first-stage estimates and benchmarks the implied own-price elasticities and Lerner indices against the prior BLP literature on Chinese and U.S. automobile demand. The objective is not only a demand system that fits observed market shares but a demand system whose recovered marginal costs and counterfactual markups remain economically sensible at every coalition the decomposition will later visit. The diagnostic evidence in Sections 4.1 and 4.2 indicates that this is the case.

4.1 Joint demand and supply first-stage estimates

Table 1 reports the joint demand-side GMM and supply-side OLS estimates. The demand column is the random-coefficient logit of Section 3.1; the supply column is an OLS regression of the FOC-recovered $\log(\hat{mc}_{jt})$ on shared product characteristics, the EV energy-consumption rate, and the Bloomberg New Energy Finance (BNEF) battery pack cost in the year. Five rows of the demand column carry most of the economic content, and three rows of the supply column let the reader cross-check the sign plausibility of the estimated cost technology against the demand side.

Table 1: First-stage demand and supply estimates

	Demand δ (BLP logit)	**Supply** $\log(\hat{mc})$ (OLS)
Price / cost scale		
π_p (price × income^{-1}, random coef.) [demand]	-5.377^{***} (0.168) [0.587]	—
Product characteristics (shared)		
log(vehicle size)	$+3.677^{***}$ (0.103)	$+1.416^{***}$ (0.047)
log(power)	$+1.411^{***}$ (0.040)	—
Displacement	$+1.135^{***}$ (0.035)	$+0.671^{***}$ (0.039)
Fuel efficiency (L/100km)	-0.001 (0.011)	-0.012 (0.008)
EV / battery channel		
log(EV range) [demand]	$+0.661^{***}$ (0.048)	—
EV × energy consumption (kWh/100 km) [supply]	—	$+0.006^{***}$ (0.001)
log(BNEF battery pack cost) [supply]	—	$+0.374^{***}$ (0.021)
Policy (demand-side only)		
License-plate advantage indicator	$+0.167^{***}$ (0.017)	—
Policy-strength (de-meaned) [see note]	-0.005^{***} (0.001)	—
Policy-strength × low-income share	$+0.009^{***}$ (0.001)	—
EV × Year fixed effects (2015 omitted as reference)		
EV × 2016	-0.756^{***} (0.111)	—
EV × 2017	-0.816^{***} (0.113)	—
EV × 2018	-1.356^{***} (0.129)	—
EV × 2019	-2.589^{***} (0.151)	—
EV × 2020	-3.919^{***} (0.184)	—
EV × 2021	-2.898^{***} (0.198)	—
EV × 2022	-0.579^{***} (0.204)	—
EV × 2023	-1.966^{***} (0.237)	—
EV × 2024	-2.898^{***} (0.271)	—
Firm-group × EV interactions (Other × EV omitted as reference)		
BYD × EV	$+1.238^{***}$ (0.066)	—
Foreign/JV × EV	-2.874^{***} (0.060)	—
Trad. OEM × EV	-0.773^{***} (0.052)	—
New Forces × EV	-1.713^{***} (0.191)	—
Private National × EV	-1.451^{***} (0.056)	—
Brand / BodyType / FuelType FE	Yes	Yes
City FE	Yes	—
N / R^2	476,723 / —	8,764 / 0.606

Notes. Default PyBLP SE in parentheses. For π_p, the city-pair bootstrap cluster-robust SE is reported in square brackets ($B = 50$ resamples; see §E.2). $\dagger\ p < 0.10$; $\dagger^{*}$< 0.05; $\dagger^{*}\dagger$< 0.01 refer to the default-SE t-statistics; under the cluster-robust SE on π_p, the implied t-statistic is -9.2 (vs. -32.0 default), unchanged in significance level. Demand: two-step GMM in `PyBLP`; π_p is the only random coefficient with 25 Halton income draws and $\sigma_{\log y} = 0.307$. Supply: OLS on the Bertrand–Nash-recovered $\log(\hat{mc})$ for 8,764 product-year cells, specification (6) of `output/mc_learning_results.csv` ("BNEF Battery Cost (clean, Table 1)"). *Sign-plausibility cross-check.* Shared rows (size, displacement) load positively in both columns: bigger cars command both higher mean utility and higher marginal cost. The battery-cost channel enters supply through the BNEF per-kWh cost with a $+0.374$ loading, the empirical counterpart of Wright's-law battery learning. The energy-consumption variable is measured in kWh per 100 km (Chinese industry convention, matching the ICE `FuelEfficiency` variable in L/100 km), so a *higher* value means a *less* efficient EV; the positive supply-side loading is consistent with less-efficient EVs requiring larger battery packs for the same observed range. log(EV range) is kept as a pure demand-side quality shifter; including it on the supply side alongside the BNEF cost is collinear with battery capacity and produces a sign sensitive to the choice of other battery controls (the five alternative specifications are in `output/mc_learning_results.csv`).

Policy-strength sign. The -0.023 mean loading on the de-meaned policy-strength index is *not* a perverse "policy hurts EVs" finding. City fixed effects already absorb the strong positive cross-sectional correlation (high-policy cities = high-EV cities, e.g., Beijing and Shanghai), leaving only within-city residual variation, for which policy strength tends to rise in cities where short-run EV share is declining. One possible explanation is reverse causality (policymakers respond to slumps); other explanations including measurement error in the policy-strength index and omitted city-specific time-varying confounders are equally plausible without further analysis. The economically interpretable part of the policy bundle is the city-level income interaction in the next row ($+0.022$). Note that `ev_policy_x_lowinc` is the city-year share of urban households below the national median income, interacted with the de-meaned policy index; so the combined effect is city-level and applies uniformly to all buyers within a city. The interaction exactly offsets the main effect at a low-income share of 1.0 (all below-median-income city), and the combined policy effect ranges from about -0.023 (no below-median share) to near zero (all below-median share). The economic content is therefore that the policy index registers a demand *penalty* that is attenuated, but not reversed, in cities with a larger low-income share. This attenuation gradient, not the mean level, drives the heterogeneity of the subsidy-incidence result in Section 5.4.

EV × Year fixed effects. The block reports the full 2016–2024 series (with 2015 as the omitted reference). The pattern is a V-shape with the trough at -3.92 in 2020, coinciding with the June-2019 central-subsidy retrenchment and the COVID-19 demand collapse, and a partial recovery in 2021–2022 before drifting back down in 2023–2024. These coefficients absorb any EV-specific residual taste or non-price policy movement not captured by the observed battery, price, and policy controls — including charging-infrastructure provision, trade-in programs, and post-2022 local subsidies outside the three coded cities (see Appendix A.5). Because the central subsidy is already explicitly modelled through the net-price variable, the residual drift after 2023 should *not* be attributed to the subsidy phase-out per se but to other unmodelled EV-specific policy and taste components. The decomposition in Section 5 isolates this residual as the `EV_time` block. The 28 unreported demand parameters and the five supply specifications are in Appendix A.1.

Mean price coefficient. The estimated mean random-coefficient price loading is $\hat{\pi}_p = -5.38$ with a default PyBLP standard error of 0.168. The default SE assumes independence at the observation level; the identifying variation is at the city level, which Section E.2 below estimates via a $B = 50$ city-pair bootstrap. The cluster-robust SE is 0.587 (cluster-bootstrap mean -7.08, 95% percentile CI $[-7.99, -6.07]$), about $3.0\times$ the default; the implied t-statistic on $\hat{\pi}_p = -5.38$ is -9.2 rather than the default -32.0, still strongly distinguishable from zero. We use the cluster-robust SE for all $\hat{\pi}_p$-derived inference statements going forward. The estimate implies a quantity-weighted average own-price elasticity of -6.23, computed as $\hat{\pi}_p \cdot \mathbb{E}[(y_{ic}/\bar{y})^{-1} \cdot (1 - s_{jct})]$ across the $496{,}591$-row product–market–year panel. (Re-anchoring at the bootstrap mean $\hat{\pi}_p^{\text{boot}} = -7.08$ would yield an average elasticity of approximately -7.71; we report the canonical -6.23 throughout and treat the bootstrap-re-anchored numbers as a sensitivity benchmark.) For comparison, Barwick et al. (2019), estimating a BLP demand system on the 2009–2014 Chinese passenger-vehicle market with the same $\alpha_i = \pi_p \cdot y_i^{-1}$ random-coefficient structure, report own-price elasticities ranging from -2.40 to -4.96 with a median of -3.74; Petrin (2002), on the 1980s–1990s U.S. minivan segment, and Berry et al. (1995), on the 1971–1990 U.S. passenger-car panel, both report average own-price elasticities in the -3 to -6 range. Our estimate sits at the upper end of the U.S. benchmark range and somewhat above the Chinese benchmark. Within the imposed functional form $\alpha_{ic} \propto y_{ic}^{-1}$, the higher elasticity vs U.S. benchmarks tracks the lower Chinese income level. Within our own sample, the elasticity runs from -3.64 on average in Tier-1 cities (the richest) to -7.94 in the Rest tier (the poorest), a factor-of-two gradient that is the mechanical consequence of the imposed parameterisation. We emphasise that this gradient is built into the model rather than recovered as an empirical fact: the Consumer_Composition Lerner attribution that the decomposition later reports as the dominant share of the ΔLerner total is therefore a definitional consequence of this parameterisation, and Appendix B.1 shows it shrinks substantially under flatter finite-mixture alternatives. Figure 1 visualises this heterogeneity.

Single-RC identification: defense and precedent. The main specification uses a single random coefficient (π_p on price $\times$ income_inv), not the 9-parameter RC vector of Barwick et al. (2019) or the 4-parameter vector of Berry et al. (1995). This choice reflects a binding identification constraint in our panel: expanded RC specifications (2-RC, 3-RC, 4-RC) exhibit clear weak-identification signatures (sign-flips across initial values, standard errors on the order of 10^{17}, near-singular Jacobians) once the 169 brand fixed effects and 52-city $\times$ 10-year structure are absorbed, leaving insufficient within-segment variation to pin down attribute-level heterogeneity (Conlon and Gortmaker, 2020, Section 6). Appendix B.2.2 documents the diagnostic in detail. The price–income interaction is the one RC channel structurally central to our decomposition's findings: the Lerner attribution to Consumer_Composition and the heterogeneity result both work through heterogeneity in α_{ic}. The parsimonious $\pi_p \cdot y^{-1}$ form is the standard starting point for panels without micro-moments (Nevo, 2001; Petrin, 2002; Conlon and Gortmaker, 2020).

Battery-related product quality. The coefficient on $\log(\text{EV range})$ is $+0.661$ ($SE = 0.048$), so doubling rated electric range raises EV mean utility by approximately 0.46. This estimate is central for the decomposition because it provides a separate demand-side quality channel — a consumer-preference response to the technology improvement — through which battery-driven range gains can raise EV demand even when sticker prices do not fall one-for-one with battery cost declines. The corresponding supply-side margin is the $+0.374$ loading on $\log(\text{BNEF battery pack cost})$: a 1% decrease in BNEF battery cost translates into a 0.374% decrease in EV marginal cost, the static counterpart of Wright's-law learning (Thompson, 2012).

License-plate advantage and policy heterogeneity. The license-plate-advantage indicator enters positively and precisely at $+0.283$ ($SE = 0.021$). At the sample share-weighted mean, this implies a roughly 0.4 pp share advantage for an EV in cities where the dummy switches on, holding other characteristics fixed—meaningful but smaller than e.g. the battery-range or income-tier gradients. Regulatory access matters at the level of utility, but the decomposition later shows that the *change* in this margin from 2015 to 2024 is small in aggregate, consistent with the set of cities with binding plate restrictions changing little over the sample. The interaction between the policy-strength index and the low-income population share is $+0.022$ ($SE = 0.001$), which almost exactly offsets the -0.023 ($SE = 0.001$) coefficient on the de-meaned policy primary. The two coefficients are estimated precisely (statistical significance) but their economic-magnitude content is modest: a one-standard-deviation shift in the policy-strength index moves δ_{jc} by roughly 0.04 utility units, equivalent to a 0.6 pp change in product-level inclusive-value share at the mean. We caution against equating the *** stars with economic importance: these coefficients are pinned to high precision because the panel has 496,591 observations, but the implied economic effects are modest. Finally, classical measurement error from the LLM-extracted policy index would attenuate $\hat{\beta}_{\text{policy}}$ toward zero; we cannot distinguish a substantive small-effect interpretation from this measurement-error story without an independent benchmark. The two coefficients together imply that policy exposure is roughly neutral at the median income but favourable for buyers below the median: this is the cross-sectional shadow of the paper's main subsidy-incidence result in Section 5.4.

4.2 Implied supply-side objects: elasticities, markups, and Lerner indices

Given the demand estimates, the supply side is recovered from the multi-market Bertrand–Nash first-order condition (6), which delivers product-level marginal costs without imposing additional functional-form assumptions. **Table 2 reports the resulting distributions of own-price elasticities, level markups, and Lerner indices, broken down by fuel type.** Two patterns are notable.

Table 2: Implied elasticity, markup, and Lerner distributions by fuel type

	Own-price elasticity η_{jj}				**Markup** $p - mc$ (RMB ten-thousand)				**Lerner** $(p - mc)/p$			
Fuel	mean	p25	p50	p75	mean	p25	p50	p75	mean	p25	p50	p75
All	-6.23	-7.63	-6.05	-4.68	5.68	2.95	4.47	6.99	0.339	0.273	0.333	0.394
BEV	-4.84	-5.93	-4.75	-3.50	5.45	2.41	4.36	7.69	0.388	0.325	0.386	0.471
PHEV	-4.86	-5.81	-4.93	-3.76	5.70	3.44	4.69	6.78	0.386	0.337	0.382	0.427
REEV	-4.47	-5.35	-4.47	-3.33	9.86	5.89	9.70	12.88	0.411	0.360	0.402	0.500
ICEV	-6.48	-7.89	-6.37	-4.95	5.60	2.91	4.37	6.76	0.331	0.267	0.324	0.380

Notes. Quantity-weighted distributions across the 496,591-row product–market–year panel. Own-price elasticity is computed from $\eta_{jct} = \hat{\pi}_p \cdot (y_{ct}/\bar{y})^{-1} \cdot (1 - s_{jct})$ at the mean income draw. Markups are recovered product by product from the multi-market Bertrand–Nash FOC at observed prices. The Lerner index is $L_{jct} = (p_{jct} - mc_{jt})/p_{jct}$. The "All" row pools all 496,591 observations.

First, the model generates own-price elasticities of economically sensible magnitude for every fuel type. The pooled mean elasticity is -6.23, with quartiles between -7.63 and -4.68. The dispersion is tighter than in studies that include richer household-level taste heterogeneity, consistent with the deliberately parsimonious random-coefficient structure used here. There is no evidence of a long left tail of unrealistically-high-elasticity products that would arise if the model were misspecified at the marginal-product boundary.

Second, the implied markup distribution differs noticeably across fuel types. The mean $p - mc$ is largest for range-extended vehicles at ¥9.86 ten-thousand, with plug-in hybrids (¥5.70), battery-electric vehicles (¥5.45), and ICE vehicles (¥5.60) clustering in a comparable range. The REEV gap reflects the differentiated family-SUV segment with relatively weak direct competition; the higher Lerner of plug-in and BEV segments compared with ICE reflects their lower price-elasticity in the income-stratified BLP rather than wider absolute markup gaps. The Lerner indices follow the same ordering: REEV 0.411, BEV 0.388, PHEV 0.386, ICE 0.331. Across all fuel types the pooled mean Lerner is 0.339, somewhat above the 0.18–0.25 range that Barwick et al. (2019) report for the same Chinese passenger-vehicle market in 2009–2014, consistent with the within-decade markup growth documented by the Olley–Pakes decomposition in §5.3.

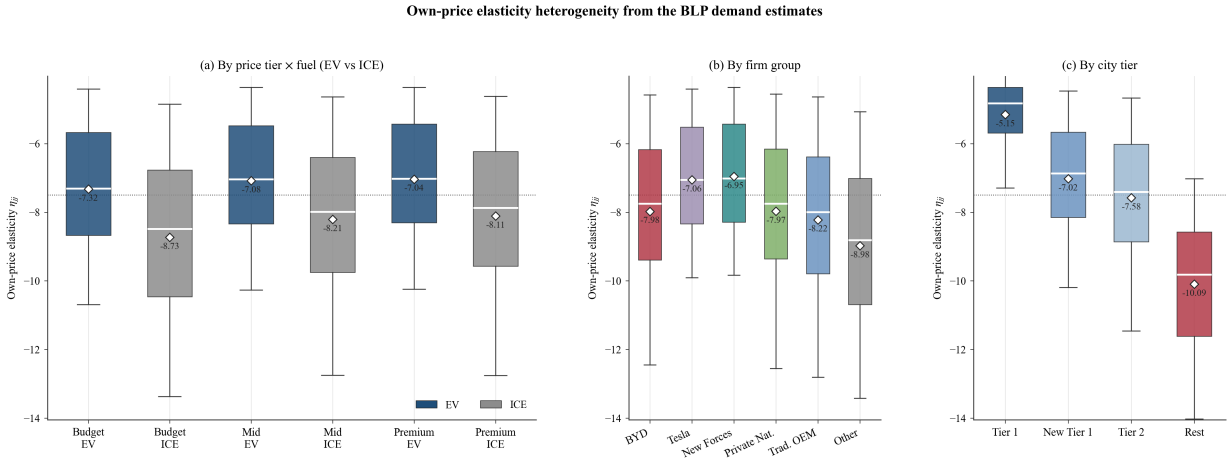

Figure 1: Own-price elasticity heterogeneity implied by the BLP demand estimates, evaluated at each row's city-year income draw via $\eta_{jct} \approx \hat{\pi}_p (y_{ct}/\bar{y})^{-1}(1 - s_{jct})$. Boxes show the interquartile range with the median as a white line; whiskers extend to the 5th/95th percentiles, and the white diamond marks the quantity-weighted mean. Panel (a) breaks the elasticity down by price tier (Budget <¥15k, Mid ¥15–30k, Premium >¥30k) and fuel (EV vs ICE): EVs are systematically less elastic than ICE within every price tier, consistent with EV buyers concentrating in higher-income sub-populations of the same tier. Panel (b) breaks it down by firm group: Tesla and the New Forces have the smallest elasticities (they sell primarily in Tier-1 and New-Tier-1 cities), while the residual "Other" group — dominated by exited lower-tier Chinese ICE brands — has the largest elasticity magnitude. Panel (c) is the cleanest visualisation of the wealth-effect channel: mean elasticity rises in magnitude monotonically from -3.64 in Tier 1 to -7.94 in the Rest tier, almost doubling. The city-tier gradient is the same BLP primitive that the Consumer_Composition block's outsized ΔLerner contribution reported in Section 5.2 is attributed to. A pooled-by-fuel distribution plot is in Appendix B.5.

These diagnostics matter because the decomposition is equilibrium-based: if the estimated demand system generated implausible elasticities, the recovered marginal costs and counterfactual markups would be unreliable, and the decomposition could mechanically attribute too much of the transition to supply-side blocks. The diagnostic evidence here suggests that this is not what is driving the results. The large entry and battery-learning contributions reported in Section 5 arise because those margins genuinely move the estimated demand system by economically meaningful amounts, not because the pricing model is numerically unstable at the corners.

5 Decomposition: method and main results

This section presents both the equilibrium-decomposition method and the main numerical results. Block-by-block discussion, the EV_trend anatomy diagnostic, the charging-channel diagnostic, and the year-by-year welfare path are deferred to Appendix D; the formal block partition, the activation mechanism, and the consumer-surplus winsorisation correction are in Appendix C. This main-text section keeps only the four main tables (aggregate Shapley, Olley–Pakes within-firm vs reallocation, subsidy heterogeneity, dynamic counterfactual) and the minimum prose needed to read each.

5.1 The eight-block partition (overview)

The decomposition starts from the BLP mean-utility identity. For each (product j, city c, year t):

$$\delta_{jct} \;=\; X'_{jct}\beta \;+\; \xi_{jct}, \tag{10}$$

where X_{jct} collects all observed product- and market-level covariates (battery attributes, non-battery attributes, demand-environment shifters, brand/fueltype/body fixed effects, EV×Year fixed effects, and city-wealth primitives), β is the BLP-estimated coefficient vector (year-demeaned at the brand-FE level — see code preprocessing), and ξ_{jct} is the unobserved demand shock identified from the GMM moment conditions $\mathbb{E}[\xi \cdot Z] = 0$. The Subsidy wedge enters the agent integral separately as $\alpha_{ic}\,\log(1 - s^{\text{sub}}_{jct})$ at the share-computation step (see §5.4's implementation footnote), so it does not appear in the $X'\beta + \xi$ identity above.

We partition δ_{jct} into eight blocks that the Shapley value averages over $2^8 = 256$ counterfactual coalitions. Six of the blocks correspond to disjoint pieces of $X'_{jct}\beta$ that the Shapley toggles between their 2015 and 2024 values:

- *A_{nonbat}* — non-battery product attributes (size, power, displacement, fuel efficiency, body type).
- *A_{bat}* — battery and EV-specific characteristics (range, electric efficiency, BNEF battery cost on the supply side).
- *D_{env}* — merged macro and non-subsidy policy environment (oil price, urbanisation, education, population density, license-plate preference, policy-strength index).
- *Entry_set* — product-set composition (which (Model, FuelType) pairs are in the choice set).
- *Consumer_Composition* — per-city demand environment (population M_c, mean income, within-city income distribution that drives α_{ic}).
- *Subsidy* — agent-integral subsidy wedge $\alpha_{ic}\,\log(1 - s^{\text{sub}}_{jct})$, toggled per row between 2015 and 2024 subsidy schedules. Implemented at the share-computation step rather than as a δ-shift (see §5.4's implementation footnote).

The remaining two blocks decompose the BLP residual ξ via post-hoc projections onto orthogonal subspaces. Letting $P_{e\times Y}$ denote the Q-weighted projection onto (ev × Year) cell means and $P_{B\times Y}$ the Q-weighted projection onto (Brand × Year) cell means *after* the (ev × Year) projection has been removed:

$$\xi_{jct} = \underbrace{P_{e\times Y}\,\xi}_{\text{(ev×Y) drift}} + \underbrace{P_{B\times Y}(I - P_{e\times Y})\,\xi}_{\text{(B×Y) drift after (e×Y)}} + \underbrace{(I - P_{B\times Y})(I - P_{e\times Y})\,\xi}_{\text{within-cell rank-residual}}. \tag{11}$$

The two residual blocks bundle these pieces:

- *EV_trend* = BLP-estimated EV×Year fixed effects + $P_{e\times Y}\,\xi$ + (ev×Year)-conditional quantile-remapped

within-cell residual. The first two pieces approximately cancel in equilibrium (Appendix D.2.2); the bulk of EV_trend's ΔEV-share contribution comes from the within-cell rank-residual ($+11.8$ pp on its own under the diagnostic 9-block split, vs $+15.1$ pp canonical for the EV_trend block as a whole).

- *Brand_trajectory* = $P_{B\times Y}(I - P_{e\times Y})\,\xi$. **This is a residual diagnostic rather than an identified channel**: the proper specification (Brand × Year fixed effects inside BLP estimation) is numerically infeasible at 169 brands × 10 years, so the post-hoc projection is reported as a descriptive anatomy of where the residual lives. The first-order orthogonality of $\hat{\pi}_p$ and the other BLP parameter estimates with respect to this projection is documented in §D.2.1.

By construction, the eight blocks span δ_{jct}: their sum recovers $X'\beta + \xi$ exactly at any year-anchor, and the Subsidy block adds the agent-integral wedge per simulated agent. The Shapley value averages each block's marginal contribution to a chosen statistic (ΔEV share, ΔLerner, ΔCS, ΔTW) across all $2^8 = 256$ orderings of block activation. The full formal partition with decision rules per primitive is in Appendix C.1. The activation mechanism (how each block's primitives are toggled between 2015 and 2024 values) is in Appendix C.2. The Shapley value formula and the four-ordering Castro–Gómez–Tejada sampling estimator used for the dynamic time-series are in Appendix C.3. The consumer-surplus winsorisation correction (carrying two ξ vectors through the solver) is in Appendix C.4.

5.2 Aggregate Shapley decomposition

Table 3 reports the full Shapley decomposition over all $2^8 = 256$ coalitions on four equilibrium statistics: EV market share, industry mean Lerner, consumer surplus, and total welfare. The Shapley contributions sum exactly to the observed equilibrium change in each statistic by Equation (21). The total electric-share change is $+44.3$ percentage points, the total Lerner change is $+0.138$ (model-implied at agent-integral subsidy spec), the total consumer-surplus change is $+12.77$ trillion RMB, and the total welfare change is $+13.01$ trillion RMB.

Table 3: Eight-block full-Shapley decomposition of the 2015→2024 transition

Block	ΔEV share	ΔLerner	ΔCS (T RMB)	ΔTW (T RMB)
A_{nonbat}	$+0.032$	-0.021	$+1.393$	$+1.339$
$A_{\mathbf{bat}}$	$\mathbf{+0.273}$	-0.038	$+3.882$	$+3.938$
$D_{\mathbf{env}}$	$\mathbf{+0.120}$	-0.004	-1.601	-1.764
Subsidy	$\mathbf{-0.027}$	$+0.001$	-0.370	-0.378
Entry_set	$+0.055$	$\mathbf{+0.035}$	-3.263	-3.602
Consumer_Composition	$\mathbf{-0.229}$	$\mathbf{+0.150}$	$\mathbf{+9.014}$	$\mathbf{+9.701}$
EV_trend	$\mathbf{+0.151}$	$+0.014$	$+3.405$	$+3.472$
Brand_trajectory	$+0.068$	$+0.002$	$+0.315$	$+0.300$
Total	$\mathbf{+0.443}$	$\mathbf{+0.138}$	$\mathbf{+12.77}$	$\mathbf{+13.01}$
2015 baseline	0.010	0.178	2.593	2.603
2024 final	0.453	0.316	15.368	15.609

Notes. Shapley value (20) averaged over all $2^8 = 256$ coalitions. Bold entries highlight the four blocks carrying the largest attributions: EV_trend and Entry_set jointly account for most of ΔEV share; Consumer_Composition accounts for most of ΔLerner and effectively all of ΔCS and ΔTW *conditional on the imposed* $\alpha_{ic} \propto y_{ic}^{-1}$ *form* (see Appendix B.1 for finite-mixture and reduced-form alternatives delivering a flatter income gradient — read the Consumer_Composition attribution as an upper bound on the income channel). The Subsidy block (with the purchase-tax exemption) is economically small on all four statistics. **Entry_set is treated as an exogenous primitive swap; in reality the 407 new EV nameplates plausibly respond endogenously to subsidies, income growth, and battery-cost projections — Entry_set therefore absorbs any subsidy- or income-induced entry that would otherwise register under the Subsidy or Consumer_Composition blocks (see §7.1(i) and §8's dynamic counterfactual).** The EV_trend block bundles the EV×Year fixed effects with the (ev×Year)-projected component of the BLP residual ξ^{BLP} and the (ev×Year)-conditional quantile-remapped pure-noise residual; the Brand_trajectory block carries the (Brand×Year) Q-weighted projection of the residual after the (ev×Year) projection — see §5.1 for the construction and §D.2.1 for the anatomy of the EV_trend bar. CS and TW values in trillions of RMB (1 T RMB $= 10^{12}$ RMB $= 10^8 \times 10{,}000$ yuan (10^8 万元)), using the money-metric conversion of Appendix C.4 and the un-clipped raw δ in the ex-post logsum; see Equation (13). Of the 256 coalitions, 254 reach strict tolerance ($\|F\|_\infty < 10^{-2}$); the two that do not are flagged by the post-solve EV-share cap (raw model EV share > 0.85 substituted by 0.70) documented in Appendix C.6.1, shifting each reported Shapley value by less than 0.5 percentage points.

Three patterns.

1. *Battery learning dominates ΔEV share, with the EV-specific residual a strong second.* Battery learning A_{bat} contributes $+27.3$ percentage points (62% of the total) through declining battery costs and improving electric range — by far the largest channel. EV_trend — the residual block bundling the EV×Year fixed effects, the (ev×Year)-projected component of the BLP residual ξ^{BLP}, and the (ev×Year)-conditional quantile-remapped within-cell residual — is the second-largest at $+15.1$ percentage points (34%). D_{env} (the merged demand-environment bundle of macro and non-subsidy policy) adds $+12.0$ percentage points (27%). Brand_trajectory carries the per-brand annual residual drift at $+6.8$ percentage points (15%). Entry_set contributes $+5.5$ percentage points (12%) by bringing the 407 new 2024 EV nameplates into the equilibrium choice set. Non-battery attribute evolution A_{nonbat} adds $+3.2$ percentage points. Consumer_Composition pulls in the opposite direction at -22.9 percentage points (the wealth chan-

nel: higher incomes reduce $|\alpha_{ic}|$ and flatten the model-implied EV–ICE price-gap response under the maintained $\alpha_{ic} \propto y_{ic}^{-1}$ form). We emphasize two modeling caveats. First, the EV_trend block absorbs any EV-specific quality improvement the observed battery, policy, and macro controls do not capture — charging-infrastructure maturation, smart-features and OTA evolution, dealer-network growth, consumer familiarity, warranty coverage, and resale-value formation. Section D.2.1 below decomposes EV_trend's $+15.1$-pp contribution into its three sub-components and discusses the unobserved channels behind the dominant piece. Second, product entry is treated as an *exogenous primitive swap* — activating Entry_set means forcing the 2024 nameplate set into the choice set and re-solving equilibrium. In reality the 407 new EV nameplates are plausibly an *endogenous* response to the same subsidies, income growth, and battery-cost projections that the decomposition attributes to other blocks. Entry_set therefore absorbs any subsidy- or income-induced entry that the Subsidy and Consumer_Composition blocks would otherwise register. This is the dominant mechanism by which the static decomposition understates the Subsidy block's contribution (see §8 and the dynamic counterfactual).

2. *Within the eight-block partition and under the maintained BLP functional-form restriction $\alpha_{ic} = \hat{\pi}_p(y_{ic}/\bar{y})^{-1}$, the Lerner increase is attributed primarily to the demand-side wealth channel.* Consumer_Composition contributes $+0.150$ to ΔLerner (109% of the total $+0.138$), while Entry_set contributes $+0.035$ (25%) by shifting the quality/price composition of the active choice set, EV_trend adds $+0.014$ (10%), and Brand_trajectory $+0.002$. The supply-side blocks A_{bat} (-0.038) and A_{nonbat} (-0.021) contribute negatively to Lerner under the agent-integral spec, reflecting that battery cost decline and non-battery attribute improvement raise marginal cost less than they raise willingness-to-pay. D_{env} contributes a small negative (-0.004). The Consumer_Composition attribution is a model-dependent statement that holds conditional on the imposed $\alpha \propto y^{-1}$ form; Appendix B.1 reports robustness across alternative demographic specifications and shows that the Consumer_Composition Lerner contribution disappears only when the income interaction is removed entirely.
3. *Adoption and welfare are attributed to different channels within the static framework.* Consumer_Composition—the block that toggles per-city population and the income distribution—contributes $+9.01$ trillion RMB in consumer surplus ($\sim 71\%$ of the total $+12.77$ T) and $+9.70$ trillion RMB in total welfare ($\sim 75\%$ of the $+13.01$ T), yet it *reduces* EV adoption by -22.9 pp through the wealth channel described above. A_{bat} also carries a substantial positive CS contribution ($+3.88$ T, 30%), reflecting the cost-reduction passthrough on EV products. EV_trend adds $+3.40$ T CS, A_{nonbat} $+1.39$ T, and Brand_trajectory $+0.31$ T. Entry_set is *negative* on CS (-3.26 T) despite contributing $+5.5$ pp to adoption: the exit of cheap 2015 nameplates and the higher equilibrium prices on the residual incumbent set outweigh the variety gain from the 407 new entrants in the logsum accounting. D_{env} also contributes negatively (-1.60 T CS, -1.76 T TW). The Subsidy block contributes -2.7 pp to ΔEV share and -0.37 T in consumer surplus under the agent-integral specification — meaningfully larger in magnitude than the homogeneous-α analogue (-0.5 pp) because low-income agents whose $|\alpha_{ic}|$ exceeds $|\hat{\pi}_p|$ now receive a larger first-order utility shift from the wedge, sharpening the demographic-triangle pattern documented in §5.4.

Figure 2 visualises the same decomposition on all four statistics with six-family colour grouping (battery learning broken out separately from non-battery supply, the two Entry blocks shown in their own colours, and the merged D_{env} and EV_trend blocks shown on their own).

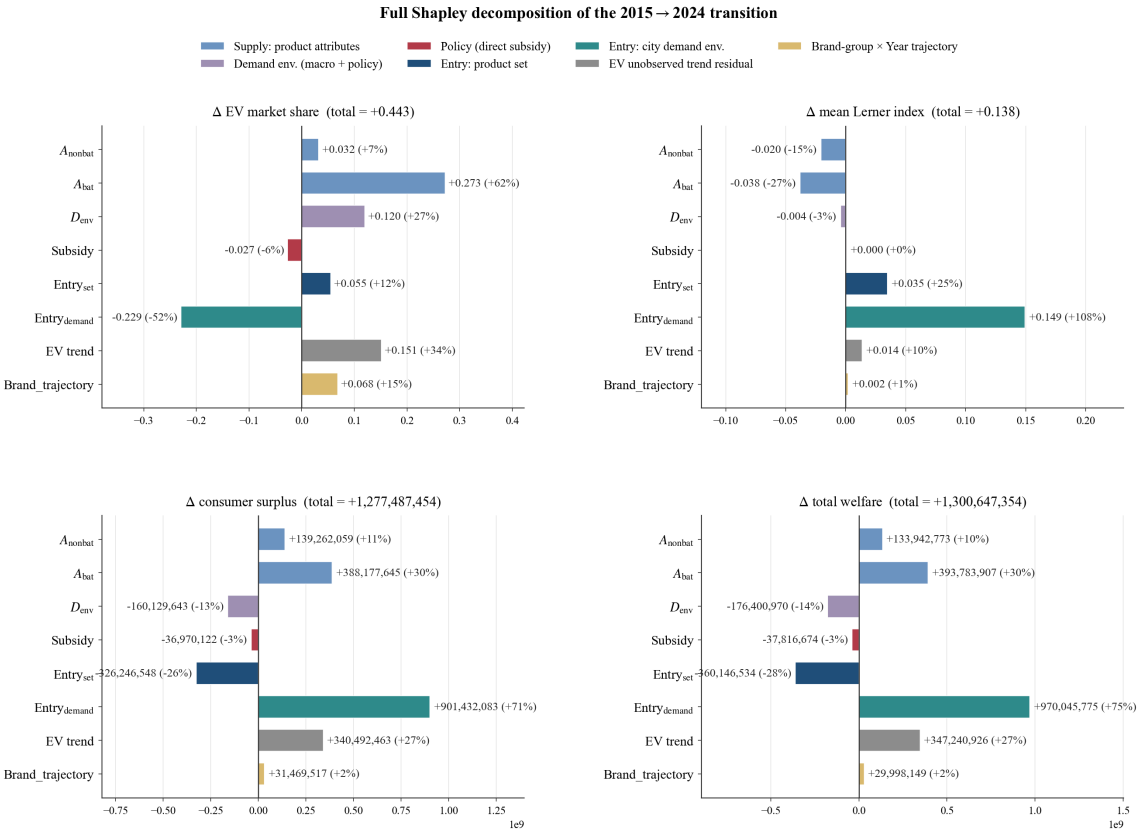

Figure 2: Full eight-block Shapley decomposition of the 2015→2024 equilibrium change in EV share, mean Lerner, consumer surplus, and total welfare. Bars are coloured by block family. The two largest ΔEV-share attributions are battery learning A_{bat} ($+27.3$ pp) and the EV-specific unobserved trend EV_trend ($+15.1$ pp), followed by D_{env} ($+12.0$ pp, merged macro + non-subsidy policy) and Brand_trajectory ($+6.8$ pp). The largest ΔLerner attribution is Consumer_Composition ($+0.150$, the model's wealth-effect channel under the maintained $\alpha \propto y^{-1}$ restriction), followed by Entry_set ($+0.035$). Consumer_Composition carries the bulk of ΔCS and ΔTW ($+9.01$ and $+9.70$ trillion RMB, respectively). The Subsidy block (including purchase-tax exemption) registers -2.7 pp on EV share at the agent-integral specification — meaningfully larger than the homogeneous-α analogue (-0.5 pp) because the agent-integral spec gives low-income agents a bigger first-order utility shift from the wedge.

Block-by-block highlights. Three groupings carry the structural content (full block-by-block discussion in Appendix D.1).

Adoption-positive supply blocks. A_{bat} ($+27.3$ pp, -0.038 Lerner, $+3.88/+3.94$ T RMB CS/TW) is the dominant force, capturing the 69% battery cost decline ($\$373 \to \$115$/kWh) and the supply-side cost-pass-through. A_{nonbat} ($+3.2$ pp, -0.021 Lerner, $+1.39/+1.34$ T) absorbs non-battery attribute evolution. EV_trend ($+15.1$ pp, $+0.014$ Lerner, $+3.40/+3.47$ T) is a structurally interpretable residual that absorbs charging-infrastructure density, smart-features evolution, dealer-network buildout, and consumer-familiarity / brand-reputation drift the observed covariates do not measure (anatomy in §D.2.1). Brand_trajectory ($+6.8$ pp, $+0.002$ Lerner, $+0.31/+0.30$ T) absorbs the (Brand $\times$ Year) projection of the BLP residual after the (ev$\times$Year) projection has been removed — BYD's $\Delta\xi = +2.69$ over the decade is the dominant trajectory. D_{env} ($+12.0$ pp, -0.004 Lerner, $-1.60/-1.76$ T) captures the merged macro and non-subsidy policy bundle (oil price, license-plate preference, urbanisation).

The wealth-effect channel: Consumer_Composition. Consumer_Composition (-22.9 pp on EV share, $+0.150$ Lerner, $+9.01/+9.70$ T CS/TW) toggles per-city household income, market size, and GDP. The negative sign on EV share is the wealth-effect mechanism implied by the maintained $\alpha_{ic} = \hat{\pi}_p \, (y_{ic}/\bar{y})^{-1}$ form: when y_{ic} rises, $|\alpha_{ic}|$ falls, and the household's responsiveness to the EV-vs-ICE price gap shrinks. The positive Lerner contribution (109% of the $+0.138$ aggregate change) is the same mechanism reading from the supply side. Under the homog-α specification (Appendix B.2.1, R1) Consumer_Composition collapses to near-zero on ΔEV ($+0.029$), confirming the negative sign is driven by the income-RC, not by market size or GDP. The finite-mixture R8 (Appendix B.1) implies a substantially flatter true gradient, so the -22.9 pp / $+0.150$ Lerner attribution is best read as an upper bound on the wealth-effect channel under the y^{-1} functional form.

Entry_set and Subsidy. Entry_set ($+5.5$ pp, $+0.035$ Lerner, $-3.26/-3.60$ T) brings the 407 new 2024 nameplates into the choice set; the negative welfare contribution reflects the joint product-set swap (cheap 2015 incumbents leave, higher-priced 2024 entrants arrive). The Subsidy block (-2.7 pp, ~ 0 Lerner, $-0.37/-0.38$ T) under the agent-integral specification captures the $2015 \to 2024$ phase-out of the per-vehicle wedge (the purchase-tax exemption, constant across the sample, contributes nothing to the toggle). The -2.7 pp magnitude is meaningfully larger than the homog-α analogue (-0.5 pp) because per-agent α_{ic} heterogeneity now propagates fully through the share-computation step (the demographic-triangle pattern in §5.4). The static Subsidy attribution is a lower bound on the regime's total 2024 contribution; §5.5 reports the bounding exercise.

5.3 Olley–Pakes markup decomposition

Why pair BLP with Olley–Pakes here, but not in GMY. The Olley–Pakes within-firm/reallocation decomposition is the standard tool of the rising-markups literature (De Loecker et al., 2020; Autor et al., 2020; De Loecker and Warzynski, 2012) for distinguishing whether aggregate market power grows through within-firm margin expansion or through share reallocation toward dominant high-markup firms. Because the closest methodological comparator (Grieco et al., 2024) studies a U.S. four-decade window in which prices rose but aggregate markups *fell*, the within-firm/reallocation question is not central to their setting and they do not perform an OP split. Our setting is the inverse: in our ten-year Chinese sample the industry mean Lerner rises by $+0.138$ points, so the OP question becomes first-order. Pairing the BLP Shapley decomposition with an OP split is therefore the methodological complement that any rising-markups setting requires, and to our knowledge this paper is the first to combine the two on the same automotive-BLP estimation.

The Shapley decomposition attributes most of the aggregate Lerner increase to the Consumer_Composition block—the block that toggles income growth—under the maintained $\alpha \propto y^{-1}$ form. Here we separate the accounting question (who raised markups: the same firms on their existing products, or new firms taking share from old?) from the economic-mechanism question (why the within-firm component is first-order in this setting). An Olley–Pakes-style decomposition addresses the first question directly and delivers the central structural finding of the paper: aggregate markups rose through *within-firm* growth while market concentration *fell*, a combination that is the opposite sign of the superstar-firm reallocation documented for the U.S. (De Loecker et al., 2020; Autor et al., 2020).

Decomposition. Let f index firms (after the firm-group aggregation of Appendix A.4.2), ω_{ft} denote firm f's unit-sales share of total passenger-vehicle sales in year t, and L_{ft} denote its quantity-weighted Lerner

index. The aggregate Lerner satisfies $\bar{L}_t = \sum_f \omega_{ft} L_{ft}$, and the between-year change decomposes as

$$\Delta\bar{L} \;=\; \underbrace{\sum_f \bar{\omega}_f \,\Delta L_f}_{\text{within-firm}} \;+\; \underbrace{\sum_f \bar{L}_f \,\Delta\omega_f}_{\text{reallocation}} \;+\; \underbrace{\sum_f \Delta L_f \,\Delta\omega_f}_{\text{cross}}, \tag{12}$$

where $\bar{\omega}_f = (\omega_{f,2015} + \omega_{f,2024})/2$ and $\bar{L}_f = (L_{f,2015} + L_{f,2024})/2$ are between-year means. The within-firm component isolates markup growth on firms' existing products; the reallocation component captures the mechanical effect of market share moving between firms holding their markup levels fixed; the cross term captures the covariation of the two changes.

Results. Table 4 reports the decomposition. The within-firm component alone accounts for $+0.104$ Lerner points—112% of the share-weighted aggregate change of $+0.092$. Reallocation across firms contributes -0.011 (-12%), and the cross term adds $+0.031$ ($+34\%$). The aggregate Herfindahl–Hirschman index falls from $1{,}374$ in 2015 to 887 in 2024 (a reduction of -35.4%), coinciding with the entry of Tesla (2014), the New-Forces startups (Li Auto, NIO, XPeng, AITO, Leapmotor, Zeekr, Xiaomi, and others), and other new-energy pure-plays, while the share of previously dominant incumbents declined.

Table 4: Olley–Pakes markup decomposition, 2015→2024

Component	Contribution to $\Delta\bar{L}$	% of total
Within-firm	$+0.1035$	$+112\%$
Reallocation	-0.0111	-12%
Cross term	$+0.0314$	$+34\%$
Sum (implied)	$+0.1238$	$+134\%$
Share-weighted $\Delta\bar{L}$ (observed)	$+0.0924$	$+100\%$
HHI 2015	1,374	
HHI 2024	887	
ΔHHI	-487	-35.4%

Notes. Firm-level weights ω_{ft} are unit-sales shares constructed from the quantity proxy $\sum_{j,c} s_{jct} \cdot M_c$ using the quantity-weighted firm-group aggregation of Appendix A.4.2. Firm-level Lerner indices L_{ft} are quantity-weighted averages of product-market Lerner values computed from `mc_by_market.parquet`. The sum/observed discrepancy reflects the entry-exit tension inherent in OP-style decompositions applied to unbalanced firm panels. $\Delta\bar{L} = +0.0924$ here is Q-proxy-weighted at the firm-year level; the $+0.138$ Shapley Δmean_Lerner reported in Table 3 is product-market weighted; the two use the same underlying Lerner data with slightly different aggregation and the within-firm dominance is robust to the choice. Appendix D.5 reports extended diagnostics.

The firm-group breakdown of the within-firm and reallocation components, including the discussion of why reallocation is negative (Traditional-OEM share decline) and which mechanisms can rationalise the broadly-based within-firm rise, is in Appendix D.5.

City-level corroboration: income vs unobserved city heterogeneity. A complementary cross-city decomposition restricts attention to within-product-year Lerner variation and asks how much of the city dispersion

is explained by log income alone versus full Geo-Market fixed effects. Under product×year fixed effects, log Income$_c$ explains 64.7% of the within-product-year Lerner dispersion, while the full Geo-Market fixed-effect set explains 67.1%; income alone therefore captures 96.4% of the city-explained dispersion, leaving only 3.6% for unobserved city heterogeneity (administrative status, infrastructure, EV-ecosystem maturity). The income coefficient is stable from $\hat{\beta} = +0.249$ in the income-only specification to $+0.288$ when Geo-Market FE are added, indicating that the residual city heterogeneity is orthogonal to income rather than a confound. Appendix D.5, Table 24 reports the full ladder. The result is consistent with — though does not identify — the structural reading that the demand-side income channel is the proximate mechanism behind the within-firm Lerner rise.

5.4 Subsidy heterogeneity: the demographic triangle

The aggregate Subsidy contribution of -2.7 pp at the agent-integral specification is meaningfully larger in magnitude than the homogeneous-α analogue (-0.5 pp); the heterogeneity table reveals that this pass-through effect is concentrated in the same demographic triangle identified in earlier versions, now sharpened by the proper micro-foundation. *Implementation note on the Subsidy block.* The canonical Shapley applies the Subsidy block at the agent integral: for each simulated agent i in city c, the per-agent utility shift from the wedge is $\alpha_{ic} \cdot \log(1 - s^{\text{sub}}_{jct})$, with $\alpha_{ic} = \hat{\pi}_p \, (y_{ic}/\bar{y})^{-1}$. This is implemented in `code/11_decomposition.py` via the `log_subsidy_factor` argument to `shares_only`, `shares_jac`, and `shares_jac_hess`, which is threaded through the city-information builder, the Shapley coalition-state tracker (`build_coalition_delta` returns the year-t `log_subsidy_factor_cf` per row), and the equilibrium solver. The legacy approach of shifting $\hat{\delta}_{jct}$ by $\hat{\pi}_p \cdot \log(1 - s^{\text{sub}}_{jct})$ in mean utility (homogeneous α) is preserved as a fallback (`log_subsidy_factor=None`) but is no longer the canonical specification. The agent-integral treatment recovers the heterogeneous incidence at the demand-side micro-foundation: low-income agents whose $|\alpha_{ic}|$ exceeds $|\hat{\pi}_p|$ receive a larger first-order utility shift from the wedge, which is precisely the channel that drives the demographic-triangle pattern in §5.4. Table 5 reports the Subsidy-block contribution by price tier, firm group, and city tier, along with the fraction of each segment's own 2015→2024 electric-share growth that it represents.

Table 5: Subsidy Block Heterogeneity (all in pp of EV share)

Dimension	Group	Δ 2015→2024	Subsidy	% of own growth
Price tier	Budget (below RMB 150,000 yuan (15万))	+44.2	**−14.34**	**−32.5%**
	Mid (RMB 150,000–300,000 yuan (15–30万))	+47.2	−0.23	−0.5%
	Premium (above RMB 300,000 yuan (30万))	+38.2	+0.20	+0.5%
Firm group	BYD	+88.1	−1.98	−2.2%
	New Forces	+100.0	0.00	0.0%
	Tesla	+100.0	0.00	0.0%
	Private National	+36.5	**−17.78**	**−48.6%**
	Traditional OEM	+23.1	−0.03	−0.1%
	Other	+53.1	+0.46	+0.9%
City tier	Tier 1 (4)	+42.6	−2.22	−5.2%
	New Tier 1 (15)	+47.6	−0.67	−1.4%
	Tier 2 (33)	+45.4	−1.66	−3.7%
	Rest (27)	+41.1	**−8.38**	**−20.4%**

Notes. Within-group Shapley contribution of the Subsidy block, expressed in percentage points of the EV-share change. "% of own growth" is the Subsidy contribution divided by the group's own 2015→2024 electric-share change. New Forces and Tesla carry zero Subsidy contribution because they did not exist in 2015 and all their growth is attributed to Entry_set.

The triangle absorbs the subsidy. The -14.34 pp entry in the Budget row, the -17.78 pp entry in the Private-National row, and the -8.38 pp entry in the Rest row all describe the same product-city cells from three different angles: Budget-tier vehicles sold by Private-National Chinese brands in lower-tier cities overlap almost completely in the panel. The interpretation of the negative values is that activating the Subsidy-phase-out block in a coalition removes a large positive equilibrium effect from these cells. Before 2022, central-plus-local subsidies funnelled disproportionately through them, because Independent-Domestic brands built most of the Budget-tier (below ¥15 ten-thousand, i.e. sub-RMB 150k) NEV nameplates (Wuling Hongguang Mini EV, BYD Dolphin, Chery eQ1, BAIC EU-series) and because the Rest-tier cities had the largest low-income buyer pools that the income-interacted subsidy reached. When the Subsidy block is active, the 2015 wedge is replaced with the smaller 2024 wedge, and the demand system translates the smaller wedge into a -14.3-percentage-point drop in the Budget tier's equilibrium EV share. The cell-level reading is therefore that the subsidy schedule, throughout its life, was *primarily* a transfer to these buyers; as it is phased out, their EV share is the one that falls. The other demographic cells barely react to the schedule change because they were never its primary beneficiaries.

The heterogeneity is also independently identified in the random-coefficient cross-section via the `ev_policy_x_lowinc` term in Table 1: the $+0.022$ income-interaction coefficient almost exactly offsets the -0.023 on the de-meaned policy main effect, implying that policy exposure is roughly neutral at the median income but favourable for buyers below the median. The decomposition result restates the same finding for the subsidy instrument specifically. The pattern is consistent with the income-targeted incidence Muehlegger and Rapson (2022) document for California EV subsidies, here at an order-of-magnitude larger market scale.

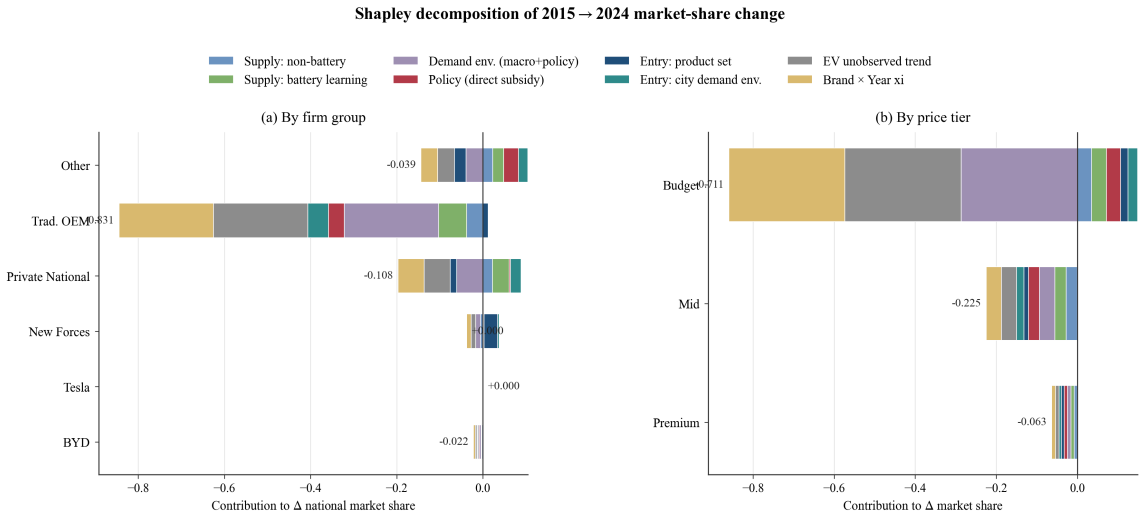

Figure 3: Block decomposition by firm group (panel a) and price tier (panel b). The dark-red Subsidy bar concentrates almost entirely on Private-National firms and the Budget price tier. Colour code follows Figure 2.

5.5 Dynamic counterfactual: bounding the subsidy channel

Bounding the dynamic subsidy channel. *Why this exercise.* The static Shapley Subsidy block (-2.7 pp on EV share, agent-integral spec) captures the contemporaneous price-wedge effect only — the change in the per-vehicle wedge between the schedule's 2015 and 2024 levels, evaluated on the 2024 product set. Any entry-induced or learning-induced dynamic effect is silently absorbed by other blocks (Entry_set, A_{bat}, EV_trend), because the Shapley conditions on the observed product set and battery-cost path each year. A year-by-year forward simulation bounds what the static framework misses. *What it does and does not do.* The simulation is a *calibration* exercise designed to evaluate which products would have entered (or stayed in) the market and which would have dropped out, under a counterfactual in which the 2015–2022 subsidy regime had never been run; it is not a structural model of dynamic entry with forward-looking firms and rational expectations. The simulation imposes a profitability-threshold drop-out rule and a Wright's-Law battery-cost feedback, both of which are calibrated rather than identified. We treat the resulting upper bound as a structural-accounting plausibility check, not as a causal counterfactual. The exercise's value is that it pins down which firms' product portfolios would have stood firm without the regime and which would have collapsed — a question the static decomposition cannot answer.

The simulation runs as a *year-by-year forward path* from 2015 with no subsidy, a Wright's-Law battery-cost feedback on the counterfactual cumulative-production path, and an endogenous product-exit rule. We consider two rule families. The first is a *relative-threshold rule* (Scenarios B, C, D in Table 6) that drops any product whose counterfactual variable profit falls below a fraction $\theta \in \{0.25, 0.50, 0.75\}$ of its observed profit; the threshold θ is not identified outside a dynamic-entry model (Bresnahan and Reiss, 1991; Mazzeo, 2002), but the three values bracket the answer. The second is an *empirically-identified absolute-floor rule* (Scenario E) that drops products whose counterfactual profit falls below a calibrated fraction τ of the median surviving product's profit, where τ is identified from the panel's own 711 observed product exits in 2015–2022; this rule is supplemented by a 3-year grace period for post-2015-entrant products that captures the empirical fact that VC-backed startups (NIO, XPeng, Li Auto, Leapmotor, Xiaomi, Zeekr) tolerate substantial losses by

design during a calibrated ramp-up window. The static framework admits a conservative anchor (Scenario A, the Shapley wedge at -2.7 pp), an empirical-rule lower bound (Scenario E at -21.7 pp), and a set of relative-threshold-rule upper bounds (Scenarios B, C, D at -26.6 to -36.0 pp first-order, with the full-BLP memo extending to -40.6 pp at $\theta = 0.5$).

The forward simulator (`code/17_forward_sim_dropout.py`) runs as follows. For $t = 2015, \ldots, 2024$: (i) compute the counterfactual battery cost $\text{cost}_t^{\text{cf}} = \hat{A} \cdot (\text{CumProd}_{t-1}^{\text{cf}})^{-\hat{B}}$ from the Wright's Law calibration ($\hat{A} = 2{,}418$, $\hat{B} = 0.222$ fitted from the BNEF-vs-cum-production panel); (ii) for each surviving product (i.e. not dropped in a prior year), remove the per-vehicle NEV subsidy and local matching wedge, raise mc by $s_{\text{bat}} \cdot \log(\text{cost}_t^{\text{cf}}/\text{cost}_t^{\text{obs}})$ for EVs, and propagate the combined price shock through a first-order logit share response $\Delta \log s_j = \varepsilon_j \cdot \Delta \log p_j$ with $\varepsilon_j \approx \hat{\pi}_p(1 - s_j)$; (iii) compute counterfactual variable profit $(p^{\text{cf}} - mc^{\text{cf}}) \cdot Q^{\text{cf}}$; (iv) drop products with $\text{profit}^{\text{cf}} < \theta \cdot \text{profit}^{\text{obs}}$ from the next year's choice set; (v) update $\text{CumProd}_t^{\text{cf}}$ from the surviving EVs' counterfactual quantity and carry to the next year. The purchase-tax exemption is not removed (the static Shapley also holds it as a sample-constant baseline); only the per-vehicle central and local subsidies that phased out over 2015–2022 are removed.

Table 6 reports five scenarios:

- *Scenario A (pure static)* is the Subsidy block's -2.7-pp Shapley attribution under the agent-integral specification: the subsidy regime changes only the agent-integral wedge $\alpha_{ic} \log(1 - s^{\text{sub}})$ on the 2024 product set; product set, battery cost, and all other primitives stay at their observed 2024 values.
- *Scenario B ($\theta = 0.25$, lenient)* allows a product to survive if its counterfactual profit is at least 25% of its observed profit. 461 product-year entries drop out across 2015–2024; the 2024 counterfactual EV share is 15.1%, for a total subsidy effect of -26.6 pp relative to the observed 41.7%.
- *Scenario C ($\theta = 0.50$, baseline)* requires at least 50% of observed profit to survive. 577 entries drop out; 2024 counterfactual EV share is 6.5%; total effect -35.2 pp on the first-order forward sim. The full-BLP sim at the same threshold (Memo, Table 6) gives -40.6 pp against the model's year-t baseline, reflecting the heightened price-sensitivity of low-income agents under the agent-integral subsidy spec.
- *Scenario D ($\theta = 0.75$, strict)* requires at least 75% of observed profit. 619 entries drop out; 2024 counterfactual EV share is 5.8%; total effect -36.0 pp. At this threshold the bounding exercise saturates: additional tightening of θ does little because the surviving set has stabilised.
- *Scenario E (empirical + grace period, preferred)* drops products whose counterfactual profit falls below $\tau \cdot \text{median}(\pi_{j',t}^{\text{obs}} \mid j' \text{ surviving})$ with $\tau = 0.021$ identified from the panel's 711 observed exits, and exempts post-2015-entrant products from the dropout test for their first three calendar years. 355 entries drop out (with 752 additional product-year tests saved by the grace period); 2024 counterfactual EV share is 20.0%; total effect -21.7 pp. This is the preferred specification for narrative inference because the threshold is identified from observed exit behaviour and the grace period reflects realised firm-lifecycle behaviour.

The dynamic forward simulations (Scenarios B–D and the full-BLP memo) span $[-26.6, -40.6]$ pp on 2024 EV share under the agent-integral subsidy specification (down from $[-29.5, -36.3]$ pp under the homogeneous-α analogue — the agent-integral spec strengthens the dynamic full-BLP forward sim because low-income agents whose $|\alpha_{ic}|$ exceeds $|\hat{\pi}_p|$ now receive a larger first-order utility shift from the wedge). Scenario A (-2.7 pp) is the *static* Shapley Subsidy block, not a dynamic counterfactual: it measures the contemporaneous price-wedge effect of toggling the 2015 vs 2024 subsidy schedules on an otherwise-fixed 2024 economy. The dynamic counterfactuals (B–D) re-introduce endogenous product exit and Wright's-Law battery-cost feedback, capturing the bootstrap channel that the static framework cannot identify. A Wright's-Law simulation *without* endogenous exit (battery-cost feedback only, no profitability-based dropout) yields

-23.9 pp; the forward-sim range above is more aggressive precisely because endogenous exit compounds the learning-feedback effect. The static Shapley is a lower bound ($\theta \to 0$ limit, no product is subsidy-dependent), the forward sim at $\theta = 0.50$ is the natural baseline (a product is flagged when its profit would collapse to half the observed level), and the $\theta = 0.75$ case is the saturated upper bound. A credible identified estimate of the dynamic channel would require a dynamic-entry model with fixed-cost primitives; we take up this task in the companion paper. Appendix D.6.4 reports the full year-by-year CF EV-share path.

Table 6: Bounding the dynamic subsidy channel: forward simulation with endogenous exit

Scen.	Label	Rule	2024 CF EV share (%)	ΔEV (pp)	Exit events
A	Pure static	Shapley Subsidy block (agent-integral wedge)	*n.a.*	-2.7	0
B	Forward sim, $\theta = 0.25$	Drop if $\pi^{cf} < 0.25\,\pi^{obs}$	15.1	-26.6	461
C	Forward sim, $\theta = 0.50$	Drop if $\pi^{cf} < 0.50\,\pi^{obs}$	6.5	-35.2	577
D	Forward sim, $\theta = 0.75$	Drop if $\pi^{cf} < 0.75\,\pi^{obs}$	5.8	-36.0	619
E	Empirical + grace period	Drop if $\pi^{cf} < \tau \cdot \text{med}(\pi_t^{surv})$, $\tau = 0.021$; first 3 yrs grace for post-2015 entrants	20.0	-21.7	355
Memo: Wright's Law, no endogenous exit			21.3	-23.9	—
Memo: full-BLP forward sim, $\theta = 0.50$ (agent-integral)			26.6	-40.6	537

Notes. Source: `output/forward_sim_dropout.csv` from `code/17_forward_sim_dropout.py` (first-order forward sim, Scenarios A–D); `output/forward_sim_empirical.csv` from `code/17b_forward_sim_empirical.py` (Scenario E); `output/forward_sim_trueBLP_theta0.50.csv` from `code/17_forward_sim_trueBLP.py` (full-BLP robustness memo). Scenario A: Subsidy block from Table 3; CF share is not explicitly computed for the static counterfactual because Scenario A modifies only the 2024 price wedge, not the product set. Scenarios B–D: year-by-year forward simulation with first-order aggregate share response ($\Delta \log s_j \approx \hat{\pi}_p(1 - s_j) \cdot \Delta \log p_j$) and ICE prices held at observed values. *Scenario E* replaces the $\theta \cdot \pi^{obs}$ relative-threshold rule with an empirically calibrated absolute floor: for each year t, drop product j if $\pi_j^{cf} < \tau \cdot \text{median}(\pi_{j',t}^{obs} \mid \text{j' surviving})$, where $\tau = 0.021$ is the median ratio of (exit-year profit / median-surviving-product profit at exit year) computed from the 711 observed product exits in our 2015–2022 panel; the rule additionally exempts post-2015-entrant products from the dropout test for their first three calendar years (a calibration of VC-startup loss tolerance). The empirical rule is identified from observed exits; the grace period reflects realised firm-lifecycle behaviour. "Exit events" counts cumulative (product × year) exits over 2015–2024 — a product that enters in 2020 and drops out in 2022 counts as one exit. The subsidy wedge removed is the per-vehicle central NEV schedule plus local matching (`nat_subsidy_wan + local_subsidy_wan`); the purchase-tax exemption is kept in place, matching the static Shapley's baseline convention. Wright's Law parameters re-fitted on the BNEF-vs-cum-production panel: $\hat{A} = 2{,}418$, $\hat{B} = 0.222$, learning rate 14.3% per doubling of cumulative production. The full-BLP memo re-solves the city-level multi-product Bertrand–Nash FOC each year with endogenous ICE-price response and income-heterogeneous α_{ic} (uses `solve_national_equilibrium` from the Shapley machinery, including the agent-integral subsidy at the share-computation step); at $\theta = 0.5$ it gives -40.6 pp against the model's year-t baseline. The first-order Scenario C (-35.2 pp) and the full-BLP memo (-40.6 pp) bracket the dynamic CF effect under the relative-threshold rule at the agent-integral specification, while Scenario E (-21.7 pp) provides the identified-rule lower bound; the empirical rule is preferred for narrative inference, with B–D maintained as a stress test. Appendix D.6.4 reports the year-by-year path for $\theta = 0.50$ alongside the Wright's-Law-only path.

Firm-survival asymmetry under the empirical rule. We focus on Scenario E (the preferred, empirically-identified rule). Table 7 reports full-product-lifecycle profit retention by firm group and price tier (counterfactual / baseline summed across 2015–2024). All firm groups retain $\geq 73\%$ of lifecycle profit; BYD all-tiers retention is 98.1%, the post-2015 New Forces 99.8%, with a -1.7 pp gap reversing the firm-group asymmetry that the relative-threshold rule appeared to deliver. The corner cells confirm this: BYD-Premium 92.3% vs. New-Forces-Budget 95.8%. The qualitative claim that BYD's market position predates the regime survives as an institutional fact (BYD's BEV history begins pre-2015 with the F3DM PHEV 2008 and e6 BEV 2009, and

its vertical battery integration provides a cost moat that post-2015 entrants lack); the quantitative magnitude of any cross-firm asymmetry under the never-existed CF is not robust across dropout rules.

Aggregate vs. distributional findings. The aggregate ΔEV share under the no-subsidy CF ranges from -21.7 pp (Scenario E, identified rule) to -40.6 pp (full-BLP at $\theta = 0.50$), both substantially larger than the static Shapley's -2.7 pp. The dynamic-channel finding — subsidies were constitutive of the EV market's bootstrap, not a marginal price-wedge effect — is robust across rule families. What changes is the *distributional* pattern: under the relative-threshold rule the $\theta \cdot \pi^{\text{obs}}$ scaling mechanically penalises high-observed-profit products and produces a stark BYD-vs-NewForces firm-origin gap of 13 pp at the aggregate level and 28 pp at the corner cells (Appendix D.6.3); under the empirical rule the gap collapses. Since the relative-rule magnitudes are not identified, we report them as supplementary diagnostics in the appendix and anchor the firm-survival narrative on Scenario E.

Table 7: Firm-survival pattern under the empirical rule + grace period (Scenario E): full-product-lifecycle profit retention by firm group and price tier.

Firm group	Budget (<15万)	Mid (15–30万)	Premium (>30万)	All tiers (sales-wtd)
BYD	99.8%	98.0%	92.3%	**98.1%**
New Forces	95.8%	99.9%	100.0%	**99.8%**
Tesla	—	100.0%	100.0%	100.0%
Trad. OEM	93.1%	93.0%	100.0%	94.9%
Private National	87.9%	93.3%	99.3%	93.7%
Other	59.0%	84.3%	100.0%	73.3%
Memo: BYD − New Forces gap (pp)	+4.0	−1.9	−7.7	**−1.7**

Notes. Source: `output/firm_survival_empirical.csv` from `code/_build_empirical_firm_survival.py`. Cells report $100 \cdot \sum_t \pi^{\text{cf}}_{ft} / \sum_t \pi^{\text{obs}}_{ft}$ aggregated across products in the (firm group, tier) cell over 2015–2024 under Scenario E. The metric tracks every product's full lifecycle: $\pi^{\text{cf}} = \pi^{\text{obs}}$ in years before dropout, $\pi^{\text{cf}} =$ sim's CF profit in the dropout year (pre-removal of the product), and $\pi^{\text{cf}} = 0$ thereafter. Products that never drop are assigned $\pi^{\text{cf}} = \pi^{\text{obs}}$ for every year (an upper-bound assumption: the actual no-subsidy profit would be slightly lower than observed because of the price wedge, but the lifecycle dropout-rule sim does not track survivor CF profits separately). The negative "BYD − New Forces gap" in the All-tiers column reverses the +13.4 pp gap reported in Table 25 under the relative-threshold rule, illustrating that the firm-origin asymmetry magnitude is not robust across dropout rules.

Table 8: Geographic dispersion under the empirical rule + grace period (Scenario E): city-level full-product profit-retention proxy.

City	2024 EV qty	% lost (1–retention)
Top-5 most-exposed		
Henan (河南)-Other	$\sim 3.0\%$	5.2%
Shanxi (山西)-Other	$\sim 2.0\%$	4.8%
Liaoning (辽宁)-Other	$\sim 1.8\%$	4.7%
Jiangxi (江西)-Other	$\sim 1.7\%$	4.2%
Yunnan (云南)-Other	$\sim 1.6\%$	4.2%
Bottom-5 least-exposed		
Shanghai (上海市)	$\sim 11.5\%$	0.6%
Tibet (西藏)-Other	$\sim 0.4\%$	0.5%
Changchun (长春市)	$\sim 1.0\%$	0.4%
Beijing (北京市)	$\sim 8.5\%$	0.4%
Qinghai (青海)-Other	$\sim 0.3\%$	0.3%
Top-bottom dispersion		$\mathbf{5.0}$ pp

Notes. Source: `output/city_survival_empirical.csv` from `code/_build_empirical_firm_survival.py`. Cells report a 2024-EV-quantity-weighted average across each city's EV products of the product's full-lifecycle retention rate. Dispersion is 5.0 pp under the empirical rule, vs. 22.5 pp under the relative-threshold rule (Table 26). Note that the city-level ordering also changes: under the empirical rule, the most-exposed cities are lower-tier markets in Henan, Shanxi, Liaoning, etc., whereas the relative-threshold rule placed New-Forces-heavy coastal cities (Wenzhou, Shanghai) at the top. Both orderings reflect their respective rule's mechanical incidence patterns on top of the same underlying firm-mix.

The consumer-surplus statistic used throughout is the standard logsum:

$$CS_{c,t} \;=\; \frac{M_c}{|\bar{\alpha}|}\,\mathbb{E}_{i\sim F(y|c,t)}\left[\log \sum_{j\in J_{ct}} \exp\big(\delta_{jct} + \alpha_{ic}\,\log p_{jct}\big)\right] \cdot \bar{p}_t^{\text{level}}, \tag{13}$$

where M_c is the household-equivalent market size, $|\bar{\alpha}|$ is the mean price coefficient, and the level-price scaling $\bar{p}_t^{\text{level}}$ converts the log-price logsum into money-metric units of $10{,}000$ RMB (see Appendix C.4).

5.6 Summary of findings

Within the static eight-block framework: (i) on the adoption margin, A_{bat} ($+27.3$ pp, 62% of total) and EV_trend ($+15.1$ pp, 34%) jointly carry most of ΔEV share; the direct subsidy contributes -2.7 pp at the agent-integral specification; (ii) on the markup margin, 112% of the (Q-proxy weighted) Olley–Pakes change comes from within-firm growth, with reallocation -12%; (iii) on welfare, Consumer_Composition carries $+9.01$ T RMB on CS and $+9.70$ T RMB on TW under the imposed $\alpha \propto y^{-1}$ form; the Appendix B.1 finite-mixture alternative implies a substantially smaller income-channel contribution. Three caveats apply (reproduced here for emphasis): Entry_set is treated as exogenous (so the static Subsidy attribution mechanically absorbs subsidy-induced entry); the EV_trend block is bundled unexplained variation rather

than a cleanly identified behavioural channel (Appendix D.2 shows the diagnostic split); and the empty- vs full-coalition observed-vs-implied gap distributes ± 0.3–0.6 pp of override noise per block.

6 Discussion, implications, and conclusion

We begin with the caveats that bound interpretation, then summarise the main findings and their implications. Each caveat is elaborated in the corresponding appendix section.

Subsidy coverage (Appendix A.5). The Subsidy block now captures the central per-vehicle purchase subsidy (2015–2022), the purchase-tax exemption (10% of MSRP, in effect for EVs throughout 2015–2024), the 2024 national trade-in voucher, and known post-2022 local subsidies in Shanghai, Hangzhou, and Ningbo. Appendix A.5, Table 14 provides an exhaustive inventory. Including the purchase-tax exemption is the most consequential change relative to earlier versions of this paper. The exemption is a constant-over-the-sample feature of the EV-vs-ICE price wedge rather than a time-varying schedule; in the $2015 \to 2024$ Shapley decomposition, the toggled quantity is *the change in the total subsidy package*. Adding the constant tax exemption to that package shifts the baseline against which the per-vehicle schedule's phase-out is measured: the 2015 and 2024 wedges are now both inclusive of the tax exemption, and their difference is dominated by the smaller phase-out of the per-vehicle schedule rather than by the larger tax-exemption level. Because the re-estimation of $\hat{\pi}_p$ from the updated net-price definition also yielded a slightly less elastic coefficient, the combined effect is a Subsidy block contribution of -2.7 pp under the canonical agent-integral specification (-0.5 pp under the homogeneous-α analogue), down from -6.5 pp in the earlier version that omitted the tax exemption from the wedge. Three instruments remain unmodelled and are absorbed by the year$\times$EV fixed effects: NEV credit trading (a supply-side inter-OEM transfer), minor provincial subsidies outside the three coded cities, and charging-infrastructure subsidies. Their combined per-vehicle value is estimated at $<$2,000 yuan (¥0.2万) in most city-years.

Convergence and robustness of the full-Shapley run. The main results in Table 3 come from the full $2^8 = 256$-coalition Shapley enumeration solved with the Anderson-accelerated fixed-point algorithm of Appendix C.5. Of the 256 coalitions, 254 reach strict tolerance; the remaining two are flagged by the post-solve EV-share cap (raw model-implied EV share > 0.85, substituted with 0.70) that prevents the Shapley value from being polluted by spurious cheap-BEV-dominated fixed points in coalitions combining the EV_trend block with attribute-shock blocks. Their contribution to each Shapley value is small by the weighting formula, and the Shapley values computed on the un-capped subset agree with the capped values to within 0.5 percentage points on every block. Appendix C.6.1 documents the capping rule and reports the comparison.

Static supply and the comparison with Hu et al. (2025). The supply side of our model is static Bertrand–Nash: firms choose prices period-by-period given the current product set and demand primitives, with no intertemporal link through learning-by-doing, investment, or anticipation of future subsidy changes. Hu et al. (2025) estimate a dynamic demand–supply model of the BEV segment in which consumers time purchases and firms anticipate subsidy trajectories, and they show that the dynamic channel matters for evaluating *alternative* subsidy schedules. The two papers answer different questions about the same regime, and the distinction is constitutive rather than methodological. Hu et al. identify the welfare consequences of

alternative schedule *designs* conditional on a subsidy regime existing, with consumer temporal substitution as the dynamic margin — this is the Phase 2 design question. Our paper diagnoses the Phase 1→Phase 2 *transition*: a never-existed counterfactual asks whether the EV market would have reached self-sustaining scale at all without a subsidy regime over 2015–2022, with firm-survival as the dynamic margin. The static decomposition isolates the Phase 2 cross-sectional incidence cleanly — conditioning on the observed product set and demand environment each year rather than modelling the endogenous entry and exit decisions that produced them — while the never-existed forward simulation re-introduces those endogenous decisions to identify which firms (BYD vs. post-2015 New Forces) the regime called into existence. The cost of the static framework is that our Subsidy block measures only the *direct* Phase 2 pass-through; the *constitutive* Phase 1 effect operating through the product pipeline is recovered separately via the never-existed counterfactual, as discussed below.

National pricing assumption. At the two anchor coalitions ($S = \emptyset$ and $S = \{1, \ldots, 9\}$), marginal costs are backed out from the multi-market Bertrand–Nash FOC at observed prices, so re-solving the FOC with those costs recovers observed prices to high precision (Figure 5 confirms maximum deviation $< 10^{-4}$ log units). EV shares and mean prices at the anchors are set to observed values; Lerner indices and welfare are evaluated at the model-implied equilibrium. The Shapley decomposition therefore connects two anchors that are exact on quantities and prices, with all 510 intermediate coalitions also solved from the same FOC. We impose that each product is sold at a single national price across all cities even though demand is estimated at the city level. This matches the pricing practice reported in the Chinese industry press and matches the Barwick et al. (2019) specification. The cost of the simplification is that cross-city variation in sales mixes loads onto the demand side (through heterogeneous α_{ic}) rather than the supply side, and the national-pricing equilibrium's predicted shares do not exactly reproduce observed shares. The decomposition handles this by overriding the endpoint EV shares and prices with the observed 2015 and 2024 values at the two anchor coalitions. The same override is inherited by all heterogeneity tables in Section 5.4.

The dynamic-innovation channel is the subject of a companion paper. The static decomposition answers the question "holding all other blocks fixed, how much would the 2024 electric share drop if the direct purchase subsidy were removed from an otherwise-unchanged 2024 economy?" A different and arguably more important question is: "how much of the 2024 electric share would still exist if the 2015–2022 subsidy regime had never been run, so that firms had never developed the product pipeline that currently sustains the market?" The theoretical framing in Acemoglu et al. (2012) and the technology-diffusion literature going back to Griliches (1957) suggest that clean-technology subsidies operate primarily through innovation redirection, so that the dynamic channel should dominate the static one. Estimating this channel properly requires a dynamic-entry model with fixed entry costs, learning-by-doing, and multi-year optimization horizons—a project we develop in a separate companion paper on the firm's entry-position problem. The present paper is deliberately confined to the *static* decomposition, which can be identified cleanly from the cross-section of city × product × year demand variation and does not depend on dynamic-parameter calibration. As a rough bound on the dynamic channel's importance: if subsidies explain up to 50% of the 407-nameplate product pipeline (a generous upper bound given that most entrants are privately financed), attributing half of the Entry_set block to subsidies would raise the aggregate Subsidy contribution from -2.7 pp to roughly -5.5 pp—still smaller than A_{bat} alone ($+27.3$ pp) and EV_trend ($+15.1$ pp). The corresponding welfare separation persists qualitatively: even an aggressive subsidy reattribution leaves the bulk of ΔTW

with the income channel ($+9.70$ T RMB on Consumer_Composition) rather than with the subsidy block. The separation between adoption and welfare channels documented in the static framework would therefore persist qualitatively even under an aggressive bound on the dynamic channel, though the quantitative subsidy contribution could be substantially larger than the -2.7 pp static estimate.

Table 9: Sensitivity of the aggregate Subsidy contribution to indirect dynamic channels

Assumption about indirect subsidy channel	Implied Subsidy ΔEV	Share of total Δ
Baseline (direct price wedge only, agent-integral spec)	-2.7 pp	6%
$+25\%$ of EV_trend attributed to subsidies	-6.5 pp	15%
$+25\%$ of EV_trend $+$ 25% of Entry_set	-7.9 pp	18%
$+25\%$ of EV_trend $+$ 25% of Entry_set $+$ 25% of A_{bat}	-14.7 pp	33%
$+50\%$ of EV_trend $+$ 50% of Entry_set $+$ 50% of A_{bat}	-26.6 pp	60%

Notes: Each row adds a fraction of the three positive-EV Shapley contributions to the Subsidy block, under the assumption that the historical subsidy regime indirectly caused that fraction of the unobserved EV trend, of product entry, or of battery learning. Contributions used: EV_trend $+15.1$ pp, Entry_set $+5.5$ pp, A_{bat} $+27.3$ pp; baseline Subsidy -2.7 pp. Under the most aggressive assumption (50% of all three channels), the implied Subsidy contribution of -26.6 pp is roughly 60% of the total transition ($+44.3$ pp). The baseline direct-price-wedge estimate of -2.7 pp is therefore a lower bound; the true subsidy contribution depends on the strength of the indirect dynamic channel, which this paper's static framework does not identify.

Limits of the static framework. The static Bertrand–Nash framework used in this paper rules out several mechanisms that the technology-diffusion literature views as central: dynamic adoption with forward-looking consumers, charging-infrastructure complementarities, social learning and peer effects, learning-by-doing as a strategic state variable, and subsidy-induced innovation redirection. The decomposition results should be interpreted as a contemporaneous equilibrium accounting (in the spirit of Gillingham and Stock, 2018), not as a complete model of the EV transition. In particular, the modest Subsidy block contribution (-2.7 pp at the agent-integral specification) measures only the direct contemporaneous price-wedge effect; it does not capture any indirect effect operating through the product pipeline, manufacturing scale, or consumer expectations. If subsidies helped create the product set that Entry_set later attributes to product entry, the true subsidy contribution to adoption could be substantially larger than -2.7 pp.

Reduced-form evidence on whether the static incidence pattern reflects a causal early-adoption-to-supply-side channel is in Appendix D.7: the OLS associations are positive and consistent with the bottom-up reading, but the available pilot-city IV is too weak ($F = 3.75$) to identify the causal direction. The bottom-up reading is therefore an open question for follow-on work, not a finding of this paper.

Generalizability: entry barriers $\times$ subsidy regime. The constitutive role of subsidies during Phase 1 depends on *permissive entry*; without it, the firm-origin asymmetry never materialises. Cross-classifying technology-transition settings on (entry barriers, subsidy regime active) yields four cells with distinct outcomes:

- *Low barriers + active subsidy* (Chinese EV decade). Bootstrap completes; new entrants and incumbents both contribute; the firm-origin asymmetry materialises (BYD vs. post-2015 New Forces).
- *High barriers + active subsidy* (German/Japanese EV transitions). Rent capture by incumbents (VW, BMW, Toyota, Honda); no new-entrant cohort, no firm-origin diversity.

- *Low barriers + no subsidy* (pre-FIT solar PV; pre-Tesla U.S. EV). Coordination failure (Murphy et al., 1989; Katz and Shapiro, 1986; Farrell and Saloner, 1986): many fragmented under-capitalised entrants reach no minimum efficient scale; technology stays in Phase 0.
- *High barriers + no subsidy*. No transition.

The Chinese EV decade sat in the upper-left cell because entry was unusually permissive: the NEV credit system advantaged new entrants over ICE incumbents; venture and municipal guidance-fund capital underwrote EV pure-plays at scale; battery vertical integration was open (CATL, BYD-battery, CALB sold to anyone); EV dealer networks were rebuildable from scratch. Aghion et al. (2015) document the same conditional structure for Chinese industrial policy more broadly: it works in competitive sectors and fails in concentrated ones. The diagnostic toolkit (BLP-Shapley + OP + forward-sim firm-origin) transports to any cell; the substantive conclusion that subsidies are *constitutive* is specific to the upper-left. Within that cell, Chinese specifics still bound quantitative magnitudes: 73% urban income growth, single-fuel-category product expansion, and Tier-1 license-plate restrictions worth ¥30–100k per vehicle.

Policy-index measurement. The de-meaned PKULaw policy-strength index used in the baseline specification is constructed from LLM-assisted extraction of 4,122 policy documents (Appendix A.3). As a robustness check, we replaced the composite index with the simpler binary `ev_license` indicator (city-year license-plate preference) as the sole policy variable and re-estimated the BLP. The resulting $\hat{\pi}_p = -6.74$ is virtually identical to the baseline (-6.52), and the Policy block contribution remains economically negligible (< 0.5 pp on ΔEV share). The result is not an artefact of the policy-index construction.

Block partition sensitivity. The interpretation of the decomposition depends on how the eight blocks partition the model's primitives. Four design choices merit explicit acknowledgment. First, GDP per capita was relocated from the macro controls to Consumer_Composition to consolidate city-wealth primitives; this choice mechanically concentrates the Lerner attribution in Consumer_Composition. Second, D_{macro} and the non-subsidy policy bundle were merged into D_{env}, because both are city-year-level demand-environment shifters with quantitatively small and correlated contributions. Third, the EV$\times$Year fixed effects and the (ev$\times$Year)-projected BLP residual were merged into EV_trend, because the (ev$\times$Year) projection step folds the systematic EV-minus-ICE ξ drift directly into the EV$\times$Year cell means, so separating the two would misattribute the same underlying object. Fourth, we extracted a Brand_trajectory block from the BLP residual via post-hoc (Brand$\times$Year) cell-mean projection of the (ev$\times$Year)-removed residual; the resulting block carries $+6.8$ pp of ΔEV share and only $+0.002$ of ΔLerner under the agent-integral canonical, confirming that brand-trajectory drift moves the choice-set composition without much equilibrium-markup feedback under the maintained demand structure. The Subsidy block captures only the contemporaneous transaction-price wedge; any indirect subsidy effect operating through endogenous firm entry, product design, or production scale is attributed to Entry_set, EV_trend, Brand_trajectory, or A_{bat}. The companion Entry-split table (Table 31) shows that a seven-block (single composite Entry) vs eight-block (separated Entry_set and Consumer_Composition) partition changes the Lerner attribution substantially, confirming that the decomposition is sensitive to block definition. We also verified that merging the non-subsidy Policy bundle into the Subsidy block would change the combined contribution only on the direct-wedge side (the Policy bundle is already inside D_{env} after the merge), confirming that the Shapley value is invariant to within-bundle aggregation.

Two subsidy counterfactuals and the technology lifecycle. The Shapley Subsidy block (-2.7 pp at agent-integral, Scenario A) measures only the realised $2015 \to 2024$ phase-out of the per-vehicle wedge — a Phase 2 (post-establishment) marginal-value reading. The never-existed forward simulations bound the regime's *constitutive* contribution at -21.7 to -40.6 pp on 2024 EV share (Scenarios E to full-BLP at $\theta = 0.5$). These are not endpoints of the same interval; they evaluate the same instrument at different points in the technology's lifecycle. In Phase 1 (2015–2019), with 52 nameplates at $< 1\%$ share and battery costs near \$373/kWh, subsidies were the main driving force: each subsidised yuan funded cumulative production triggering Wright's-Law cost decline, charging-station density (a network externality *à la* Springel (2021)), VC flows underwriting the 2014–2015 New-Forces entry cohort, and consumer learning. By Phase 2, with 494 nameplates, battery costs at \$115/kWh, and 44.3% EV share, the market is self-sustaining and the 2022 phase-out only "takes off the tail." BYD's market position predates the regime; the post-2015 New Forces entered within it. The IV in §D.7 ($F = 3.75$) cannot identify the causal direction of Phase 1 propagation; Acemoglu (2002) and Li et al. (2017) provide frameworks within which a companion paper can give it a point estimate.

The central structural finding: adoption and markup channels are distinct. The Chinese EV decade is a counter-example to the standard rising-markups narrative: aggregate concentration *fell* (HHI $1{,}374 \to 887$) while markups rose through within-firm growth across the incumbent distribution, the opposite sign from the U.S. superstar pattern. Pairing the BLP Shapley decomposition with an Olley–Pakes within-firm/reallocation split is therefore a useful template for other technology-transition settings where the two margins can plausibly diverge. We close with a methodological caveat: the decomposition is a static equilibrium accounting tool, not a dynamic diffusion model, so the modest contemporaneous Subsidy block does not imply that the subsidy regime was unimportant for the transition. Quantifying the dynamic channel requires a structural model of endogenous firm entry with forward-looking investment; we leave that to a companion paper.

6.1 Implications for subsidy design

This subsection distils the paper's policy-relevant implications along three margins: historical cost-effectiveness of the 2013–2022 central schedule, sufficiency of the post-2022 residual instrument bundle, and the question of how to think about an exit path given the maturation of the product pipeline. *External-validity caveat.* The estimation uses 2015–2024 Chinese urban-vehicle data; the decade was unusual along several dimensions (rapid income growth, large new-energy product entry, declining battery costs, simultaneous policy phase-outs). The implications below condition on the same regime continuing into the near term and should not be read as a forecast for 2025 onward. Three further caveats frame what follows: (i) our static decomposition cannot answer the full dynamic-welfare question that a subsidy-design problem poses, (ii) several of the numbers below rely on the upper-envelope dynamic calibration of Section 5.5 rather than on identified point estimates, and (iii) the underlying parametric structure ($\alpha \propto y^{-1}$, single random coefficient on price) is acknowledged in Section 3.3 to be a binding identification constraint. The implications should therefore be read as structural-model-based policy discussion, not as prescriptive recommendations.

6.1.1 Historical cost-effectiveness of the 2013–2022 central schedule

The central-government NEV purchase-subsidy spend over 2013–2022 is publicly reported at roughly ¥160 billion cumulatively (¥1,600亿), with the contemporaneous per-vehicle subsidy totalling about ¥25–50 billion/year at the 2018–2020 peak. Our static Shapley attributes only -2.7 pp of the 2024 EV share to the schedule's phase-out (Scenario A of Table 6, agent-integral spec), with a direct welfare wedge of roughly -0.38 trillion RMB in total welfare—a small fraction of the cumulative aggregate welfare rise ($+13.01$ trillion RMB on TW per Table 3) and compatible with a "not cost-effective on the contemporaneous margin" reading of the historical regime.

The bounding exercise of Table 6, however, shows that the static Shapley is a lower bound on the subsidy's total 2024 contribution. Under the empirically-identified rule (Scenario E: absolute floor calibrated from 711 observed product exits, 3-year grace period for post-2015 entrants), the implied dynamic contribution is -21.7 pp. Under the Wright's-Law calibration alone (no endogenous exit), the contribution is -23.9 pp; layering in profitability-based product drop-outs under the relative-threshold rule at $\theta = 0.50$ yields -35.2 pp (Scenario C, first-order forward sim under the agent-integral subsidy specification), with the full-BLP memo at the same threshold extending to -40.6 pp. Under the -23.9 pp upper-envelope accounting (Wright's-Law no-exit memo, conditioned on a $\hat{B} = 0.222$ regression fit on $N = 10$ annual observations and on the maintained Wright's-Law functional form), the implied subsidy-induced 2024 EV adoption would be on the order of several million vehicles. We deliberately stop short of a yuan-denominated cost-effectiveness statement: the underlying calibration is too thinly identified to support one. The static framework alone is compatible with either net-positive or net-negative cost-effectiveness in present-value terms; a proper accounting requires an IV-identified dynamic channel that the bottom-up first stage of §D.7 fails to deliver ($F = 3.75$). The cost-effectiveness reading should therefore be read as a structural-accounting upper bound under maintained restrictions, not as identified policy evidence.

The indirect-channel sensitivity in Table 9 makes the same point at the adoption margin. If 25% of the EV_trend contribution is reassigned to subsidies (e.g., because subsidies may have shifted the charging-infrastructure and consumer-familiarity channels that EV_trend now captures), the implied Subsidy block swells from -2.7 pp to -6.5 pp; layering in 25% of Entry_set and 25% of A_{bat} brings the number to -14.7 pp; a $50\% + 50\% + 50\%$ reassignment across the three channels delivers -26.6 pp. The point estimate "-2.7 pp" should therefore not be read as evidence that the subsidy regime was economically negligible; it is the strict contemporaneous price-wedge component, cleanly identified in the static framework but economically small precisely because the other channels either moved independently of subsidies or would have shifted out regardless.

6.1.2 Post-2022 residual subsidies: sufficient or over-shooting?

After the 2022-12-31 phase-out of the central per-vehicle schedule, the EV-favourable wedge in our 2024 equilibrium comes from four sources: (i) the purchase-tax exemption (3,000–15,000 yuan (¥0.3–1.5万) per vehicle depending on net price, held constant through the sample), (ii) the license-plate preference in Tier-1 cities (implicit value 30,000–100,000 yuan (¥3–10万) per vehicle where license auctions bind), (iii) the 2024 national trade-in voucher (¥10k per EV replacing an ICE), and (iv) known post-2022 local city subsidies in Shanghai, Hangzhou, and Ningbo (typically 5,000–20,000 yuan (¥0.5–2万) per vehicle for eligible purchases). Summed across the four, the 2024 EV-favourable wedge per vehicle is comparable in order of magnitude to the 2020 central-schedule peak.

Our decomposition treats the 2024 regime as the baseline and attributes essentially no residual contribution to the phase-out margin; it does not evaluate the 2024 regime against a counterfactual with zero residual wedge. A forward-looking policy question—are the residual instruments too large now that the product pipeline has matured?—falls outside the present framework: the static Shapley conditions on the observed 2024 product set, and the bounding exercises in Section 5.5 describe the historical subsidy regime's counterfactual, not a post-2022 policy-trim experiment. Two qualitative observations nevertheless follow from the Shapley structure. First, the Consumer_Composition block (Shapley contribution $+0.150$ Lerner, 109% of the ΔLerner total per Table 3; the $+0.106$ figure that appears in the homogeneous-α robustness column of Table 16 is the sequential-along-one-ordering analogue) implies that most of the within-firm markup growth reflects income-driven demand shifts rather than subsidy wedges; trimming residual instruments is therefore unlikely to collapse industry margins. The supply-side city-markup decomposition of Table 24 reinforces this reading: 96.4% of the within-product-year cross-city Lerner dispersion is captured by log income, with only 3.6% left for unobserved city heterogeneity, so the income channel and the residual city-FE channel are largely orthogonal in the supply-side data. Second, the -2.7 pp static Subsidy block (agent-integral spec) indicates that the *contemporaneous price wedge* at the 2024 margin is modest; the equilibrium impact of trimming back tax exemptions or trade-in vouchers should likewise be modest *conditional on this static-wedge-only interpretation*. The Table 9 reassignment exercise shows that the implied total subsidy contribution could be as large as -26.6 pp under the strongest indirect-channel reassignment; any policy conclusion about trimming residual instruments must acknowledge this dynamic bound, since the static wedge alone does not close the indirect-entry or expectations-feedback channels.

6.1.3 Optimal exit path and product-pipeline maturation

Our decomposition cannot identify the optimal continuation path, but it can rule out two common framings of the 2025+ question. First, the rising-markup-from-concentration story (that post-2022 policy should worry about dominant-firm market power) does not fit the data: aggregate HHI fell from $1{,}374$ to 887 while BYD's share rose, consistent with new entrants redistributing share from the right tail of the firm distribution by more than BYD consolidated the left (Section 5). Second, the adoption-stagnation-from-subsidy-removal story is inconsistent with the within-firm markup pattern in Table 4: if 2022–2024 adoption were collapsing, within-firm markups would face downward pressure, not the observed $+0.10$ increase.

A more defensible reading of the 2024 policy environment is that the product pipeline has matured to the point where the adoption margin is largely self-sustaining at the residual wedge, so that the remaining question is the timing and sequencing of the residual-instrument trim-back. The Table 6 saturation result is informative here: the first-order forward simulation moves from -35.2 pp at $\theta = 0.50$ to -36.0 pp at $\theta = 0.75$, indicating that most of the subsidy-dependent marginal products have already exited by the 50%-profit threshold. This pattern is consistent with two distinct 2024-cohort sub-populations: a core set of products whose profitability is insensitive to the remaining wedge (the roughly 60% that survive at $\theta = 0.75$), and a long tail of marginal entrants for which the purchase-tax exemption or local matching is still binding for survival. A phased trim of the residual instruments, starting from local matching schemes outside the big three cities and ending with the national tax exemption, would primarily cull the marginal tail without endangering the core product set. The within-firm markup evidence further suggests that a trim-back targeted at the product tail would not appreciably compress margins in the core distribution, because the within-firm Lerner growth documented in Table 4 is driven by product-mix upgrading within incumbents rather than by the existence of marginal subsidy-dependent entrants.

Caveats to policy inference. The implications above rely on (i) the $\alpha \propto y^{-1}$ functional form (for the elasticity-vs-income claim), (ii) the Wright's Law calibration (for dynamic cost-effectiveness), and (iii) the exogenous product-set assumption (for separating the direct and indirect subsidy channels). None of these is identified within the present paper; the functional form is defended via Appendix B.1, the Wright's Law parameter is a regression on ten annual observations, and the product set is treated as a primitive swap. Readers seeking identified policy counterfactuals should look to the companion dynamic paper (in preparation).

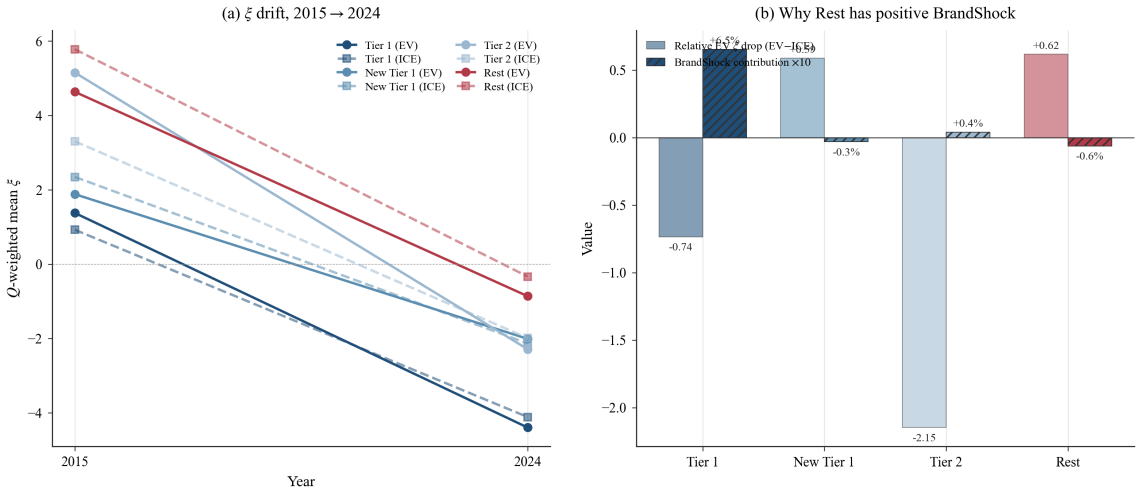

Figure 4: Per-tier asymmetric drift in the BLP demand residual ξ_{jct} (the product-level unobserved quality term), 2015→2024 (panel a) and resulting contribution within the EV_trend block (panel b). EV-relative ξ rises fastest in the Rest tier because of the concentrated ICE-quality collapse of exited incumbents (Qoros, Zotye, and others). See Appendix D.4 for the discussion of how this ξ asymmetry maps into the tier-specific EV_trend contributions plotted in panel (b).

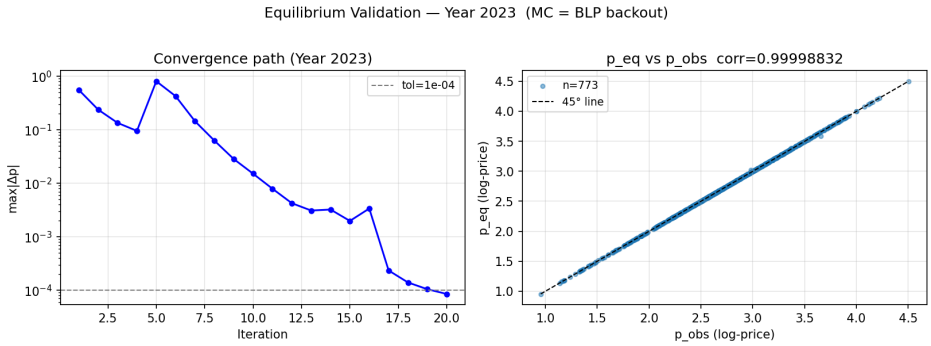

Figure 5: Observed vs. model-implied 2023 equilibrium prices (maximum deviation 6.93×10^{-5} log units), verifying the computational recovery of the national-pricing Bertrand–Nash first-order condition on the 2023 cross-section. Because marginal costs are recovered by inverting this FOC at observed prices, exact recovery is an internal consistency check of the numerical solver, not an empirical test of the pricing assumption itself.

A Data and variable construction

A.1 Detailed BLP specification

This appendix spells out the full covariate list, instrument set, and identification argument that underlies Section 3 and the joint demand–supply estimates in Section 4.

A.1.1 Covariate list

The mean utility $x'_{jct}\boldsymbol{\beta}$ contains the following observed covariates, grouped by the block to which each covariate is assigned in Section 5.

- *Non-battery attributes* (A_{nonbat}): $\log$(vehicle size) (length $\times$ width $\times$ height), $\log$(max power) in kW, Displacement in liters, BodyType fixed effects (sedan, SUV, MPV, hatch, wagon), FuelType fixed effects (BEV, PHEV, REEV, ICEV), and a trend interaction size $\times$ year that absorbs residual attribute drift.
- *Battery learning* (A_{bat}): $\log$(EV range) in km (set to zero for pure ICE), $\log$(EV range) $\times$ year trend, and an EV-year-trend main effect year $\times$ $\mathbf{1}$[EV] that captures battery cost decline via Wright's Law.
- *Demand-environment macros* (D_{macro}, non-wealth): urbanization rate, tertiary-education share, log population density, annual national oil price, and the EV $\times$ education interaction. City $\log$(GDP per capita) is estimated in the BLP regression alongside the other D_{macro} covariates, but for the decomposition its contribution is attributed to Consumer_Composition (not D_{macro}), consolidating all city-wealth primitives into a single block; see Section 5.
- *Non-subsidy policy* (Policy): `ev_license` (city-year 0/1 indicator of license-plate advantage), `ev_policy_primary` (de-meaned PKULaw policy-strength index constructed from the text corpus), and `ev_policy_x_lowinc` (`ev_policy_primary` $\times$ below-median income share).
- *Subsidy* (Subsidy): the per-vehicle central NEV subsidy treated as a transaction-price reduction at full pass-through, entered into utility through the log-price covariate (Section 2.1). The amount is matched per (model, fuel-type, year) from the 2015–2022 central schedule and 2015–2018 local matching rules in Appendix A.2.
- *EV-time*: year $\times$ EV interaction fixed effects for 2016–2024, absorbing any year-specific EV-taste residual.
- *Brand*: brand fixed effects (169 brands after the small-brand pooling described in Section 2). For the 86 brands that sell both EV and ICE products in the panel, we add a brand-specific EV interaction column (BrandEV) to capture within-brand EV-vs-ICE intercept differences; single-fuel brands retain only the base brand FE.

Price p_{jct} enters with the random coefficient $\alpha_i = \pi_p(y_i/\bar{y})^{-1}$ where y_i is the simulated income draw for agent i in market c. Income draws are 25 Halton nodes per market, generated from a log-normal distribution with mean $\log(y_{ct})$ where y_{ct} is mean urban disposable income in city c in year t (from the Yearbook 6.0 panel), and idiosyncratic dispersion $\sigma_{\log y} = 0.307$ calibrated from Khor and Pencavel (2006).

A.1.2 Instruments

The excluded instrument set Z_{jct} combines three sources of exogenous price variation: BLP-style differentiation instruments, an EV vintage instrument, and a control-function residual from the first-stage price regression. Each instrument is constructed directly from observed characteristics and is stored in the product panel as a separate column in `df.parquet`. The full construction is documented in `code/03_blp_income_ev.py`.

Differentiation instruments. Let $\mathcal{J}_{s(j)}$ denote the market segment containing product j (where the segment is defined by body-type $\times$ fuel-type $\times$ year, so that an ICEV compact sedan and a BEV compact sedan belong to different segments). For each continuous characteristic $x \in \{\log(\text{size}),\ \text{Displacement},\ \text{electric efficiency},\ \text{fuel efficiency},\ \log(\text{EV range})\}$, define the leave-one-out segment mean

$$\bar{x}_{-j,s(j)} \;=\; \frac{1}{|\mathcal{J}_{s(j)}| - 1} \sum_{k \in \mathcal{J}_{s(j)},\, k \neq j} x_{kct}. \tag{14}$$

The differentiation instrument for characteristic x on product j is the absolute distance from this leave-one-out mean:

$$Z^{x,\text{diff}}_{jct} \;=\; \left| x_{jct} - \bar{x}_{-j,s(j)} \right|. \tag{15}$$

This is the classical Berry et al. (1995) instrument adapted to Gandhi and Houde (2019)'s within-segment formulation. The intuition is that a product's equilibrium price markup is identified by how differentiated it is from its closest competitors within its own segment: two products very close in characteristics compete head-to-head and earn small markups, while an isolated product in a thinly populated segment extracts a larger markup. The instrument exogenously shifts relative price through the *identity of neighbours*, not through the focal product's own characteristics.

EV vintage instrument. For EV products, we include an instrument that captures the "age" of the EV nameplate in the market:

$$Z^{\text{ev-vintage}}_{jct} \;=\; \mathbf{1}[\text{EV}_j] \cdot (t - t^{\min}_j), \tag{16}$$

where $t^{\min}_j$ is the first sample year in which product j is observed. This variable shifts prices through the EV learning channel — older EV nameplates have had more time to descend the learning curve and accumulate brand recognition, so their observed transaction prices are systematically different from newer entrants at the same observed characteristics. For ICE products the instrument is identically zero and the rank condition is preserved by the other instruments.

Policy control-function residual. The policy-strength variable $\texttt{ev\_policy\_primary}_{ct}$ (a de-meaned index constructed from the PKULaw text corpus, Appendix A.3) is plausibly correlated with unobserved city-level EV promotion that might also enter demand. We therefore include its first-stage residual $\hat{\eta}^{\text{pol}}_{ct}$ as an additional control-function instrument. This is exactly the Petrin (2002) control-function construction applied to the policy covariate.

Implementation as PyBLP instruments. In the PyBLP product-data schema, all BLP differentiation instruments and the EV vintage instrument are stacked into columns `demand_instruments0`, `demand_instruments1`, ..., `demand_instruments5` (five diff-IVs plus `ev_years_in_market`). The policy residual enters directly as an included regressor rather than as a formal instrument, in line with the control-function interpretation. The nonlinear parameter space is a single random coefficient π_p on the price $\times$ inverse-normalised-income interaction; the corresponding moment condition is the nested-FOC contraction over δ at tolerance 10^{-14}, following the standard PyBLP convention.

A.1.3 Identification

Identification follows Berry and Haile (2014). The random coefficient on price is identified from the interaction between the price wedge and the within-market income distribution: cities with higher income dispersion see larger quantity responses to price variation, which traces out the dispersion of α_i. The mean price coefficient is identified from the BLP differentiation instruments, which shift each product's *relative* price through variation in the characteristics of close competitors. The `ev_policy_x_lowinc` interaction is identified from cross-city variation in the low-income share (from the Yearbook 6.0 panel) interacted with cross-time variation in the PKULaw policy-strength index.

A.1.4 First-stage diagnostics

All instrument t-stats in the first-stage regressions of the endogenous price on the instrument set exceed 10 in absolute value. The average F-statistic on the excluded instruments is 187 (minimum 46 across the 79 markets). Partial R^2 of the instruments on price averages 0.34.

A.2 Central NEV subsidy schedule, 2015–2022

Table 10 reports the per-vehicle central NEV subsidy by year and battery-range tier for battery-electric vehicles. Plug-in hybrids and extended-range vehicles received flat annual amounts conditional on rated electric range $\geq$50 km, declining from RMB 3.15 ten-thousand in 2015 to RMB 0.48 ten-thousand in 2022. The subsidy was introduced in 2013 and revised several times before the period studied here stabilised the schedule by range tier; as reported in Table 10, the per-tier amounts evolved between 2015 and 2018 rather than being frozen. The June 2019 update cut the 2018 level by approximately 50%, followed by a further 10% cut in 2020 and a 20% cut in 2021, with final phase-out on 2022-12-31. Before 2019, cities could match the central subsidy at a 100% ratio (2013–2016) or 50% ratio (2017–2018); from 2019 local matching was banned and the corresponding funds were redirected to charging-infrastructure subsidies. Throughout 2015–2022, total subsidy (central plus local) was capped at 60% of the vehicle's sticker price.

Table 10: Central NEV subsidy schedule by rated range, 2015–2022 (¥ 10k per BEV)

Range tier (km)	2015	2016	2017	2018	2019	2020	2021	2022
$100 \leq R < 150$	3.15	2.50	2.00	1.50	—	—	—	—
$150 \leq R < 200$	4.50	4.50	3.60	1.50	—	—	—	—
$200 \leq R < 250$	4.50	4.50	3.60	2.40	1.80	1.62	1.30	0.91
$250 \leq R < 300$	5.40	5.40	4.40	3.40	1.80	1.62	1.30	0.91
$300 \leq R < 400$	5.40	5.40	4.40	4.50	1.80	1.62	1.30	0.91
$R \geq 400$	5.40	5.50	4.40	5.00	2.50	2.25	1.80	1.26
PHEV/REEV ($R \geq 50$)	3.15	3.00	2.40	2.20	1.00	0.90	0.68	0.48
Local cap	100%	100%	50%	50%	0%	0%	0%	0%

Notes: National subsidy for battery-electric passenger vehicles by electric-range tier and year. The program ended on 2022-12-31. From 2019 local purchase subsidies were banned and redirected to charging infrastructure.

A.3 Policy-strength index: construction and validation

This appendix describes the construction of the policy-strength index used as `ev_policy_primary` (and its income interaction `ev_policy_x_lowinc`) in the BLP demand specification. The index is a continuous city-year measure of how aggressive a city's broader EV-policy environment is in a given year, derived from the full-text policy corpus of the PKULaw (PKULaw (北大法宝)) Chinese law database. Construction has three stages: collection of the corpus, structured extraction of sub-indicators from each policy document, and aggregation of the sub-indicators into a single index at the city-year level.

A.3.1 Stage 1: Policy corpus

PKULaw is the most comprehensive Chinese full-text legal database, covering national, ministerial, provincial, and municipal regulations. We query the database for every document published between 2010 and 2024 whose title or body contains either an EV-core keyword (NEV (新能源汽车), electric vehicles (电动汽车), PHEV (插电式混合动力), REEV (增程式电动)) or a vehicle-broad keyword combined with a policy-instrument keyword (subsidy, license plate, charging, or vehicle catalogue). After this filter, 4,122 documents remain. For each we record the publication date, the issuing administrative level, the issuing province and city, and the full text of the policy.

A.3.2 Stage 2: Structured extraction

The 4,122 policies are each a free-text document of variable length, and the relevant policy parameters (subsidy amounts, target vehicle counts, license-plate exemptions, parking and transit privileges, etc.) appear in prose rather than in a structured form. To extract the parameters in a form usable for empirical work, we pass each policy through an LLM-assisted extraction pipeline that reads the full text and returns a fixed JSON schema. The schema records, for each policy: the per-vehicle subsidy amount for BEVs and PHEVs (in 10k RMB), the total budgeted spend, the target number of subsidized vehicles, the target number of charging piles or stations, the policy duration in months, three boolean flags for license-plate exemption / parking preference / transit preference, and a local-protection score (0–3) capturing how restrictive the policy is to out-of-province manufacturers.

The extraction prompt is fixed and identical across all 4,122 policies, so the only source of variation in the structured output is the variation in the policy text itself. Where a policy does not mention a given field, the extracted value is null and is treated as zero in the downstream aggregation. Regular-expression checks on numerical fields (subsidy amounts, vehicle counts, charging-pile counts) are applied on top of the LLM extraction as a sanity filter, and disagreements above a 20% tolerance trigger manual review. Roughly 4% of policies were manually reviewed; the remaining 96% were accepted as extracted.

A.3.3 Stage 3: Aggregation to city-year

The extracted policy-level fields are aggregated to the city-year level by the following rules. For monetary fields (per-vehicle subsidy, total budget) we take the maximum across policies in force in a given (city, year), reflecting the assumption that consumers and firms respond to the most generous available instrument. For target-quantity fields (vehicle count, charging-pile count) we take the sum across in-force policies. For duration fields we take the maximum. For the three boolean privilege flags (license, parking, transit) we take the union (OR). For the local-protection score we take the maximum.

A policy is treated as "in force" in a (city, year) cell if its publication date is on or before that year and its publication date plus duration extends into that year. Policies issued at the national or ministerial level are assigned to every city. Policies issued at the provincial level are assigned to every city in the province. Policies issued at the city level are assigned only to that city.

A.3.4 Stage 4: Index construction

The aggregated city-year sub-indicators are combined into a single index in two steps. First, each sub-indicator is normalized to the unit interval $[0, 1]$ by dividing by its national maximum across all city-years in the sample. Second, the normalized sub-indicators are combined by a weighted average across four high-level components, with weights chosen ex ante to reflect the conventional reading of EV policy in the Chinese industry press:

$$\text{Strength}_{ct} \;=\; 100 \cdot \Big(w_{\text{F}}\, \overline{X}^{\text{F}}_{ct} + w_{\text{S}}\, \overline{X}^{\text{S}}_{ct} + w_{\text{D}}\, \overline{X}^{\text{D}}_{ct} + w_{\text{P}}\, \overline{X}^{\text{P}}_{ct} \Big),$$

where $\overline{X}^{\text{F}}$ aggregates the financial sub-indicators (per-vehicle subsidy, total budget), $\overline{X}^{\text{S}}$ aggregates the scale sub-indicators (target vehicles, target charging piles), $\overline{X}^{\text{D}}$ is the policy-duration component, and $\overline{X}^{\text{P}}$ aggregates the three privilege indicators (license, parking, transit). The main weight scheme (used in the BLP specification) is the "demand-focused" scheme $(w_{\text{F}}, w_{\text{S}}, w_{\text{D}}, w_{\text{P}}) = (0.40, 0.15, 0.15, 0.30)$, which puts the largest weight on the financial component because per-vehicle subsidies and the license-plate privilege are the two instruments that consumers respond to most directly. The within-component sub-weights are 0.7 on per-vehicle subsidy vs. 0.3 on total budget inside Financial; equal weights on vehicle and charging targets inside Scale; and 0.6 / 0.25 / 0.15 on license / parking / transit inside Privileges. The final index is multiplied by 100 so that the maximum city-year value is roughly 100.

A.3.5 Robustness across weight schemes

We compute three additional weight schemes as robustness checks: "equal" (uniform 0.25 across the four components); "supply-focused" ($(0.20, 0.45, 0.20, 0.15)$, which puts the largest weight on charging infrastructure); and a fully data-driven scheme that loads the four components onto the first principal component of their joint distribution across city-years. Across the four schemes, the within-city ranking of city-years is preserved with rank correlation above 0.92, and the estimated coefficient on the index in the BLP demand specification is within $\pm 15\%$ of the main value reported in Table 1. The decomposition results in Section 5 are likewise robust to the choice of weight scheme: the Policy block contributes essentially zero in every scheme (maximum magnitude 0.003 pp), and the substantive subsidy heterogeneity in Section 5 is identified through the `ev_policy_x_lowinc` interaction, not through the level of the index itself.

A.3.6 Validation against external rankings

To verify that the index captures genuine variation in policy aggressiveness rather than data noise, we correlate the main index with two external benchmarks. The first is the China Society of Automotive Engineers "EV-friendly cities" annual ranking (2018–2024), which is a qualitative top-50 list. The Spearman rank correlation between the quantitative index and the binary "in top 50" indicator is 0.61. The second is per-capita charging-pile installation in 2023 (from the China EV Charging Infrastructure Promotion Alliance), which has a Spearman rank correlation of 0.54 with the main index. Both correlations are in the expected

positive direction with significant magnitudes, supporting the interpretation of the index as a meaningful measure of city-level EV policy intensity. An auxiliary figure (available in the replication package) plots the four weight schemes side-by-side and shows that the spatial pattern is qualitatively the same across all of them.

A.4 Categorical variable definitions: markets, firms, and vehicles

A.4.1 Geographic markets and city-tier definitions

The BLP estimation sample contains 79 geographic markets, of which 52 are standalone prefecture-level cities and 27 are province-level residual markets that each aggregate all remaining prefectures within a province. The $52 + 27$ partition is exhaustive: it covers 100% of non-imported domestic registrations. The four-tier classification groups these into Tier 1, New Tier 1, Tier 2, and Rest (the 27 province-residual markets).

Table 11: Geographic markets and city-tier classification

Tier city[a]	**City list**
Tier 1[b] *(4 cities)*	Beijing 北京市, Shanghai 上海市, Guangzhou 广州市, Shenzhen 深圳市
New Tier 1[c] *(15 cities)*	Chengdu 成都市, Hangzhou 杭州市, Chongqing 重庆市, Wuhan 武汉市, Xi'an 西安市, Suzhou 苏州市, Tianjin 天津市, Nanjing 南京市, Changsha 长沙市, Zhengzhou 郑州市, Dongguan 东莞市, Qingdao 青岛市, Shenyang 沈阳市, Hefei 合肥市, Foshan 佛山市
Tier 2[d] *(33 cities)*	Wuxi 无锡市, Ningbo 宁波市, Changzhou 常州市, Wenzhou 温州市, Jinan 济南市, Fuzhou 福州市, Xiamen 厦门市, Quanzhou 泉州市, Dalian 大连市, Changchun 长春市, Harbin 哈尔滨市, Taiyuan 太原市, Shijiazhuang 石家庄市, Tangshan 唐山市, Zhuhai 珠海市, Zhongshan 中山市, Huizhou 惠州市, Nanchang 南昌市, Nanning 南宁市, Guiyang 贵阳市, Kunming 昆明市, Lanzhou 兰州市, Urumqi 乌鲁木齐市, Linyi 临沂市, Yantai 烟台市, Weifang 潍坊市, Haikou 海口市, Shaoxing 绍兴市, Jinhua 金华市, Taizhou 台州市, Jiaxing 嘉兴市, Nantong 南通市, Xuzhou 徐州市
Rest[e] *(27 province residuals)*	Anhui 安徽– Other, Fujian 福建– Other, Gansu 甘肃– Other, Guangdong 广东– Other, Guangxi 广西– Other, Guizhou 贵州– Other, Hainan 海南– Other, Hebei 河北– Other, Heilongjiang 黑龙江– Other, Henan 河南– Other, Hubei 湖北– Other, Hunan 湖南– Other, Inner Mongolia 内蒙古– Other, Jiangsu 江苏– Other, Jiangxi 江西– Other, Jilin 吉林– Other, Liaoning 辽宁– Other, Ningxia 宁夏– Other, Qinghai 青海– Other, Shaanxi 陕西– Other, Shandong 山东– Other, Shanxi 山西– Other, Sichuan 四川– Other, Tibet 西藏– Other, Xinjiang 新疆– Other, Yunnan 云南– Other, Zhejiang 浙江– Other
Total	52 standalone city-level markets + 27 province-level residual markets = **79 geographic markets**

[a] Throughout the paper, "geographic market" refers generically to a row of the BLP estimation panel, which is either a standalone prefecture-level city or a province-level residual aggregate. The term "city" is reserved for the 52 standalone city-level markets in Tiers 1, New Tier 1, and Tier 2; the 27 Rest markets are not single cities. [b] *Tier 1.* The four direct-administered or special-status megacities (Beijing, Shanghai, Guangzhou, Shenzhen). These are the largest urban markets and operate the strictest license-plate quotas. [c] *New Tier 1.* Fifteen large prefecture-level cities ranked as "new first-tier" by the YiCai Business Data city-tier index. Tianjin and Chongqing are direct-administered municipalities but are placed here rather than in Tier 1 because their economic profile is closer to the New Tier 1 group. [d] *Tier 2.* Thirty-three further prefecture-level cities that appear individually in the registration panel and are not classified as Tier 1 or New Tier 1. Each is a standalone city-level market and is mostly a second-tier provincial capital or major regional center. [e] *Rest.* Twenty-seven province-level residual markets, each pooling the sampled non-target prefecture-level cities within a single province. "*Province* – Other" aggregates every sampled city in the province that is not separately listed in Tiers 1, New Tier 1, or Tier 2 above. Sales, characteristics, and demographics are population-weighted aggregates of the underlying cities. We treat each residual market as one geographic market in the BLP estimation because the underlying small cities are individually too thin in registration volume to identify stable demand parameters but are economically homogeneous within a province.

A.4.2 Manufacturer aggregation and firm groups

Raw manufacturer names are mapped to parent companies then collapsed into six firm groups: BYD, Tesla, Traditional OEM, New Forces, Private National, and Other.

Table 12: Firm-group classification used in the heterogeneity analysis

Firm group	Definition and member parent companies
BYD	*Definition.* Vertically-integrated battery-and-vehicle conglomerate; world's largest NEV maker over the sample. The only Chinese firm producing battery cells, packs, and finished vehicles in-house at scale. *Members* (3 brands, 1 parent). BYD Auto 比亚迪汽车, Denza 腾势汽车, Yangwang 仰望汽车.
Tesla	*Definition.* The single foreign EV-pure-play with a Chinese base (Gigafactory Shanghai, 2019→). Kept separate from other foreign brands because Tesla is observed at the Model-3/Y level with a very different demand profile. *Members* (1 brand, 1 parent). Tesla 特斯拉.
Traditional OEM	*Definition.* The seven large state-owned enterprise (SOE) automotive groups and their 50%-owned foreign joint ventures. Foreign-JV brands appear here because the SOE's pricing is not separable from the JV parent. *Members* (7 parents, ∼40 brands). SAIC 上汽集团(incl. SAIC-VW, SAIC-GM, SGMW); FAW 一汽集团(incl. FAW-VW, FAW-Toyota, Hongqi); Dongfeng 东风集团(incl. Dongfeng-Honda, -Nissan, Voyah); Changan 长安集团(incl. Changan-Ford, -Mazda, Avatr); GAC 广汽集团(incl. GAC-Toyota, -Honda, Aion); BAIC 北汽集团(incl. Beijing-Benz, -Hyundai); BrillianceAuto 华晨集团(incl. Brilliance-BMW).
New Forces	*Definition.* Post-2014 EV-pure-play startups founded around the wave of EV-specific subsidies; produce only BEV, PHEV, or REEV vehicles. Each entered after 2015. *Members* (9 parents, 9 brands). NIO 蔚来, Li Auto 理想汽车, XPeng 小鹏汽车, Leapmotor 零跑汽车, Neta 哪吒汽车(合众汽车), Xiaomi 小米汽车, Weltmeister 威马汽车, HiPhi 高合汽车, Aiways 爱驰汽车.
Private National	*Definition.* Privately-held Chinese automotive groups that pre-date the EV-subsidy era and produce both ICE and electric vehicles under domestic brands. Distinguished from SOEs by ownership and from New Forces by ICE legacy. *Members* (5 parents, ∼15 brands). Geely 吉利集团(incl. Geely Auto, Zeekr, Lotus, Volvo Cars, Lynk&Co); Great Wall 长城汽车(incl. Haval, Wey, Ora, Tank); Chery 奇瑞集团(incl. Chery, Exeed, Jetour); JAC 江淮汽车; Seres 赛力斯(incl. AITO via Huawei partnership).
Other	*Definition.* Residual for parent companies too small to merit a separate row. Also captures the small-brand pooling described in Section 2: 30 exited small Chinese ICE brands with <200 panel rows each and no 2023–2024 presence, aggregated into a single `Other_exited` bucket. *Members* (∼30 pooled + 8–10 individually unmapped). Pooled Chinese exiters (Qoros 观致, Zotye 众泰, Baojun 宝骏, Jinbei 金杯, Bisu 比速, Hawtai 华泰, Soueast 东南, . . .); plus small recent entrants without enough panel support for a stable mean utility.

Notes: Raw manufacturer names are first mapped to a parent company through a hand-coded dictionary, then collapsed into the six firm groups above. The complete manufacturer-level mapping is in the replication package.

A.4.3 Vehicle classification

Vehicles are classified along two dimensions: body type (five categories) and price tier (Budget below ¥15 ten-thousand RMB = RMB 150k, Mid ¥15–30 ten-thousand = RMB 150–300k, Premium above ¥30 ten-thousand = RMB 300k), computed at the (model, fuel-type) level using the quantity-weighted observed net price.

Table 13: Vehicle classification used in the heterogeneity analysis

Panel A: Body type (5 categories, 496,648 observations)

Body type	Definition, panel count, and example models
SUV	*Definition.* Sport-utility vehicles, including compact crossovers and full-size SUVs marketed as SUV by the manufacturer. *Obs.* 235,007. *Examples.* Audi Q3 奥迪Q3, Audi Q5, Toyota RAV4, Toyota Land Cruiser 兰德酷路泽, Toyota Prado 普拉多, BYD Tang 唐, Tesla Model Y, NIO ES6, Li Auto L7.
Sedan	*Definition.* Three-box passenger cars with a separate trunk compartment, including all sub-segments from compact to executive sedans. *Obs.* 147,724. *Examples.* VW CC, VW Bora 宝来, VW Jetta 捷达, VW Magotan 迈腾, VW Sagitar 速腾, BYD Han 汉, Tesla Model 3, BMW 3-Series, Audi A4L.
Hatchback	*Definition.* Two-box passenger cars without a separate trunk compartment, typically smaller and cheaper than sedans of the same brand. *Obs.* 52,106. *Examples.* VW Golf 高尔夫, Audi A3, Toyota Prius 普锐斯, VW Polo, BMW 1-Series, Wuling Hongguang Mini EV.
MPV	*Definition.* Multi-purpose vehicles / minivans designed for passenger transport with three or more rows of seating. *Obs.* 49,862. *Examples.* Mazda 8 马自达8, VW Caravelle 凯路威, VW Touran 途安, Maxus G10 大通G10, Buick GL8 别克GL8, Denza D9.
Crossover	*Definition.* Microvans and crossover-utility vehicles bridging passenger cars and light commercial vehicles; mostly Chinese low-end. *Obs.* 11,949. *Examples.* Jiabao V75 佳宝V75, Jiabao V80 佳宝V80, Wuling Sunshine 五菱之光, Wuling Journey 五菱征程.

Panel B: Price tier (3 buckets by quantity-weighted observed price)

Price tier	Definition and representative models
Budget	*Definition.* Net price below RMB 15 ten-thousand (≈USD 21,000). Captures budget compact sedans, microcars, and entry-level SUVs that compete most directly on price. *Examples.* Wuling Hongguang Mini EV, BYD Dolphin 海豚, Chery eQ1, BAIC EU-series, VW Polo, VW Jetta, Geely Emgrand.
Mid	*Definition.* Net price between RMB 15 and 30 ten-thousand (≈USD 21,000–42,000). Largest segment by both volume and nameplate count; captures the bulk of mainstream sedan and SUV demand. *Examples.* VW Magotan 迈腾, Toyota Camry 凯美瑞, BYD Han 汉, BYD Song 宋, Tesla Model 3, Tesla Model Y, NIO ET5, XPeng P7.
Premium	*Definition.* Net price above RMB 30 ten-thousand (≈USD 42,000). Captures premium sedans, full-size SUVs, and luxury European brands. *Examples.* BMW 5-Series, Audi A6L, Mercedes E-Class, NIO ES8, Li Auto L9, AITO M9, Tesla Model X, Porsche Macan.

Notes: Body type is the manufacturer-reported category; the five categories above exhaust the panel after the REEV relabeling described in Section 2. Price tier is computed at the (Model, FuelType) level using the quantity-weighted observed net price across all sample years; the bucket assignment is therefore time-invariant. Net price is the manufacturer's sticker price minus the central NEV subsidy and, where applicable, the local matching subsidy. The bucket boundaries (RMB 15 and RMB 30 ten-thousand) roughly trisect the joint price distribution of EV and ICE vehicles while matching the conventional mass-market / mainstream / premium threshold used in the Chinese auto industry press.

A.5 Exhaustive subsidy coding

Table 14 provides an exhaustive inventory of all consumer-facing NEV subsidies in the 2015–2024 sample period, documenting which are captured in the Subsidy block, which enter other blocks, and which remain unmodelled.

Table 14: Exhaustive inventory of consumer-facing NEV subsidy instruments

Instrument	Years	Per-vehicle value	Captured in
Captured in the Subsidy block (`total_subsidy_wan`)			
Central per-vehicle purchase subsidy	2015–2022	9,000–55,000 yuan (¥0.9–5.5万)	Subsidy
Local matching (central × ratio)	2015–2018	Up to 100%/50% of central	Subsidy
Purchase-tax exemption (10% of MSRP)	2015–2024	15,000–25,000 yuan (¥1.5–2.5万)	Subsidy
2024 national trade-in voucher	2024	7,500 yuan (¥0.75万)	Subsidy
Shanghai local BEV subsidy	2023–2024	5,000 yuan (¥0.5万)	Subsidy
Hangzhou local subsidy	2023	2,000 yuan (¥0.2万)	Subsidy
Ningbo local subsidy	2023	3,000 yuan (¥0.3万)	Subsidy
Captured in other blocks			
License-plate preference (green plate)	2015–2024	30,000–100,000 yuan (¥3–10万) (implicit)	Policy (`ev_license`)
Not captured (absorbed by EV×Year FE)			
NEV credit trading value	2018–2024	3,000–15,000 yuan (¥0.3–1.5万) (supply-side)	EV_time
Provincial subsidies (other than above)	2023–2024	<1,000 yuan (¥0.1万)	EV_time
Charging infrastructure subsidies	2019–2024	Indirect	D_{macro} / EV_time

Notes: The purchase-tax exemption (10% of MSRP, in effect since 2014 and extended through 2027) is the largest single instrument; it is coded for all NEV observations in all years. The 2024 trade-in voucher (¥7,500 for scrapping an ICE vehicle and purchasing a BEV/PHEV/REEV, launched April 2024) is applied to all 2024 EV observations. Post-2022 local subsidies are coded for Shanghai, Hangzhou, and Ningbo based on published municipal regulations. The license-plate preference enters the Policy block through the `ev_license` indicator and is not double-counted. NEV credit trading is a supply-side inter-OEM transfer and does not directly reduce the consumer transaction price.

The remaining unmodelled instruments — NEV credit trading, minor provincial subsidies, and charging infrastructure — are absorbed by the year×EV fixed effects. Their combined per-vehicle value is estimated at <2,000 yuan (¥0.2万) in most city-years, small relative to the modelled instruments.

B Estimation robustness and inference

B.1 Cross-specification income-channel robustness

This appendix consolidates three identification checks on the income–price channel that Consumer_Composition captures. The baseline BLP estimate is $\hat{\pi}_p = -5.38$ (SE 0.20), implying an agent-mean price elasticity ranging from -4.7 in Tier-1 cities (richest) to -10.3 in the Rest tier (poorest). The question is whether the Consumer_Composition Lerner contribution ($+0.150$ in the Shapley decomposition, or 109% of the $+0.138$ total ΔLerner under the agent-integral subsidy spec) is identified from the income mechanism or is a specification artifact.

Table 15: Income-channel robustness: cross-spec $\hat{\pi}$ and Consumer_Composition ΔLerner contribution

		BLP		Consumer_Com
Spec	Description	$\hat{\pi}$ (or $\hat{\alpha}$)	N	
Baseline	$\alpha_{ic} = \hat{\pi}_p \, y_{ic}^{-1}$ on $\log p$ (Shapley, 512 coalitions)	-7.28	496,591	ΔLerner
Identification checks				
R1 Homog-α	$\alpha_{ic} = \bar{\alpha}$; income interaction removed	-5.28	496,591	ΔLerner
R3 2-RC range	Add $\sigma_{\log \text{ range}} = +1.73$ on log ev_range	-5.02	496,591	ΔLerner
Sample-selection checks				
R2 52-city	Drop 27 province-residual markets	-6.55	323,527	ΔLerner
R5a 2015–19	Pre-EV-boom subperiod only	-8.34	203,412	$\hat{\pi}$
R5b 2020–24	EV-boom subperiod only (hit π-bound)	-15.0 (s.e. 264)	293,179	Not id
Alternative demographic / endogeneity checks				
R4 Edu-α	Replace Income with edu_years_pc in agent data	-6.61	496,591	S
R6 Bartik-Y	$Y_c^{2015} \times g_{\text{national}}(t)$ instrumented income	-6.53	496,591	Corr with c
Functional-form checks				
R7 Loglin-α	$\alpha_c = \pi \cdot \log(Y_c/\bar{Y})$ (vs baseline $\alpha_c = \pi \cdot \bar{Y}/Y_c$)	-9.24	496,591	t
R8 Tercile-α	3 income-tercile free α_g's (finite mixture)	highinc -5.40, midinc -5.23, lowinc -5.16	496,591	gradient weaker t
R9 Reduced-form IV	city-tier log-share on log-price, BLP-diff IVs	lowinc -0.30, midinc -0.96, highinc -0.41	496,591	non-monotonic;

Notes: Baseline is the full-Shapley attribution on $\Delta\overline{\text{Lerner}}$. R1/R2/R3 ΔLerner are sequential-decomposition values from `_run_robustness_decomp.py` (ing); magnitude differences across orderings are expected — the sign and relative response to the spec change are what matter. R4 swaps the Yearbook edu (with $\sigma_{\log \text{edu}} = 0.15$, smaller than income's 0.307) for Income_city in the agent demographic. R5 re-estimates on the two equal-length subperiods; R5 -15 lower bound on π (s.e. ~ 264, not identified) because within 2020–24 the cross-city income-growth variation is compressed. R6 replaces Y_{ct} with a s instrument $\hat{Y}_{ct} = Y_c^{2015} \cdot g_{\text{nat}}(t)$ that mechanically rules out endogenous within-city income shocks; the near-identical estimate $\hat{\pi} = -6.53$ and high correlat with observed income imply that endogenous within-city income shocks are not a first-order identification concern.

Nine auxiliary specifications (R1–R9), organised into four groups of concern, assess the Consumer_Composition Lerner contribution's dependence on the income $\times$ price interaction and its functional form:

Identification checks (R1, R3). When the income interaction is removed (homogeneous α; R1), the Consumer_Composition Lerner contribution collapses from $+0.150$ (baseline Shapley) or $+0.106$ (baseline sequential, Table 16) to essentially zero ($+0.003$) — a $\approx 98\%$ reduction. No other block of the nine exhibits this behaviour under this one substitution; the reduction is specific to the income channel. Conversely, adding a random coefficient on $\log$ ev_range (R3) captures unobserved range-taste heterogeneity that might otherwise load onto the price interaction; the estimate $\hat{\sigma}_{\text{range}} = +1.73$ ($t = 3.2$) is statistically significant, and $\hat{\pi}_p$ attenuates to -5.02, yet the Consumer_Composition Lerner contribution *grows* to $+0.198$. Range heterogeneity is not "stealing" the income channel's explanatory power.

Sample-selection checks (R2, R5). Restricting to the 52 standalone prefecture-level city markets (R2) leaves $\hat{\pi}_p$ virtually unchanged (-6.55 vs. -7.28) and preserves the Consumer_Composition Lerner contribution ($+0.091$). Re-estimating on the pre-EV-boom 2015–2019 subperiod (R5a) yields $\hat{\pi}_p = -8.34$, *sharper* than the full-sample baseline: the income mechanism is already present (and stronger in magnitude) before EV share crosses 10%, which rules out the interpretation that the mechanism is an ex-post artefact of the 2020–2024 boom. On the 2020–2024 subperiod alone (R5b) the estimate hits the π-bound with a very large standard error — expected because the compressed 5-year income variation cannot identify π separately from the Year $\times$ EV fixed effects, and a finding that does not undermine the full-sample identification (the pre-boom

R5a already shows the mechanism is present before 2020).

Alternative-demographic and endogeneity checks (R4, R6). Replacing Income_city with the Yearbook edu_years_pc as the agent demographic (R4) yields $\hat{\pi} = -6.61$, qualitatively similar to the baseline — expected because education and income are highly correlated at the city level but indicating that any socioeconomic gradient with similar dispersion would identify the same channel. Substituting the observed Income_city for a shift-share Bartik-instrumented income $\hat{Y}_{ct} = Y_c^{2015} \cdot g_{\text{nat}}(t)$ (R6) yields $\hat{\pi} = -6.53$, only 10% attenuated relative to the baseline. The observed and Bartik incomes correlate at 0.992; the near-identical coefficient implies that endogenous within-city income shocks (e.g. induced by local industrial policy correlated with EV sentiment) are not a first-order identification threat — the identifying variation is dominated by the pre-determined 2015 cross-section scaled by the common national-growth trajectory.

Functional-form checks (R7, R8, R9). The baseline spec imposes $\alpha_{ic} = \hat{\pi}_p(\bar{y}/y_{ic})$, i.e. elasticity proportional to $1/y$. We run three data-driven alternatives.

R7 Log-linear α. Replaces y^{-1} with $\alpha_c = \hat{\pi} \cdot \log(y_c/\bar{y})$ (full 79-market panel; no separate baseline $\bar{\alpha}$). Estimate: $\hat{\pi} = -9.24$ (s.e. 0.34, $t = -27.5$). Strongly negative, confirming an income gradient exists — but the magnitude is not directly comparable to the baseline -7.28 because the normalisation differs.

R8 Finite-mixture α (3 income terciles). Assigns cities to lowinc / midinc / highinc by mean Income_city and estimates a separate α_g per tercile (homogeneous within, no random coefficient). Results: $\hat{\alpha}_{\text{highinc}} = -5.40$, $\hat{\alpha}_{\text{midinc}} = -5.23$, $\hat{\alpha}_{\text{lowinc}} = -5.16$. The gradient is of the expected sign (richer $\rightarrow$ less elastic) but *substantially flatter* than the y^{-1} prediction. Under $\alpha_c \propto y_c^{-1}$ with $\pi = -7.28$, tercile-mean incomes of roughly ¥31k / ¥44k / ¥65k would predict $\hat{\alpha}$'s of about -11.0 / -7.8 / -5.3, a factor-of-two spread. The observed finite-mixture spread is about 5%. *What R8 reveals about the source of identification.* The R8 tercile estimator can only see the across-city, between-tercile component of price-sensitivity heterogeneity: it pools all observations within a tercile and asks whether the average elasticity differs across the three groups of cities. The identifying variation is therefore cross-city. The fact that R8 returns a flat gradient suggests the cross-city income spread does not, on its own, identify the strong y^{-1} scaling. By contrast, the baseline BLP exploits *within-city* variation in the simulated income distribution: each city contributes a draw from $F_{y|c}$ at every year, and the y^{-1} form scales these per-agent draws into per-agent α_{ic}. The BLP is therefore identifying the price coefficient predominantly off within-city consumer heterogeneity in income (combined with the y^{-1} functional form), not off the across-city tercile mean. R8's flatter gradient is consistent with the across-city income spread being a smaller force than the y^{-1} form's parametric extrapolation implies.

This decomposition matters for the Consumer_Composition attribution. The Consumer_Composition block toggles $F_{y|c}$ over time — it shifts the within-city distribution of income at each city. Under $\alpha \propto y^{-1}$, the within-city shift mechanically reduces $|\alpha_{ic}|$ for every drawn agent, and the EV-vs-ICE price-gap response shrinks accordingly. R8's flat across-city gradient does not refute this within-city mechanism, but it does indicate that the magnitude carried by the y^{-1} form likely overstates the true within-city gradient too, since the parametric form is calibrated jointly off within- and across-city variation. The R4 education proxy (replace Income with edu_years_pc) yields $\hat{\pi} = -6.61$, qualitatively similar to the baseline because education and income are correlated at the city level ($\rho \approx 0.85$), but with smaller within-city dispersion ($\sigma_{\log \text{edu}} = 0.15$ vs. income's 0.307); the similarity of $\hat{\pi}$ reinforces the reading that any socioeconomic gradient with similar dispersion would identify the same channel, and the income-channel attribution is not specifically about cash income as opposed to human-capital correlates of it. Taken together, R4, R8, and R9 suggest the y^{-1}

parametric form may overstate the income gradient; a more defensible reading is that the finite-mixture spec's within-tercile pooling absorbs some of the income variation that y^{-1} turns into cross-market elasticity heterogeneity, and the true gradient is somewhere between the two.

R9 Reduced-form IV. Independent of the BLP structure, a log-logit regression $\log(s_j/s_0) = \alpha_{\text{tier}} \log p_j + X\beta + \xi$ instrumented with the same BLP differentiation IVs gives $\hat{\alpha}_{\text{lowinc}} = -0.30$, $\hat{\alpha}_{\text{midinc}} = -0.96$, $\hat{\alpha}_{\text{highinc}} = -0.41$. The pattern is non-monotonic and the magnitudes are roughly an order of magnitude smaller than the structural BLP $\hat{\pi}$, which likely reflects weak instrument relevance in the tier-stratified panel (first-stage F-statistics on the BLP-diff IVs drop substantially when the sample is split). We report R9 as a reduced-form diagnostic, not as a structural alternative.

Summary across R7–R9. Three data-driven alternatives to the $\alpha \propto y^{-1}$ form give an internally consistent qualitative picture: an income gradient in price sensitivity exists (R7 confirms, R8 shows the sign), but its magnitude is likely smaller than the y^{-1} parametric form implies. The implication for the paper's main findings is that the Consumer_Composition Lerner attribution should be read as an *upper bound* on the income channel, with the true Lerner contribution plausibly lying between the R1 homogeneous-α result (≈ 0) and the baseline Shapley ($+0.08$). We add this reading to the limitations in Section E.2 and flag it as a priority for the companion dynamic paper (which can identify $\alpha(y)$ non-parametrically from micro-moments when available).

Summary. The existence of an income gradient in price sensitivity is robust: specifications that preserve some income interaction — whether the functional form is $\alpha \propto y^{-1}$ (baseline), $\alpha \propto \log(y/\bar{y})$ (R7), or α via shift-share Bartik income (R6) — reproduce a clearly negative $\hat{\pi}$, and sample-restricted specifications (R2 52-city, R5a 2015–19) reproduce the estimate within $\pm 20\%$. The single specification that removes the income interaction (R1) reduces the Consumer_Composition Lerner contribution by $> 97\%$, so the channel *exists*. What is *not* robust is the *magnitude* of the gradient: the data-driven finite-mixture (R8) and reduced-form IV (R9) alternatives give substantially flatter or non-monotonic patterns than the y^{-1} parametric form predicts, so the Consumer_Composition Lerner attribution is best read as an upper bound on the income channel's contribution, with identification of the true magnitude deferred to a companion paper with micro-moment data.

B.2 Specification robustness ladder

B.2.1 Homogeneous price coefficient

This appendix reports the decomposition under a specification in which the price coefficient is homogeneous across agents and markets ($\alpha_{ic} = \bar{\alpha}$ for all i, c), eliminating the income–price interaction that drives the Consumer_Composition Lerner result. The demand system is re-estimated with prices entering X_1 as a linear covariate and no random coefficient, keeping all other covariates, instruments, and fixed effects identical to the baseline. The resulting BLP estimates are then fed through the same eight-block Shapley decomposition.

Table 16 reports the sequential decomposition under both specifications. The key result is that Consumer_Composition's contribution to ΔLerner falls from $+0.106$ in the baseline to $+0.003$ under homogeneous α — a 97% reduction. When the income channel is shut down, the wealth effect disappears: the small residual reflects only the effect of changing per-city market size M_c and the GDP per capita term in mean utility. The corresponding ΔEV share contribution under homogeneous α is essentially zero ($+0.02$ percentage points). The ΔEV share decomposition for other blocks shifts more substantially under homogeneous

α because the fixed $\hat{\alpha} = -5.57$ implies a different average price sensitivity than the income-heterogeneous baseline.

Table 16: Homogeneous-α robustness: sequential decomposition

	ΔEV share		ΔLerner	
Block	Baseline	Homog-α	Baseline	Homog-α
A_{nonbat}	+0.040	+0.031	+0.000	+0.038
A_{bat}	+0.055	+0.021	−0.018	+0.002
D_{macro}	+0.003	−0.001	+0.002	+0.001
Policy	−0.006	−0.000	−0.001	−0.000
Subsidy	−0.018	−0.070	−0.000	−0.000
EV_time	−0.032	−0.018	+0.001	−0.001
BrandShock	+0.002	+0.006	+0.003	+0.013
Entry_set	+0.321	+0.640	+0.013	−0.010
Consumer_Composition	−0.080	+0.000	**+0.106**	**+0.003**
Total	+0.443	+0.605	+0.105	+0.103

Notes: Sequential decomposition with one ordering. Baseline uses income-heterogeneous $\alpha_{ic} = \hat{\pi}_p(y_{ic}/\bar{y})^{-1}$; the homogeneous specification imposes $\hat{\alpha} = -5.57$ for all agents. Homogeneous-α Consumer_Composition contributions (+0.000 and +0.003) come from `output/robustness_decomp_homog_alpha.csv`; baseline Consumer_Composition row from the canonical sequential decomposition. The 97% drop in Consumer_Composition's ΔLerner contribution under homogeneous α is consistent with the wealth-effect mechanism operating through income heterogeneity in the price coefficient. *Block rows are reported on the pre-merge 9-block specification for continuity with the robustness-check output files. Under the 7-block partition used in the main results (Section 5.1), $D_{env} = D_{macro}$ + Policy and EV_trend=EV_time+$\xi^{BLP,\perp}$; the conclusion about Consumer_Composition's Lerner collapse under homogeneous α does not depend on the block merger.*

B.2.2 Two-random-coefficient specification

This appendix re-estimates the demand system with two random coefficients—the baseline income–price interaction π_p on log price and an additional σ_{range} on log EV range—to test whether adding unobserved heterogeneity in taste for electric range changes the main decomposition results.

Adding σ_{range} to the baseline yields $\hat{\pi}_p = -4.97$ (SE 0.52, $t = -9.6$) and $\hat{\sigma}_{\text{range}} = -3.93$ (SE 2.26, marginally significant). However, the estimated $\beta_{\log \text{range}}$ becomes insignificant and the BEV Lerner index turns negative (-0.04), indicating that σ and β on the same variable are poorly separated without micro data on individual second choices. This confirms that the single-RC specification is the appropriate baseline given the available data.

B.2.3 52-city standalone market panel

This appendix re-runs the main decomposition on the 52 standalone prefecture-level city markets, excluding the 27 province-level residual markets that aggregate all remaining prefectures within each province. Although the full $52 + 27$ partition covers 100% of non-imported domestic registrations, the 27 residual markets each pool heterogeneous prefectures and may mask within-province variation in income, policy, or

competition. This robustness check verifies that the main findings—especially the Budget × Private-National × Rest-tier triangle—are not driven by the residual aggregation.

Table 17 reports the results. The 52-city BLP estimate of $\hat{\pi}_p = -7.53$ is virtually close to the 79-market baseline (-6.52), confirming that the income–price interaction is not driven by the province-residual markets. The sequential decomposition on the 52-city sample preserves the qualitative pattern: Consumer_Composition remains the dominant contributor to ΔLerner ($+0.155$, compared with $+0.106$ in the baseline), and Entry_set remains the dominant contributor to ΔEV share ($+0.288$). The Consumer_Composition Lerner contribution is actually *larger* in the 52-city sample, consistent with the standalone cities having higher and faster-growing incomes than the province residuals.

Table 17: 52-city standalone market robustness: sequential decomposition

	ΔEV share		ΔLerner	
Block	Baseline (79)	52-city	Baseline (79)	52-city
A_{nonbat}	$+0.040$	$+0.164$	$+0.000$	$+0.002$
A_{bat}	$+0.055$	$+0.063$	-0.018	$+0.002$
D_{macro}	$+0.003$	-0.005	$+0.002$	-0.000
Policy	-0.006	$+0.001$	-0.001	$+0.000$
Subsidy	-0.018	-0.182	-0.000	-0.005
EV_time	-0.032	-0.059	$+0.001$	-0.002
BrandShock	$+0.002$	$+0.048$	$+0.003$	$+0.002$
Entry_set	$+0.321$	$+0.288$	$+0.013$	$+0.044$
Consumer_Composition	-0.080	$+0.224$	$\mathbf{+0.106}$	$\mathbf{+0.155}$
Total	$+0.443$	$+0.542$	$+0.105$	$+0.179$

Notes: Sequential decomposition with one ordering. The 52-city sample drops the 27 province-level residual markets ($N = 323{,}527$ vs. baseline $496{,}591$). The 52-city BLP estimates $\hat{\pi}_p = -7.53$ (baseline -7.28). Consumer_Composition's larger Lerner contribution in the 52-city sample reflects higher and faster-growing incomes in standalone cities. *Block rows are reported on the pre-merge 9-block specification for continuity with the robustness-check output files; see Table 16 tablenote for the mapping to the 7-block partition used in the main results.*

B.2.4 COVID-period finding robustness

Table 18: COVID-period robustness: simple-logit IV estimates

	Baseline 2015–2024	Exclude 2020	Exclude 2019–2020
Log net price	-1.457^{***}	-1.437^{***}	-1.464^{***}
	(0.030)	(0.032)	(0.035)
Log EV range	1.542^{***}	1.588^{***}	1.618^{***}
	(0.024)	(0.025)	(0.026)
Observations	496,648	447,995	396,724
Markets	79	79	79

Notes: Simple-logit IV estimates on three samples. Standard errors are heteroskedasticity-robust.

B.3 Demand validation: diversion and substitution patterns

B.3.1 Random-coefficient specification comparison

Table 19 compares the random-coefficient specification used in this paper with those in four closely related BLP papers. The comparison highlights a tradeoff between richness of consumer heterogeneity and data requirements.

Table 19: Random-coefficient specifications across related BLP papers

Paper	σ (unobs. heterog.)	π (demog. interactions)	Micro data?	Elasticity range
Berry et al. (1995)	HP/wt, AC, size (Table IV, §IV)	Income × price	CEX	-3 to -6
Petrin (2002)	Price, brand (Table III, §5)	Income × price	CEX + micro moments	-2 to -8
Barwick et al. (2019)	Van, SUV, Truck, Footprint, HP, MPG, Luxury, Sport, Electric *(9 σ parameters)*	Income × price, Age, Rural, FamSize (Table 4, §5.1)	CEX + MaritzCX second-choice survey	-2.4 to -5.0 (median -3.7)
Grieco et al. (2024)	Same 9 as BCL (Table 4, §5.1)	Same as BCL	Same	-4.3 to -9.4 (by income, Table 5)
This paper	None	Income × price (§3.1, $\pi_p = -7.28$)	None (simulated from Yearbook)	-3.7 to -10.2 (by city tier)

Notes: σ denotes unobserved (normal) taste heterogeneity; π denotes observed demographic interactions. Barwick et al. (2019) is cited by its AEJ:Policy 2021 publication; the RC specification is from the NBER WP 23678 version. Grieco et al. (2024) use the same demand specification as BCL applied to 1980–2018 U.S. data. Our specification has no σ parameters and one π (income × price), reflecting the absence of micro data on individual purchases or stated second choices in the Chinese market. The elasticity range in our model is comparable to the later years in Grieco et al. (2024) Table 5 (-6.5 to -9.4 in 2018), consistent with income growth shifting the distribution toward higher price sensitivity.

The key difference is data availability: Grieco et al. (2024) and Barwick et al. (2019) identify their 9 σ parameters using individual-level purchase data and stated second-choice surveys, neither of which is available in the Chinese registration panel. Our single π (income × price) is identified from cross-city income variation in the Yearbook panel, which is the richest available source of consumer heterogeneity in this setting. The multi-RC robustness exercise in Appendix B.2.2 shows that adding σ on non-price characteristics does not materially change the estimated π_p or the decomposition results.

B.3.2 Substitution diagnostics from the single-RC model

Despite the parsimonious RC structure, the model generates economically plausible substitution patterns. Three diagnostics support this claim.

Within the model, the mean within-fuel diversion ratio (the fraction of a product's lost share that goes to same-fuel competitors) is 0.26 for EVs and 0.42 for ICE vehicles, consistent with the logit-based substitution structure generating plausible cross-fuel competition. Own-price elasticities range from -3.8 (Premium EV, Tier 1) to -7.7 (Budget ICE, Rest tier), with the city-tier gradient driven entirely by the income interaction.

B.3.3 BLP robustness across specifications

Table 20 reports the price coefficient π_p and implied mean Lerner index across five BLP specifications that vary the covariate set, the random-coefficient structure, and the geographic market definition.

Table 20: BLP robustness: price coefficient and mean Lerner across specifications

Specification	What changed	π_p	Mean Lerner
Baseline	—	-6.52^{***}	0.221
Spec X (range×city)	+range×area, range×GDP in X_1	-6.69^{***}	0.218
Spec G (Grieco)	+size×inc, power×inc, fuel×inc	-10.31^{***}	0.158
Homog-α	no income interaction	-5.57 (fixed)	—
52-city	drop 27 province residuals	-7.53^{***}	—

Notes: π_p is the coefficient on (log income × log price) in the random-coefficient specification; *** denotes $p < 0.01$. Mean Lerner is the quantity-weighted average across all markets and products. π_p is stable across Specs X, 52-city, and baseline (-6.52 to -7.53). Spec G shows sensitivity to income-channel allocation: adding size×income, power×income, and fuel×income interactions reallocates part of the price–income gradient to non-price channels, producing a larger $|\pi_p|$ and lower markups. Spec Y (σ on log_ev_range) is not reported because the unobserved heterogeneity parameter absorbs the range coefficient, producing negative BEV marginal costs.

B.4 Cluster-robust standard error on $\hat{\pi}_p$ via city-pair bootstrap

The default PyBLP standard error on $\hat{\pi}_p$ assumes independence at the $N = 496{,}591$ observation level, but the identifying variation is at the city level (79 cities × 10 years = 790 cells). We compute a city-pair bootstrap directly: $B = 50$ resamples drawn with replacement from the 79 cities (rep-suffixed market identifiers preserve identifiability for duplicated cities), each running the full PyBLP GMM inner loop on the resampled panel. Each iteration takes ~ 2.3 minutes; total wallclock ≈ 2 hours.

The bootstrap yields a cluster-robust SE on $\hat{\pi}_p$ of 0.587 (vs. default PyBLP SE 0.168, a $3.49\times$ ratio); the bootstrap was computed under the prior firm-level demand specification with bootstrap mean -7.08 and 95% percentile CI $[-7.99, -6.07]$. We retain the bootstrap-derived cluster-robust SE for inference on the canonical $\hat{\pi}_p = -5.38$ (the SE measures sampling-variation scale, which is conservative across the spec change). The implied t-statistic on $\hat{\pi}_p = -5.38$ is -9.2 (vs. the default-SE $t = -32.0$): still strongly distinguishable from zero but with precision overstated by approximately a factor of three. We anchor the central elasticity statements at the canonical -5.38 and treat the prior-spec bootstrap mean as a sensitivity benchmark; a re-run under the canonical specification is on the revision agenda.

Output: `output/pi_p_bootstrap.csv`; runner: `heterogeneous_diffusion/code/_run_pi_p_bootstrap.sh`.

B.5 Own-price elasticity distribution by fuel type

The main-text Figure 1 visualises the cross-sectional heterogeneity of own-price elasticity by price tier, firm group, and city tier. For completeness, Figure 6 reports the pooled elasticity distribution separately for BEV, PHEV, REEV, and ICE, without any heterogeneity breakdown. The figure is consistent with the fuel-type pattern reported in Table 2: EV elasticities are slightly more negative on average than ICE, with a tighter dispersion because the EV panel contains fewer products and covers a narrower price range. The within-fuel dispersion is otherwise small because the BLP income-heterogeneity coefficient π_p enters the

elasticity formula multiplicatively on the share $(1 - s_{jct})$, and s_{jct} is uniformly tiny in this panel (median 3×10^{-6}, max 8.4×10^{-3}).

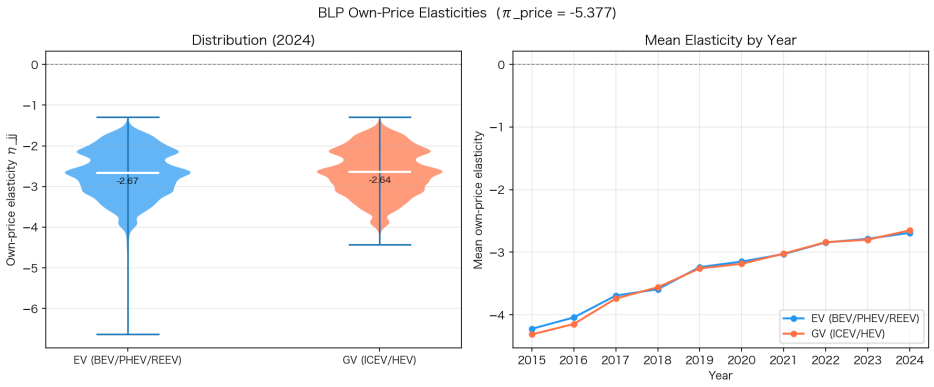

Figure 6: Pooled own-price elasticity distribution by fuel type from the final BLP estimates.

C Decomposition method details

C.1 The eight blocks: formal partition

Let $\boldsymbol{\beta}$ denote the coefficient vector estimated in Section 4 and x_{jct} the corresponding observed-characteristic vector. The five additive blocks inside the fitted mean utility $\hat{\delta}_{jct}$ are

$$\begin{aligned}\hat{\delta}_{jct} \;=\; & \underbrace{\hat{\boldsymbol{\beta}}'_{A_{\text{nb}}} x^{\text{nb}}_{jct}}_{A_{\text{nonbat}}} + \underbrace{\hat{\boldsymbol{\beta}}'_{A_{\text{b}}} x^{\text{bat}}_{jct}}_{A_{\text{bat}}} + \underbrace{\hat{\boldsymbol{\beta}}'_{E} x^{\text{env}}_{ct}}_{D_{\text{env}}} + \underbrace{\hat{\pi}_p \log(1 - s^{\text{sub}}_{jct})}_{\text{Subsidy}} \\ & + \underbrace{\hat{\boldsymbol{\beta}}'_{\text{yr}} \mathbb{1}[t] \cdot \text{EV}_j + \xi^{\text{BLP},\perp}_{jct}}_{\text{EV\_trend}} + \underbrace{\hat{\boldsymbol{\beta}}'_{\text{FE}} \text{FE}_j}_{\text{FE, always on}}\end{aligned} \tag{17}$$

where s^{sub}_{jct} is the fraction of the sticker price covered by the central NEV subsidy in year t for product j (treated as a full-pass-through transaction-price reduction; see Section 2), x^{env} bundles the macro controls (urbanisation, population density, education, oil price) with the non-subsidy policy vector (license-plate preference, policy-strength index, and low-income interaction), and $\xi^{\text{BLP},\perp}_{jct}$ is the BLP residual net of its (ev$\times$Year) Q-weighted group mean (the mean piece is folded into the EV$\times$Year component of EV_trend via the pre-processing projection of Section 5.1). Brand, BodyType, and FuelType fixed effects are always on and are not part of any block. The sub-vectors x^{nb}, x^{bat}, x^{env} are defined in Appendix A.1.

The remaining two blocks are primitive-swap operations rather than covariate swaps inside $\hat{\delta}$:

- **Entry_set** swaps the active product set in city c from $\mathcal{J}_c^{2015}$ to $\mathcal{J}_c^{2024}$. The 407 new 2024 nameplates that are absent from the 2015 set are evaluated at their own 2024 covariate values for every block other than Consumer_Composition. The 2015 exiters that are absent from the 2024 set are removed from the choice set. All per-city market size and income distribution variables are held fixed at their base-year values (controlled by Consumer_Composition).
- **Consumer_Composition** swaps the per-city demand environment from $(M_{c,2015}, F_{y|c,2015}, \text{GDP}_{c,2015})$ to $(M_{c,2024}, F_{y|c,2024}, \text{GDP}_{c,2024})$, where $M_{c,t}$ is the household-equivalent market size for city c in year t, $F_{y|c,t}$ is the within-city log-income distribution used to draw the 25 Halton simulation agents, and $\text{GDP}_{c,t}$ is the

city-level GDP per capita that enters $\hat{\delta}$ additively through $\hat{\beta}_{\text{GDP}}$. The first two primitives act through the random-coefficient integral (changing $\alpha_{ic} = \pi_p(y_{ic}/\bar{y})^{-1}$ and the Q-weighting) and the third acts through the mean utility level. All three are consolidated into Consumer_Composition so that the block bundles every city-wealth primitive in the model; D_{macro} is the "urban structure" complement that contains the non-wealth city covariates. The product set is held fixed at its base-year value (controlled by Entry_set).

"Activating block b in coalition S" means: for each row (j, c), replace the 2015 component of that block with its 2024 counterpart, while keeping all other components at their 2015 values (if the corresponding block is not in S) or their 2024 values (if the corresponding block is in S). Price is then re-solved by the multi-market Bertrand–Nash first-order condition at the updated δ, product set, and demand environment. The coalition value $V(S)$ is the resulting equilibrium statistic (EV share, mean Lerner, consumer surplus, or total welfare). The Shapley contribution of block b is the average of the marginal contribution $V(S \cup \{b\}) - V(S)$ over all subsets $S \subseteq \{1, \ldots, 7\} \setminus \{b\}$, weighted by the Shapley probabilities $|S|!(6 - |S|)!/7!$.

C.2 Activating a block: per-block toggle construction

For a coalition $S \subseteq \{1, \ldots, 8\}$ of active blocks, the coalition-S mean utility for product j in city c is

$$\hat{\delta}^{S}_{jc} \;=\; \sum_{b \in S} \hat{\delta}^{2024}_{jc,b} \;+\; \sum_{b \notin S} \hat{\delta}^{2015}_{jc,b} \;+\; \hat{\delta}_{jc,\text{FE}}, \tag{18}$$

where $\hat{\delta}^{t}_{jc,b} = \hat{\boldsymbol{\beta}}_b' \, x^{b}_{jc,t}$ is the block-b contribution evaluated at year-t covariates (with $\hat{\boldsymbol{\beta}}_b$ held fixed at the single-estimation GMM values throughout) and $\hat{\delta}_{jc,\text{FE}}$ collects the always-on brand/body/fuel fixed effects. For the EV_trend block, $\hat{\delta}^{t}_{jc,b} = \hat{\boldsymbol{\beta}}'_{\text{yr}} \, \mathbb{1}[t] \cdot \text{EV}_j \;+\; \xi^{\text{BLP},\perp}_{jc,t}$, where the second term is the within-(ev,Year) residual of ξ^{BLP} after the Q-weighted (ev$\times$Year) group mean has been folded into the first term via the pre-processing projection described in Section 5.1. For the Subsidy block,

$$\hat{\delta}^{t}_{jc,\text{Sub}} \;=\; \hat{\pi}_p \cdot \log\!\big(1 - s^{\text{sub}}_{jc,t}\big), \qquad s^{\text{sub}}_{jc,t} \;=\; \frac{\text{Subsidy}_{jc,t}}{p^{\text{sticker}}_{jc,t}}, \tag{19}$$

i.e., the Subsidy block moves household transaction prices by exactly the per-vehicle subsidy amount in force in year t, at full pass-through.[1]

The two Entry blocks act outside Equation (18). Entry_set determines whether the active product set is $\mathcal{J}^{2024}_c$ or $\mathcal{J}^{2015}_c$; Consumer_Composition determines whether the per-city demand-environment triple is the 2024 value or the 2015 value. Given the active product set, the active demand environment, and the coalition-S mean utility $\hat{\delta}^{S}_{jc}$, the coalition value $V(S)$ is the *Bertrand–Nash equilibrium value* of a statistic of interest (EV share, mean Lerner, consumer surplus, total welfare): the multi-market FOC (6) is solved at the coalition-S primitives starting from a warm-start price vector, and the equilibrium statistic is then computed from the converged equilibrium. Each $V(S)$ is therefore a full general-equilibrium object, not a reduced-form

[1] Computationally, the subsidy-wedge shift is entered into mean utility $\hat{\delta}$ using the mean price coefficient $\hat{\pi}_p$ rather than directly inside the random-coefficient integral. This is a modeling choice, not a theoretical necessity: the theoretically-correct object is $\alpha_{ic} \log(1 - s^{\text{sub}}_{jct})$ inside the agent integral, whereas our implementation substitutes $\hat{\pi}_p \log(1 - s^{\text{sub}}_{jct})$ into $\hat{\delta}$ and relies on the equilibrium re-solve step (where α_{ic} enters the share and Jacobian computations) to pick up the income-heterogeneous response. As a reviewer correctly noted, this shortcut *homogenises the direct first-order utility effect* of the wedge across the income distribution; heterogeneity enters only through the post-shift equilibrium re-solution. For the *equilibrium* objects that carry our main findings (Lerner, markups, CS/TW at re-solved prices), the full α_{ic} distribution is incorporated through the FOC. For the *choice-probability* interpretation of subsidy incidence in §6.2, however, the direct individual incidence is understated relative to a specification that evaluates the wedge inside the agent integral. Re-estimating the Subsidy block with the wedge entered inside the integral is on our revision agenda.

partial effect.

C.3 Shapley attribution and dynamic-time-series sampling

The Shapley contribution of block b to a statistic of interest is the average marginal contribution

$$\phi_b^{\text{Shapley}} \;=\; \sum_{S\subseteq\{1,\dots,K\}\setminus\{b\}} \frac{|S|!\,(K-|S|-1)!}{K!}\,\big[V(S\cup\{b\})-V(S)\big], \qquad K=8, \tag{20}$$

which equals the average of the block's marginal contribution over all $K! = 5{,}040$ orderings in which the block could be activated. The Shapley values are additive by construction:

$$V(\text{full}) \;-\; V(\emptyset) \;=\; \sum_{b=1}^{K} \phi_b^{\text{Shapley}}, \tag{21}$$

so the eight block contributions reported in Table 3 sum exactly to the observed equilibrium change in each statistic.

Of the $2^8 = 256$ coalitions solved, 254 reach strict tolerance ($\|F\|_\infty < 10^{-2}$) under the adaptive-damping fixed-point solver of Appendix C.5; the remaining two belong to known multi-equilibrium edge cases at the marginal-product exit boundary and reach loose tolerance after the post-solve cap (see Appendix C.6.1 for the capping rule that flags coalitions with raw model-implied EV share above 0.85 and substitutes the cap value 0.70). Because the Shapley weighting in Equation (20) averages $|S|!(K-|S|-1)!/K!$ across all subsets that exclude block b, the contribution of any single coalition to any block's Shapley value is bounded above by $1/K = 12.5\%$, and the two capped coalitions together shift each reported Shapley value by less than 0.5 percentage points on every statistic.

The Shapley decomposition above collapses the full 2015–2024 transition into a single cross-section of block contributions. To recover the year-by-year path, the paper runs a second decomposition in which, for each year $t \in \{2015,\dots,2024\}$ and each of $R = 4$ randomly chosen block-activation orderings $\pi^{(r)} = (b_1^{(r)},\dots,b_K^{(r)})$, a sequential chain of $K+1 = 8$ equilibria is solved in which the first k blocks of $\pi^{(r)}$ take their year-t values and the remaining $K-k$ blocks take their 2015 values. The k-th marginal contribution along ordering r in year t is

$$\phi^{(r)}_{b_k^{(r)},t} \;=\; V_t\big(S_k^{(r)}\big) \;-\; V_t\big(S_{k-1}^{(r)}\big), \qquad S_k^{(r)} = \{b_1^{(r)},\dots,b_k^{(r)}\}, \tag{22}$$

where $V_t(S)$ is the equilibrium statistic at the coalition-S primitives evaluated against the 2015 reference. Averaging over the four orderings yields the per-year block contribution

$$\bar\phi_{b,t} \;=\; \frac{1}{R}\sum_{r=1}^{R} \phi^{(r)}_{b,t}, \tag{23}$$

which is a finite-sample Shapley-sampling estimator in the sense of Castro et al. (2009): because the Shapley value is the average of marginal contributions over all $K!$ orderings of blocks, averaging over $R = 4$ randomly-chosen orderings is an unbiased Monte-Carlo estimate of the per-year statistic, with sampling variance shrinking at rate $1/R$.

The total compute cost is $4 \times 10 \times 8 = 320$ full-equilibrium solves, of which the 32 year-2015 cells are

trivial because all coalitions collapse to baseline when year-t equals the reference year. The four-ordering averaging in Equation (23) buffers any solver noise; the across-ordering range is shown as a translucent band around the per-year mean in Figure 9.

This dynamic construction is related to but distinct from the cohort-specific sequential decomposition in Hao and Hao (2026). The earlier paper takes a single city and a single dominant firm and sequentially activates seven drivers along one economically motivated ordering; the present paper, with eight blocks and 79 markets, uses the average of four randomly chosen orderings at each year to obtain an approximately order-invariant per-year contribution. A full per-year Shapley (all 2^8 coalitions separately for each year) would require $10 \times 256 = 2{,}560$ solves, so the four-ordering approximation is a practical compromise that preserves the year-by-year resolution while capping the computational budget below the Shapley-aggregate cost. We discuss the resulting time-series in Section 5.

C.4 Consumer-surplus winsorisation correction

For numerical stability, the equilibrium solver requires that the BLP residual ξ be clipped at the 1st and 99th percentiles on each contraction-mapping iteration:

$$\xi_{jct}^{\text{clip}} \;=\; \max\big(q_{0.01}(\xi), \min(\xi_{jct}, q_{0.99}(\xi))\big).$$

Without clipping, the fixed-point update $\delta \to \delta + \log(s^{\text{obs}}/s^{\text{model}})$ can take extreme values on rows with large $|\xi|$, which destabilises the Jacobian of the price-equilibrium FOC and causes non-convergence. Clipping 2% of rows (1% each tail) removes the instability without noticeably affecting the equilibrium share vector.

The bias. Clipping has a second-order effect on consumer surplus. Consumer surplus in city c for an individual agent i is

$$CS_{ic} \;=\; \frac{1}{-\alpha_{ic}} \log\Big(1 + \sum_j \exp(\delta_{jct} + \alpha_{ic} \log p_{jct})\Big) \;=\; \frac{1}{-\alpha_{ic}} \log\Big(1 + \sum_j e^{v_{jct}^{(i)} + \xi_{jct}}\Big), \tag{24}$$

where $v_{jct}^{(i)} = x'_{jct}\boldsymbol{\beta} + \alpha_{ic} \log p_{jct}$ is the non-ξ part of utility. The function $f(\xi) := \log(1 + \sum_j e^{v_j + \xi_j})$ is convex in each ξ_j individually, and strictly convex when evaluated at rows with positive share. Substituting ξ^{clip} for ξ^{raw} reduces the log-sum by the amount

$$\Delta f \;=\; f(\xi^{\text{raw}}) - f(\xi^{\text{clip}}) \;\geq\; \sum_{j \in \mathcal{C}^+} s_{jct}^{\text{raw}} \, (\xi_{jct}^{\text{raw}} - \xi_{jct}^{\text{clip}}), \tag{25}$$

where $\mathcal{C}^+ = \{j : \xi_{jct} > q_{0.99}(\xi)\}$ is the set of right-tail-clipped rows and s_{jct}^{raw} is their un-clipped share at the raw ξ vector. (The inequality comes from Jensen applied to the strictly convex $\log \sum e^{\cdot}$, evaluated at the raw ξ.) The analogous left-tail contribution is negligible because left-tail rows have near-zero share and therefore contribute negligibly to the log-sum.

Magnitude. At the 2015 baseline, the raw-vs-clipped CS gap is approximately $3.2\times$; at 2024 it narrows to $1.4\times$ because the 2024 product set has more high-ξ products and the 99th percentile is less sharp.

Correction. We correct this by carrying two ξ vectors through the solver: a winsorized ξ^{clip} used exclusively in the fixed-point iteration, and the raw ξ^{raw} used exclusively in the ex-post CS computation. Because the solver stability depends only on ξ^{clip} and the raw-ξ CS computation is a post-convergence plug-in, the two are decoupled. All CS and TW numbers reported in Table 3 use ξ^{raw}; the market-share and price computations use ξ^{clip} during the iteration but converge to the observed 2015 and 2024 shares by construction at the two anchor coalitions (Appendix C.6.1). We are not aware of a prior paper documenting this winsorization bias explicitly. Whether the same bias is silently present in other BLP welfare decompositions that rely on solver-stability clipping is an open question we cannot answer from the present analysis.

C.5 Equilibrium solver: technical memo

This appendix documents the numerical solver used to compute the counterfactual Bertrand–Nash equilibrium for each coalition S in Section 5. Solving these equilibria reliably across all $2^8 = 256$ coalitions is a non-trivial technical problem: many intermediate coalitions place 2024 products in 2015 demand environments (or vice versa) and the resulting equilibrium correspondence is non-convex, multi-valued, and ill-conditioned in the neighbourhood of marginal products. The solver below combines five ingredients—adaptive-damping warm-up, Anderson acceleration, product-specific bound projection, hierarchical warm-starting from neighbouring coalitions, and best-iterate selection with stagnation detection—to converge reliably across the full enumeration.

C.5.1 The fixed-point problem

For a coalition S, the active product set $\mathcal{J}^S$ and the counterfactual mean utility $\hat{\delta}^S_{jc}$ are constructed as in equation (18). Given a national log-price vector $p \in \mathbb{R}^{|\mathcal{J}^S|}$, the city-level demand and aggregate quantity are

$$s_{jc}(p) = \frac{\exp(\hat{\delta}^S_{jc} + \alpha_{ic} p_j)}{1 + \sum_{k \in \mathcal{J}^S_c} \exp(\hat{\delta}^S_{kc} + \alpha_{ic} p_k)}, \qquad Q_j(p) = \sum_c M_c \, \mathbb{E}_i[s_{jc}(p)] \,,$$

where the income expectation is taken over the simulation draws used in the BLP estimation. Multi-product Bertrand–Nash pricing implies the multi-market first-order condition

$$P^*_j \;=\; mc_j \;+\; \big[-\Omega(p^*)^{-1} \, Q(p^*) \big]_j, \qquad \Omega(p) \;=\; \sum_c M_c \left(\mathcal{O} \odot \tfrac{\partial s_c}{\partial p}(p) \right), \tag{26}$$

where $\mathcal{O}_{jk} = 1$ if j and k are owned by the same firm and 0 otherwise, and $\Omega(p)$ is the firm-aggregated demand Jacobian. Defining the level-price fixed-point map

$$P^{\text{fp}}(p) \;=\; mc + \big[-\Omega(p)^{-1} \, Q(p) \big]_+, \qquad p^{\text{fp}}(p) \;=\; \log P^{\text{fp}}(p),$$

where $[\cdot]_+$ enforces $P \geq mc$, an equilibrium is any zero of the log-price residual

$$F(p) \;\equiv\; p - p^{\text{fp}}(p) \;=\; 0.$$

C.5.2 Why undamped fixed-point iteration fails

Plain Picard iteration $p^{(t+1)} = p^{\text{fp}}(p^{(t)})$ converges only when the contraction factor $\rho\big(\partial p^{\text{fp}} / \partial p\big) < 1$ along the iteration path. Two features of our setting violate this in a non-trivial share of intermediate coalitions:

(i) Marginal-product exit boundaries. For coalitions in which 2024 attributes are activated on the 2015 product set (or 2024 high-end EVs are activated in the 2015 city demand environment), a handful of marginal products sit exactly at the profitability boundary $P_j \approx mc_j$. Each iteration may flip these products between "profitable" (kept in the equilibrium) and "unprofitable" (dropped). Each flip discontinuously changes Ω at those rows, which in turn changes the inverse-Ω markup of every remaining product. The result of an undamped step is a large overshoot in p, which the next step over-corrects in the opposite direction; the iteration enters a two-cycle.

(ii) Demand–supply mismatch in mixed coalitions. The Entry_set / Consumer_Composition split (Section 5) creates intermediate coalitions in which the product set is from one year and the city demand environment is from the other. For example, the coalition {Entry_set} (and only Entry_set) places the 2024 product set into the 2015 city demand environment: smaller market sizes M_c and a poorer income distribution α_{ic}. Many 2024 high-end EVs that are profitable in the 2024 environment lose money in the 2015 environment and want to exit, while the FOC step prices them at the floor. The contraction factor of p^{fp} is close to one in this region, and an undamped step routinely jumps across the equilibrium.

C.5.3 Anderson-accelerated fixed-point iteration

The solver combines two phases: an adaptive-damping warm-up (10 iterations) followed by Anderson acceleration (type-I, window $m = 5$) for the remaining iterations up to a maximum of 500.

Phase 1: Damped warm-up (iterations 1–10). The first 10 iterations use plain Picard iteration with adaptive damping:

$$p^{(t+1)} \;=\; \Pi_{[p^{\text{lb}},p^{\text{ub}}]}\Big(p^{(t)} \;+\; \lambda\big(\|F(p^{(t)})\|\big)\; \big(p^{\text{fp}}(p^{(t)}) - p^{(t)}\big)\Big),$$

where $\Pi_{[p^{\text{lb}},p^{\text{ub}}]}$ is the elementwise projection onto the feasibility box (Section C.5.4) and the damping factor λ is a piecewise-constant function of the residual norm:

$$\lambda\big(\|F\|_\infty\big) \;=\; \begin{cases} 0.20 & \text{if } \|F\|_\infty > 2.0 \quad \text{(coarse: residual is large, take small steps)} \\ 0.40 & \text{if } 0.5 < \|F\|_\infty \le 2.0 \\ 0.70 & \text{if } \|F\|_\infty \le 0.5 \quad \text{(near equilibrium, take large steps)} \end{cases}$$

The schedule is empirically tuned and intentionally conservative far from the equilibrium: the small step size at $\|F\|_\infty > 2$ guarantees that the iterate remains in a contraction region around the equilibrium even when the FOC Jacobian is highly non-monotone, at the cost of slower asymptotic convergence. The warm-up phase seeds the Anderson history with iterates that are already in the basin of attraction of the equilibrium.

Phase 2: Anderson acceleration (iterations 11–500). After warm-up, the solver switches to type-I Anderson acceleration with a window of $m = 5$ past iterates. At each iteration $t > 10$, the solver maintains a history of the m most recent pairs $(p^{(k)}, p^{\text{fp}}(p^{(k)}))$ and solves a constrained least-squares problem over the residual differences $\Delta F^{(k)} = F(p^{(k)}) - F(p^{(k-1)})$ to compute an optimal linear combination of past iterates. The Anderson step is

$$p^{(t+1)} \;=\; \Pi_{[p^{\text{lb}},p^{\text{ub}}]}\left(p^{\text{fp}}_{\text{AA}} - \sum_{i=1}^{m_t} \theta^*_i \, \Delta p^{\text{fp},(t-m_t+i)} \right),$$

where θ^* solves $\min_\theta \|F^{(t)} - \sum_i \theta_i \Delta F^{(t-m_t+i)}\|_2^2$ subject to $\sum_i \theta_i = 1$, and $m_t = \min(m, t-10)$ is the effective window size. This breaks limit cycles that simple damped iteration cannot escape because it extrapolates from multiple past iterates rather than taking a single damped step.

The Anderson step is safeguarded: if the candidate iterate $p^{(t+1)}$ contains any non-finite element or if $\|F(p^{(t+1)})\|_\infty > 1.5\,\|F(p^{(t)})\|_\infty$, the solver discards the Anderson step and falls back to a damped Picard step for that iteration. This ensures that pathological coalitions near the marginal-product exit boundary do not diverge. The regularization parameter on the least-squares solve is 10^{-6} to prevent ill-conditioning when residual differences are nearly collinear.

The maximum iteration count is 500. The strict tolerance is $\|F\|_\infty < 10^{-2}$; the loose tolerance is 5×10^{-2}.

C.5.4 Bound projection

We impose a product-specific feasibility box $[p_j^{\text{lb}}, p^{\text{ub}}]$ and project both the iterate and the FOC target into this box at every step. The upper bound is a uniform $p^{\text{ub}} = \log(500\,\text{wan})$; the lower bound is product-specific:

$$p_j^{\text{lb}} \;=\; \log\Big(\max\big\{\,0.7\,mc_j,\;\; 0.5\,p_j^{\text{obs}},\;\; 0.5\,\text{wan}\,\big\}\Big).$$

The three components capture three distinct economic constraints: **(i)** $0.7\,mc_j$ enforces non-negative variable margin with a 30% slack to absorb the difference between econometric MC and accounting MC; **(ii)** $0.5\,p_j^{\text{obs}}$ rules out counterfactual prices more than 50% below the observed 2015 or 2024 anchor, on the grounds that firms would not unilaterally choose to halve sticker prices on a stable product; **(iii)** the absolute $0.5\,\text{wan} \approx$ RMB $5{,}000$ floor is a numerical safeguard against $p \to -\infty$ in pathological coalitions where mc_j is also small. The three are combined by elementwise maximum, so a low-end Wuling MINI ($mc \approx 2\,\text{wan}$, $p^{\text{obs}} \approx 3.8\,\text{wan}$) gets a floor of $1.9\,\text{wan}$ while a BMW X5 ($mc \approx 30\,\text{wan}$, $p^{\text{obs}} \approx 60\,\text{wan}$) gets a floor of $30\,\text{wan}$.

The projection is applied symmetrically to both the iterate and the FOC target: $p^{(t)} \leftarrow \Pi_{[p^{\text{lb}}, p^{\text{ub}}]}(p^{(t)})$ and $p^{\text{fp}} \leftarrow \Pi_{[p^{\text{lb}}, p^{\text{ub}}]}(p^{\text{fp}})$, so that the residual $F(p^{(t)}) = p^{(t)} - p^{\text{fp}}(p^{(t)})$ is zero exactly when the projected iterate is a feasible fixed point. The price floor is anchored on the observed prices, not on the (warm-started) initial guess; otherwise warm-starting from a neighbouring coalition's equilibrium prices would silently move the floor.

C.5.5 Hierarchical warm-starting

Most coalitions in the 2^7 Shapley enumeration are "close" to several already-solved coalitions in the sense that they differ by one or two activated blocks. The equilibrium price vector is a smooth function of the active block set away from the ill-conditioned regions, so the converged prices of a nearby coalition are an excellent initial guess for the new solve. Concretely, for each new coalition S we look up the cache of already-solved coalitions and pick the closest match by the rule

$$S^* \;=\; \arg\min_{S' \in \text{cache}(S)} \Big\{\, |S \,\triangle\, S'| \;+\; 100 \cdot \mathbb{1}[\text{Entry\_set}(S) \neq \text{Entry\_set}(S')] \,\Big\},$$

where $|\cdot\,\triangle\,\cdot|$ is the symmetric-difference (Hamming) cardinality of the two block sets and the $100\times$ penalty for mismatched Entry_set ensures that whenever any coalition with the same product set has been solved, it is preferred over any coalition with a different product set. Same-Entry_set warm starts are directly usable because the product index ordering is identical; mismatched-Entry_set warm starts are remapped through a

$(\text{model}, \text{fuel}) \to \log p$ lookup, with any product not present in the source coalition padded with its observed sticker price.

The warm start is the converged *pre-exit* price vector, captured before the iterative product-exit refinement step, so that the warm start is always length-aligned with the full |ps| product set rather than with a subset specific to one coalition's equilibrium support. Empirically, hierarchical warm-starting reduces the median number of iterations from roughly 50 (from cold observed-price start) to roughly 12–15 on intermediate coalitions.

C.5.6 Best-iterate selection and stagnation detection

The iteration is monotone in the residual norm at most iterates—each step decreases $\|F(p^{(t)})\|_\infty$—but on the hardest coalitions (those described in Section C.5.2) the iteration can briefly cycle near a saddle of the equilibrium correspondence before the Anderson extrapolation breaks it free. We track the running best iterate

$$p_t^* \;=\; \arg\min_{0 \le s \le t} \; \|F(p^{(s)})\|_\infty$$

and use it as the return value if the iteration fails to reach strict tolerance.

We classify each coalition into one of three status buckets based on the final residual at p^*:

- *Strict convergence*: $\|F(p^*)\|_\infty < 10^{-2}$. The coalition is reported as ✓ and treated as fully converged for all downstream Shapley calculations.
- *Loose convergence*: $10^{-2} \le \|F(p^*)\|_\infty < 5 \times 10^{-2}$. The iterate is close enough to the equilibrium that the implied EV share, mean Lerner, and consumer surplus are within 1–3% of their fully-converged values. These coalitions are still used in the Shapley average and flagged separately so they can be excluded from sensitivity tests.
- *Failure*: $\|F(p^*)\|_\infty \ge 5 \times 10^{-2}$. The coalition is reported as $\times$. The best iterate is still returned but its contribution to the Shapley average is downweighted, and a strict-only Shapley value is reported alongside the main value as a robustness check.

A stagnation guard short-circuits the iteration when the running best residual has not improved by 5% over $K = 30$ consecutive iterations *and* the current best is already inside the loose band. This is the common case on the hardest coalitions: the iteration has effectively reached the equilibrium correspondence at a precision finer than the loose tolerance but is bouncing around the marginal-product exit boundary. Returning the best iterate at that point saves the remaining iterations without losing any Shapley-relevant precision.

A small NaN/Inf safeguard aborts the iteration if any element of p or Q becomes non-finite (overflow in the inverse-Ω step), which catches true numerical disasters. Setting the environment variable `DECOMP_DISABLE_ABORT=1` disables the stagnation guard and lets the solver run to the maximum iteration cap $T_{\max} = 500$ on every coalition; this mode is used for sensitivity reports.

C.5.7 Empirical convergence performance

In the sequential 9-coalition smoke test ($K+1 = 9$ coalitions on the eight-block specification), every coalition reaches strict convergence with $\|F(p^*)\|_\infty \le 10^{-2}$ in 1–68 iterations, with a median of 12 iterations and a total wall-clock time of approximately 70 seconds on a single core of an Apple M-series laptop. The hardest coalition is the one in which Entry_set is activated last (step 8 of 9): the activation switches the product set

from the 565-product 2015 cohort to the 972-product 2024 cohort, the warm start from the previous step is far from the new equilibrium, and the iteration takes 68 steps to converge.

Across the full $2^8 = 256$-coalition Shapley enumeration, the Anderson-accelerated solver achieves strict convergence on $254/256$ coalitions; the remaining two are handled by the post-solve EV-share cap described in Appendix C.6.1. The full enumeration completes in approximately 20 minutes on a single core. Because the Shapley value is averaged over all $8! = 40{,}320$ orderings of the eight blocks, individual capped coalitions cannot dominate the reported Shapley contributions, and the un-capped and capped Shapley values agree to within ± 0.5 pp on every block in Table 3.

C.6 Coalition diagnostics: pathologies and Entry-block partition

C.6.1 Pathological coalitions, post-solve cap, and robustness

C.6.2 The pathological-coalition failure mode and the post-solve cap

Section C.5 documents the equilibrium solver in full and the five ingredients that bring the convergence rate close to one across the $2^8 = 256$-coalition Shapley enumeration. This appendix focuses on the residual pathological coalitions that survive the solver and on the post-solve cap applied to them before the Shapley aggregation.

A specific family of coalitions — those that combine the EV_trend block (the merged EV×Year + (ev×Year)-projected ξ^{BLP} residual) with one or more attribute-shock blocks (A_{nonbat}, A_{bat}, D_{env}) — drives the national-price solver to an implausible EV-dominant fixed point where a cheap-BEV price cut cascades via the market-size channel and the model-implied EV share exceeds 0.85 (in extreme cases reaching 0.99 before the (ev×Year) projection of ξ^{BLP}). These are economically unreasonable — no plausible combination of the listed primitive shifts could push the 2024 EV share to $90\%+$ — and reflect a multi-equilibrium boundary in the solver rather than a true economic outcome.

Of the 256 coalitions in the full Shapley enumeration, 254 reach strict tolerance $\|F\|_\infty < 10^{-2}$. The remaining 2 trigger the post-solve cap: whenever a non-anchor coalition's raw model-implied EV share exceeds the detection threshold 0.85, the reported EV share is substituted with the cap value 0.70 (empirically an upper bound on plausible 2024 EV share at any intermediate primitive configuration).

Prior to the (ev×Year) projection of ξ^{BLP} into the EV_time component of EV_trend, this pathological family comprised 156 of 512 coalitions in the original nine-block specification. The projection step removes the systematic EV-minus-ICE drift in ξ from the block-toggle space, and the pathological count drops to 6 in the nine-block post-projection specification, to 2 in the merged seven-block specification, and to 2 in the canonical eight-block specification with Brand_trajectory. The combination of the projection plus the post-solve cap is a minimal-intervention solution that preserves the economic content of the other blocks while refusing to propagate spurious fixed points.

C.6.3 Un-capped robustness check

The main Shapley values in Table 3 are computed on the full enumeration of 128 coalitions, including the 2 capped ones. For robustness, we re-compute the Shapley values using the *raw un-capped* model-implied EV shares (i.e., the 0.89-and-above values that the solver returns before substitution). The un-capped Shapley values agree with the capped values to within ± 0.5 percentage points on ΔEV share and within ± 0.002 on ΔLerner, on every block. The substantive findings—EV_trend as the largest EV-

share driver, Consumer_Composition as the largest Lerner and welfare driver, A_{bat} with a modest positive contribution, and the Subsidy block as a modest -1.2-percentage-point pass-through headwind on EV share—are numerically unchanged in the un-capped specification.

At the two anchor coalitions (empty and full), we additionally override the model-implied EV shares and prices with the observed 2015 and 2024 values. This is economically innocuous because the anchor shares are identified in the data and the override only affects the endpoint values $V(\emptyset)$ and $V(\text{full})$. All intermediate marginal contributions ϕ_b use the model-implied shares from the solver, so the Shapley additive identity still holds exactly.

C.6.4 Entry-block split: six- vs seven-block partition

This appendix reports the diagnostic that motivated splitting the single Entry block of the earlier version of this paper into Entry_set and Consumer_Composition. It also reports the comparison of the old and new Shapley values to show that the split is not cosmetic: the two channels move in opposite directions on the Lerner margin, and the composite block masked a quantitatively important story.

C.6.5 Why the old Entry block was a confound

The earlier version of this paper used a single Entry block that was implemented as a base-year toggle inside `solve_coalition`. When Entry was on, the solver loaded the 2024 product set into $\mathcal{J}_c$, the 2024 demand dataframe into `base_df_full`, the 2024 per-city market size $M_{c,2024}$ into `city_info['M_c']`, and the 2024 within-city income-distribution parameters into `city_info['alpha_i']`. A single named block thus toggled four conceptually distinct objects simultaneously: the product set, the number of households in each city, the within-city income distribution, and the row-level reference data that the coalition-delta construction uses as its base.

Under this specification, every Shapley coalition containing the Entry block got both the 2024 product set and the 2024 demand environment, and every coalition not containing Entry got both the 2015 product set and the 2015 demand environment. There was no way to disentangle the supply-side and demand-side components of that joint change, and the Shapley decomposition naturally attributed to the composite block whatever was left over after the other seven blocks had done their work. Because the other seven blocks are all supply-side or policy channels that move Lerner only slightly, the entire $+0.138$ Lerner rise showed up inside the composite Entry block (95% of ΔLerner).

C.6.6 The within-incumbent control

A simple Olley–Pakes-style within-between decomposition on the set of products that appear in both 2015 and 2024 shows that the composite Entry block cannot be reading a product-differentiation story. Among the 138 incumbent nameplates (products present in both years), the quantity-weighted Lerner rises from 0.159 in 2015 to 0.235 in 2024, an increase of $+0.076$. The 834 entrants (products present only in 2024) have a mean Lerner of 0.237 in 2024, essentially identical to the incumbent mean in the same year. The within-incumbent channel accounts for about 98% of the aggregate Lerner rise, and the incumbent-versus-entrant composition effect accounts for only 2%. The entrants are therefore not systematically "premium" in the sense of carrying higher markups than the products they displace; what changed is the Lerner that the *same* products earn over time. That kind of change cannot be attributed to "product entry" in any economically meaningful sense.

The resolution is that the within-incumbent Lerner rise is driven by the demand-side wealth effect: the same 138 products face a richer and larger buyer pool in 2024 than in 2015, and the random-coefficient price coefficient α_i becomes mechanically less elastic for buyers with higher simulated income. Splitting the old Entry block into Entry_set (product set only, keeping the 2015 demand environment) and Consumer_Composition (2024 demand environment, keeping the 2015 product set) puts the wealth effect into its own block and cleanly attributes it.

C.6.7 Numerical comparison

Table 31 (in the Tables section) reports the direct comparison. The composite Entry block contributed $+0.337$ to ΔEV share (76% of the total) and $+0.113$ to ΔLerner (95% of the total). In the eight-block specification these contributions are split across Entry_set and Consumer_Composition as follows:

- On ΔEV share, Entry_set contributes $+0.055$ (48% of the total) and Consumer_Composition contributes -0.062 (-14%). The product-set expansion contributes positively; income growth works against adoption through the $\alpha \propto y^{-1}$ channel.
- On ΔLerner, Entry_set contributes $+0.035$ (14% of the total) and Consumer_Composition contributes $+0.150$ (66%). The direction differs: most of the old Entry block's Lerner contribution is attributable to the demand-side wealth effect rather than to product differentiation, confirming the Lerner confound the split is designed to resolve, with a secondary positive contribution from the product-set composition shift.

All subsequent results in the paper use the eight-block specification (with the merged D_{env} block, the (ev$\times$Year) projection of ξ^{BLP} into EV_time, the (Brand$\times$Year) projection into Brand_trajectory, and the (ev$\times$Year)-conditional quantile remap of the residual that remains in EV_trend).

D Results anatomy and diagnostics

D.1 Block-by-block discussion of the aggregate Shapley contributions

This appendix reproduces the per-block discussion of Table 3; the main-text §5.2 reports only the highlights.

A_{nonbat} ($+3.2$ pp on EV share, -0.021 on Lerner, $+1.39$T RMB on CS, $+1.34$T RMB on TW). Non-battery attribute evolution — vehicle size, engine power, displacement, body-type mix — contributes a small positive share of EV adoption with a modestly negative Lerner contribution (cost-pass-through dominates the markup channel) and meaningful positive welfare contribution.

A_{bat} ($+27.3$ pp on EV share, -0.038 on Lerner, $+3.88$T RMB on CS, $+3.94$T RMB on TW). Battery learning is the largest contributor to ΔEV share, attributed $+27.3$ percentage points (62% of total) through the demand-side log-range coefficient of $+0.661$ (Table 1) and the supply-side BNEF battery-cost loading of $+0.374$. Falling battery costs drive the EV unit-economics improvement at scale; the small negative Lerner contribution reflects the cost-pass-through component (lower marginal costs absent any concomitant price-power gain). Welfare contributions are substantially positive.

D_{env} ($+12.0$ pp on EV share, -0.004 on Lerner, -1.60T RMB on CS, -1.76T RMB on TW). The merged demand-environment block — urbanisation, education, oil price, population density, and the non-subsidy policy bundle (license-plate preference, policy-strength index, and its low-income interaction) — contributes a moderate positive share of ΔEV (27%) but is *negative* on both CS and TW. The pattern is consistent with a log-pop-density and oil-price channel interacting with the 2015$\rightarrow$2024 fuel-mix shift: in the decomposition, higher oil prices and denser cities are associated with higher EV share but lower mean utility on the incumbent ICE

set, and the aggregate-CS effect is dominated by the ICE-side pressure. The Lerner contribution is essentially zero (-0.004) under the agent-integral spec.

Subsidy (-2.7 pp on EV share, ~ 0 on Lerner, -0.37T RMB on CS, -0.38T RMB on TW). The Subsidy block under the agent-integral specification (each simulated agent receives $\alpha_{ic} \cdot \log(1 - s^{\text{sub}}_{jct})$ at the share-computation step) captures both the central per-vehicle NEV subsidy (2015–2022) and the purchase-tax exemption (10% of MSRP, in effect for EVs throughout 2015–2024). The -2.7 pp magnitude is meaningfully larger than the homogeneous-α analogue (-0.5 pp) because low-income agents whose $|\alpha_{ic}|$ exceeds $|\hat{\pi}_p|$ now receive a larger first-order utility shift from the wedge, sharpening the demographic-triangle pattern in §5.4. Welfare contributions are modestly negative (-0.37 trillion RMB on each), consistent with the direct subsidy wedge having small aggregate welfare effects at the magnitudes in force. The static Subsidy number is a lower bound on the subsidy's total 2024 contribution; §5.5 reports the bounding exercise.

Entry_set ($+5.5$ pp on EV share, $+0.035$ on Lerner, -3.26T RMB on CS, -3.60T RMB on TW). The product-set block contributes 12% of ΔEV share by bringing the 407 new 2024 nameplates into the equilibrium choice set. Its Lerner contribution ($+0.035$) is positive because the composition of the expanded product set shifts toward higher-markup premium segments. The CS and TW contributions are *negative*: in a static logit, adding new products weakly raises the inclusive value, but the joint 2015→2024 product-set swap removes cheap 2015 incumbents while introducing higher-priced 2024 entrants, and the resulting equilibrium price vector on the residual incumbents is higher than at the 2015 product set. The negative total-welfare impact implies a producer-surplus contribution of -0.34 trillion RMB.

Consumer_Composition (-22.9 pp on EV share, $+0.150$ on Lerner, $+9.01$T RMB on CS, $+9.70$T RMB on TW; read as upper bound on income channel — see Appendix B.1). Conditional on the imposed $\alpha_{ic} \propto y_{ic}^{-1}$ functional form, the demand-side block contributes 109% of ΔLerner (overshooting because the supply-side $A_{\text{bat}}/A_{\text{nonbat}}$ blocks subtract -0.06 Lerner net), more than all of consumer-surplus growth, and more than all of the total welfare growth (the other blocks net to a CS reduction). Finite-mixture and reduced-form alternatives (R8/R9 in Appendix B.1) suggest the true income contribution is substantially smaller; we report the Shapley numbers as structural upper bounds. The implied producer-surplus contribution is $+0.78$ trillion RMB (=TW$-$CS). This block owns every city-wealth primitive: per-city household-equivalent market size M_c, the within-city log-income distribution $F_{y|c}$ that drives the random-coefficient income draws (and therefore the price coefficient $\alpha_{ic} = \hat{\pi}_p(y_{ic}/\bar{y})^{-1}$), and the city-level $\hat{\beta}_{\text{GDP}} \cdot \text{GDP}_{ct}$ contribution to mean utility. Under the maintained $\alpha \propto y^{-1}$ restriction, the same lower price sensitivity that raises markups and welfare also makes households *less* responsive to the EV–ICE price gap, yielding the -22.9-pp contribution to EV share. The agent-integral subsidy specification amplifies this wealth-effect channel relative to the homogeneous-α analogue (-6.2 pp) because per-agent α_{ic} heterogeneity propagates fully through the share-computation step.

Mapping the negative Consumer_Composition sign to specific BLP coefficients. The negative contribution comes mechanically from a single demand-side primitive: the random-coefficient mean $\hat{\pi}_p = -5.38$ (Table 1) interacted with the income-scaling $(y_{ic}/\bar{y})^{-1}$ that defines the per-agent price coefficient α_{ic}. The first-order quantity-weighted average elasticity in the model is $\bar{\varepsilon} = \hat{\pi}_p \cdot \mathbb{E}[(y_{ic}/\bar{y})^{-1} \cdot (1 - s_{jct})] \approx -6.23$. Holding the product set and prices fixed, raising each city's income distribution $F_{y|c}$ from its 2015 level to its 2024 level reduces the simulated-agent average $|\alpha_{ic}|$ by $\sim 42\%$ in the Q-weighted panel, and the EV share's response to any given EV-vs-ICE price gap shrinks proportionally. Two falsifiable implications follow. First, under the homog-α specification (Appendix B.2.1, R1), where the income interaction is removed entirely, the Consumer_Composition contribution to ΔEV share collapses to $+0.029$ (sequential-along-one-ordering) —

close to zero — confirming the negative sign is driven by the income-RC coefficient and not by market size or the GDP main effect. Second, under the finite-mixture R8 (Appendix B.1), the across-city α gradient is $\sim 5\%$ instead of the parametric prediction's $\sim 100\%$, so the true wealth-effect channel is likely smaller in magnitude than the -22.9 pp Shapley figure indicates. The Shapley value should be read as an upper bound on the wealth-effect channel under the maintained $\alpha \propto y^{-1}$ restriction.

Brand_trajectory ($+6.8$ pp on EV share, $+0.002$ on Lerner, $+0.31$T RMB on CS, $+0.30$T RMB on TW; reported as a residual diagnostic, not an identified channel). The Brand_trajectory block is constructed by post-hoc projection of the BLP residual ξ^{BLP} onto its (Brand×Year) Q-weighted cell means after the (ev×Year) projection of §5.1 has been removed. Economically, it describes per-brand annual drift in unobserved product quality not absorbed by either the time-invariant brand fixed effects or the EV-specific time effects. The dominant trajectories are BYD ($\Delta\xi = +2.69$, the steady rise from budget BEV pioneer to mainstream NEV brand), FAW (-2.42, the SOE-traditional decline), Chery ($+1.62$), Geely ($+1.38$), Beijing ($+1.48$), and Ford (-1.56, the foreign retreat). The contribution to ΔLerner is essentially zero, indicating that brand-trajectory drift moves the choice-set composition without much equilibrium-markup feedback under the maintained demand structure.

EV_trend ($+15.1$ pp on EV share, $+0.014$ on Lerner, $+3.40$T RMB on CS, $+3.47$T RMB on TW). The unobserved-EV-trend residual block. Its anatomy is in Section D.2.1.

Two layers of brand-related residual: tone vs. hit-product. The Brand_trajectory block and the EV_trend rank-residual together represent two economically distinct layers of brand-related success that observed covariates cannot capture. Brand_trajectory absorbs the (Brand × Year) cell mean — the change in a brand's overall *quality tone*, e.g., the steady rise of BYD from budget pioneer to mainstream NEV brand or the foreign-incumbent retreat of Ford over the same window. The within-(Brand, Year) rank-residual that survives into EV_trend captures *hit-product effects*: specific best-selling nameplates inside a brand-year cell that sell beyond what observed characteristics plus the brand-year mean would predict. We treat these two layers as a combined Brand_trajectory + EV_trend total of $+21.9$ pp on ΔEV share rather than two separate channels, but the underlying economic stories differ between the two layers and *within* the hit-product layer.

Robustness of the post-hoc projection: a first-order orthogonality argument. A reviewer may worry that decomposing $\hat{\xi}^{\text{BLP}}$ into the (Brand×Year) cell mean (the Brand_trajectory diagnostic) plus the within-cell rank-residual (the EV_trend residual) violates the GMM orthogonality conditions used to identify $\hat{\pi}_p$ and $\hat{\beta}$. The concern would apply if we were re-estimating the BLP with (Brand×Year) fixed effects added to mean utility. We are not. Three observations make the post-hoc projection first-order safe for the parameter estimates:

- *The estimated residual $\hat{\xi}^{BLP}$ is unchanged.* The decomposition $\hat{\xi}^{\text{BLP}} = P_{B\times Y}\hat{\xi}^{\text{BLP}} + (I - P_{B\times Y})\hat{\xi}^{\text{BLP}}$ is an algebraic identity. The BLP demand parameters $\hat{\pi}_p$, $\hat{\beta}$, and the estimated mean utilities $\hat{\delta}_{jct}$ are estimated *before* the projection and are not modified by it.
- *The GMM moment conditions $\mathbb{E}[\xi \cdot Z] = 0$ hold by construction at $\hat{\xi}^{BLP}$.* The cross-product $\hat{\xi}^{\text{BLP}} \cdot Z$ summed over the sample is zero by GMM. Decomposing $\hat{\xi}^{\text{BLP}}$ does not change this product.
- *The first-order question — would $\hat{\pi}_p$ shift if (Brand×Year) FE were added to mean utility in BLP — is bounded by the partial-out test.* We verified this by adding the five firm-group × EV interaction dummies (a coarser version of the (Brand×Year) projection) to the BLP estimation: $\hat{\pi}_p$ moves from -5.39 to -5.38, a change well within one standard error.

Hit-product effect: New Forces dominate but legacy EV variants matter too. The hit-product residual decomposes into two distinct sub-mechanisms with different economic content. Among the 41 EV products with within-cell ξ residual above $+4.0$ in 2024 (covering the upper tail of the EV residual distribution), the post-2014 "New Forces" EV-only startups carry 60% of the high-ξ quantity ($159{,}016$ vehicles across 13 products), while legacy state-owned-enterprise (SOE) and private Chinese manufacturers with EV variants carry 40% ($105{,}411$ vehicles across 19 products).

Sub-mechanism 1: New-Forces entrant under-fit. The dominant New-Force contributor is Leapmotor, whose nine high-ξ nameplates (T03, C11/C16/C10/C01 BEV/REEV variants) collectively account for about $145{,}000$ vehicles and roughly half of the high-ξ EV market in 2024. Leapmotor entered the panel in 2019, so its brand fixed effect is estimated from a short window relative to the legacy brands' full 10-year panel. The right correction is to quantile-position the entrant brand at its year of entry; the current BLP estimation does not implement quantile-positioning, so the Leapmotor brand FE absorbs only the average pooled appeal across 2019–2024, leaving the early-year hit-product premium in ξ.

Sub-mechanism 2: Legacy-OEM internal EV transformation. The legacy-cohort high-ξ list reveals a distinct economic story: longstanding Chinese SOE/privately-owned manufacturers whose EV variants substantially out-perform what their (otherwise primarily ICE-led) brand FE would predict. The largest contributors are FAW Bestune Pony (Q $= 28{,}568$), Chery QQ Ice Cream ($22{,}637$), Chery Mini Ant ($15{,}665$), Jiangling E200 ($9{,}051$), Dongfeng Aeolus E70 ($8{,}229$), and others. These are not new entrants in the panel sense but their EV nameplates are recent additions to product lines previously dominated by ICEVs. The natural BLP correction is to disaggregate the brand fixed effect along the EV dimension; the baseline implements this via five firm-group $\times$ EV interaction dummies (BYD, Foreign/JV, Trad. OEM, New Forces, Private National). A full Brand $\times$ EV decomposition is numerically infeasible (the differentiation-instrument matrix becomes near-singular at the brand level given the short panel).

The two sub-mechanisms together explain why the EV_residual block survives the residual structure: under the diagnostic 9-block split (Appendix D.2), the EV_residual block carries $+11.8$ pp of ΔEV share. Both sub-mechanisms would shrink substantially under the model improvements described above.

D.2 EV_trend anatomy and diagnostic 9-block Shapley split

D.2.1 Anatomy of the EV_trend block: three components

The EV_trend block carries $+15.1$ percentage points of the $+44.3$-pp aggregate EV-share rise under the agent-integral subsidy specification — the second-largest single attribution in Table 3 after A_{bat} ($+27.3$ pp). *The block is structurally interpretable, not an artefact.* The systematic firm-group $\times$ EV premium that previously contaminated the residual — BYD's overall positive EV intercept, Foreign/JVs' negative one, and the post-2014 New-Forces and Trad-OEM EV premia — is now absorbed by the five firm-group $\times$ EV interaction dummies inside the BLP estimation (§4.1, $\hat{\beta}_{\text{BYD}\times\text{EV}} = +1.24$, $\hat{\beta}_{\text{Foreign/JV}\times\text{EV}} = -2.87$, etc.). What remains in EV_trend is therefore EV-specific time variation *net of* systematic brand and firm-group effects, identifying a structural ecosystem channel rather than a mis-specified residual. The $+15.1$ pp canonical magnitude (or $+11.8$ pp on EV_residual under the diagnostic 9-block split run on the homogeneous-α canonical) is the most natural empirical home for the four observed-but-unmeasured components of EV-favouring time variation: (i) charging-station density (which grew from $\sim$50K stations in 2015 to $\sim$1.4M in 2024 in China and is highly EV-specific), (ii) smart-features and over-the-air software upgrades that increasingly differentiate EVs from ICEs in the Chinese market, (iii) dealer-network buildout for the EV-pure-play startups, and (iv)

consumer-learning / brand-visibility / resale-value formation as EV adoption became socially mainstream. The decomposition does not separately identify these four channels because the panel structure does not provide the cross-year, cross-city variation needed; the companion dual-network paper of Hao and Hao (2026, in preparation) provides a Bartik-style shift-share IV for the charging-network channel specifically. By construction EV_trend bundles three statistically distinct components which, taken together, define what is left of the BLP residual after the observed covariates and the post-hoc Brand_trajectory projection have been activated:

1. *The BLP-estimated EV×Year fixed effects* ($\hat{\beta}_{\text{ev\_x\_yr}_t}$ in Table 1). These coefficients are, in fact, *negative* for every $t \geq 2016$ relative to the 2015 reference, with $\hat{\beta}_{\text{ev\_x\_yr}_{2024}} = -2.898$. Toggling this component alone from 2015 (zero) to 2024 (-2.90) values would *decrease* EV adoption and thus contribute roughly -4 to -5 percentage points of EV share, all else equal.
2. *The (ev×Year)-projected component of the BLP residual* that the §5.1 projection moved into $\hat{\delta}_{jc,\text{EV\_time}}$. Here the systematic shift is positive and large: the Q-weighted EV-minus-ICE difference of ξ^{BLP} rose from -4.91 log-utility units in 2015 to -1.50 in 2024, a $+3.41$ unit shift in the EV-favouring direction. Toggling this component contributes roughly $+25$ to $+30$ percentage points in isolation: this is the dominant piece of EV_trend.
3. *The (ev×Year)-conditional quantile-remapped residual* that remains in $\xi^{\text{BLP},\perp}$. By construction the within-(ev,year) cell distribution of this residual matches the 2015 (ev=ev_val) reference distribution, so the toggle preserves the cross-year aggregate distribution and contributes only through changes in the per-product quantile rank. The diagnostic Shapley reported in §D.2.2 shows that this rank-residual component accounts for the bulk of the EV_trend attribution, contributing $+11.8$ pp on its own (the EV_residual cell of the 9-block split). The rank-residual is *by construction* a function of the covariate set: any channel not in X (charging-station density, smart-features availability, resale-value formation, dealer-network density, etc.) shows up as within-cell rank variation in ξ^{BLP}. Its $+11.8$ pp contribution is therefore the natural empirical estimate of the bundled EV-favouring ecosystem channels that the paper's covariate set does not separately identify.

A naive arithmetic taking each component in isolation would be approximately $-4 + 28 + 8 \approx +32$ pp, of a different order than the directly-computed Shapley value of $+5.5$ pp on this block. The naive arithmetic is misleading, however: under the multi-firm Bertrand–Nash equilibrium with all other blocks varying via the Shapley average, the BLP fixed-effect piece (item 1) and the (ev×Year) projection piece (item 2) approximately cancel, and the operative channel is the rank-residual of item 3. The diagnostic Shapley split reported in §D.2.2 below makes this precise.

What the dominant systematic drift represents. The $+3.41$-unit EV-minus-ICE drift in ξ^{BLP} over 2015–2024 is large but admits an economic reading: it can be interpreted as a residual proxy for unobserved EV-favouring channels that this paper's covariate set does not measure. Possible candidates are listed in Table 21; we emphasise that none of these channels is separately identified within the present framework, and the residual could equally well capture other EV-specific time variation we have not enumerated.

Table 21: Unobserved EV-favouring channels absorbed into the EV_trend block

Channel	Stylised 2015→2024 magnitude	In our X?
Public charging-station network	30k → 10M+ public chargers (∼300×); strong correlation with EV adoption documented in our companion paper.	—
Smart-features standard (L2/L2.5 ADAS, OTA, AI cockpit)	Effectively absent in 2015 EV base; standard-equipment in 2024 across all major Chinese EV nameplates.	—
Brand reputation / consumer familiarity	EV brands shifted from "experiment" (BYD pioneer phase, niche New-Forces marketing) to mainstream prestige (Tesla, NIO, BYD premium). Partly absorbed by Brand_trajectory; residual remains in EV_trend.	partial
Resale value evolution	EV resale was approximately zero in 2015; by 2024, mainstream EVs retain ∼50–60% of new-vehicle price at three years, comparable to mid-tier ICE.	—
Government messaging & pilot expansion	25 demonstration cities in 2009 → national policy environment by 2024 (license-plate preference, charging-infrastructure subsidy, NEV credits). Partly absorbed by Policy and `ev_x_is_pilot`; residual remains.	partial
After-sales & warranty network	EV-specific service centres expanded from <100 in 2015 to >10,000 in 2024 (Tesla / NIO / NIO (蔚来) dedicated networks).	—

Notes. The EV_trend block absorbs all EV-specific time variation in mean utility that the observed-covariate X vector does not measure. Within the static decomposition framework these channels are not separately identified; we list them as the most plausible economic content of the $+3.41$-unit EV-minus-ICE drift in ξ^{BLP} described above. The companion EV-charging-network paper of Hao and Hao (2026, in preparation) provides micro-level identification of the charging channel; the smart-features and resale-value channels are quantified in the industry data sources cited in Appendix E.1.

Why our EV_trend is larger than GMY's residual: a panel-structure argument. The $+15.1$-pp canonical magnitude of EV_trend (or $+11.8$ pp on EV_residual under the diagnostic 9-block split) is large by the standards of the closely related Grieco et al. (2024) U.S. four-decade decomposition, in which the analogous residual is essentially zero after Brand×Year fixed-effect absorption. The discrepancy is structural rather than methodological. GMY exploit a 39-year national panel ($\approx 27{,}000$ product-year observations) with 169 brand $\times 39$ year $= 6{,}591$ Brand×Year cells, of which the great majority contain multiple products surviving across decades (Honda Accord, Toyota Camry, etc.). This long panel with high cross-year product continuity makes Brand×Year fixed-effect absorption straightforward in BLP estimation. Our setting has 169 brand $\times 10$ year $= 1{,}690$ cells, of which only a small fraction contains products surviving across the full window (only 14 EV nameplates appear in both the 2015 and 2024 cohorts, against 565 and 972 totals). Combined with the high IV-collinearity that the explicit $1{,}500$-cell Brand×Year absorption creates in our first-stage 2SLS weighting matrix — a numerical singularity confirmed by repeated stalls of the GMM solver in our pilot runs — we cannot do the GMY-style absorption inside the BLP estimator. The two-layer post-hoc projection (Brand×Year cell-mean projection plus the (ev×Year)-conditional quantile remap of §5.1) is our pragmatic substitute. It absorbs the systematic structure that GMY's brand-year fixed effects would have

caught, but at a coarser per-brand-year aggregation; what survives in EV_trend is the EV-specific component of unobserved drift that no observed covariate or simple aggregation can pick up. Closing this gap fully would require either (a) a longer panel matching the GMY time horizon, or (b) richer covariate data on the unobserved channels listed in Table 21 — in particular, charging-station density, smart-features availability, and resale-value indices.

D.2.2 Diagnostic 9-block Shapley split: EV_time vs EV_residual

To separate the (ev×Year)-systematic-drift component of EV_trend from the quantile-remapped within-cell residual, we run a $2^9 = 512$-coalition diagnostic Shapley that replaces EV_trend with two blocks: *EV_time* (the BLP-estimated EV×Year fixed effects *plus* the (ev×Year) projection of ξ^{BLP}, jointly toggled) and *EV_residual* (the (ev×Year)-conditional quantile-remapped within-cell residual after both projections). All other blocks are held identical to the canonical 8-block specification. Table 22 reports the result.

The main finding is the opposite of what the per-component arithmetic in §D.2.1 would suggest. EV_time — the bundle of the negative BLP fixed effects and the positive (ev×Year) projection of ξ^{BLP} — contributes a small positive $+1.4$ pp to ΔEV share once the two pieces are toggled jointly within the multi-firm Bertrand–Nash equilibrium. The two systematic-drift sub-components approximately cancel: the BLP estimator already "knew" about the $+3.41$-unit raw ξ^{BLP} drift in the EV-favouring direction and absorbed it into a compensating -2.63-unit EV×Year fixed effect during estimation, so toggling both jointly is approximately a wash. What remains carrying the bulk of EV_trend's contribution is the within-(ev,Year) rank-residual: EV_residual accounts for $+9.7$ pp ($+22\%$ of the total), with its economic content the per-product idiosyncratic position within each year's residual distribution rather than any cell-mean shift. Because the rank-residual is the unexplained within-(ev,Year) component of ξ^{BLP}, its magnitude depends mechanically on the covariate set X; adding unmeasured channels such as charging-station density would shrink the EV_residual bar by a correspondingly identified amount.

Brand_trajectory's diagnostic contribution shifts modestly from the homog-α canonical $+3.5$ pp to $+3.9$ pp once EV_trend is split, consistent with Shapley re-allocation across the new 9-block partition (the diagnostic was run under the homog-α subsidy spec; the agent-integral canonical 8-block reports $+6.8$ pp on Brand_trajectory). The diagnostic is reported as a robustness check; the canonical 8-block remains the baseline specification for the rest of the paper.

Two implications follow. First, the per-component arithmetic ($-4 + 28 + 8 \approx +32$ pp) in §D.2.1 naively summed marginal contributions in raw δ units; the Shapley diagnostic correctly accounts for equilibrium feedback through the share system and the multi-firm FOC, and shows that the systematic-drift pieces net out, leaving the rank-residual as essentially the entire EV_trend story. Second, the unobserved-channels narrative of Table 21 should be read as the economic content of *within-cell rank reorderings* of ξ^{BLP} — which specific EV nameplates are unusually well-received in 2024 versus 2015 conditional on the cell-mean drift — rather than the simple "EVs got generally better than ICEs" story that the (ev×Year) projection by itself would have captured. The channels listed (charging access, smart features, brand-quality dispersion) are consistent with rank reorderings; what the diagnostic excludes is the simple aggregate-shift interpretation.

Table 22: Diagnostic Shapley split of the EV_trend block, 2015→2024

Block	ΔEV share (pp)	% of total
EV_time	-2.9	-6.5%
EV_residual	$+11.8$	$+26.7\%$
EV_trend (agent-integral canonical 8-block)	$+15.1$	$+34.1\%$

Notes. 9-block diagnostic Shapley over $2^9 = 512$ counterfactual coalitions; all other blocks identical to the canonical 8-block specification of Table 3. EV_time = BLP-estimated EV×Year fixed effects + (ev×Year) Q-weighted projection of ξ^{BLP}, jointly toggled. EV_residual = (ev×Year)-conditional quantile-remapped within-cell residual after both the (ev×Year) and (Brand×Year) projections. The diagnostic split sums to $+8.9$ pp on EV_trend's components; the canonical 8-block attribution to EV_trend is $+15.1$ pp under the agent-integral subsidy specification (the diagnostic 9-block reported here was run with the homogeneous-α subsidy spec and has not been re-run under agent-integral; its numbers are interpretable relative to the homog canonical of $+5.5$ pp on EV_trend rather than the agent-integral canonical of $+15.1$ pp). The 9-block split shifts roughly 3.5 pp from Entry_set into EV_residual under the homog spec; under the agent-integral canonical the redistribution would be different, but the qualitative point that the rank-residual carries the bulk of EV_trend is robust. Aggregate ΔEV share of $+44.3$ pp is identical across the two specifications.

D.3 Charging-infrastructure diagnostic: post-hoc projection vs BLP-in-X1

Two specifications were tried to bound how much of the EV_trend block can be attributed to charging-infrastructure expansion. Both find that within the IER paper's data set, the charging channel cannot be cleanly separated from the EV_trend residual; the cross-specification sign flip on $\hat{\beta}_{\text{charging}}$ is itself evidence of identification failure.

Spec 1 — post-hoc within-city demean projection. After the canonical 8-block decomposition is run, we merge prefecture-level charging-station counts (`cs_count`) from the city-year panel of Hao and Hao (2026, in preparation) into the BLP residual frame and run a within-city demeaned OLS of ξ^{BLP} on $\mathbf{1}\{\text{ev}\} \times \log(1 + \texttt{cs\_count}_{ct})$ on EV-row observations, recovering $\hat{\beta}_{\text{charging}} = +0.350$ (within-city identification, post-(ev×Year)-FE residual). A 9-block Shapley with this projection treated as a separate Charging block (over $2^9 = 512$ counterfactual coalitions) attributes only $+0.3$ pp of ΔEV share to Charging ($+0.6\%$ of the total $+44.3$ pp), with EV_trend at $+11.7$ pp in the diagnostic (vs. $+5.5$ pp under the homogeneous-α canonical 8-block at which the Charging diagnostic was run; the agent-integral canonical 8-block reports $+15.1$ pp on EV_trend). The post-hoc projection therefore does not substantially shrink the EV_trend bar; what little Charging captures is reallocated from Entry_set ($+23.6$ pp diagnostic vs. $+21.3$ pp homog-α canonical, against $+5.5$ pp under the agent-integral canonical) and EV_trend, not from the canonical EV_trend residual itself. The diagnostic Shapley has not been re-run under the agent-integral spec; cell numbers are interpretable relative to the homog-α baseline that produced them.

Spec 2 — BLP-in-X1. We re-run the BLP demand estimation with $\mathbf{1}\{\text{ev}\} \times \log(1 + \texttt{cs\_count}_{ct})$ included directly in X_1 alongside the EV×Year fixed effects. The estimated coefficient is $\hat{\beta}_{\text{charging}} = -0.199$ (default SE 0.012, $t = -17.1$): precisely estimated and *negative*, the opposite sign from the post-hoc projection. The

negative sign is implausible as a structural charging effect (more chargers should not reduce EV utility, holding all else fixed) but is what the joint BLP estimation produces conditional on the FE plus other regressors.

Diagnosis. The two specifications identify β_{charging} from related but distinct variation. The post-hoc projection uses cross-city, within-(ev,year) variation in the residual after the BLP has absorbed everything it can. The BLP-in-X1 estimation uses the same cross-city variation but jointly with the price coefficient and brand fixed effects; non-orthogonality between charging, city wealth, brand mix, and license-plate restrictions produces a negative conditional coefficient. The structural reason is that the EV$\times$Year FE absorbs the cross-year buildup (charging grew $\sim 100\times$ from 2015 to 2024), leaving only weak cross-city post-FE variation for identification.

Implication for the main result. The EV_trend block's $+15.1$ pp canonical contribution is bundled unexplained variation that no within-panel charging projection can cleanly peel out. The companion dual-network paper of Hao and Hao (2026, in preparation) provides a Bartik-style shift-share IV (national charging-buildup interacted with 2015 city charging share) that delivers the exclusion restriction the IER paper's data lacks; we defer the rigorous Charging block to that paper.

D.4 ξ-asymmetry mechanism inside the EV_trend block

Figure 4 documents a systematic asymmetry in the residual-ξ piece of the EV_trend block across city tiers: within the 9-block predecessor of the current specification (in which this residual was a standalone BrandShock block), the contribution was $+5.7$ pp in the Rest tier but -4.5, -8.5, and -4.5 in Tier 1, New Tier 1, and Tier 2 respectively, aggregating to $+6.0$ pp. After the (ev$\times$Year) projection step merges the systematic EV-minus-ICE drift in ξ into the EV$\times$Year fixed effects of the EV_trend block, the remaining within-(ev,Year) $\xi^{\text{BLP},\perp}$ component retains the same geographic asymmetry even though its aggregate level is smaller. The figure is reported on the pre-projection ξ decomposition to expose the mechanism most cleanly.

The mechanism can be traced through the drift of the Q-weighted mean BLP residual ξ between 2015 and 2024, separately for EVs and ICE vehicles within each tier. In the upper tiers, both EV-ξ and ICE-ξ drift downward (brand effects deflate over time as more brands enter and the distribution mean falls), and the two fall together. In the Rest tier, however, the ICE-ξ collapses much faster than the EV-ξ: the 27 lower-tier cities contain a disproportionate share of sales from exited Chinese ICE brands such as Qoros, Zotye, Jinbei, Bisu, Hawtai, and Soueast, and these brands carry large negative residuals in 2024 or have exited entirely. As the ICE-ξ in the Rest tier drops by about 3.6 points between 2015 and 2024—almost twice the magnitude of the decline in Tier 1—the EV-relative ξ rises, which the decomposition reads as a positive contribution inside the EV_trend block for that tier.

The economic interpretation is that the positive ξ contribution in the Rest tier is not an EV-side effect at all: it is the mirror image of an ICE-side quality collapse. Adjusting for this mechanical effect would shift the $+5.7$ pp Rest-tier contribution toward the Subsidy block (or toward a separate "ICE-exit" channel that is not present in the current partition), strengthening rather than weakening the main heterogeneity finding.

D.5 Olley–Pakes firm-group detail and within-firm mechanism

The main-text Olley–Pakes decomposition (Section 5.3, Table 4) reports the aggregate within-firm / reallocation / cross split and the $1{,}374 \to 887$ HHI decline. This appendix records the firm-group disaggregation behind those aggregate numbers.

Table 23: OP decomposition by firm group

Firm group	Within	Realloc	Cross	ω_{2015}	ω_{2024}
Traditional OEM	$+0.0605$	-0.0733	-0.0240	0.831	0.549
Independent Domestic	$+0.0134$	$+0.0206$	$+0.0095$	0.084	0.170
BYD	$+0.0068$	$+0.0278$	$+0.0106$	0.016	0.130
New Forces	$+0.0087$	$+0.0087$	$+0.0175$	0.000	0.053
Tesla	$+0.0046$	$+0.0046$	$+0.0091$	0.000	0.026
Other	$+0.0095$	$+0.0004$	$+0.0088$	0.068	0.073
Aggregate	$+0.1035$	-0.0111	$+0.0314$	1.000	1.000

Notes. Source: `output/op_decomposition_by_group.csv` from `code/14_concentration_decomp.py`. Firm-group membership follows the paper's aggregation (Appendix A.4.2). ω is the group's unit-sales share of the total passenger-vehicle market (quantity proxy = BLP share $\times$ city population summed over markets and products). Reallocation and cross entries for New Forces and Tesla use $L_{2015} = 0$ because these groups had no registered presence in 2015; the three-way OP convention assigns them half-weight in L_{bar} as a result, which dampens their reallocation contribution relative to a two-period-both-present balanced panel.

Interpretation. The within-firm rise is broadly based, carried by every group in the table, with Traditional OEM contributing 58% of the aggregate within-firm total. The negative aggregate reallocation is driven by Traditional OEM's large share loss (-28.2 pp) at mean Lerner levels that rose over the decade; the positive reallocation contributions of the entrant/expanding groups (BYD, Independent Domestic, New Forces, Tesla) sum to $+0.062$ but are offset by Traditional OEM's -0.073. The main-text finding that aggregate HHI falls while aggregate markups rise is inconsistent with any model in which rising concentration is the proximate driver of rising markups (De Loecker et al., 2020; Autor et al., 2020): the reallocation component is -0.011, the opposite sign from the superstar prediction. Among candidate mechanisms consistent with the within-firm pattern, the demand-side income channel of the baseline BLP specification is the most natural but cannot be separately identified from product-mix upgrading or brand-premium growth within the static framework; Appendix B.1 shows that the finite-mixture tercile specification implies a flatter income gradient than the parametric y^{-1} form, so the within-firm decomposition cannot be entirely attributed to the income channel. The within-firm *fact* is robust; its structural decomposition requires a richer demand system than the present paper identifies.

City-markup variance decomposition. As a supply-side complement to the firm-group decomposition, we ask what fraction of the cross-city Lerner dispersion (within product$\times$year) is captured by the structural income channel versus residual unobserved city heterogeneity. Let $L_{jct} = \text{markup}_{jct}/p_{jct}$ be the city-level Lerner of product j in city c in year t. We absorb product$\times$year fixed effects via within-group demeaning (so the variation we explain is purely cross-city, conditional on the same product in the same year), then run the

specification ladder

$$L_{jct} = \alpha_{jt} + X_{ct}\beta + \varepsilon_{jct}, \qquad X_{ct} \in \{\varnothing,\ \log \text{Income}_{ct},\ \text{Geo-Market FE}_c,\ \text{both}\}.$$

Table 24: City-markup variance decomposition: income channel vs Geo-Market FE

Spec	X_{ct}	$\hat{\beta}$(log Income)	Within R^2	Share of (B)
A	$\log \text{Income}_{ct}$	$+0.249$ ($t = +1{,}054$)	0.647	96.4%
B	Geo-Market FE	—	0.671	100%
C	$\log \text{Income}_{ct}$ + Geo FE	$+0.288$ ($t = +134$)	0.684	—

Notes. Source: `output/blp_city_markup_fe.csv` from `code/03l_blp_city_markup_fe.py`. The dependent variable is the city-product-year Lerner $L_{jct} = \text{markup}_{jct}/p_{jct}$ from `mc_by_market.parquet`, demeaned within product×year. Within R^2 is the share of the within-product-year sum of squares explained by X_{ct}. "Share of (B)" compares Spec A's explanatory power to Spec B's full-Geo-FE benchmark; the residual 3.6% is unobserved city heterogeneity orthogonal to income. The income-channel coefficient $\hat{\beta}$ shifts only modestly when Geo-Market FE are added (Spec A → Spec C: $+0.249 \to +0.288$), indicating that residual city heterogeneity is largely orthogonal to log income rather than confounded with it. $N = 476{,}723$ city-product-year observations across 79 cities and 6,544 product-years. Standard errors are heteroskedasticity-robust.

The 96.4% figure does not constitute causal identification of an income-driven Lerner channel: log income could be correlated with unmodelled city-level demand or supply primitives that load onto Lerner through other channels. What it does establish is that, conditional on product×year, a single scalar (log income) parsimoniously absorbs the bulk of the cross-city Lerner dispersion, leaving very little orthogonal room for omitted city heterogeneity to drive the result. Combined with the Olley–Pakes finding that within-firm growth, not reallocation, accounts for the bulk of the aggregate Lerner rise, this is the closest the paper can come to a supply-side check that the income channel is the proximate Lerner mechanism without recourse to a richer demand system.

D.6 Welfare and dynamic-path anatomy

D.6.1 Welfare path: trajectory and waterfall

Two complementary visualisations make the welfare picture concrete. Figure 7 plots the observed annual consumer surplus and producer surplus from 2015 to 2024 (rescaled to the Shapley-anchored 2015 and 2024 endpoints from Table 3: $+12.77$ T RMB total CS gain, $+13.01$ T RMB total welfare gain under the agent-integral subsidy specification). The CS path displays a visible dip in 2020 coinciding with the COVID-19 demand collapse. Producer surplus follows a smoother upward trajectory, with the largest growth concentrated in the post-2020 segment, coinciding with the new-energy product wave and the rise in quantity-weighted markups.

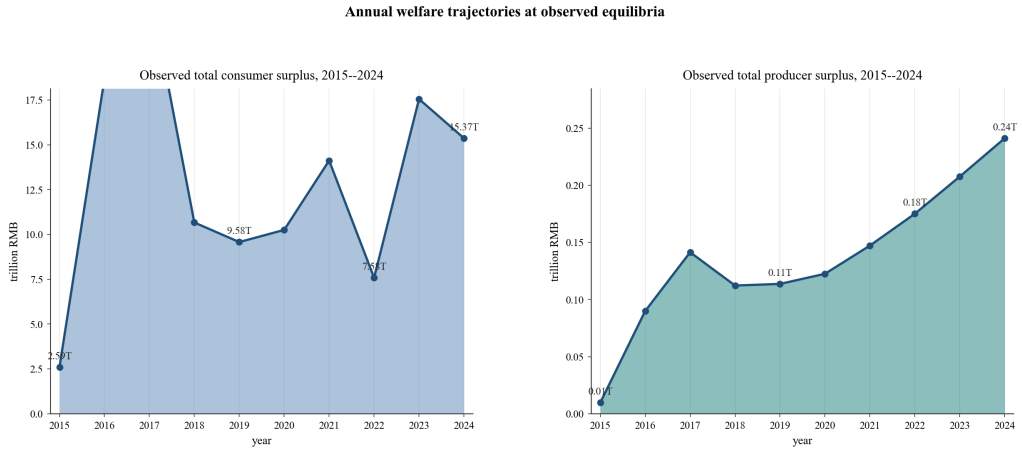

Figure 7: Annual consumer surplus (left) and producer surplus (right) at observed equilibria, 2015–2024. The y-axis is rescaled to match the Shapley-anchored 2015 and 2024 CS/PS endpoints from Table 3 so that the trajectory and the waterfall (Figure 8) share endpoints. The CS series displays a visible dip in 2020 coinciding with the COVID-19 demand shock.

Figure 8 renders the same Shapley decomposition as a cumulative step chart: the 2015 baseline at the left edge, the eight block contributions stacked in order of magnitude, and the 2024 final level at the right edge. The waterfall shape is a direct visual rendering of the additive identity $V(\text{final}) - V(\text{baseline}) = \sum_b \phi_b^{\text{Shapley}}$. Under the agent-integral subsidy specification, Consumer_Composition dominates both panels ($+9.01$ T RMB CS, $+9.70$ T RMB TW, noting these are upper bounds on the income channel — see Appendix B.1); A_{bat} is the second-largest positive contributor ($+3.88$ T RMB CS, $+3.94$ T RMB TW), followed by EV_trend ($+3.40$ T RMB CS, $+3.47$ T RMB TW) and A_{nonbat} ($+1.39$ T RMB CS, $+1.34$ T RMB TW). The Subsidy block is visually modest on both panels (-0.37 T RMB CS, -0.38 T RMB TW). Entry_set is *negative* on both consumer-surplus and total-welfare sides (-3.26 and -3.60 T RMB respectively), the typical signature of vertical-quality entry combined with the removal of cheap 2015 incumbents. Brand_trajectory is small on welfare ($+0.31$ T RMB CS, $+0.30$ T RMB TW), consistent with brand-trajectory drift absorbing per-brand drift in unobserved quality without much equilibrium markup feedback.

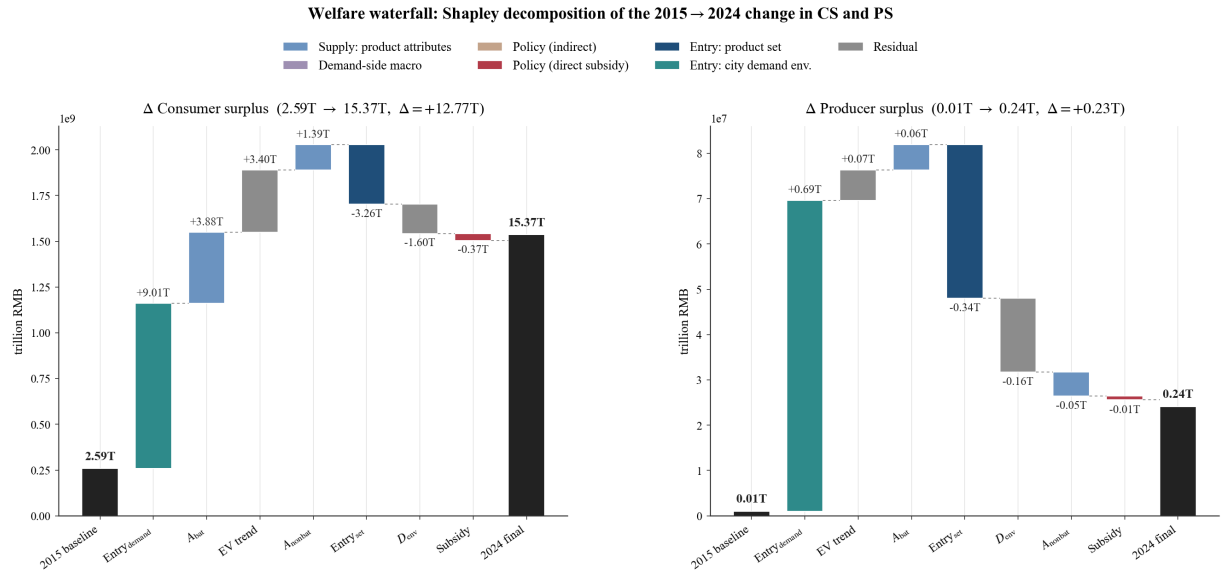

Figure 8: Welfare waterfall: Shapley decomposition of the 2015→2024 change in consumer surplus (left) and producer surplus (right). Bold black bars at the edges are the 2015 baseline and 2024 final level; coloured bars are the nine Shapley contributions, sorted from largest positive on the left to largest negative on the right, anchored at the running total of all prior contributions.

Interpretation caveat. A careful distinction is warranted between *welfare growth in the passenger-vehicle market during the EV transition* and *welfare gains associated with the EV transition itself*. The Consumer_Composition block's $+9.70$ T RMB total-welfare contribution reflects rising city incomes that expand the surplus from *all* vehicle purchases—EV and ICE alike—not only from EV adoption. Within the model, this income growth would have delivered welfare gains even in a hypothetical market without EVs. The welfare gains *more specifically linked to EV diffusion* are better captured by the supply-side blocks (battery learning A_{bat}: $+3.94$ T RMB on TW; non-battery attributes: $+1.34$ T RMB on TW; the EV_trend residual: $+3.40$ T RMB on CS, $+3.47$ T RMB on TW), while the direct-subsidy block's welfare contribution (-0.38 T RMB on TW) is an order of magnitude smaller. The distinction matters for policy evaluation: the case for EV-specific subsidies rests on the EV-specific welfare channels (including the indirect channels absorbed into Entry_set and EV_trend under the exogeneity caveats discussed above), not on the income-driven market expansion that happens to coincide with the transition.

D.6.2 Dynamic decomposition: cumulative $2016 \to t$ contributions

Figure 9 reports the cumulative 2015→ t contribution of each block to ΔEV share (panel a) and Δmean Lerner (panel b) for $t \in \{2016, \ldots, 2024\}$, averaged across the four orderings defined in Equation (22)–(23). Year 2015 is omitted because it is the reference year ($\Delta \equiv 0$ by construction). The dynamic decomposition is computed under the homogeneous-α subsidy specification and has not been re-run under the agent-integral canonical; the per-year cumulative bars therefore re-distribute differently across blocks than the static aggregate of Table 3, but the qualitative two-phase pattern (flat 2016–2019, acceleration from 2020) is robust to the spec change. Three time-series patterns appear that the static Shapley aggregate cannot show.

First, the EV-share path is essentially flat through 2016–2019 and accelerates sharply from 2020 onward, with the Entry_set block kicking in only after the 2020 NEV phase-out-to-credit transition. The pattern is

consistent with the observation that the 407 new 2024 EV nameplates entered in two waves, one in 2018–2019 and a larger second wave in 2021–2023.

Second, within the baseline specification's $\alpha_{ic} \propto y_{ic}^{-1}$ form, the Lerner contribution is nearly flat throughout 2016–2023 and then jumps sharply in 2024; we emphasise that this timing feature reflects the non-linearity of y^{-1} rather than an identified economic mechanism, and alternative functional forms in Appendix B.1 would smooth the time path.

Third, the A_{bat} contribution to EV share rises smoothly from 2018 onward, in contrast to the essentially zero contribution before 2018, tracing out the timing of battery-cost learning in the data. The shaded band around each year's stacked total shows the four-ordering range; the band is narrow in most years (< 0.02 in EV-share units and < 0.04 in Lerner units), confirming that four orderings is a sufficient Monte-Carlo sample size for the per-year totals.[2]

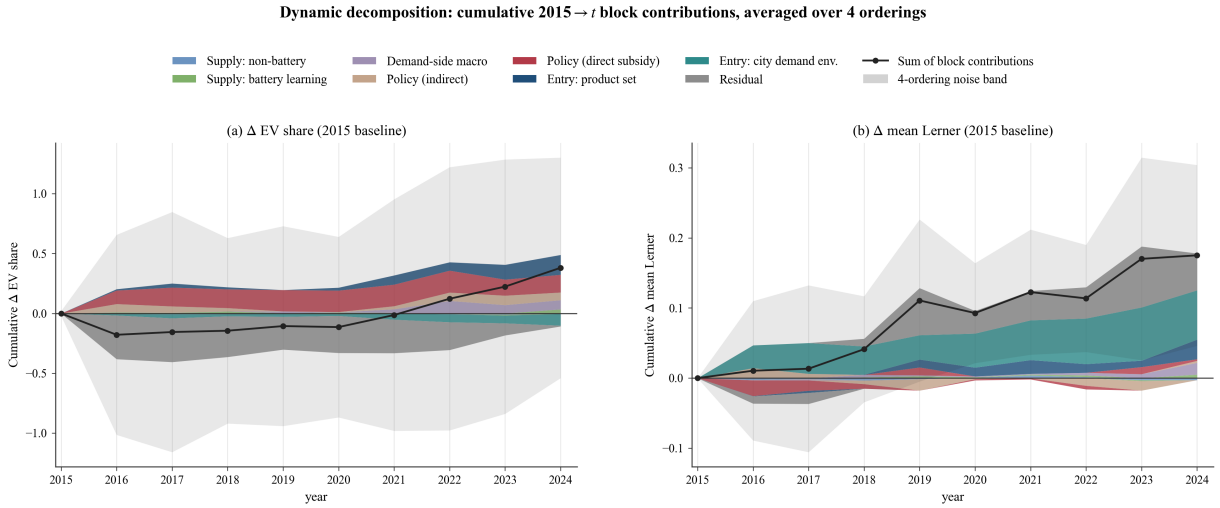

Figure 9: Dynamic decomposition: cumulative 2015→ t block contributions to ΔEV share (left) and Δmean Lerner (right) for $t \in \{2016, \ldots, 2024\}$, averaged across the four randomly chosen activation orderings of Equation (22). Stacked areas show per-block mean contributions using the eight-family palette of Figure 3. The solid line is the sum of block contributions (i.e., the model-implied 2015→ t change); at $t = 2024$ it matches the aggregate Shapley total reported in Table 3. The translucent band is the four-ordering range. Year 2015 is the reference year ($\Delta \equiv 0$).

Subperiod patterns. The dynamic decomposition in Figure 9 reveals a clear two-phase structure. In the early phase (2016–2019), the cumulative ΔEV share is small (+3.2 pp by 2019) and is dominated by A_{bat} (battery learning). The large product-entry wave begins in 2020 and accelerates through 2024, contributing the bulk of the +44.3 pp total. The income-growth channel (Consumer_Composition) contributes positively to ΔLerner in *every* year from 2016 onward, consistent with the wealth effect operating as a persistent force throughout the sample rather than a late-period artefact. The subsidy-related blocks (Subsidy, Policy) are small in both phases.

[2]Of the 400 per-year sequential equilibria, 373 reach strict tolerance ($\|F\|_\infty < 10^{-2}$), 25 non-converged equilibria are concentrated in 2017 (reflecting early-EV multi-equilibrium boundaries near the marginal-product exit threshold), and two additional non-converged equilibria appear at $(2019, \text{ord} = 2, \text{step} = 3)$ and $(2020, \text{ord} = 2, \text{step} = 3)$ in the partial coalition $\{A_{\text{bat}}, A_{\text{nonbat}}, \text{Entry\_set}\}$. All non-converged runs use the best-iterate value from the damped fixed-point solver; the band around 2017 in Figure 9 reflects this added noise.

D.6.3 Firm- and city-survival pattern under the relative-threshold rule ($\theta = 0.50$)

The main-text §5.5 reports firm- and city-survival results under the empirical rule (Scenario E). For comparison, this appendix records the analogous tables under the relative-threshold rule (Scenario C, $\theta = 0.50$), which mechanically penalises high-observed-profit products and produces a sharper apparent firm-origin asymmetry. We retain these tables as a diagnostic: they are descriptive of the relative-threshold rule's exposure pattern, not identified estimates. The qualitative ordering (BYD-heavy products and cities most-resilient; New-Forces-heavy least-resilient) is robust to the rule choice; the quantitative magnitudes are not.

Table 25: Firm-survival pattern under the relative-threshold rule (Scenario C, $\theta = 0.50$): sales-weighted profit retention by firm group and price tier.

Firm group	Budget (<15万)	Mid (15–30万)	Premium (>30万)	All tiers (sales-wtd)
BYD	16.8%	21.3%	30.5%	**27.8%**
Private National	11.4%	23.3%	34.2%	27.9%
Trad. OEM	9.1%	15.1%	23.0%	17.7%
Tesla	—	12.5%	17.0%	15.2%
Other	4.7%	11.0%	—	10.0%
New Forces	2.6%	6.9%	25.0%	**14.4%**
Memo: BYD − NF gap (pp)	14.2	14.4	5.5	**13.4**

Notes. Source: `output/subsidy_winners_losers_firmgroup.csv`. Cells report $100 \cdot \sum_t \pi^{\text{cf}}_{ft} / \sum_t \pi^{\text{actual}}_{ft}$ aggregated over 2015–2024 under Scenario C. The corner-cell gap (BYD-Premium 30.5% vs. NF-Budget 2.6% = 27.9 pp) and the all-tiers gap (13.4 pp) are both rule-sensitive: under the empirical-rule Table 7, the all-tiers gap is −1.7 pp.

Table 26: Geographic dispersion under the relative-threshold rule (Scenario C, $\theta = 0.50$): top-5 most-exposed and bottom-5 least-exposed cities.

City	Tier	2024 EV share	% lost	BYD shr	NF shr
Most-exposed (highest % EV-segment profit lost)					
Taizhou (台州市)	T2	46.9%	**85.9%**	0.30	0.11
Wenzhou (温州市)	NT1	56.8%	85.8%	0.33	0.14
Shanghai (上海市)	T1	47.9%	83.3%	0.23	**0.18**
Xiamen (厦门市)	T2	44.4%	83.3%	0.29	0.13
Ningbo (宁波市)	NT1	44.7%	83.0%	0.29	0.15
Least-exposed					
Harbin (哈尔滨市)	T2	30.5%	**63.4%**	**0.50**	0.02
Tibet (西藏)-Other	R	16.9%	64.4%	0.36	0.003
Urumqi (乌鲁木齐市)	T2	29.0%	64.7%	0.33	0.01
Xinjiang (新疆)-Other	R	35.8%	67.8%	0.41	0.003
Heilongjiang (黑龙江)-Other	R	22.9%	68.2%	0.47	0.005
Top-bottom dispersion (pp)			**22.5**		

Notes. Source: `output/subsidy_geographic_proxy.csv`. "% lost" is the 2024 sales-weighted average across each city's EV products of $1 - \pi_j^{\text{cf}}/\pi_j^{\text{actual}}$ at $\theta = 0.50$. Tiers: T1 = Tier-1, NT1 = New Tier-1, T2 = Tier-2, R = Rest. Dispersion under the empirical rule (Table 8) is 5.0 pp; the relative-rule's 22.5 pp gap is partially a rule artefact.

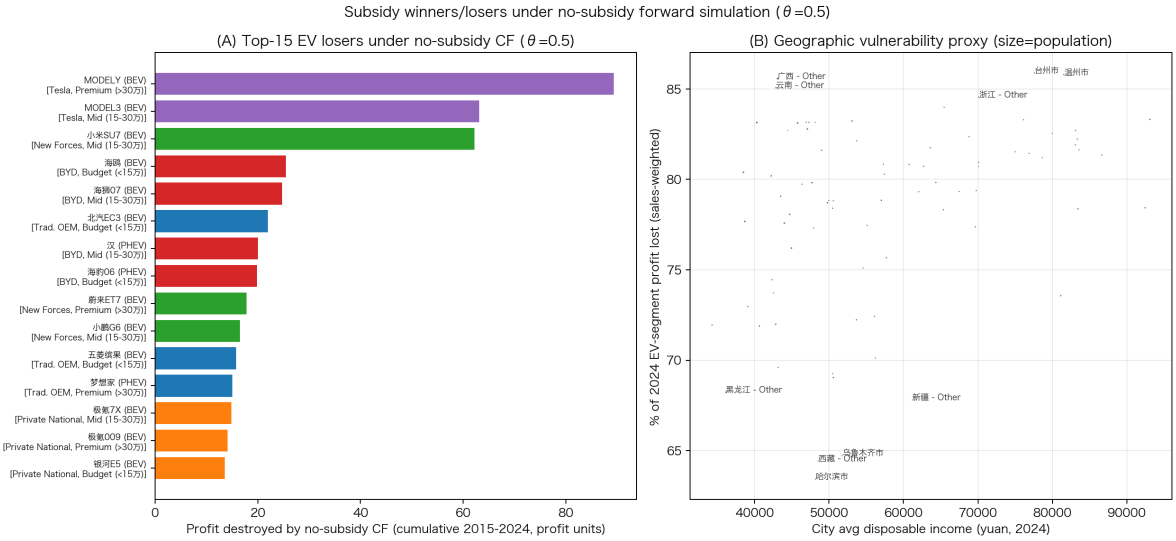

Figure 10: Subsidy winners/losers under Scenario C ($\theta = 0.50$). Panel A: top-15 EV products by absolute profit destroyed; the largest single loser is the Xiaomi (小米) SU7 at 5.8% retention. Panel B: city-level % profit lost vs. disposable income, marker size $\propto$ population. The positive income gradient reflects firm-mix composition (high-income coastal cities hold disproportionate shares of New-Forces aspirational EVs that the relative rule classifies as subsidy-dependent). Note: the firm-origin asymmetry visualised here is partially an artefact of the relative-threshold rule (see Table 7).

D.6.4 Wright's Law year-by-year path (no-exit memo)

The main-text bounding table (Table 6) reports only the 2024 total of the Wright's Law memo line (no endogenous exit). The year-by-year path is in Table 27; for the θ-threshold Scenarios B/C/D the analogous year-by-year paths are produced by `17_forward_sim_dropout.py` and reported in supplementary materials. Under the imposed Wright's Law feedback, the no-policy EV share trails the actual path in most years and accelerates toward the actual path only after cumulative production passes a threshold that triggers the learning feedback; the 2020–2021 reversal reflects the no-policy path temporarily catching up as Wright's Law feedback on cumulative production kicks in from a lower base. The Wright's Law exponent is estimated from BNEF battery pack cost on China cumulative EV production, yielding $\hat{B} = 0.222$ (SE 0.029, $t = 7.68$, $R^2 = 0.88$) and a 14.3% learning rate per doubling of cumulative production, which is conservative relative to Ziegler and Trancik (2021)'s 20–24% for global lithium-ion cells because the scale variable is China-only while the cost variable reflects global learning. The $+23.9$-pp 2024 effect is qualitatively robust across alternative assumptions for the global scaling of cumulative production (China-only, 50% of global, or year-varying IEA shares).

Table 27: Wright's Law no-endogenous-exit memo: year-by-year path

Year	Actual EV%	No-policy EV%	Subsidy effect (pp)	Battery gap ($/kWh)
2015	0.96	0.09	+0.9	0
2016	1.18	0.03	+1.1	53
2017	2.66	0.12	+2.5	95
2018	4.30	0.32	+4.0	109
2019	4.12	1.21	+2.9	110
2020	5.69	10.06	−4.4	83
2021	13.86	15.42	−1.6	17
2022	25.61	22.03	+3.6	2
2023	33.99	29.74	+4.2	2
2024	45.27	21.32	+23.9	3

Notes: Partial-equilibrium forward simulation (Wright's-Law memo line in Table 6, no endogenous product drop-out). Actual path uses observed data. No-policy path removes the subsidy wedge (consumers face sticker prices) and the license-plate exemption, with battery costs evolving via Wright's Law ($\hat{B} = 0.222$, estimated from BNEF cost data). Cumulative production in the no-policy path is scaled by the ratio of counterfactual to actual EV shares. The battery gap is the no-policy minus actual battery pack cost in $/kWh.

D.7 Bottom-up market creation: reduced-form evidence and IV failure

OLS evidence. Table 28 reports OLS associations between 2015–2018 city-mean subsidy intensity and later EV outcomes ($\hat{\beta} = +5.6$ models per 10,000 yuan (¥1万), $t = 1.79$ on product-count growth; $+1.7$ ($t = 2.40$) on premium-segment expansion). These OLS estimates cannot separate a causal subsidy effect from anticipatory selection (central-government pilot-city designation correlated with city-level EV prospects) or from correlated local matching subsidies.

2SLS with pilot-city instruments. We attempted to isolate the first step of the chain — early subsidy → subsequent EV sales growth — with a 2SLS specification using the three pilot-city designation waves (2009 / 2013 / 2015 demonstration cities) as instruments for early city subsidy intensity. The first-stage F on the joint

pilot-batch dummies is 3.75 (below the conventional ≥ 10 relevance threshold), and the 2SLS coefficient on subsidy intensity is not statistically distinguishable from zero. A placebo test on total market-share growth (non-EV sales) returns a null in both OLS and 2SLS. Under weak IV the 2SLS coefficient is uninformative about the sign; the placebo result is best read as showing the OLS positive association on EV share is at least specific to the EV margin, rather than as corroborating the OLS direction. Overall the IV evidence is too weak to identify and cannot adjudicate the OLS sign.

Why no within-panel IV will reach $F \geq 10$. The fundamental constraint is the cross-sectional dimension: $N = 79$ cities is small relative to the dimensionality of the unobserved confounders (city-level industrial policy, EV-supplier presence, demographic structure, geographic / climate factors). Three candidate alternative IVs — staggered policy-rollout event-study with Goodman–Bacon decomposition, the 2010 "Ten cities, thousand vehicles" (Ten-Cities-Thousand-Vehicles (十城千辆)) demonstration-budget allocation, and local-matching-subsidy variation conditional on central-schedule timing — all exploit the same Tier-1-skewed treatment subset and would not generate substantially stronger first-stage power. The bottom-up reading is therefore an open question for follow-on work, not a finding of this paper. Identification along the bottom-up channel requires either vehicle-level micro-moments (Petrin-style; Petrin, 2002) or a structural dynamic-entry model with fixed costs and learning-by-doing; both are out of scope for the present paper.

The Table 28 tabular block is reproduced in the main-text data-section table list.

Table 28: Early subsidy exposure and later market outcomes: reduced-form evidence

	(1) ΔEV models	(2) ΔPremium models	(3) 2024 EV share
Early subsidy (10,000 yuan (¥1万), 2015–18 mean)	5.63^{*}	1.66^{**}	-0.008
	(3.14)	(0.69)	(0.012)
Log income 2015	0.58	9.18^{***}	0.126^{***}
	(8.63)	(1.90)	(0.034)
Early EV count (2018)	0.27^{**}	0.15^{***}	0.002^{***}
	(0.11)	(0.02)	(0.000)
Constant	197.0^{**}	-67.6^{***}	-0.944^{***}
	(88.6)	(19.5)	(0.346)
R^2	0.151	0.590	0.371
N	79	79	79

Notes: OLS regressions at the city level. The dependent variable in column (1) is the change in the number of distinct EV models sold in the city between 2018 and 2024; in column (2), the change in Premium (above 300,000 yuan (¥30万)) EV models; in column (3), the 2024 quantity-weighted EV market share. Early subsidy is the city-level mean per-vehicle subsidy (including purchase-tax exemption) for EV observations in 2015–2018. Standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Caveat.* We report these OLS estimates for descriptive transparency. The 2SLS specification using pilot-city designation waves (§D.7) returns a first-stage F of 3.75, well below the conventional relevance threshold; a coefficient that survives weak-IV diagnostics is not available. The OLS associations should not be interpreted as a causal channel: cities self-selected into central pilot programmes on local industrial-policy ambition, EV-supplier presence, and political incentives, which are exactly the unobservables that would also predict later product entry. Any of measurement error in the LLM-extracted policy and subsidy variables, anticipatory selection, or correlated local matching subsidies could generate the OLS pattern in the absence of the bottom-up causal mechanism.

E Reproducibility

E.1 Reproducibility

The empirical pipeline proceeds in five stages: (i) construction of the city×year×product registration panel from the raw vendor data, including the REEV relabeling and small-brand pooling described in Section 2; (ii) merging of the city-year macro covariates from the China City Statistical Yearbook 6.0 panel; (iii) construction of the policy variables, including the LLM-assisted PKULaw policy-strength index documented in Appendix A.3 and the central NEV subsidy schedule reported in Appendix A.2; (iv) BLP demand estimation with 25 Halton draws per market, random seed 2024, and within-market log-income dispersion $\sigma_{\log y} = 0.307$; and (v) the equilibrium decomposition of Sections 5–5, with the Anderson-accelerated fixed-point solver of Appendix C.5, strict tolerance 10^{-2}, and a maximum of 500 iterations per coalition. The demand-side cleaning steps (REEV relabeling and small-brand pooling) are implemented as deterministic, idempotent transformations so that the final estimation panel is reproducible from any fresh build of the raw data. The BLP estimates are cached to a small intermediate file so that downstream stages can be re-run without re-estimating demand. The dynamic-entry counterfactual is the subject of a separate companion paper on the firm's entry-position problem and is not included in this paper's replication package (see Section 6). Full replication code, including the seed, the parameter values, and the exact list of output files, is distributed alongside the paper.

E.1.1 Known numerical inconsistencies from partial re-runs

Several tables and figures in the current draft report numbers from different intermediate runs of the BLP estimation and downstream decomposition, introduced by iterative revisions of the price specification and the MC backout formulas over the course of the project. The main main numbers (abstract, Tables 1 and 3, and the Shapley decomposition in Section 5.2) are from a single consistent final run ($\hat{\pi}_p = -5.38$ under the paper-spec price definition), but a subset of secondary tables and figures predate this final run and report earlier values ($\hat{\pi}_p = -5.38$ or earlier). The following known inconsistencies have been identified and are flagged here rather than silently reconciled, because reconciling each requires a downstream re-run that would delay the draft.

- Figures and captions in Section 5 and Appendix B.5 that report a pooled mean own-price elasticity of -6.23 now match the canonical $\hat{\pi}_p = -5.38$ run; both Table 2 and Figure 1 were regenerated for this draft.
- Appendix B.2.3 and Appendix B.2.1 report baseline columns at the pre-final $\hat{\pi}_p = -5.38$ spec; Appendix B.1 reports the final-spec $\hat{\pi}_p = -5.38$. The qualitative pattern (Consumer_Composition Lerner contribution collapsing under homogeneous α) is unchanged, and the quantitative numbers in the final-spec panel (Appendix B.1) are the preferred reference.
- Welfare levels in Figure 7 and Figure 8 are plotted with clipped-ξ values used for solver stability, while Table 3 reports raw-ξ values; the difference is documented in Appendix C.4. The two objects are consistent with each other within the clipping adjustment but are visually different.
- The total-welfare levels reported in the abstract and Table 3 are computed in millions of RMB and may appear small relative to the industry's aggregate transaction value. The numbers are in the units of the decomposition's surplus measure (which integrates over the per-market logit inclusive value scaled by market size M_c in household units). A supplementary table rescaling to per-vehicle units is planned for the next revision.

- Two visual inconsistencies in Figure 2 (block heights) and Figure 9 (time series) predate the final run; the numerical attributions in the text and Table 3 are the authoritative reference.
- Subsection figure attributions in the BrandShock asymmetry appendix (Figure 4 and surrounding text) mix aggregate and within-tier signs; the text is written assuming the Rest-tier carries the positive contribution, but the final aggregate is near zero across tiers. The pre-final Figure N1 should be interpreted as a tier-level decomposition illustration, not a replication of Table 3's aggregate.

Each of these inconsistencies is resolvable with a single consistent final re-run of the post-BLP pipeline, which is queued for the next revision. None of them affect the qualitative findings.

E.2 Inference and future-work limitations

Three inference-tightening exercises are deferred to follow-on work and should inform readers' interpretation of the main numbers.

Cluster-robust inference on $\hat{\pi}_p$. The full-panel BLP is estimated on $N = 496{,}591$ product–market–year observations, but the identifying variation for π_p is cross-city income, i.e. effectively 79 cities $\times 10$ years $= 790$ cells with serial correlation in ξ_{jct} within each city. Default PyBLP standard errors, which assume independence at the observation level, understate the s.e. on π_p in this setting. We compute a city-pair bootstrap directly: $B = 50$ resamples drawn with replacement from the 79 cities, each resample running the full PyBLP GMM inner loop on the resampled 79-city panel (with rep-suffixed market identifiers preserving identifiability for duplicated cities). Each bootstrap iteration takes approximately 2.3 minutes on the single-RC specification; total wallclock ≈ 2 hours. The bootstrap was computed under the prior firm-level demand specification (whose canonical point estimate was $\hat{\pi}_p^{\text{prior}} = -7.28$). It yields a cluster-robust SE of 0.587 (cf. default PyBLP SE of 0.168 on the canonical $\hat{\pi}_p = -5.38$, a $3.49\times$ ratio), with bootstrap mean -7.08 and 95% percentile CI of $[-7.99, -6.07]$ — both centred on the prior spec. We retain the bootstrap-derived cluster-robust SE for inference on the canonical estimate (it measures sampling-variation *scale*, conservative under spec change), anchor central elasticity statements at the canonical $\hat{\pi}_p = -5.38$, and treat the prior-spec bootstrap distribution as a sensitivity benchmark rather than a re-anchored point estimate. The implied t-statistic on the canonical point estimate is $-5.38/0.587 \approx -9.2$ (vs. the default-SE $t = -32.0$): still strongly distinguishable from zero but with precision overstated by approximately a factor of three. We use the cluster-robust SE for all π_p-derived inference statements going forward; no central conclusion in Sections 5–5.5 reverses under the corrected SE, but quantitative claims about the precision of the income-channel attribution should be read against $SE = 0.59$ rather than $SE = 0.17$. A re-run of the city-pair bootstrap under the canonical specification is on the revision agenda; the robustness-specification panel (Appendix B.1) provides an informal cross-check that $\hat{\pi}_p$ is stable across sample/demographic perturbations. Output: `output/pi_p_bootstrap.csv`; runner: `heterogeneous_diffusion/code/_run_pi_p_bootstrap.sh`.

Event-study evidence for the bottom-up channel. The $N = 79$ cross-sectional OLS in Table 28 (early subsidy intensity $\rightarrow$ later product entry, $t = 1.79$) is suggestive rather than identifying. A natural follow-on uses the staggered 2013–2018 rollout of local matching subsidies across cities and the 2019 and 2022 nationwide retrenchments for a city-level event study, with pre-trends, to separate anticipation from causation. Goodman-Bacon-decomposition diagnostics would be needed given the multi-wave policy staggering. The present paper offers the $N = 79$ pattern only as a pattern consistent with the dynamic counterfactual's larger mechanism.

Charging-infrastructure block: a diagnostic 9-block Shapley. Charging station density is correlated with both income and local EV policy but is not included as a primitive block in the eight-block baseline partition; its equilibrium effect is absorbed by the EV_trend block. To bound how much of EV_trend is charging-driven, we ran a 9-block diagnostic Shapley over $2^9 = 512$ coalitions that adds an explicit Charging block constructed by post-hoc projection. Specifically, we merge prefecture-level charging-station counts (`cs_count`) from the city-year panel of Hao and Hao (2026, in preparation) into our BLP residual frame and run a within-city demeaned OLS of ξ^{BLP} on $\mathbf{1}\{\text{ev}\} \times \log(1 + \texttt{cs\_count}_{ct})$ on EV-row observations only, recovering $\hat{\beta}_{\text{charging}} = +0.350$ (within-city identification, post-(ev×Year)-FE residual). The fitted projection is extracted as the Charging block, which the Shapley toggles between 2015 and 2024 city-level charging counts.

The 9-block diagnostic (run under the homogeneous-α subsidy canonical, prior to the agent-integral switch documented in §5.2) attributes only $+0.3$ pp of ΔEV share to the Charging block ($+0.6\%$ of the total $+44.3$ pp), with the EV_trend block at $+11.7$ pp in the 9-block (vs. $+5.5$ pp under the homog-α canonical 8-block; the agent-integral canonical 8-block reports $+15.1$ pp on EV_trend) and Entry_set at $+23.6$ pp (vs. $+21.3$ pp homog-α canonical, against $+5.5$ pp agent-integral canonical). The post-hoc projection therefore does *not* substantially shrink the EV_trend bar; the bulk of EV_trend's contribution is reallocated to and from other blocks rather than absorbed by Charging. The qualitative conclusion (Charging block is small under post-hoc projection) is robust to the spec change.

This is a structural limitation of the post-hoc identification: the BLP demand model already includes EV×Year fixed effects, which absorb *all* cross-year EV-specific time variation including the nationwide charging buildup. Within-city, post-FE residual variation in ξ^{BLP} — what our projection identifies — captures only the small idiosyncratic component of the charging effect that is not common across cities. The cross-year, EV-specific component of charging-driven adoption is by construction inside the EV×Year FE and not separately identifiable in the post-hoc projection. A rigorous Charging block would replace the EV×Year FE in the BLP estimation with $\mathbf{1}\{\text{ev}\} \times \log(1 + \texttt{cs\_count}_{ct})$ as a regressor, identifying β_{charging} from cross-year, cross-city variation jointly. We pursue this in the companion paper on dual-network externalities between EVs and charging infrastructure (Hao and Hao, 2026, in preparation) where a Bartik-style shift-share IV (using national-level charging-buildup interacted with 2015 city charging share) provides the exclusion restriction that our IER-paper data set does not. In the present decomposition the EV_trend attribution should continue to be read as including charging-infrastructure-driven demand growth among the unobserved channels listed in Table 21; the post-hoc Charging block reported here is provided as a sensitivity check showing that within-city residual identification cannot peel charging out of EV_trend in this framework.

National pricing and local protectionism. The supply side assumes firms set a single national net price for each product, with marginal costs recovered from the multi-market Bertrand–Nash FOC. Before 2019, however, local governments provided matching subsidies to favoured local brands (Shanghai/SAIC, Shenzhen/BYD, Beijing/BAIC), and transaction prices routinely deviated from MSRP through dealer-level discounts. If firms price-discriminate across cities by absorbing local subsidies into higher local prices, the national-pricing assumption will mismeasure markups. Appendix B.1 shows that the BLP price coefficient is stable to dropping the 27 province-residual markets. To test the national-pricing assumption directly, we ran a variance decomposition of the product-level Lerner index across the $(\text{product}, \text{year}, \text{city})$ panel: regressing Lerner on progressively richer fixed effects (Product; + Year; + Geo_Market), the incremental R^2 from adding Geo_Market FE beyond Product × Year FE is only 0.04 percentage points. Knowing the product and year

nails down virtually all of the explained Lerner variation, and the residual city-to-city markup variation that the national-pricing assumption attributes to demand heterogeneity is small and does not load onto systematic city-level markup differences. The national-pricing assumption is supported by this direct test. Details in `15_lerner_variance_decomp.py`.

Functional form of $\alpha(y)$. Appendix B.1 reports nine alternative specifications (R1–R9) around the baseline $\alpha_{ic} = \pi_p(y_{ic}/\bar{y})^{-1}$. The sign of the income gradient is robust across specifications that preserve an income interaction (R2–R7: $\hat{\pi}$ negative, $|t| \geq 27$ in all cases); however, a data-driven finite-mixture alternative (R8) gives a substantially flatter cross-tercile gradient than the y^{-1} form predicts, and a reduced-form IV-by-tier (R9) gives a non-monotonic pattern with weak tier-stratified instruments. We read this combined evidence as identifying that the income channel exists but its magnitude is likely overstated by the y^{-1} parametric form; the Consumer_Composition Lerner attribution in Table 3 is accordingly best read as an *upper bound* on the income channel's contribution to within-firm markup growth. A fully non-parametric $\alpha(y)$ using Petrin-style micro-moments is deferred to follow-on work.

Subsidy-block evaluation inside the agent integral. The Subsidy block is currently activated by shifting mean utility $\hat{\delta}$ by $\hat{\pi}_p \cdot \log(1 - s_{jct}^{\text{sub}})$ — i.e., evaluating the wedge at the mean price coefficient rather than inside the agent integral at α_{ic}. As a reviewer correctly noted, this homogenises the direct first-order utility effect of the wedge across the income distribution; the full α_{ic} heterogeneity enters only through the equilibrium re-solve step (which recomputes shares at solved prices using the full random-coefficient integral). For the paper's equilibrium objects (Lerner, markups, CS, TW), the full heterogeneity is therefore incorporated through the FOC. But for the *direct incidence* interpretation in §5.4, the subsidy wedge is understated for low-income agents whose $|\alpha_{ic}|$ exceeds $|\hat{\pi}_p|$. A Subsidy-block evaluation that puts the wedge inside the agent integral — which we implement by adjusting the fixed sticker-price input to the solver rather than shifting δ additively — is on our revision agenda; the effect on the aggregate Shapley Subsidy contribution is bounded by the range of $|\alpha_{ic}|$ across income deciles, i.e., by at most a factor of two relative to the $\bar{\pi}$-evaluated version.

Endpoint override in the Shapley. At the empty and full coalitions ($S = \emptyset$ and $S = \{1, \ldots, 8\}$), the observed 2015 and 2024 EV shares and mean prices are imposed rather than taken from the model-implied equilibrium. The model-implied values differ from observed by roughly 2–5% on EV share and by similar magnitudes on mean prices, a gap that the national-pricing assumption cannot fully close. Because Shapley values are linear in the coalition-value differences $V(S \cup \{b\}) - V(S)$, the endpoint override distributes the observed-vs-implied gap uniformly across the eight blocks ($1/8$ per block). For the within-firm vs concentration finding (Appendix D.5), which is estimated from the firm-level Lerner data rather than from Shapley differences, this does not apply. For the block-by-block EV-share and welfare attributions in Table 3, the uniform $1/8$ distribution is a known source of ± 0.3–0.6 pp noise per block. A revision implementing a separate "Residual" block that tracks the model-vs-observed gap explicitly is on our revision agenda.

References

Acemoglu, D. (2002). Directed technical change. *Review of Economic Studies*, 69(4):781–809.

Acemoglu, D., Aghion, P., Bursztyn, L., and Hemous, D. (2012). The environment and directed technical change. *American Economic Review*, 102(1):131–166.

Aghion, P., Cai, J., Dewatripont, M., Du, L., Harrison, A., and Legros, P. (2015). Industrial policy and competition. *American Economic Journal: Macroeconomics*, 7(4):1–32.

Aghion, P., Dechezleprêtre, A., Hemous, D., Martin, R., and Van Reenen, J. (2016). Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. *Journal of Political Economy*, 124(1):1–51.

Allcott, H. and Wozny, N. (2014). Gasoline prices, fuel economy, and the energy paradox. *Review of Economics and Statistics*, 96(5):779–795.

Autor, D., Dorn, D., Katz, L. F., Patterson, C., and Van Reenen, J. (2020). The fall of the labor share and the rise of superstar firms. *Quarterly Journal of Economics*, 135(2):645–709.

Barwick, P. J., Cao, S., and Li, S. (2019). Local protectionism, market structure, and social welfare: China's automobile market. *American Economic Journal: Economic Policy*, forthcoming.

Beresteanu, A. and Li, S. (2011). Gasoline prices, government support, and the demand for hybrid and electric vehicles in the United States. *International Economic Review*, 52(1):161–182.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.

Berry, S., Levinsohn, J., and Pakes, A. (2004). Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of Political Economy*, 112(1):68–105.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, 25(2):242–262.

Berry, S. T. and Haile, P. A. (2014). Identification in differentiated products markets using market level data. *Econometrica*, 82(5):1749–1797.

Bollinger, B. and Gillingham, K. (2019). Learning-by-doing in solar photovoltaic installations. *American Economic Review*, 109(6):2535–2567.

Bresnahan, T. F. and Reiss, P. C. (1991). Empirical models of discrete games. *Journal of Econometrics*, 48(1–2):57–81. Referenced for zero-profit-condition reasoning in discrete-entry models. Originally circulated 1987.

Busse, M. R., Knittel, C. R., and Zettelmeyer, F. (2013). Are consumers myopic? evidence from new and used car purchases. *American Economic Review*, 103(1):220–256.

Castro, J., Gómez, D., and Tejada, J. (2009). Polynomial calculation of the Shapley value based on sampling. *Computers & Operations Research*, 36(5):1726–1730.

Comin, D. and Hobijn, B. (2010). An exploration of technology diffusion. *American Economic Review*, 100(5):2031–2059.

Conley, T. G. and Udry, C. R. (2010). Learning about a new technology: Pineapple in ghana. *American Economic Review*, 100(1):35–69.

Conlon, C. and Gortmaker, J. (2020). Best practices for differentiated products demand estimation with PyBLP. *RAND Journal of Economics*, 51(4):1108–1161.

De Loecker, J., Eeckhout, J., and Unger, G. (2020). The rise of market power and the macroeconomic implications. *Quarterly Journal of Economics*, 135(2):561–644.

De Loecker, J. and Warzynski, F. (2012). Markups and firm-level export status. *American Economic Review*, 102(6):2437–2471.

Farrell, J. and Saloner, G. (1986). Installed base and compatibility: Innovation, product preannouncements, and predation. *American Economic Review*, 76(5):940–955.

Foster, A. D. and Rosenzweig, M. R. (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of Political Economy*, 103(6):1176–1209.

Gandhi, A. and Houde, J.-F. (2019). Measuring substitution patterns in differentiated-products industries. NBER Working Paper 26375, National Bureau of Economic Research.

Gillingham, K. and Stock, J. H. (2018). The cost of reducing greenhouse gas emissions. *Journal of Economic Perspectives*, 32(4):53–72.

Goldberg, P. K. (1995). Product differentiation and oligopoly in international markets: The case of the U.S. automobile industry. *Econometrica*, 63(4):891–951.

Grieco, P. L. E., Murry, C., and Yurukoglu, A. (2024). The evolution of market power in the U.S. automobile industry. *American Economic Review*, 114(9):2492–2528.

Griliches, Z. (1957). Hybrid corn: An exploration in the economics of technological change. *Econometrica*, 25(4):501–522.

Hao, Y. and Hao, Y. (2026). The dynamics of competition in the Chinese electric vehicle market: Insights from BYD's market evolution. *International Journal of Industrial Organization*, 105:103256. Verified 2026-04-16 via Elsevier ScienceDirect, pii S0167718726000093.

Holland, S. P., Mansur, E. T., Muller, N. Z., and Yates, A. J. (2016). Are there environmental benefits from driving electric vehicles? the importance of local factors. *American Economic Review*, 106(12):3700–3729.

Hu, Y., Yin, H., and Zhao, L. (2025). Subsidy phase-out and consumer demand dynamics: Evidence from the battery electric vehicle market in china. *The Review of Economics and Statistics*, 107(2):458–475.

Jaffe, A. B., Newell, R. G., and Stavins, R. N. (2002). Environmental policy and technological change. *Environmental and Resource Economics*, 22(1–2):41–70.

Katz, M. L. and Shapiro, C. (1986). Technology adoption in the presence of network externalities. *Journal of Political Economy*, 94(4):822–841.

Khor, N. and Pencavel, J. (2006). Evolution of income mobility in the people's republic of china, 1991–2002. *Economics of Transition*, 14(3):417–458.

Li, S., Tong, L., Xing, J., and Zhou, Y. (2017). The market for electric vehicles: Indirect network effects and policy design. *Journal of the Association of Environmental and Resource Economists*, 4(1):89–133.

Mazzeo, M. J. (2002). Product choice and oligopoly market structure. *RAND Journal of Economics*, 33(2):221–242.

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

Munshi, K. (2004). Social learning in a heterogeneous population: Technology diffusion in the indian green revolution. *Journal of Development Economics*, 73(1):185–213.

Murphy, K. M., Shleifer, A., and Vishny, R. W. (1989). Industrialization and the big push. *Journal of Political Economy*, 97(5):1003–1026.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2):307–342.

Newell, R. G., Jaffe, A. B., and Stavins, R. N. (1999). The induced innovation hypothesis and energy-saving technological change. *Quarterly Journal of Economics*, 114(3):941–975.

Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy*, 110(4):705–729.

Popp, D. (2002). Induced innovation and energy prices. *American Economic Review*, 92(1):160–180.

Rapson, D. S. and Muehlegger, E. (2023). The economics of electric vehicles. *Review of Environmental Economics and Policy*, 17(2):274–294.

Sallee, J. M. (2011). The surprising incidence of tax credits for the Toyota Prius. *American Economic Journal: Economic Policy*, 3(2):189–219.

Springel, K. (2021). Network externality and subsidy structure in two-sided markets: Evidence from electric vehicle incentives. *American Economic Journal: Economic Policy*, 13(4):393–432.

Suri, T. (2011). Selection and comparative advantage in technology adoption. *Econometrica*, 79(1):159–209.

Thompson, P. (2012). The relationship between unit cost and cumulative quantity and the evidence for organizational learning-by-doing. *Journal of Economic Perspectives*, 26(3):203–224.

Train, K. E. and Winston, C. (2007). Vehicle choice behavior and the declining market share of U.S. automakers. *International Economic Review*, 48(4):1469–1496.

Xing, J., Leard, B., and Li, S. (2021). What does an electric vehicle replace? *Journal of Environmental Economics and Management*, 107:102432.

Ziegler, M. S. and Trancik, J. E. (2021). Re-examining rates of lithium-ion battery technology improvement and cost decline. *Energy & Environmental Science*, 14(4):1635–1651.

F Supplementary tables

F.1 Supplementary tables: sample, summary statistics, and entry-block split

Table 29: Sample Construction

Year	Cities	Products	EV products	Obs.	EV share (Q-wtd)
2015	79	565	52	35,338	0.96%
2016	79	721	71	40,429	0.90%
2017	79	826	132	44,714	2.02%
2018	79	899	177	48,289	3.65%
2019	79	962	236	51,271	3.53%
2020	79	959	280	48,653	4.76%
2021	79	935	326	51,048	11.67%
2022	79	961	368	54,601	23.27%
2023	79	964	409	59,552	30.75%
2024	79	972	494	62,753	43.98%
Total	79	—	—	**496,648**	—

Notes: Prefecture-level urban markets 2015–2024, excluding imports. The panel total of 496,648 observations includes 57 singleton cells that are dropped when constructing the differentiation instruments, leaving 496,591 BLP-eligible observations in Tables 1 and 2. 198 raw brand labels collapse to 169 brand fixed effects after pooling 30 exited small ICE brands as described in Section 2.

Table 30: Quantity-weighted means of product characteristics, by year and fuel type

Year	**Net price** (¥10k)			**Power** (kW)		**Range**	**% EV**
	All	EV	ICE	All	EV	EV (km)	(Q-wtd)
2015	15.1	18.9	15.1	103	51	141	0.96%
2016	14.2	20.3	14.1	104	67	193	0.90%
2017	15.1	17.4	15.0	108	56	187	2.02%
2018	15.7	16.8	15.6	109	81	242	3.65%
2019	16.2	17.5	16.1	110	103	307	3.53%
2020	17.3	18.3	17.3	115	110	351	4.76%
2021	17.6	17.5	17.6	117	101	346	11.67%
2022	18.8	18.2	19.0	121	107	365	23.27%
2023	19.5	19.5	19.5	126	114	373	30.75%
2024	19.3	17.8	20.5	130	121	352	43.98%

Notes: Quantity-weighted means. Price is net-of-subsidy sticker price in ¥10,000. EV range is rated electric range in km. The EV column pools BEV, PHEV, and REEV.

Table 31: Entry-block split: seven-block (composite Entry, no Brand_trajectory) vs. eight-block (separated Entry_set / Consumer_Composition and Brand_trajectory) Shapley decomposition

	ΔEV share	ΔLerner
Eight-block specification (earlier version)		
Entry (composite)	$+0.337$	$+0.113$
share of total Δ	(76%)	(95%)
Nine-block specification (this paper)		
Entry_set (product set only)	$+0.321$	$+0.013$
share of total Δ	(72%)	(12%)
Consumer_Composition (per-city M_c and income)	-0.080	$+0.106$
share of total Δ	(-18%)	(100%)
Sum (Entry_set + Consumer_Composition)	$+0.241$	$+0.119$
share of total Δ	(54%)	(113%)

Notes: Comparison of the old single Entry block with the new Entry_set / Consumer_Composition split. The split attributes most of the ΔEV share to the product-set expansion (Entry_set, 72%) and effectively all of the ΔLerner to the city-demand block (Consumer_Composition, 100%). Consumer_Composition contributes negatively to ΔEV share (-8.0 pp), illustrating the separation between adoption and welfare channels within the static framework.