The Trade-Off Between Decarbonization and Reshoring

Evidence from the French EV Environmental Score Reform *

Clément Malgouyres† Thierry Mayer‡ Lewin Nolden§

Preliminary
August 2025

Abstract

Governments increasingly tie clean-technology subsidies to social and industrial aims, hoping to accelerate decarbonization while nurturing domestic production. In 2024, a French reform embedded a life-cycle environmental score in its electric-vehicle (EV) purchase bonus (based on estimated emissions during the production process). Linking exhaustive French registration records to assembly plants and battery origins, we exploit the sharp induced eligibility shock in a difference-in-differences design. Treated EV models experience 60% decline in sales relative to EVs remaining eligible. We combine these estimates with survey-based stated preferences in order to estimate a nested-logit demand system. Model-based counterfactuals suggest that the reform reduced total EV penetration by 0.9 pp, implying slower fleet decarbonization. The reform cut both public spending and consumer surplus and raised tank-to-wheel emissions enough to offset manufacturing-stage CO2 gains. A budget-neutral variant that recycles the fiscal savings into larger bonuses for eligible models would restore the lost electrification, and would cost about €4 500 of consumer surplus per reshored battery while turning the net environmental gains positive.

Keywords: car industry, subsidies, environment

JEL Codes: H23, Q58, L62, D12, L52

*The authors thank Bertrand-Olivier Ducreux for his help with the registration data collected by Ademe. We thank Philippe Aghion, Laurent Bach, Antonin Bergeaud, Paul Dutronc, Gabrielle Fack, Etienne Fize, Dimitrije Ruzic, Timothy Van Zandt and Farid Toubal for their comments as well as audience at Collège de France, Paris-Dauphine, HEC and IPP seminars.

†CREST-CNRS, IPP-PSE, clement.malgouyres@ensae.fr

‡Sciences Po, CEPII, and CEPR, thierry.mayer@sciencespo.fr

§University of Oxford (from October 2025), lewin.nolden@trinity.ox.ac.uk

1 Introduction

Passenger cars are a natural focal point for climate policy: they alone account for roughly 10% of global transport CO_2 emissions (Sen and Miller, 2023). In many advanced economies, this proportion is higher and has grown during recent decades.[1] To hit mid-century decarbonization targets, more than twenty-five countries (and many sub-national jurisdictions) have announced that sales of new Internal Combustion Engine (ICE) cars must cease by 2035 or earlier. These phase-out pledges are reinforced by generous purchase incentives that narrow the price gap between battery-electric vehicles (EVs) and their internal combustion engine (ICE) counterparts.

The transition, however, is not merely environmental—it is also industrial. Batteries and final assembly represent a large part of the value added of a vehicle. If, as is the case for ICEs, these stages tend to cluster, the shift to EVs risks to be accompanied by a large spatial shift in production. China's rapid ascent sharpens that concern: By 2023, China already represented 62% of the world's EVs and nearly 80% of battery cells, far above its share in the traditional ICE market. Policymakers have therefore begun to tie EV subsidies to localisation or domestic-content rules, the most prominent example being the U.S. Inflation Reduction Act (Allcott et al., 2024).

Then, there is the CO_2 emissions trade-off. Building an EV typically embeds 2 to 4 tons of CO_2-equivalent more than building a comparable ICE car, mainly because of battery manufacturing and aluminum-heavy bodies (IEA, 2024). On the road, the ordering reverses: EVs emit virtually no CO_2, whereas tank-to-wheel emissions of ICE cars are roughly 120–180 grams of CO_2 per kilometer.[2] The "break-even" distance at which an EV's life-cycle emissions fall below those of an ICE is therefore around 30,000–50,000 km, depending on how clean the electricity mix is (IEA, 2024). Policies that screen EV models on life-cycle (including the manufacturing-phase) emissions curb high-emitting production today but this can be at the price of slowing down the long-run decarbonization of the fleet if they restrict the set of eligible EVs to such an extent that electrification of the fleet stalls.

France's 2024 reform illustrates this tension well. Only EVs for which manufacturing emissions fall beneath a fixed life-cycle threshold continue to qualify for the purchase bonus, a rule that instantly disqualified models accounting for about 30% of EV sales in 2023. Because life-cycle intensity varies strongly with the location of assembly and battery production and (to a lesser extent) with transport costs, the measure also acts as an filter that favors regional supply chains. This paper asks whether such Life-Cycle-Assessment (LCA)-based incentives can *simultaneously* speed up fleet decarbonization by steering demand toward lower-emission production while favoring domestic producers, and at what economic cost.

We assemble a new monthly panel that links every passenger-car registration in France to its place of final assembly and—when the vehicle is electric—to the country (or region) in which its battery cells were produced. Two sharp policy changes in January 2024 provide quasi-experimental

[1] For example, while France's overall greenhouse-gas emissions fell by 20% between 1990 and 2019, transport's emissions stagnated and its share of overall emissions went up from 27% to 32% over the same period. Within that sector, passenger cars represent over 50% of France's transport-sector CO2 footprint in recent years (CITEPA, 2023), implying that passenger cars represent at least 16% of overall emissions.

[2] The official WLTP sales weighted average is about XXX g CO_2 per kilometer in France, but recent works in France and the EU (European Commission, 2024) show that WLTP norms tend to underestimate the real world emissions, with real world emissions exceeding the WLTP norm by 20% for gasoline and diesel engines and 350% for plug-in hybrids.

price variation: first, and most importantly, the life-cycle "environmental-score" filter mentioned above disqualified a large share of EV models from the French purchase "bonus". Second, the pre-existing CO_2 "malus" on internal-combustion cars (ICEs) was increased for high-emitting models. By comparing within-segment market shares of treated vs untreated vehicles before and after each reform in a difference-in-differences design, we are able to assess how responsive demand is to policy-induced price changes.

To translate those reduced-form findings into substitution patterns usable for the whole sector, we embed those shocks in a three-level nested-logit model. Households first decide among not buying, purchasing an ICE, or purchasing an EV; conditional on the powertrain, they choose a vehicle segment and then a model. The correlation parameters that govern within-nest substitution are identified jointly from the policy variation and from "second-choice" survey moments that report whether EV buyers would remain in the EV nest were their preferred model unavailable (Allcott et al., 2024). Solving the model—either under constant monopolistic competition or under multiproduct Bertrand pricing—allows simple "exact hat-algebra" counterfactuals.

Our new dataset shows the pre-reform dominance of Chinese supply chains: in 2023, China supplied about 30 % of French EV registrations and nearly 55 % of battery cells. Our diff-in-diffs estimates confirm strong price sensitivity and dynamic specification shows no evidence of differential pre-trends, thus lending credibility to the identifying assumption. Losing the bonus cuts sales of the affected EVs by roughly 60 %, an own-price elasticity of $\varepsilon \approx -3.5$; a one-thousand-euro rise in the malus reduces sales of targeted ICE models by about 20 %, implying elasticities $\varepsilon \approx -6$.

Feeding these responses through the nested logit model allows for a quantification of the trade-offs between reshoring, decarbonization and the consumer surplus that lies at the heart of life-cycle-based industrial policy. Our results show that the environmental-score filter succeeds in redirecting demand toward European assembly plants and battery suppliers, resulting in the additional sales of about 6,000 cars with European-sourced batteries, yet at the cost of a 0.9-percentage-point fall in overall EV penetration. Consumer surplus declines by about €214 million but this is more than offset by €237 million in fiscal savings. The net CO_2 effect is (slightly) negative once higher tank-to-wheel emissions are accounted for. A budget-neutral alternative that retains the life-cycle filter but recycles the savings into a larger bonus for the remaining eligible EVs restores electrification and generates net benefits in terms of emissions. It increases the number of European-sourced battery by around 11,000 units, yet still costs roughly €4,500 in consumer surplus for every additional battery sourced in Europe.

Relationship with the literature. The first strand of work that our paper speaks to evaluates how demand-side incentives shape consumers' adoption of lower-carbon vehicles. Early quasi-experimental studies exploited a variety of policy changes and found sizable but heterogeneous elasticities of demand for hybrids and EVs—for example, Chandra et al. (2010), Klier and Linn (2015), and Yan and Eskeland (2018). More recent evidence from large-scale programs shows that elasticities remain high when incentives are targeted at lower-income households or bundled with non-price perks: see Muehlegger and Rapson (2022) for California. In comparison with Muehlegger and Rapson (2022), we tend to find larger reduced-form own price elasticities for EVs, which might reflect the fact that our policy-induced variation in price affected vehicles that present very similar attributes as far as the final consumer is concerned. Our paper is most closely related to Allcott et al. (2024) which focuses on the short-run impact of an EV incentive program (the US Inflation Reduction Act, IRA, implemented by president Biden at the end of 2022) which impacted differentially domestic and foreign car makers. We build on their approach by estimating a similar nested model and use some of their aggregate moments from survey data. We depart from them in several ways. Where Allcott et al. (2024) analyze an explicit domestic-content rule under the IRA, we examine a reform that conditions subsidies on a vehicle's cradle-to-gate CO_2 footprint and show its *de facto* domestic / European bias. Because the French reform hit a slack market without the inventory rationing that characterized the car market until mid 2023 (reference), our difference-in-differences design isolates the pure demand response to losing the bonus rather than a mix of supply and demand shocks. Moreover, we trace the provenance of both final assembly and battery cells for every registration, letting us measure how the policy redirects sourcing along the EV value chain—a stage at the heart of many industrial policy concerns (Head et al., 2024; Barwick et al., 2025). We also depart from them by adopting a somewhat different approach to identification and estimation. In the spirit of Berry et al. (2004), we show that second choice data can be used sequentially to identify the nested parameters, without making assumptions about the structural error term nor estimating the price parameters. Furthermore, applying insights from "exact hat algebra" to the nested logit, we build moments as a function of the nest parameters and observed shares only, which eases and speeds up estimation substantially.[3] Our contribution to this literature is to examine a uniquely sharp, LCA filter that removed a purchase incentive from about one-third of the French EV market overnight, thus providing clean quasi-experimental variation. We also contribute by documenting clearly the implication of the policy on the sourcing of the cars but also of the batteries—a component at the center of key industrial policy concerns.

A second, fast-growing literature studies industrial-policy dimensions of the battery electric vehicle transition. Recent models highlight how learning curves, network effects, and strategic trade considerations can justify production-side support for "green" industries (Aghion et al., 2024). Empirical work has focused on place-based subsidies for battery plants (Criscuolo et al. 2023), the reaction of location choices to demand subsidies contingent on local production (Head et al. 2024), and the export consequences of China's EV push (Barwick et al. 2025). Yet few papers quantify the trade-offs between domestic value-added and carbon abatement when purchase subsidies themselves become conditional on provenance or on LCA scores. By linking exhaustive French registration data to assembly plants and battery origins, we are able to trace how such provenance-based rules redistribute demand across firm nationalities and stages of the value chain. Our model focuses on the short to medium run impact of the policy taking location and sourcing decision by firms as given as well as overall productivity.

[3]This can be useful if one wants to use bootstrap inference to gauge the degree of statistical precision of the counterfactual (d'Haultfoeuille et al., 2014).

Organization of the paper. The paper is structured as follows. Section 2 presents the environmental score reform at the center of the paper. Section 3 shows key descriptive statistics on the dynamics of market shares among EVs and of the overall EV share around the reform. Reduced-form results are in section 4. Section 5 introduces the conceptual framework, presents the demand and supply side of the model as well as our approach to identification and estimation. Section 6 carries out a cost-benefit analysis of the reform and simulate a simple, counterfactual policy. Finally, section 7 concludes.

2 Policy background and data

2.1 Policy background

Electrification incentives for cars in France combine a demand–side incentive (bonus) for battery-electric models, a malus tax on CO_2-intensive ICEs, and several ancillary tweaks (social leasing, fleet rules). In 2024 a new environmental score (ES) fundamentally redefined bonus eligibility by imposing a life-cycle-emission threshold, while the malus schedule steepened sharply. This section documents the three building blocks—bonus/ES, malus, and ancillary changes. In the final subsection (2.1.4), we explain how we translate them into the tax-subsidy variables used in the empirical analysis and in the counterfactuals.

2.1.1 The environmental score reform: change in the main EV incentive program

The bonus before the reform. Since 2023, France's purchase subsidy has been reserved for battery-electric vehicles (BEVs) costing below €47 000 and weighing less than 2.4 t. Households earning under the 5th income decile received a €7 000 rebate; all others were eligible for €5 000. In practice, the overwhelming majority of recipients of the incentives get the €5 000 bonus with only about 10% of households getting the €7 000 (Montout and Robinet, 2024). No incentives have been available for ICE or hybrid since the 2010s and plugins (PHEV) are excluded from any bonus since 2023.

Life-cycle reform in 2024. Since January 2024, eligibility for France's ecological bonus ("bonus écologique") is determined by an *environmental score* (ES) calculated by ADEME, a public agency under the French Ministry for the Ecological Transition.[4] Manufacturers submit detailed bills of materials, and ADEME converts the data into life-cycle CO_2-equivalent emissions per vehicle. This score assesses the carbon footprint of a vehicle throughout its lifecycle, including production, assembly, and transportation. To qualify for the bonus, vehicles must achieve a minimum score of 60 out of 100.

The ES comprises two major components. The main component relates to upstream carbon footprint while a secondary one is about additional environmental factors.

- **Upstream Carbon Footprint ($\geq$ 70%)** includes emissions from:

[4] ADEME stands for *Agence de la transition écologique*, which translates to French Agency for Ecological Transition. See: `https://www.ademe.fr/` and `https://score-environnemental-bonus.ademe.fr/`.

– *Ferrous metals*: Carbon footprint from the production of ferrous metals used in the vehicle (excluding the battery).
– *Aluminium*: Emissions from aluminium production for the vehicle (excluding the battery)—see Figure 1 for the emissions factors used.
– *Other materials*: Emissions from the production of materials other than ferrous metals and aluminum used in the vehicle.
– *Battery*: Emissions related to battery production.
– *Assembly*: Emissions from intermediate transformations and vehicle assembly, depending on the country of the production site.
– *Transport*: Emissions from transporting the vehicle from the production site to the distribution site in France.

Emission factors by country and mode are shown in Table 1.

- **Additional Environmental Factors ($\leq$ 30%)**: This optional component can contribute up to 30% of the total score and considers:
  – *Use of recycled and bio-sourced materials*: Incorporation of environmentally friendly materials in the vehicle.
  – *Battery reparability*: The ease with which the battery can be repaired or replaced.

While we cannot easily construct this score as we do not have access to all the required quantities per vehicle, we can still precisely determine which cars were eligible and which were not. A ministerial document published on December 14 of 2023 listed car models qualifying for the EV incentive in early 2024. Models whose score exceeded the fixed threshold and were eligible in 2023—roughly 30% of French EV sales in 2023—lost the bonus. The list is updated quarterly.[5] Our diff-in-diffs analysis focus on car models which lost eligibility and compare them to eligible cars as of December 14, 2023. Our descriptive and counterfactual analysis, however, takes into account the expanding boundaries of the set of eligible vehicles over time—which is entirely driven by the introduction of new models and not by cars that were already available experiencing a change in eligibility status. Appendix Figure A.2.1 presents two popular Chinese-made models that lost the bonus—the Dacia Spring and the MG 4.

[5]For methodology and the current list of eligible models, see: `https://score-environnemental-bonus.ademe.fr/` and the ministerial decree: `https://www.legifrance.gouv.fr/jorf/id/JORFTEXT000048167407`.

Figure 1: Illustration of the material factors used in the ES calculation (aluminum)

$FE_{aluminium}^{site}$ sont définies suivant la zone d'implantation du site d'assemblage considéré, selon les dispositions suivantes :

Zone d'implantation du site considéré	Facteur d'émission carbone de la production d'aluminium (pur et allié) par unité de masse aluminium (pur et allié) (kg-eq CO_2/kg)
Amérique du Nord	8,5
Amérique du Sud	13,9
Europe	8,6
Chine	20,0
Japon	12,6
Conseil de Coopération du Golfe	11,4
Autre	18,5

Table 1: Selected emission factors used by ADEME.

Component	**France**	**Germany**	**China**
Ferrous metals (kg CO_2-eq./kg)	1.4	1.4	2.0
Aluminium (kg CO_2-eq./kg)	8.6	8.6	20.0
Plastics/composites (kg CO_2-eq./kg)	4.6	4.6	5.0
Battery (kg CO_2-eq./kWh)	53	53	68
Assembly (kg CO_2-eq./kg vehicle)	0.58	0.83	1.60
Maritime transport (kg CO_2-eq./t km)	0.101 (global)		
Rail transport (kg CO_2-eq./t km)	0.010	0.023	0.041
Road transport (kg CO_2-eq./t km)	0.208	0.256	0.377

2.1.2 The French feebate

The "malus" tax. France operates a tax on high tank-to-wheel CO_2-rated vehicles, including plug-in hybrids, which rises with emissions. The top bracket reaches €60,000 in 2024. Figures 2 document the schedule focusing on the lower end of the CO2 distribution where most of the sales occur. Figure 3 shows the jump between 2023 and 2024.

Figure 2: Malus schedule and volume distribution by CO_2 rating, 2022–24.

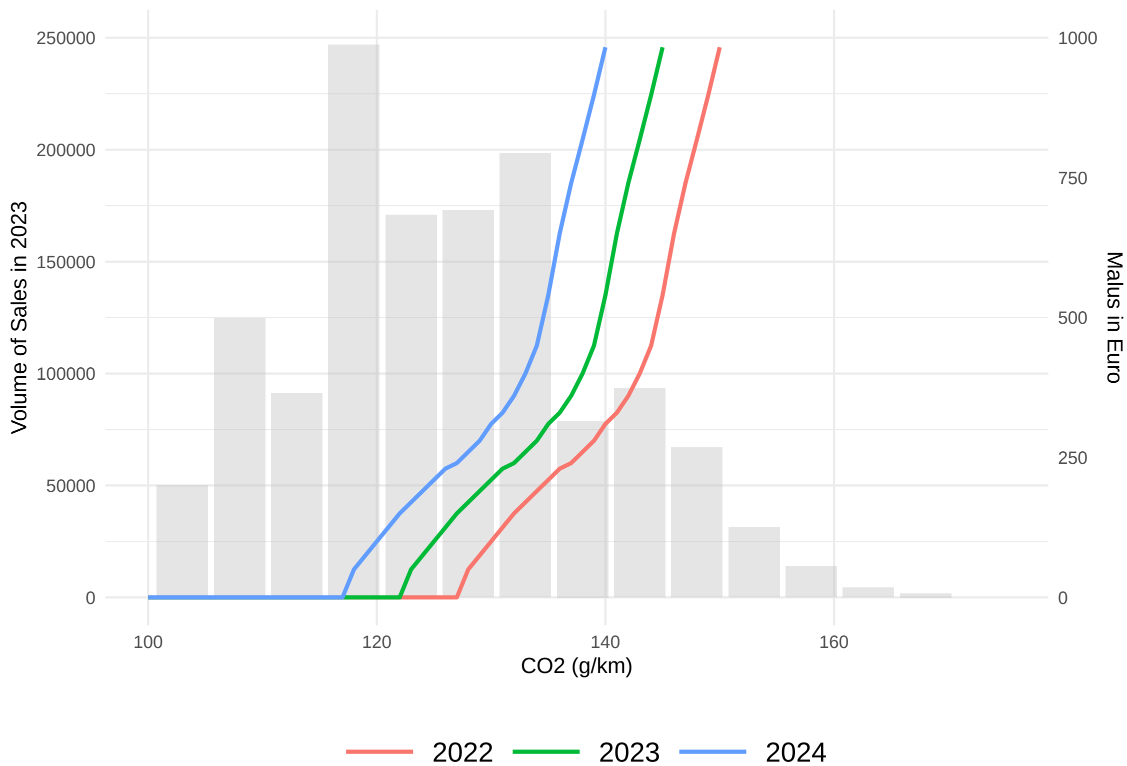

2.1.3 Contemporary policy changes

Social leasing. A temporary "leasing social" scheme subsidised long-term EV leases for low-income households. 50 000 contracts—roughly 25% of all 2023 bonus claims—were signed between January and March 2024 and the program was stopped afterward. Models eligible for the leasing social are a subset of those eligible for the main EV incentive. The cost of the program is €13 000 per vehicle. The program does not impact our estimation of the average treatment effect of losing eligibility to the EV incentive program. It matters however when we translate these effects into semi-elasticities. We take the program into account including it into the expected increase in post-tax price induced by the loss of eligibility, as detailed in section 2.1.4 (for further details see the French Senate report[6].) We cannot directly observe which purchase was subsidized by the social

[6] https://www.senat.fr/rap/a24-148-2/a24-148-212.html

Figure 3: Increase in the feebate between 2023 and 2024, weighted by 2023 sales.

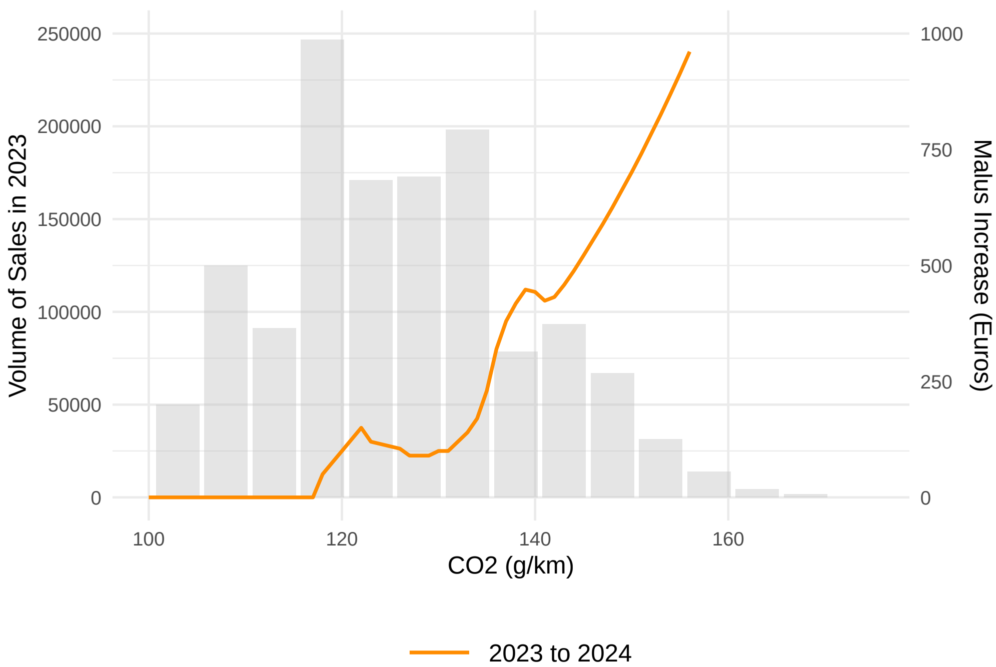

leasing. We know however that past March 2024, no new leasing contracts was offered and the program should therefore not directly impact market shares past this date.

Corporate buyers. From February 2024 onwards, corporate fleets ceased to receive the €3,000 bonus. We weight the bonus change by the 2023 split between household and corporate registrations (80% vs. 20%).

2.1.4 How large is eligibility-loss shock for EV models?

How large is the price shock for a vehicle that loses the EV purchase subsidy? While this exact magnitude has little bearing on our difference-in-differences estimates of average treatment effects, it is crucial for translating those effects into semi-elasticities and for calibrating our counterfactual simulations. We propose to bound the subsidy loss faced by models that fail the 2024 environmental-score (ES) filter. Those bounds are expected values of the loss depending on assumptions regarding the treatment of corporate fleets and the social leasing:

1. **Impacted vehicles would not have been eligible to social-leasing programme (lower bound).** In 2023, corporate fleets account for a share $\pi_c = 0.2$ of EV purchases, and household for the remainder $\pi_h = 0.8$ (Montout and Robinet, 2024). Corporate purchases receive a €3,000 bonus, while households receive €5,000 (90% of cases) or €7,000 (10%, with low enough in-

come).[7] Setting the post-reform subsidy to zero gives

$$\Delta\tau_{\text{low}} = \pi_c(3\,000/6) + \pi_h\big[0.9 \times 5\,000 + 0.1 \times 7\,000)\big] \;\approx\; 4\,260. \tag{1}$$

This version of the policy shock allows to simulate a counterfactual policy where even in the absence of the ES reform, the impacted models would *not* have been eligible to the social leasing.

2. **Including social leasing (upper bound).** The 2024 "leasing social" scheme signed about 50,000 contracts before being frozen in March.[8] Each lease embeds an implicit subsidy of roughly €13 000.[9] Removing these vehicles from the post-March data and recalculating the weighted average raises the effective bonus loss to

$$\Delta\tau_{\text{high}} = \pi_c(3\,000/6) + \pi_h\Big\{0.75\big[0.9 \times 5\,000 + 0.1 \times 7\,000\big] + 0.25 \times 13\,000\Big\} \;\approx\; 5\,900 \tag{2}$$

The fraction of low income households who could benefit from the € 7000 EV incentives is set at 20% based on first semester figures reported by the French ministry for the environment and reported in Montout and Robinet (2024). This version of the shock can be used to get the counterfactual situation where ES is not implemented and ES-impacted vehicles would have been eligible to the social leasing.

We therefore treat $\Delta\tau_{\text{low}} = 4.26\text{k}$ and $\Delta\tau_{\text{high}} = 5.90\text{k}$ as reasonable lower and upper bounds on the policy-induced price increase for disqualified EV models. More specifically, we consider the upper bound of the shock when estimating semi-elasticities empirically as this seems like the best approximation of the price shock faced by potential consumers. For our baseline simulation, however, we try to isolate the impact of the ES reform and consider the lower bound version of the shock—a scenario where the social leasing program would have excluded the ES impacted models anyway.

2.2 Data

Eligibility and sales in France. The Ministry publishes an exhaustive list of models that satisfy the 2024 life-cycle "environmental score" even though the full underlying calculations are not public. We merge that list with ADEME's monthly registration data at the CNIT level (as in Kessler et al. 2023) using TVV-CNIT correspondence. CNIT refers to "Code National d'Identification du Type" (national vehicle type identifier), which is the most granular vehicle identifier available. For example, a Renault Zoe with horsepower of 85 KW has a different CNIT identifier than a Renault Zoe with 100 KW. TVV on the other hand refers to "Type Variante Version" (type variant version),

[7] Based on ASP statistics presented in (Montout and Robinet, 2024).

[8] For comparison, that would be about 25 % of 2023 bonus claims. Budget report, French Senate, `https://www.senat.fr/rap/a24-148-2/a24-148-212.html`.

[9] Same source as footnote 8.

a vehicle identifier determined by the producer. We aggregate the data somewhat: a single observation is defined by the tuple *(model, brand, fuel-type, fiscal-power class, curb-weight, month)*, which gives us a vehicle-level panel of monthly sales.

Place of assembly and battery origin for EVs. Each vehicle is matched to its assembly plant using S&P global production database. For electric vehicles, we additionally attach the country or region where the battery cells were produced—see Appendix section A.1 for a full description of the process to determine battery country of origin. Combining these locations with weight-based emission factors from the IEA allows us to impute manufacturing CO_2 for every model.

Aggregate substitution moments from second-choice surveys. Price sensitivity parameters as well as some of the nesting parameters (for ICE) are estimated using policy variation. In order to estimate a subset of the nesting parameters, we also use US survey evidence with aggregate moments computed by Allcott et al. (2024). Specifically, the share of EV buyers who would still choose an EV if their preferred model were unavailable, and the same statistics applied to segment-EV level.

3 Descriptive evidence

Before moving to causal estimation, we present patterns of sales around the date of the reform. We begin by showing the evolution of EV market shares according to the eligibility status, place of assembly, and battery-cell origin (Figures 4–5). We then compare aggregate EV adoption in France and Germany around the reform date (Figure 6). These stylized facts motivate our reduced-form and structural analyses in subsequent sections.

EV market shares by subsidy-eligibility status. Figure 4 shows a pronounced break at the January-2024 reform. Models that were eligible for the bonus in 2023 but lost eligibility thereafter see their share shrink by roughly 60%, and the drop persists through the end of the sample. By contrast, the share of *never-eligible* models remains essentially flat.

EV market shares by country of final assembly. Chinese-assembled EVs had been gaining market share in France throughout 2023, but Figure 5 reveals that the life-cycle reform abruptly reverses that trend. European—especially French—assembly benefits most, while "Rest of World" models remain roughly constant.

EV market shares by country of battery-cell origin. Battery sourcing is even more concentrated than final assembly: China accounts for about half of all cells sold in 2023. Figure 6 indicates that the eco-score reform first flattens and then marginally reverses the upward trajectory.

Figure 4: Market shares *within* the EV segment by subsidy-eligibility status, 2023–2024

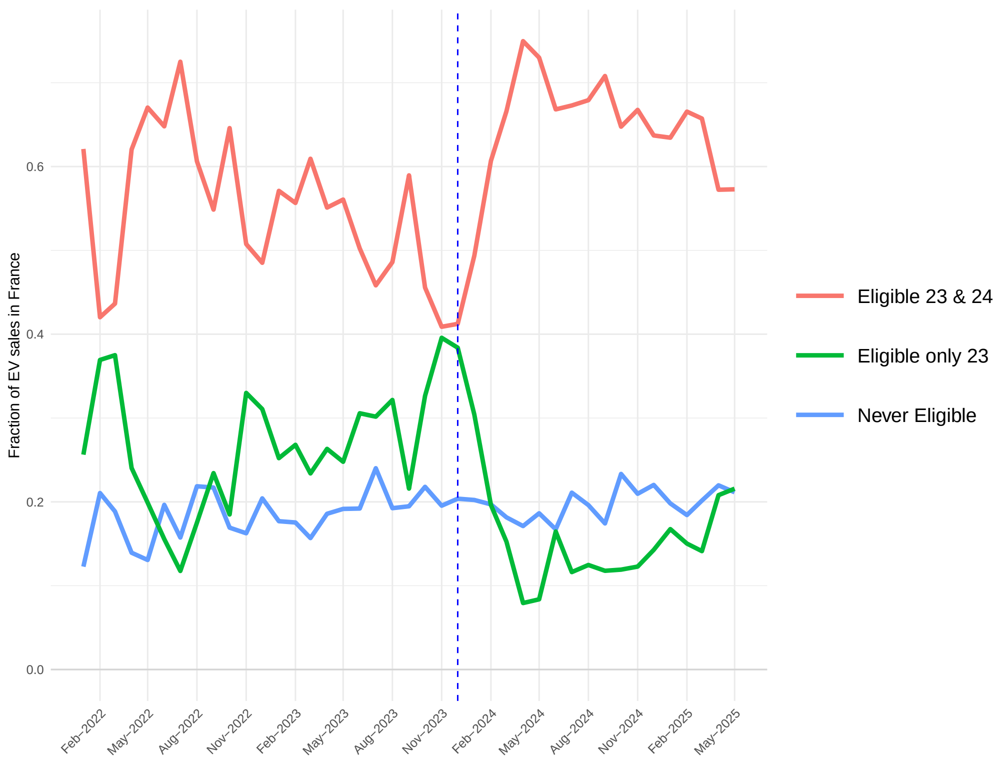

Note: Market shares within the EV segment by subsidy-eligibility status.

Figure 5: Market shares *within* the EV segment by country (or region) of vehicle assembly

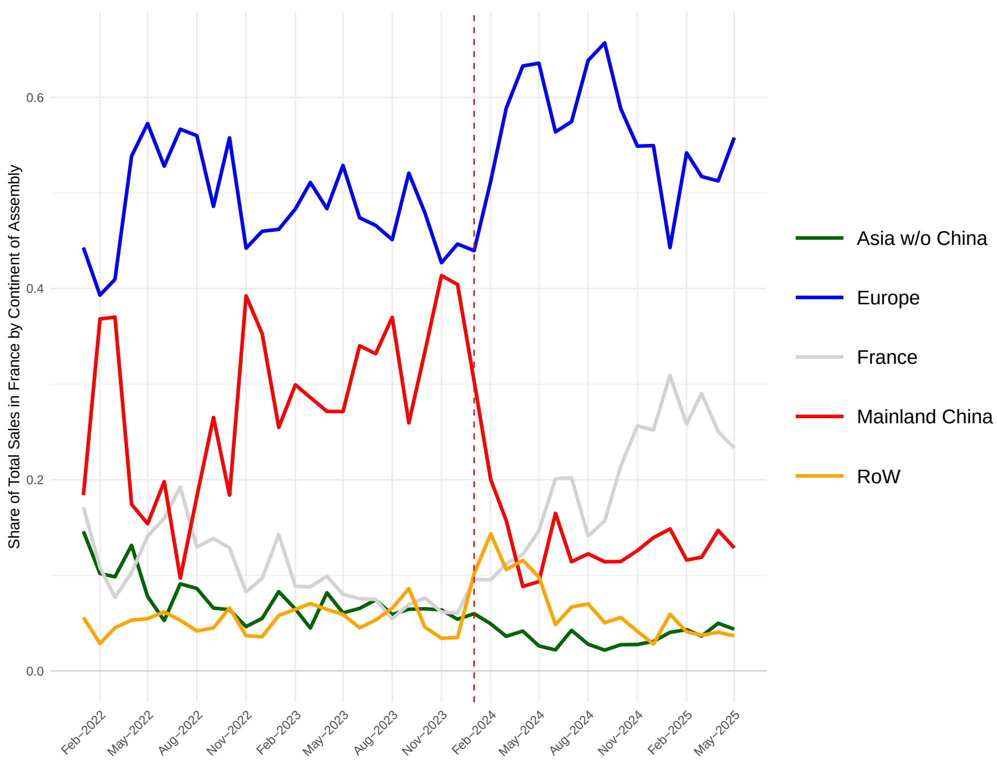

Note: Market shares within the EV segment by country (or region) of vehicle assembly. For months of 2025 data, we assume that assembly plants (which we observe until the end of 2024) are unchanged.

Figure 6: Market shares *within* the EV segment by country (or region) of battery-cell production

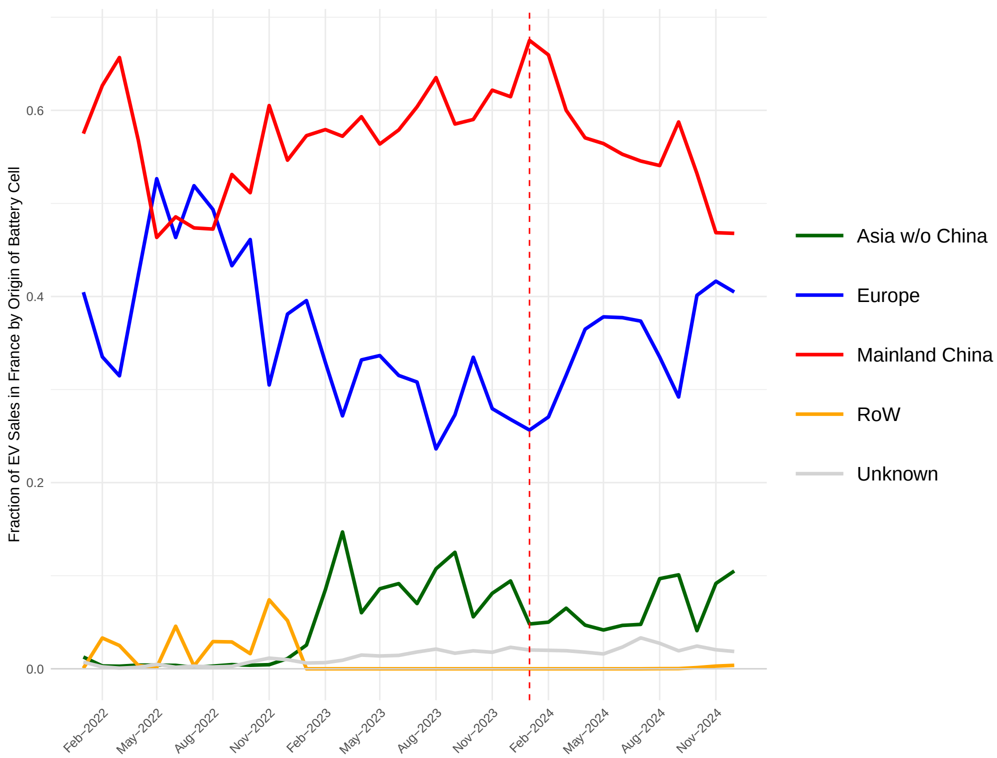

Note: Market shares within the EV segment by country (or region) of battery-cell production.

Overall EV adoption in France and Germany. While France introduced its life-cycle–based eligibility test for its EV bonus, Germany simultaneously *abolished* its purchase subsidy (both in January 2024). Registrations react sharply in Germany. The share of EVs among new cars falls by roughly one-third (figure 7). In contrast, the French decline is milder and more gradual although we do see visually a slow-down in EV adoption. A synthetic-control comparison with other EU markets confirms a significant break for Germany but finds no statistically clear break for France (see results in the online appendix section B.1), although such cross-country exercise necessarily lacks statistical power and relies on strong identifying assumption. We view this German/French comparison as a motivation for using a structural model as it shows that changes in EV incentives can impact the EV share in the overall market for new passenger cars but that reduced-form approach will not be enough to isolate and estimate precisely such effects when the reforms are not as dramatic as the German case.

Figure 7: Monthly EV market shares in France and Germany.

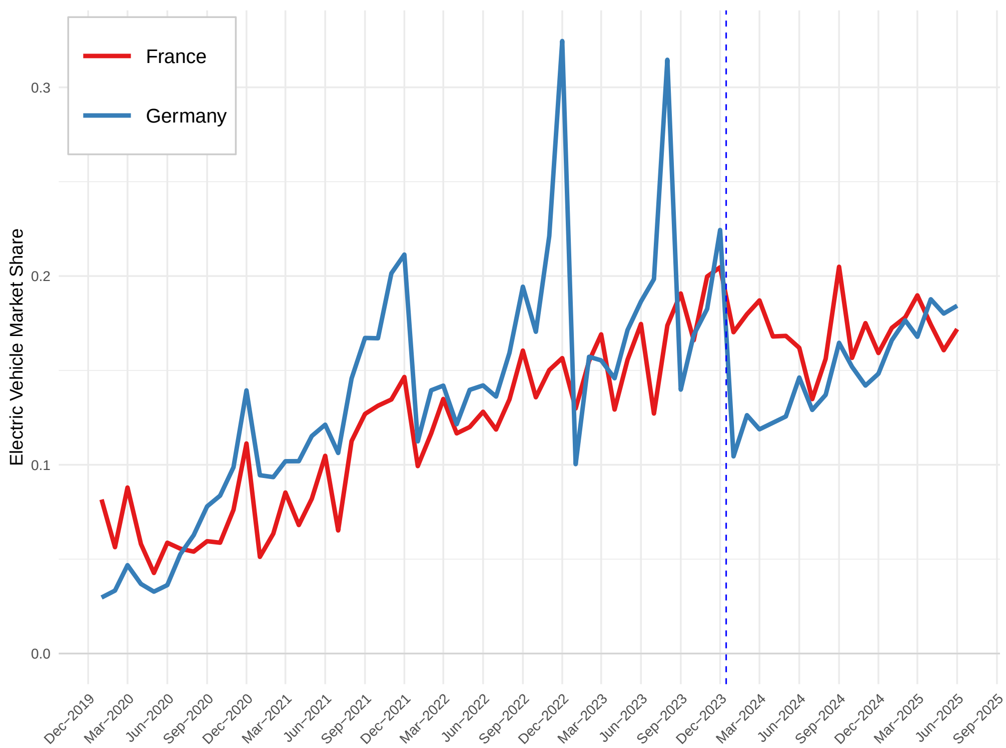

Note: Data from the EU Alternative Fuels Observatory, battery electric vehicle share among newly registered passenger cars (M1).

Overall, the 2024 reform triggered a pronounced reallocation of sales within the French EV segment—towards models that remain eligible to the subsidy—yet the aggregate EV share of total new-car registrations also turns downwards around the same date. Simple synthetic controls give an inconclusive counterfactual for France. To deal with this limitation of data-driven approach to assess overall electrification, we will rely on the structural model developed below to isolate how much of that slowdown is genuinely policy-induced.

4 Reduced-form empirical results

Having documented the raw shifts in EV shares, we next estimate a difference-in-differences specification that exploits the abrupt withdrawal of bonuses for high-LCA models to identify semi-elasticities of demand for EVs in Section 4.1. Section 4.2 presents the evidence regarding the demand for ICEs, exploiting changes in the malus.

4.1 Evidence on the reform of the EV incentive

Aggregation and econometric specification. For our analysis, we consider monthly sales grouped at the level of the *carline*, which we define as a unique combination of brand, model, and "Puissance Fiscale" (a tax-related binned measure of power).[10] This aggregation captures the variation of different trims within a car model, while reducing measurement error that might arise from changes in identifiers ("cnit") in the raw dataset over time. Appendix section B.2 presents robustness checks for different levels of aggregation. We further restrict the sample to a six months window before and after the policy change. Finally, we only consider models that had positive sales in at least one of the six months prior to the policy change.

A model is labeled as treated if it had a purchase price below 47.000 euro and a weight under 2.4 tons as of December 2023 (i.e. eligible before the reform) and if it lost its eligibility in January 2024. Overall, those represented 31% of our sample of 2023 sales. Car models that were already non-eligible in December 2023 represented around 23% and those that were eligible and retained their status throughout January 2024 made up the remaining 46%. In other words, the reform impacted about a third of the market for electric vehicles in France.

Table 2 presents further summary statistics for the electric vehicle (EV) analysis sample. Summary statistics are presented both for all electric vehicles in the sample, as well as for the subset of those that were eligible in December 2023.

We use PPML regression which has become the standard theory-consistent estimator for gravity and firm-level market share regressions (on top of accounting for issues related to 0-shares (Silva and Tenreyro, 2006; Chen and Roth, 2024), this ensures the consistency of the fixed effects with the denominator of the market share). We present both a dynamic (3) and a static (4) specification.

$$\textbf{Dynamic:}\quad s_{j|g,t} = \exp\left[\sum_{k\neq -2} \beta_k \times \mathbb{1}\{k=t\} \times \texttt{treated}_j + \underbrace{\delta_{s,g,t}}_{\text{seg}\times\text{fuel type}\times t\ \text{FE}} + \underbrace{\gamma_j}_{\text{carline FE}} + \epsilon_{j,t}\right], \quad (3)$$

$$\textbf{Static:}\quad s_{j|g,t} = \exp\left[\beta \times \mathbb{1}\{t>0\} \times \Delta\tau_j + \underbrace{\delta_{s,g,t}}_{\text{seg}\times\text{fuel type}\times t\ \text{FE}} + \underbrace{\gamma_j}_{\text{vehicle FE}} + \epsilon_{j,t}\right], \quad (4)$$

where $s_{j|g,t}$ is the share of vehicle j within fuel type $g = \{\texttt{EV},\texttt{ICE}\}$ at time t. Each carline j belongs to a segment s (Sedan, SUV, Minivan etc). The dynamic specification is an event study with treatment interacted with months-since treatment dummies. The static regression measures the impact of the

[10] This aggregates the data from the narrow definition of the "cnit".

Table 2: Summary Statistics of EV Analysis Sample

Variable	Mean	Std. Dev.	P5	P95	Obs
All models					
Number of Unique Groups	166.0	NA	NA	NA	NA
Number of Unique Treated Groups	27.0	NA	NA	NA	NA
Sales in units per month	174.0	503.7	0.0	848.1	1980
\% of 0 market share observations	15.1	NA	NA	NA	NA
Price (€)	42410.5	13542.1	29990.0	134490.0	1980
Ever Eligible models					
Number of Unique Groups	62.0	NA	NA	NA	NA
Number of Unique Treated Groups	27.0	NA	NA	NA	NA
Sales in units per month	358.1	765.3	0.0	1756.4	744
% of 0 market share observations	12.1	NA	NA	NA	NA
Price (€)	37428.9	6833.5	27308.8	45750.0	744

Note: This table reports summary statistics for the EV analysis sample, sub model by month panel. Prices are weighted by sales. Sales are measured in units; prices are filled averages over the sample period.

amount of subsidy lost ($\Delta\tau_j$), with the lower bound defined in (1). In section 5.4, we discuss the link between these double difference estimates and the demand parameters of the model we use for counterfactuals, depending on the structure of fixed effects included in the specification.

Results: diff-in-diffs for EVs. Figure 8 presents the event study graph from equation (3) where treatment is a dummy indicating eligibility loss. In March 2024, we estimate a coefficient around -1, which indicates a decrease in sales of about 60% to be compared to October 2023, our reference point. We see very little anticipation in November and December 2023, confirming that the policy change was quite unexpected. Sales of newly un-eligible carlines fall further in the next months before stabilizing. Appendix Figures B.2.4 present the same analysis with different sets of fixed effects. Appendix figures B.2.8 present the same analysis for the disaggregated sample. In appendix B.2.4, we use an alternative specification with firm (e.g. Volkswagen, Stellantis, General Motors, etc.) by segment by month fixed effects.

Table 3 presents the static results of equation (4) for all EVs in the sample across multiple sets of fixed effects. The coefficients capture the effect of a one thousand euro price increase through loss of eligibility. The coefficients are negative and significant across the specifications, indicating that losing eligibility decreased the market share of the concerned car models. A robustness check where sales are aggregated at the most granular level ("cnit" instead of carline) is presented in appendix Table B.2.1. As will be made clear in section 5.4.2, the choice of fixed effects of columns (2) and (3) allow for a clear mapping between our estimate and demand parameters of the model.

4.2 Evaluation of the impact of the feebate

Aggregation and econometric specification. For the analysis of the change in the malus, from December 2023 to January 2024, we consider observations grouped at the brand, model, and CO_2

Table 3: The impact of the loss of eligibility on new vehicle registrations

Dependent Variable:	Market Share		
Model:	(1)	(2)	(3)
Variables			
Treatment (TP)	-0.153***	-0.151***	-0.193**
	(0.042)	(0.050)	(0.085)
Fixed-effects			
Month	Yes		
Carline	Yes	Yes	Yes
Month × Segment		Yes	
Month × Segment × Firm			Yes
Fit statistics			
Observations	1,975	1,965	1,923
Pseudo R^2	0.224	0.226	0.235
Implied own-price elasticity	-4.24	-4.18	-5.35

Note: This table presents results of a static difference-in-differences regression of number of new vehicle registrations on loss of eligibility according to equation (4). Coefficients capture the effect of a 1000 euro price increase. The time period considered covers six pre- and six post-reform months. Observations are grouped at the carline (brand × model × fiscal power) level. The same regression with observations grouped at the cnit level can be found in appendix table B.2.1. Clustered (carline) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

Figure 8: Event study of loss of eligibility

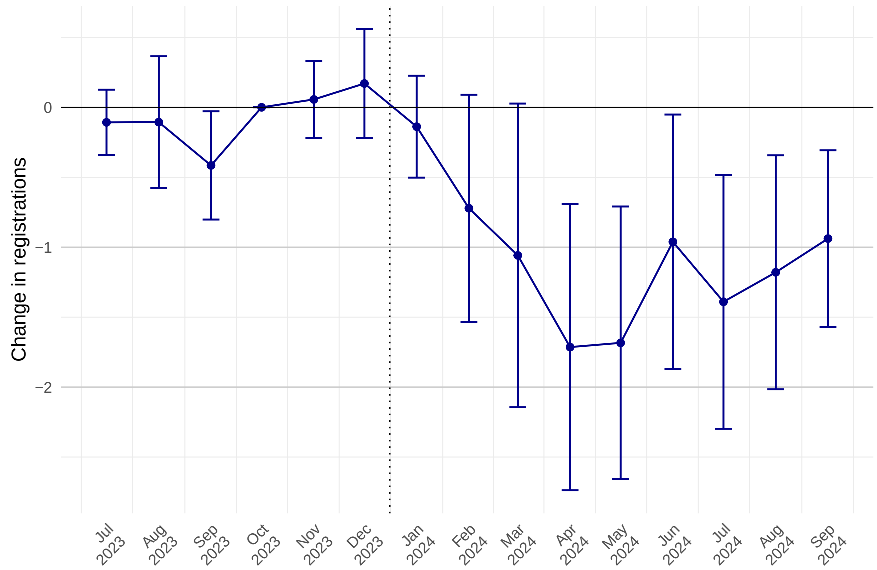

Note: This figure presents results of a dynamic difference-in-differences regression of the carline's market share in new vehicle registrations on the loss of eligibility according to equation (3). Vertical bars indicate 95% confidence intervals. Specifications with alternative fixed effects can be found in appendix B.2.4. The figures in Appendix B.2.8 present the same analysis for the disaggregated sample.

(g/km) level. Taking into account CO_2 emissions is necessary insofar as it is the primary variable determining the level of the malus. As with EVs, we restrict the sample to a six months window before and after the policy change, where we impose that at least one of the six pre-policy change months must have displayed positive sales. A final restriction that we impose is to exclude vehicles experiencing a malus increase of more than 20,000 euros. While those only represent a small fraction of sales, they are the most polluting luxury cars, which might have fundamentally different demand responses to price changes.

Table 4 presents summary statistics. It shows that the ICE analysis sample, which includes all vehicles but electric ones, has a somewhat finer grouping than the preferred aggregation of EVs, which is due to the much more differentiated supply of ICE models.

Results: diff-in-diffs for ICEs. Figure 9 presents the event study graph where treatment is a binary variable equal to one if the malus increase from 2023 to 2024 was larger than 500 euros. The average effect over the first six post reform months is -0.456 (appendix table B.3.2 reports the full static specification), which translates into a roughly 37% decrease in market share for vehicles experiencing a malus increase of more than 500 euros relative to those that do not.

Table 5 shows the benchmark results for the effect of the 2023 to 2024 increase in malus on market share. Coefficients capture the effect of a 1000 euro price increase. As in the analysis of the bonus, they are negative and significant across the three sets of fixed effects. Again, as discussed section 5.4.2, these results will allow to identify important parameters of the demand system we use to build counterfactual scenarios.

Table 4: Summary Statistics of ICE Analysis Sample

Variable	Mean	Std. Dev.	P5	P95	Obs
Diesel & Petrol models					
Number of Unique Groups	2429.0	NA	NA	NA	NA
Sales in units per month	32.4	155.1	0.0	148	25696
\% of 0 market share observations	56.4	NA	NA	NA	NA
Price (€)	28489.2	14485.6	21409.8	136850	25696

Note: This table reports summary statistics for the ICE analysis sample, sub model by month panel. Prices are weighted by sales. Registrations are measured in units; prices are filled averages over the sample period.

Figure 9: Event study of increase in malus among ICE

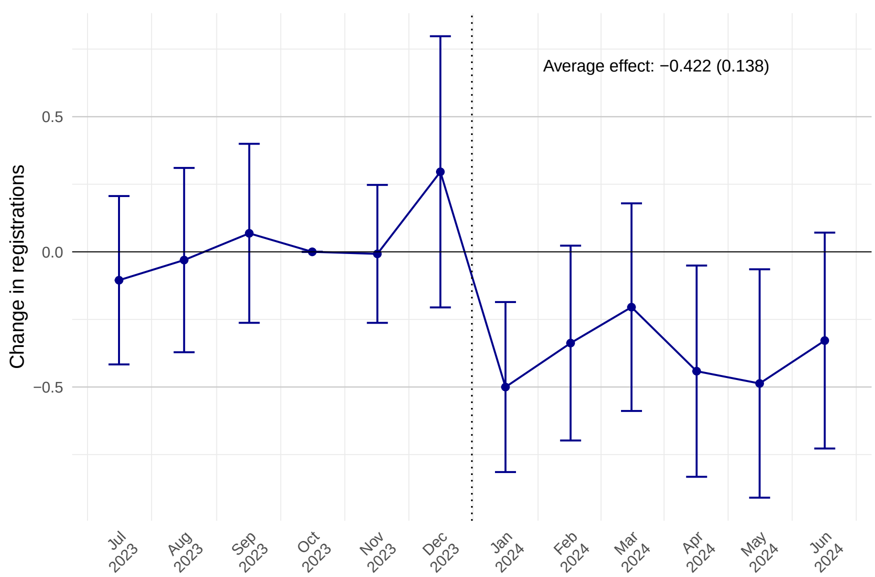

Note: This figure presents results of a dynamic difference-in-differences regression of number of new vehicle registrations on a dummy equal to one if the malus increase was more than 500 euros according to equation (3). Vertical bars indicate a 95% confidence interval. Observations are grouped at the Brand × Model × CO2 g/km level. The regression includes group and month × motorization fixed effects. Specifications with alternative fixed effects can be found in the appendix. Specifications with alternative fixed effects can be found in appendix figure B.3.4.

Table 5: The impact of the malus increase on new vehicle registrations

Dependent Variable:	Market Share		
Model:	(1)	(2)	(3)
Variables			
$\Delta\tau_j \times Post$	-0.173***	-0.168***	-0.194***
	(0.044)	(0.042)	(0.065)
Fixed-effects			
Month x Motorisation	Yes		
Carline	Yes	Yes	Yes
Month × Segment x Motorisation		Yes	
Month × Segment × Firm x Motorisation			Yes
Fit statistics			
Observations	25,696	25,685	24,908
Pseudo R^2	0.226	0.226	0.227
Implied own-price elasticity	-5.11	-4.99	-5.74

Note: This table presents results of a static difference-in-differences regression of number of new vehicle registrations on malus increase in 1000 of euros according to equation (4). The time period considered covers six pre- and six post-reform months. Observations are grouped at the brand × Model × CO_2 g/km. Appendix tables B.3.1 presents results for a regression controlling for estimated markups. Clustered (carline) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

5 Conceptual framework

To recover the deeper substitution patterns and welfare effects, we embed our semi-elasticity estimates into a three-level nested-logit framework. Section 5.1 lays out the model and section 5.2 presents the associated choice probabilities. Section 5.3 specifies the mean utility associated with each vehicle and presents our approach to identification. In particular, section 5.4 presents our sequential approach to identification. We show how second-choice survey moments pin down the nest parameters (5.4.1) and then describe how appropriately parametrized versions of our diff-in-diffs regressions identify a combination of the remaining parameters (5.4.2). We present the results in section 5.5.

5.1 Nested–logit structure

We consider a static model of car choice in the spirit of Goldberg (1995) and following a long tradition of discrete choice models (McFadden et al., 1978; Berry, 1994).

Passenger-car demand is formalized as a three–level nested logit. Consumers first choose between buying a given power-train $g \in \{\text{EV}, \text{ICE}\}$ or not buying a car; within each g they then pick a *segment* s (*SUV*, *Sedan*, *Minivan*, *Break*, *Sports*, *Other*); finally they choose a precise carline j (defined by a brand, model, and power rating combination). Figure 10 illustrates the decision tree, whose *branch parameters* $\mu_{4,g}$ (within–segment) and μ_3 (across segments) regulate substitution patterns.[11]

Figure 10: Decision tree in the two–level nested logit.

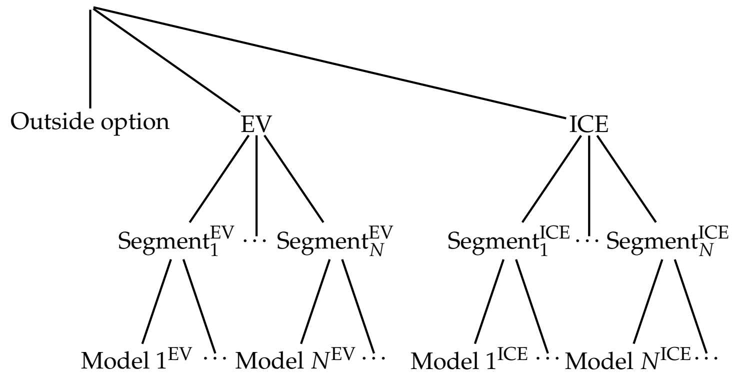

5.2 Choice probabilities

Denote by s_j the market share or choice probability for carline j among *all* options and by s_g°, $s_{s|g}^\circ$, and $s_{j|s,g}$ the conditional shares at each stage.[12] The nested-logit algebra delivers the three–way factorization

We can express the choice probability for carline j as the product of several nested conditional probabilities

$$s_j = s_g^\circ \times s_{s|g}^\circ \times s_{j|s,g}, \tag{5}$$

and the nested logit algebra delivers closed forms for each component of this factorization:

$$s_{j|s,g} = \frac{\exp(V_{jt}/\mu_{4,g})}{\sum\limits_{k\in(s,g)} \exp(V_{kt}/\mu_{4,g})}, \quad s_{s|g}^\circ = \frac{\exp\big(IV_{s|g,t}/\mu_3\big)}{\sum\limits_{r\subset g} \exp\big(IV_{r|g,t}/\mu_3\big)}, \quad \text{and} \quad s_g^\circ = \frac{\exp(IV_{g,t})}{1+\sum\limits_{h} \exp(IV_{h,t})}, \tag{6}$$

and inclusive values given by

$$IV_{s|g} = \mu_{4,g} \ln\Big[\sum_{k\in(s,g)} \exp(V_k/\mu_{4,g})\Big], \quad \text{and} \quad IV_g = \mu_3 \ln\Big[\sum_{r\subset g} \exp(IV_{r|g}/\mu_3)\Big]. \tag{7}$$

5.3 Indirect utility and the estimating equation

We now specify how we model the process that determines that drive carline j's c'es utility in month t is

$$V_{jt} = -\alpha p_{jt} + \xi_j + \xi_{jt}, \tag{8}$$

[11]When $\mu_{4,g} = \mu_3 = 1$ the nests collapse to the standard logit. Smaller values increase within-nest substitution.

[12]The subscript $\circ$ is used to signal market shares aggregated at a higher level of aggregate than the carline (segment or fuel-type).

where p_{jt} decomposes into the policy instrument (τ_{jt}), marginal cost (mc_j), and a possibly endogenous mark-up ($\tilde{p}^0_{jt}$):

$$p_{jt} = \tau_{jt} + mc_j + \tilde{p}^0_{jt}.$$

It will be useful to define the market share of j within EVs, defined as $s_{j|g,t} \equiv s_{j,t}/s^{\circ}_{g,t}$. Using (6), (7), and (8) and taking logs, we obtain that

$$\ln s_{j|g,t} = -\frac{\alpha}{\mu_{4,g}}\tau_{jt} + \frac{\xi_j - \alpha mc_j}{\mu_{4,g}} - \underbrace{\left(\frac{1}{\mu_{4,g}} - \frac{1}{\mu_3}\right) IV_{s|g,t} - \frac{IV_{g,t}}{\mu_3}}_{\text{varies at the } gst \text{ level}} - \frac{\alpha}{\mu_{4,g}}\tilde{p}^0_{jt} + \frac{\xi_{jt}}{\mu_{4,g}}. \tag{9}$$

Equation (9) highlights the difficulty of applying standard instrumental variable to identify all parameters of the model. In particular, to identify separately $\mu_{4,\text{EV}}$ and μ_3 one would need to instrument for the overall attractiveness of EV at a given time period $IV_{\text{EV},t}$. This is unattractive as it is difficult to isolate time series variation in the relative attractiveness of EV at the aggregate level that is not confounded by other possible shocks (changes in preferences, business cycle, etc). We therefore consider a specification with a rich set of fixed effects which will allow us to recover $\frac{\alpha}{\mu_{4,g}}$. The term in braces depends only on the $\langle s, g, t\rangle$ combination; it is purged with segment–fuel type–month fixed effects δ_{sgt}, whereas the time-invariant $\{\xi_j - \alpha mc_j\}$ are absorbed by model fixed effects γ_j. We can therefore write equation (9) under a form that matches the key estimation equation (4) which was fitted in section 4:

$$\ln s_{j|g,t} = -\frac{\alpha}{\mu_{4,g}}\,\tau_{jt} + \delta_{sgt} + \gamma_j + u_{jt}, \quad \text{with} \quad u_{jt} = -\frac{\alpha}{\mu_{4,g}}\tilde{p}^0_{jt} + \frac{\xi_{jt}}{\mu_{4,g}}\,. \tag{10}$$

Quasi-experimental changes in τ_{jt} (bonus withdrawals for EVs and malus hikes for ICEs) identify the key semi-elasticities $\alpha/\mu_{4,g}$. We discuss our estimation strategy for this combination of parameters in section 5.4.2. In a second step, we use U.S. second-choice data to pin down the nesting parameters $\{\mu_3, \mu_{4,\text{EV}}\}$; together these deliver α and $\mu_{4,\text{ICE}}$, completing the demand system.

5.4 Identification and estimation

The nested nature of the model allows for a sequential identification of the parameters of the model. We first show how second choice data pin down the parameters $\{\mu_3, \mu_{4,\text{EV}}\}$ and then we show how to relate our diff-in-diffs estimates to α/μ_{4g}.

5.4.1 Second choices as a source of identification in the nested logit model

We build on papers using second choice data to estimate a model of discrete choice for differentiated goods (e.g. Berry et al., 2004). We borrow two empirical moments from Allcott et al. (2024). Our procedure relies on expressing the retention rates as a function of parameters and observed shares. As such, it bypasses the need to estimate the mean utility V_j associated with each choice j. It also makes explicit the absence of dependence of these parameters and moments on the price coefficient α.

The *second-choice* data used by Allcott et al. (2024) report, for every EV buyer, which vehicle they would have purchased had their first choice been unavailable. Replicating that hypothetical removal in the model provides moments that map directly to the two unknown branch parameters μ_3 (segment level) and $\mu_{4,\text{EV}}$ (vehicle level inside the EV nest).

Inclusive values within the EV nest. Let s denote a segment inside the EV branch and V_j the deterministic utility of model j. *Inclusive values* become

$$IV_{s|\text{EV}} \;=\; \ln\left(\sum_{k\in(s,\text{EV})} e^{V_k/\mu_{4,\text{EV}}}\right), \qquad IV_{\text{EV}} \;=\; \ln\left(\sum_{r\subset\text{EV}} e^{\mu_{4,\text{EV}}/\mu_3\; IV_{r|\text{EV}}}\right). \tag{11}$$

Given (11), the three conditional choice probabilities can be rewritten as

$$s_{j|s,\text{EV}} = \frac{e^{V_j/\mu_{4,\text{EV}}}}{\sum\limits_{k\in(s,\text{EV})} e^{V_k/\mu_{4,\text{EV}}}}, \quad s^{\circ}_{s|\text{EV}} = \frac{e^{\mu_{4,\text{EV}}/\mu_3\; IV_{s|\text{EV}}}}{\sum\limits_{r\subset\text{EV}} e^{\mu_{4,\text{EV}}/\mu_3\; IV_{r|\text{EV}}}}, \quad \text{and} \quad s^{\circ}_{\text{EV}} = \frac{e^{\mu_3 IV_{\text{EV}}}}{1+\sum\limits_{g\in\{\text{EV,ICE}\}} e^{\mu_3 IV_g}}.$$

Removing a single model. Suppose model $j \in (s,\text{EV})$ is withdrawn. At the lower level, the segment's inclusive value falls by $\Delta IV_{s|\text{EV}} = \ln(1 - s_{j|s,\text{EV}}) < 0$. Propagating this shock upward yields a change in the EV nest's attractiveness:

$$\Delta IV_{\text{EV}} = \ln\left(1 - s^{\circ}_{s|\text{EV}} + s^{\circ}_{s|\text{EV}}(1 - s_{j|s,\text{EV}})^{\mu_{4,\text{EV}}/\mu_3}\right). \tag{12}$$

Because both $s_{j|s,\text{EV}}$ and $s^{\circ}_{s|\text{EV}}$ are *observed* pre-removal shares, ΔIV_{EV} is a closed-form function of the unknown pair $(\mu_3, \mu_{4,\text{EV}})$.

Retention rate moments. We can express two retention rates predicted by the nested logit structure. The first one is the fraction of j buyers that would keep on buying an EV, should j be removed from the choice set. The second one is the same concept computed at the EV $\times$ segment level.

Using equation (12), the post-removal EV share is $s^{\circ}_{\text{EV}}(\text{no } j) \;=\; s^{\circ}_{\text{EV}}\,\Delta_j/\left[\,s^{\circ}_{\text{EV}}\Delta_j + (1 - s^{\circ}_{\text{EV}})\right]$, where $\Delta_j = \exp\left(\mu_3 \Delta IV_{\text{EV}}\right)$.

The EV retention rate after removing vehicle j is therefore

$$\eta_j^{(\text{EV})} = \frac{s^{\circ}_{\text{EV}}(\text{no } j) - (s^{\circ}_{\text{EV}} - s_j)}{s_j}, \qquad 0 < \eta_j^{(\text{EV})} < 1. \tag{13}$$

The EV/segment retention rate implied by theory is

$$\eta_j^{(s,\text{EV})} = \frac{s^{\circ}_{\text{EV}}(\text{no } j)\, s_{s|\text{EV}}(\text{no } j) - (s^{\circ}_{\text{EV}}\, s_{s|\text{EV}} - s_j)}{s_j}, \tag{14}$$

with $s_{s|\texttt{EV}}(\text{no } j) = \dfrac{s_{s|\texttt{EV}}\,\Lambda_j}{s_{s|\texttt{EV}}\,\Lambda_j + 1 - s_{s|\texttt{EV}}}$ and $\Lambda_j = (1 - s_{j|\texttt{EV},s})^{\mu_4/\mu_3}$.

Averaging over all EV vehicles, we obtain:

$$X_{\text{EV}}^{m}(\mu_3, \mu_{4,\text{EV}}) = \sum_{j \in \text{EV}} s_{j|\text{EV}}\, \eta_j^{(\text{EV})}, \quad \text{and} \quad X_{s,\text{EV}}^{m}(\mu_3, \mu_{4,\text{EV}}) = \sum_{j \in \text{EV}} s_{j|\text{EV}}\, \eta_j^{(s,\text{EV})}. \tag{15}$$

As empirical counterparts to those retention rates, we take the figures reported by Allcott et al. (2024), such that $X_{\text{EV}}^{o} = 0.52$ and $X_{s,\text{EV}}^{o} = 0.33$. Matching $X_{\text{EV}}^{m} = X_{\text{EV}}^{o}$ and $X_{s,\text{EV}}^{m} = X_{s,\text{EV}}^{o}$ pins down the nest scales $\{\mu_3, \mu_{4,\text{EV}}\}$ without requiring any information on the attractiveness of each product or the price coefficient α.

Illustration. For illustrative purposes, we draw $X_{\text{EV}}^{m}(\mu_3, \mu_{4,\text{EV}})$ in figure 11 under the simplifying assumption $\mu_{4,\text{EV}} = \mu_3$: the retention rate declines sharply as μ_3 rises toward the logit limit, illustrating the moment's identifying power.

Figure 11: Model-predicted EV-on-EV retention rate as a function of μ_3 when $\mu_{4,\text{EV}} = \mu_3$.

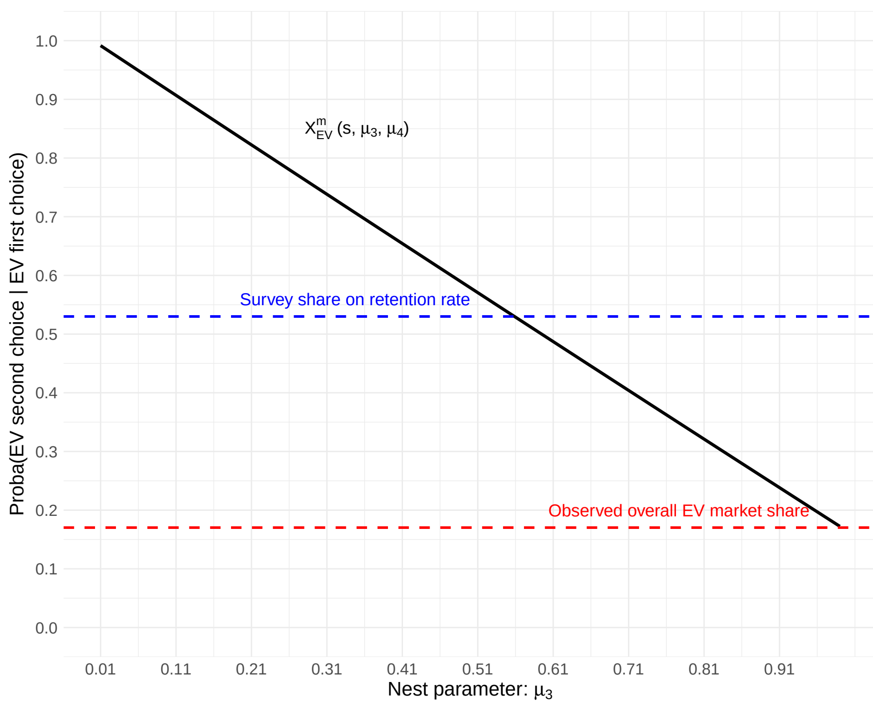

Note: This graph display the EV retention rate in single level nested logit (ICE/EV) as a function of μ_3, using French data on observed market share in 2023. The blue line represents the empirical estimate of the EV retention rate from Allcott et al. (2024). The empirical value $X_{\text{EV}}^{o} = 0.52$.

Once $\{\mu_3, \mu_{4,\text{EV}}\}$ are pinned down, the quasi-experimental estimates of $\alpha/\mu_{4,\text{EV}}$ and $\alpha/\mu_{4,\text{ICE}}$ from regression (10) deliver the full set of demand parameters $\{\alpha, \mu_3, \mu_{4,\text{EV}}, \mu_{4,\text{ICE}}\}$, closing the model used in Sections 6.

5.4.2 Mapping the diff-in-diffs estimates to the model price parameter α

Going back to equation (10), we now discuss whether, and under what assumptions, the regressions presented in section 4.1 identify the ratio $\alpha/\mu_{4,g}$.

(0) Monopolistic competition in pricing. If firms price at marginal cost (i.e. $\tilde{p}^0_{jt} = 0$ as in the perfectly competitive case) or if they do not vary conditional on firm and time fixed effect (monopolistic competition), the policy variable τ_{jt} is uncorrelated with the unobserved demand shocks ξ_{jt}; hence equation (10) can be estimated directly, provided segment-fuel type-period (gst) fixed effects are included. These coefficients are presented in columns (2) of Table 3 and Table 5 for ICE and EVs respectively.

(1) Absorbing mark-ups with firm-gst fixed effects. Under multi-product Bertrand pricing and nested logit demand, mark-ups vary at the firm $\times\, g \times s \times t$ level—see Appendix C.1.1. Introducing firm $\times\, g \times s \times t$ fixed effects absorbs those mark-up changes and breaks any spurious correlation between τ_{jt} and $\tilde{p}^0_{jt}$. These coefficients are presented in columns (3) of Table 3 and Table 5 for ICE and EVs respectively, and tend to be somewhat larger than the one including only segment-fuel type-period (gst) fixed effects. Whereas this specification is consistent with the theoretical framework (Nested logit demand with Bertrand multiproduct pricing), it throws away a lot of the variation to identify the parameter. Accordingly, we prefer the approach presented below which is theory consistent and preserves more of the variation.

(2) Iterative control for endogenous mark-ups. The parameters $\{\mu_3, \mu_{4,\text{EV}}\}$ are identified from second choice data. Starting from an initial guess for the structural parameters $\{\alpha, \mu_{4,\text{ICE}}\}$, we compute model-implied mark-ups $m_{jt}(\theta)$ where θ gathers all 4 structural parameters, add $-\alpha m_{jt}(\theta)/\mu_{4,g}$ to the dependent variable in equation (10), re-estimate equation (10), update θ with a new guess for $\alpha/\mu_{4,g}$, and repeat until convergence (appendix C.1.2). This procedure delivers a consistent semi-elasticity even when mark-ups respond endogenously to policy (see Head and Mayer, 2021; Breinlich et al., 2024).

Appendix table B.3.1 presents these results.[13] Results are very similar suggesting predicted changes in markups are small and large independent from the reform, consistent with the idea that market shares are fairly small and pass-through is close to perfect.[14]

[13] In a Nash-Bertrand equilibrium, pre-tax prices will adjust to changes in malus, through changes in markups (i.e. there's imperfect pass-through). Markups are functions of market shares, demand parameters and tax policy. Changes in markups are likely correlated with changes in tax / malus and will impact demand, thus presenting a potential threat to identification in our regression. As presented in Appendix section C.1.2, we deal with this by constructing markups based on observed market shares and assumed demand parameters. We then adjust the market share for this term assuming a given value for the within-nest price sensitivity and run the estimation using the adjusted dependent variable. Based on the results, we update the demand parameters and markups accordingly and iterate until convergence.

[14] If anything, they are slightly more negative, suggesting that endogenous changes in markups induce a positive but small bias in our baseline estimate, consistent with the notion that, on average markups have decreased slightly in response to increases in the feebate.

5.5 Estimates

The price semi-elasticities that anchor the demand system come from our difference-in-differences design: losing access to the EV purchase bonus shifts sales with an estimated semi-elasticity $\hat{\beta}_{\texttt{EV}} = -0.154$, while the 2024 increase in the CO_2-malus for ICEs implies $\hat{\beta}_{\texttt{ICE}} = -0.224$. Both coefficients stem from the specification with segment–fuel type-time and vehicle fixed effects and where we control for possible changes in markups (see Table B.3.1). Combining these coefficients with the second-choice statistics delivers the full set of nested-logit parameters collected in Table 6. The upper-nest correlation $\hat{\mu}_3 = 0.648$ indicates substantial within-fuel type substitution relative to across-fuel type substitution, and the lower-nest scale differs markedly between EVs and ICEs ($\hat{\mu}_{4,\text{EV}} = 0.569$ vs. $\hat{\mu}_{4,\text{ICE}} = 0.390$), confirming that ICE buyers view products within a segment as closer substitutes than do EV buyers.

Table 6: Estimated preference parameters

	Point estimate	
Source	Parameter	Value
Diff-in-diffs (bonus & malus)	$\hat{\beta}_{\texttt{EV}},\ \hat{\beta}_{\texttt{ICE}}$	$-0.1537,\ -0.2242$
Second choice data[a]	$\hat{\mu}_3,\ \hat{\mu}_{4,\text{EV}}$	$0.648,\ 0.569$
Combined identification	$\hat{\alpha},\ \hat{\mu}_{4,\text{ICE}}$	$0.087,\ 0.390$

[a] US new-car buyers' second-choice survey from Allcott et al. (2024), applied to French market shares.

To validate the model beyond the moments used for estimation, we compare its predictions for ICE retention behaviour—left entirely untargeted—with their empirical counterparts. As reported in Table 7, the simulated probability that an ICE buyer remains in the ICE nest when her top choice is unavailable is 0.943 (versus 0.965 in the data); conditioning on vehicle segment brings the fit even closer. The model therefore replicates second-choice substitution patterns well outside the EV calibration window, lending some more credence to the counterfactual exercises that follow.

Table 7: Untargeted retention moments (ICE nest)

	Model	**Data**
Overall ICE retention, X^m_{ICE}	0.943	0.965
Segment-specific retention, $X^m_{s,\text{ICE}}$	0.674	0.670

6 Cost benefit analysis and counterfactuals

We present the method to construct counterfactuals in terms of deviation from observed shares in section 6.1. In section 6.2, we present the results of a counterfactual policy where no environmental score reform took place.[15] In section 6.3, we carry out a cost benefit analysis, relying on additional

[15]But where impacted vehicles would have not benefited from the social leasing. See discussion in section 2.1.4

assumptions. We then implement a budget neutral alternative policy where savings are recycled to increase incentives on still eligible vehicles (section 6.4). We finally discuss the changes in producer surplus and the (limited) role of changes in markups in response to the policy (section 6.5).

6.1 Exact hat algebra for the monopolistic competition and the oligopolistic cases

Suppose the reform changes the policy instrument (bonus / malus) for model j by an amount $\Delta\tau_j$ but leaves pre-tax prices unchanged. With parameters $\theta = \{\alpha, \mu_3, \mu_4\}$ and the *observed* baseline shares $\{s_j\}$, the hat–algebra yields the counterfactual share-ratio $\widehat{s}_j \equiv s_j'/s_j$ as the product of three terms:

$$\widehat{s}_j = \underbrace{\widehat{s}_{j|s,\text{EV}}}_{\text{within segment}} \times \underbrace{\widehat{s}^{\circ}_{s|\text{EV}}}_{\text{segment within EV}} \times \underbrace{\widehat{s}^{\circ}_{\text{EV}}}_{\text{EV share in sales}},$$

with

$$\widehat{s}_{j|s,\text{EV}} = \frac{\exp\left[-(\alpha/\mu_{4,\text{EV}})\Delta\tau_j\right]}{\sum_{k\in(s,\text{EV})} s_{k|s,\text{EV}} \exp\left[-(\alpha/\mu_{4,\text{EV}})\Delta\tau_k\right]},$$

$$\widehat{s}^{\circ}_{s|\text{EV}} = \frac{\left[\sum_{j\in s,\text{EV}} s_{j|s,\text{EV}} \exp\left[-(\alpha/\mu_{4,\text{EV}})\Delta\tau_j\right]\right]^{\mu_3}}{\sum_{r\in S} s^{\circ}_{r|\text{EV}} \left[\sum_{j\in r,\text{EV}} s_{j|r,\text{EV}} \exp\left[-(\alpha/\mu_{4,\text{EV}})\Delta\tau_j\right]\right]^{\mu_3}},$$

$$\widehat{s}^{\circ}_{\text{EV}} = \frac{\sum_{r\in S} s^{\circ}_{r|\text{EV}} \left[\sum_{j\in r,\text{EV}} s_{j|r,\text{EV}} \exp\left[-(\alpha/\mu_{4,\text{EV}})\Delta\tau_j\right]\right]^{\mu_3}}{\sum_{n\in\{\text{EV},\text{ICE}\}} s^{\circ}_n \sum_{r\in S} s^{\circ}_{r|n} \left[\sum_{j\in r,n} s_{j|r,n} \exp\left[-(\alpha/\mu_{4,\text{EV}})\Delta\tau_j\right]\right]^{\mu_3}}.$$

These expressions depend only on the *baseline* market shares, the estimated demand parameters, and the vector $\Delta\tau$.

Endogenous mark-ups under multiproduct Bertrand pricing To allow pre-tax prices to adjust, we close the model with the markup formula derived in Appendix C.1. Let $m^{(0)}$ denote the baseline mark-ups. Beginning with $\Delta m^{(0)} = \mathbf{0}$,

1. Compute the predicted shares $s^{(k+1)}$ using the hat expressions above with effective price change $\Delta p^{(k)} = \Delta\tau + \Delta m^{(k)}$.
2. Feed those shares into the closed-form markup equation to obtain new mark-ups $m^{(k+1)}$; the change is $\Delta m^{(k+1)} = m^{(k+1)} - m^{(0)}$.
3. Iterate $k \to k+1$ until $\|\Delta m^{(k+1)} - \Delta m^{(k)}\| < \texttt{tol}$, typically a handful of iterations.

Convergence delivers a fixed point in which shares, mark-ups, and effective prices are jointly consistent with the demand system *and* Bertrand competition. In practice, the iterative solution moves the headline results only marginally relative to the fixed-price benchmark, consistent with the small incidence of the bonus on producers.

6.2 Counterfactual market shares in the absence of reform

Dropping the environmental-score (ES) filter. Table 8 shows what happens if the 2024 life-cycle screen is abolished while all other instruments stay in place. The composition of EV sales swings sharply away from China-assembled models toward European (and especially French) plants, yet total EV penetration falls by 1 percentage point. In other words, the ES filter is pivotal for reshoring value added but slows electrification at the margin.

Table 8: Counterfactual without the environmental-score reform (place of assembly)

Region	Cf with **no ES** reform (%)	Observed (%)	Δ (%)
Overall	17.9	16.9	-1.0
Mainland China	23.3	14.5	-8.8
Europe	51.9	57.8	5.8
Asia wo China	3.2	3.4	0.2
France	14.7	16.6	1.9
RoW	6.8	7.7	0.9

Note: Counterfactual evaluated under the *lower-bound shock* scenario (see section 2.1.4) with *fixed mark-ups* (see section 6.5 for more on varying markups). Appendix section C.2 provides details on all parameters used in the simulation.

Keeping the ES filter but canceling the 2024 CO_2 malus. Eliminating the malus while retaining the LCA threshold leaves the *within-EV* country mix essentially unchanged (Table 9); the nested structure insulates relative shares from shocks applied outside the EV nest.[16] What does change is the EV adoption rate: without the tax hike on polluting ICEs, electric cars lose ground to conventional vehicles, underscoring the malus's role as the main quantity lever.

Table 9: Counterfactual without the 2024 malus increase (assembly mix among EVs)

Region	Cf with **no MALUS** reform (%)	Observed (%)	Δ (%)
Overall	15.9	16.9	1
Mainland China	14.5	14.5	0
Europe	57.8	57.8	0
Asia wo China	3.4	3.4	0
France	16.6	16.6	0
RoW	7.7	7.7	0

Note: Counterfactual evaluated under the no 2024 malus increase (applying 2023 schedule) with *fixed mark-ups* (see section 6.5 for more on varying markups). Appendix section C.2 provides details on all parameters used in the simulation.

[16] We are transparent about this feature which is somewhat at odds with findings of more sophisticated, random coefficient models. For instance, Xing et al. (2021) tend to find that EV replace proportionally more compact ICE vehicles than large ones. Our modeling choice however is consistent with the second choice data and allows for a more transparent estimation and computation of counterfactuals.

Battery-cell sourcing. Relaxing the ES constraint has a large impact upstream in the battery supply chain (Table 10). The model predicts about a 5 pp decline in Chinese market share with a rise in EU-made battery cells of about 4 pp.

The model predicts that the reform increased EU-made battery purchases by 6 500 units. There are two countervailing forces. Holding EV market share constant, the increase in EU market share among EVs would have increased EU-made battery demand by 12 500 units, but 6000 units are lost due to shrinking EV market shares.

Table 10: Counterfactual without the environmental-score reform (battery-cell origin)

Region	Cf with **no ES** reform (%)	Observed (%)	Δ (%)
Overall	17.9	16.9	-1.0
Mainland China	61.4	56.4	-4.9
Europe	30.8	34.8	4.1
Asia wo China	5.9	6.6	0.6
RoW	0.1	0.1	0.0
Unknown	1.9	2.1	0.2

Note: Counterfactual evaluated under the *lower-bound shock* scenario (see section 2.1.4) with *fixed mark-ups* (see section 6.5 for more on varying markups). Appendix section C.2 provides details on all parameters used in the simulation.

The ES screen and the CO_2 malus act on different margins. The former reallocates demand toward lower-carbon—and more domestic—supply chains but suppresses EV uptake, whereas the latter accelerates electrification without altering the geography of value added. This result is a direct consequence of the nested logit assumption so it needs to be interpreted with caution. Overall, it suggests, however, that the two policies are complementary.

6.3 Cost benefit analysis

The model expresses every welfare component in monetary terms, permitting a like-for-like comparison between the life-cycle eligibility screen and the purchase bonus it reshapes. Two main *costs* arise. First, tighter eligibility raises average prices, cutting consumer surplus by a compensating-variation estimate of about EUR 244 million. Second, because the reform slows electrification, it increases tank-to-wheel emissions; applying a $+1.38\ \mathrm{gCO_2\,km^{-1}}$ gap over an 180 000 km lifetime and valuing carbon at 200 EUR/tCO_2 yields an external cost of roughly EUR 93 million.

Three measurable *benefits* offset those losses. Public outlays on the bonus fall by an estimated EUR 267 million—largely the mirror image of the consumer surplus loss. Second, the life-cycle screen reallocates demand toward less-carbon-intensive models, reducing manufacturing emissions by an amount worth about EUR 45 million when priced at the same social cost of carbon. Third, although harder to monetize, there are industrial pay-offs—scale economies in European battery production, learning in local assembly, and option value for future competitiveness—that add upstream value beyond what the static accounting captures.

Figure 12 assembles all quantified items. The balance of $[244 + 93 - 267] =$ EUR 70 million implies a net cost of roughly EUR 10 500 per additional EV whose battery is sourced in Europe. When the manufacturing-phase externality benefit is included, that figure falls to about EUR 2 600.

Figure 12: Decomposition of the net welfare effect of the 2024 life-cycle-based eligibility reform.

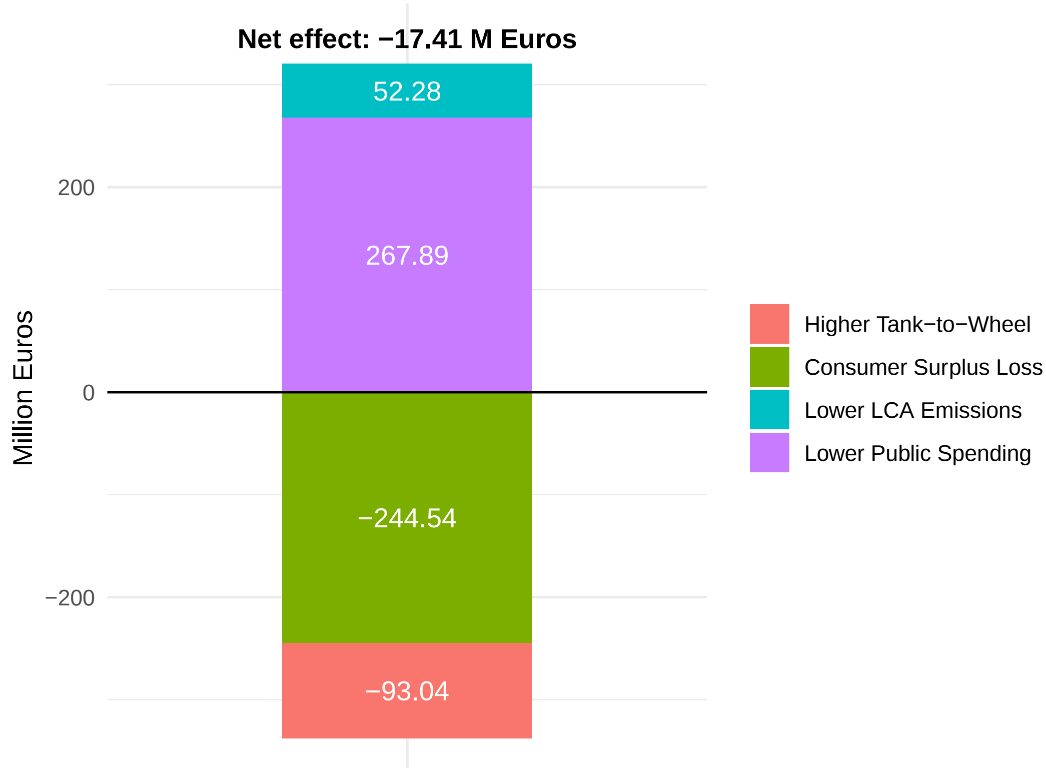

Note: Positive bars represent social benefits, negative bars social costs. Counterfactual evaluated under the *lower-bound shock* scenario (see section 2.1.4) with *fixed mark-ups* (see section 6.5 for more on varying markups). Appendix section C.2 provides details on all parameters used in the simulation.

6.4 Simple budget neutral counterfactual policy

Policy design The actual 2024 environmental-score (ES) reform cut public spending by excluding high manufacturing-phase emission models from the bonus. We consider a *budget-neutral* variant that keeps the ES eligibility filter in place but recycles the saved funds into a uniform top-up for the *remaining* eligible EVs. The revenue-neutral condition implies an average increase of $|\Delta\tau| \approx 900$ per qualifying vehicle (Fig. 13).

Figure 13: Public spending before and after re–allocating the savings from the environmental–score (ES) eligibility filter.

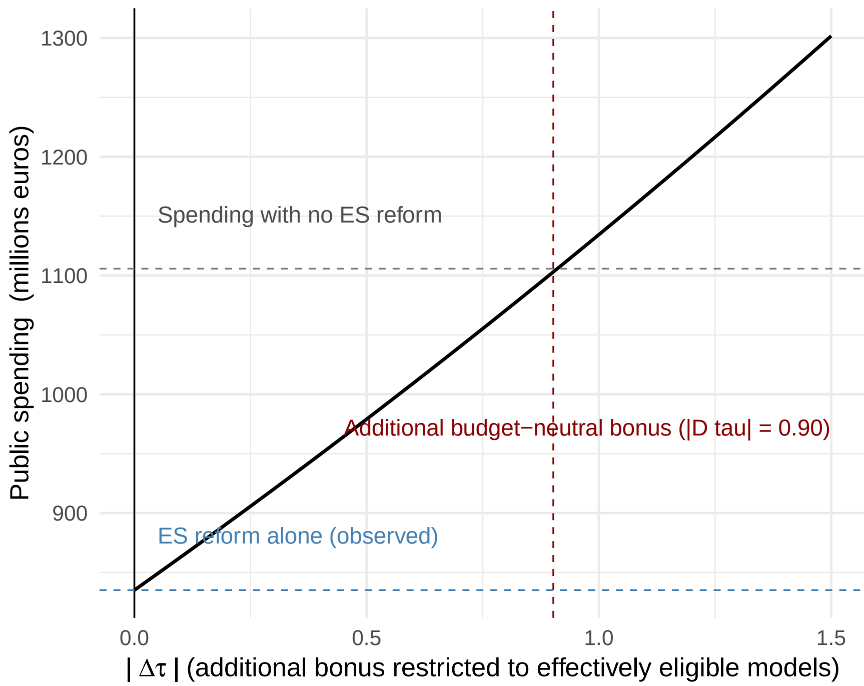

Note: Counterfactual evaluated under the *lower-bound shock* scenario (see section 2.1.4) with *fixed mark-ups* (see section 6.5 for more on varying markups). Appendix section C.2 provides details on all parameters used in the simulation.

Key outcomes. The larger bonus shifts an extra +12,500 cars toward Europe-sourced batteries while holding the aggregate EV share roughly constant. Consumer surplus (CV) declines by €61 million, so the implicit cost is $CV/\Delta battery_{EU} \approx 4\,900$ per additional "reshored" battery—about one-third of the battery's production cost. Considering the increase in tank-to-wheel emissions and manufacturing-emission benefits are included this ratio falls to about €2 600 (Fig. 14).

Environmental balance. Because the tightened eligibility still excludes high-LCA imports, the life-cycle gain per vehicle is preserved, while the higher subsidy prevents the drop in fleet electrification. As a result, the budget-neutral variant dominates the original reform on both decarbonisation and welfare grounds, albeit at a non-negligible consumer cost.

Figure 14: Decomposition of the net welfare effect for the budget-neutral variant of the ES reform.

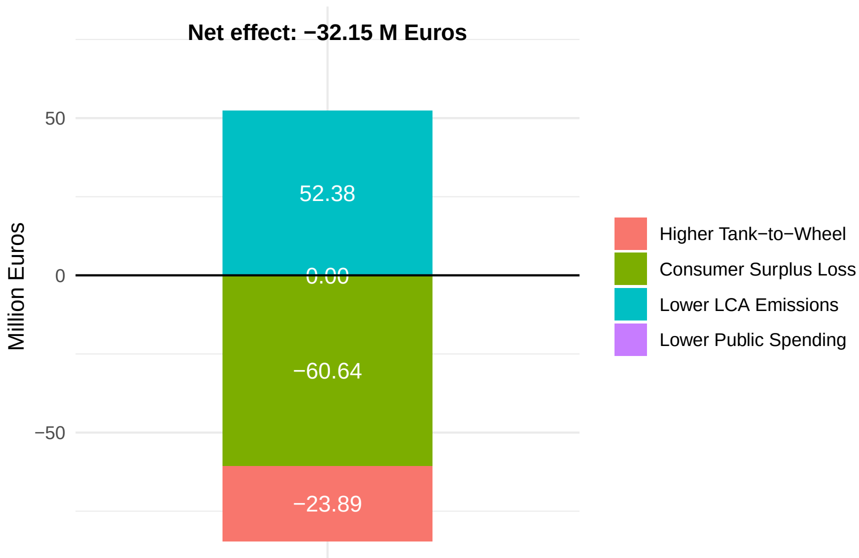

Note: Positive bars represent social benefits, negative bars social costs. Counterfactual evaluated under the *lower-bound shock* scenario (see section 2.1.4) with *fixed mark-ups* (see section 6.5 for more on varying markups). Appendix section C.2 provides details on all parameters used in the simulation.

6.5 Producer–surplus effects and variable markups

Up to this point the counter-factual experiments have held pre-tax prices constant, consistent with monopolistic competition with CRS. We now resolve every scenario under multiproduct Bertrand competition: for each firm f, $(\Theta_f \odot \Delta)\,(p_f - c_f) = -s_f$, with Θ_f the ownership matrix and Δ the Jacobian of market shares (see App. C.1). Mark-ups are re-evaluated after the policy shock and shares and prices are iterated to convergence.

As can be seen in Table 11, allowing mark-ups to adjust has a negligible impact on the reallocation away from China and on the aggregate fall in EV demand when compared with the constant-mark-up benchmark (see Table 8). Table 12 shows that the shift of battery demand from China to

Europe is virtually identical to the constant-mark-up case, confirming that price re-optimization does not materially affect sourcing patterns .

Table 11: Counterfactual *without* the environmental-score reform: EV market shares by place of assembly *with* endogenous mark-ups.

Region	Cf with **no ES** reform (%)	Observed (%)	Δ (%)
Overall	17.8	16.9	-0.9
Mainland.China	22.5	14.4	-8.1
Europe	52.1	57.8	5.7
Asia.w.o.China	3.4	3.4	0.0
France	14.9	16.6	1.7
RoW	7.0	7.7	0.7

Note: Counterfactual evaluated under the *lower-bound shock* scenario (see section 2.1.4) with *variable mark-ups* (see section 6.5 for more on varying markups). Appendix section C.2 provides details on all parameters used in the simulation.

Table 12: Counterfactual *without* the environmental-score reform: EV market shares by battery-cell origin *with* endogenous mark-ups.

Region	Cf with **no ES** reform (%)	Observed (%)	Δ (%)
Overall	17.8	16.9	-0.9
Mainland.China	53.0	47.8	-5.2
Europe	29.8	33.3	3.4
Asia.w.o.China	5.9	6.6	0.7
RoW	0.3	0.3	0.0
Unknown	10.9	11.9	1.0

Note: Counterfactual evaluated under the *lower-bound shock* scenario (see section 2.1.4) with *variable mark-ups* (see section 6.5 for more on varying markups). Appendix section C.2 provides details on all parameters used in the simulation.

Quantitatively, mark-ups change by less than €100 on average. Overall producer surplus falls only by about €7 million. The headline redistribution is internal: firms with positive change in surplus experience a ~ 152 million gains while losing firms total ~ 159 million of losses.

Because mark-ups hardly move, most of the statutory subsidy is captured by buyers rather than sellers. With prices determined by $p = c + m + \tau$, total pass-through is $\mathrm{d}p/\mathrm{d}\tau = 1 + \mathrm{d}m/\mathrm{d}\tau$. Empirically we find $\mathrm{d}m/\mathrm{d}\tau \approx 0$, implying full incidence on consumers—inline with reduced form evidence on listed prices—and confirming that the reform's welfare costs are borne almost exclusively on the demand side.

7 Conclusion

We assess France's 2024 "score environmental", i.e. life-cycle bonus–malus reform. Reduced-form estimates show registration volumes respond strongly to changes in incentives. Beyond substitution among EVs, the reform flattened the upward EV trend as our structural model attributes a 0.9 percentage-point drop in 2024 EV share to the environmental score. Our counterfactual analysis suggests that tank-to-wheel CO_2 costs exceed life-cycle gains in the baseline counterfactual. We show that a simple budget-neutral redesign would be carbon-positive while still reshoring battery demand substantially. However, the cost per battery reshored in this scenario remains fairly high at €4 600 per battery. Overall, in the current state of supply, incentives to redirect demand to EU-based producers appear fairly costly. The distributional consequences of the policy depends greatly on where the buyers of EVs that became excluded by the reform tended to sit on the income distribution. In the absence of appropriate micro-data, this is a difficult question to tackle but remains an important avenue for further research.

References

Abadie, A. (2021, June). Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature 59*(2), 391–425.

Abadie, A., A. Diamond, and J. H. and (2010). Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. *Journal of the American Statistical Association 105*(490), 493–505.

Abadie, A. and J. Gardeazabal (2003, March). The economic costs of conflict: A case study of the basque country. *American Economic Review 93*(1), 113–132.

Aghion, P., L. Barrage, D. Hémous, and E. Liu (2024). Transition to green technology along the supply chain.

Allcott, H., R. Kane, M. S. Maydanchik, J. S. Shapiro, and F. Tintelnot (2024). The effects of "buy american": Electric vehicles and the inflation reduction act. Technical report, National Bureau of Economic Research.

Barwick, P. J., H.-S. Kwon, S. Li, and N. B. Zahur (2025). Drive down the cost: Learning by doing and government policies in the global ev battery industry. Technical report, National Bureau of Economic Research.

Berry, S., J. Levinsohn, and A. Pakes (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. *The RAND Journal of Economics*, 242–262.

Breinlich, H., H. Fadinger, V. Nocke, and N. Schutz (2024). Gravity with granularity.

Chandra, A., S. Gulati, and M. Kandlikar (2010). Green drivers or free riders? an analysis of tax rebates for hybrid vehicles. *Journal of Environmental Economics and management 60*(2), 78–93.

Chen, J. and J. Roth (2024). Logs with zeros? some problems and solutions. *The Quarterly Journal of Economics 139*(2), 891–936.

CITEPA (2023). Évolution des émissions de gaz à effet de serre par secteur en France, 1990–2022. Données issues de SECTEN/CITEPA via Notre-environnement.gouv.fr.

Criscuolo, C., A. Dechezleprêtre, and G. Lalanne (2023). 7 industrial strategies for the green transition. *Sparking Europe's New Industrial Revolution*, 122.

d'Haultfoeuille, X., P. Givord, and X. Boutin (2014). The environmental effect of green taxation: the case of the french bonus/malus. *The Economic Journal 124*(578), F444–F480.

European Commission (2024). Commission report under article 12(3) of regulation (eu) 2019/631 on the evolution of the real-world co2 emissions gap for passenger cars and light commercial vehicles. COM(2024) 122 final SWD(2024) 59 final, European Commission, Brussels.

Goldberg, P. K. (1995). Product differentiation and oligopoly in international markets: The case of the us automobile industry. *Econometrica: Journal of the Econometric Society*, 891–951.

Grassi, B. (2024, November). The eu miracle: When 75 million reach high income. Bocconi University, OFCE, IGIER and CEPR.

Head, K. and T. Mayer (2021). *Poor Substitutes?: Counterfactual Methods in IO and Trade Compared*. Centre for Economic Policy Research.

Head, K., T. Mayer, M. Melitz, and C. Yang (2024). Industrial policies for multi-stage production: The battle for battery-powered vehicles. Technical report, mimeo.

Hottman, C. J., S. J. Redding, and D. E. Weinstein (2016). Quantifying the sources of firm heterogeneity. *The Quarterly Journal of Economics 131*(3), 1291–1364.

IEA (2024, April). Global EV Outlook 2024: Moving towards increased affordability. Technical report, International Energy Agency, Paris. Licence: CC BY 4.0.

Kessler, L., F. Morvillier, Q. Perrier, and K. Rucheton (2023). An ex-ante evaluation of the french car feebate. *Energy Policy 173*, 113335.

Klier, T. and J. Linn (2015). Using taxes to reduce carbon dioxide emissions rates of new passenger vehicles: evidence from France, Germany, and Sweden. *American Economic Journal: Economic Policy 7*(1), 212–42.

McFadden, D. et al. (1978). Modelling the choice of residential location.

Montout, S. and A. Robinet (2024). Le soutien au développement des véhicules électriques est-il adapté? *La note d'analyse de France Stratégie 139*(9), 1–16.

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

Sen, A. and J. Miller (2023). Vision 2050: Update on the global zero-emission vehicle transition in 2023. *International Council on Clean Transportation*.

Silva, J. S. and S. Tenreyro (2006). The log of gravity. *The Review of Economics and statistics 88*(4), 641–658.

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

Yan, S. and G. S. Eskeland (2018). Greening the vehicle fleet: Norway's CO2-Differentiated registration tax. *Journal of Environmental Economics and Management 91*, 247–262.

Online Appendix

Table of contents

A Data appendix

A.1 Determining Country of Battery Origin

In order to determine the origin of the battery of a model identified at Brand x Model level j, a two step procedure is implemented. A first data set allows us to identify the number of units of j sold in France by country of assembly. Let S^A_{jz} be the number of units of Brand x Model j sold in France and assembled in z. The share of sales of j in France that were assembled in country z then is:

$$s^A_{jz} = \frac{S^A_{jz}}{\sum_{z=1}^{Z} S^A_{jz}}$$

A second data set allows us to identify the world wide sales of model j assembled in country z by the location of battery production k. Let S^B_{jzk} be the number of units of Brand x Model j sold worldwide, assembled in country z with a battery from country k. The share of sales of j with battery from k in world wide sales of j assembled in country z is then:

$$s^B_{jzk} = \frac{S^B_{jzk}}{\sum_{k=1}^{K} S^B_{jzk}}$$

Since we only observe the world wide sales of units of j with batteries from different countries k, we approximate the share of units of j sold in France with batteries from k as follows:

$$s_{jk} = \sum_{z=1}^{Z} s^A_{jz} * s^B_{jzk}$$

A.2 Other elements

Figure A.2.1: Two models disqualified by the 2024 ES (battery and assembly in China).

B Empirical appendix

B.1 Synthetic control analysis

In order to better understand changes in the electric vehicle market in Europe we employ the synthetic control method (SCM) as developed in Abadie and Gardeazabal (2003), Abadie et al. (2010) and Abadie (2021). We analyse both the change in EV subsidies in France and the abolition of subsidies in Germany.

Data. For this analysis, we use monthly data on electric vehicle market shares in 27 European countries, provided by the European Alternative Fuels Observatory. This data is available from January 2020 until August 2024 (more requested).

Basic setup. The goal of the SCM is to generate basic counterfactuals for how the EV market would have developed absent any reform. In order to achieve this, the pre-reform trends in countries without any reforms (donor pool) are weighted in such a way as to most closely track the pre-reform trend of the country that experiences the reform. The weights are optimized for 2023 in order to align the synthetic trend with the observed market shares. Comparing the weighted trend post-reform with the observed trend then allows one to better understand whether a reform differentially impacted the affected country. The donor pool for France does not include Germany and vice versa. Furthermore, the following countries were removed from the donor pool since they themselves experience major policy changes around January 2024: Belgium, Italy and Ireland. The weights resulting from the SCM are presented in Table B.1.1. Figure B.1.7 below plots the observed and counterfactual trends for France and Germany. While initially France seems to fare better than its synthetic counterfactual, the trends quickly align again making it difficult to identify an effect. The same analysis for Germany on the other hand shows a clear difference between the observed and synthetic trend, with the observed market share underperforming its counterfactual.

Table B.1.1: Synthetic control weights

Donor	France	Germany
Netherlands	0.20	0.25
Spain	0.62	0.57
Sweden	0.18	0.18

In line with Grassi (2024), we employ two robustness exercises presented in the next section.

Leave-one-out robustness test: here we remove each non-zero weight donor one by one from the donor pool and reestimate the weights. No matter which of the positively weighted donors is removed, the overall picture remains largely unchanged.

In-country-placebo robustness test: here we plot the difference between the observed and counterfactual trends for all countries in the donor pool. While the gap between the observed French and synthetic French trend crosses into the interquartile range of other European countries, the German gap stays below it.

Figure B.1.7: Synthetic control results and robustness checks for France and Germany.

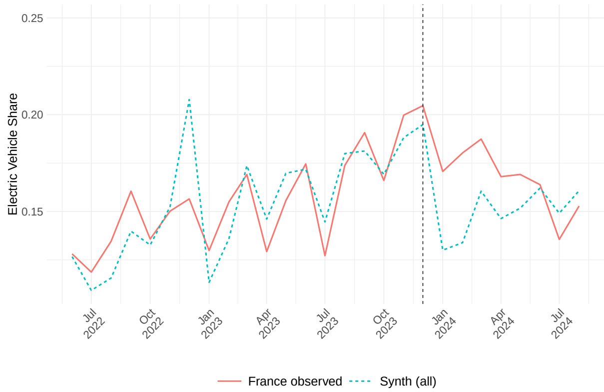

Synthetic control: France

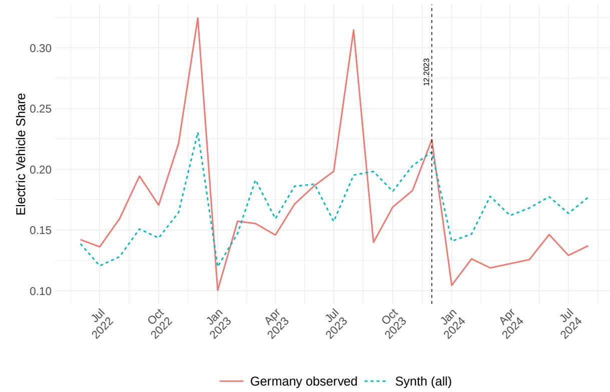

Synthetic control: Germany

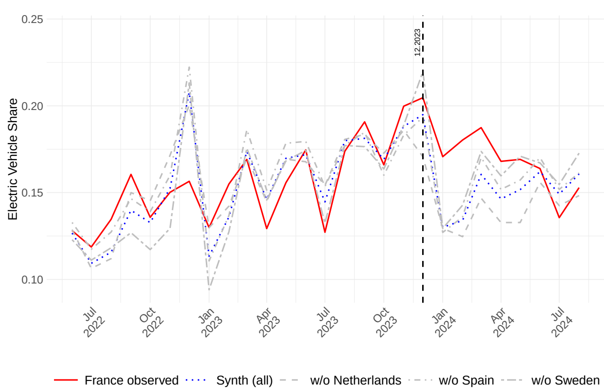

Leave-one-out: France

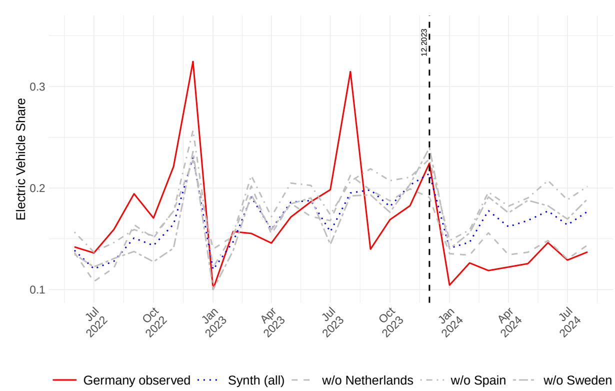

Leave-one-out: Germany

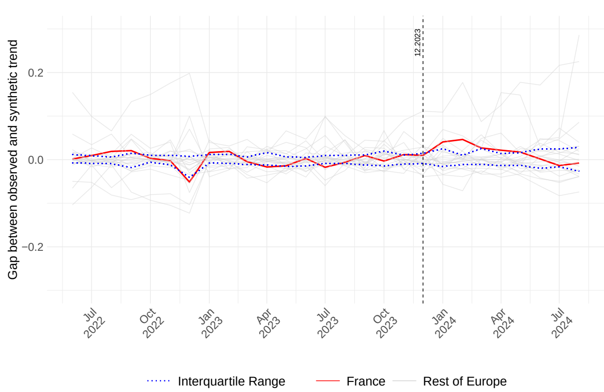

Placebo test: France

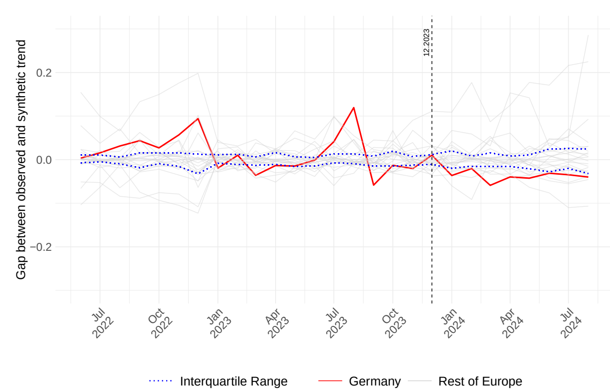

Placebo test: Germany

B.2 Bonus: Robustness and Further Event Studies

B.2.1 Brand x Model x PF aggregation

Figure B.2.4: Event study of loss of eligibility among all EVs

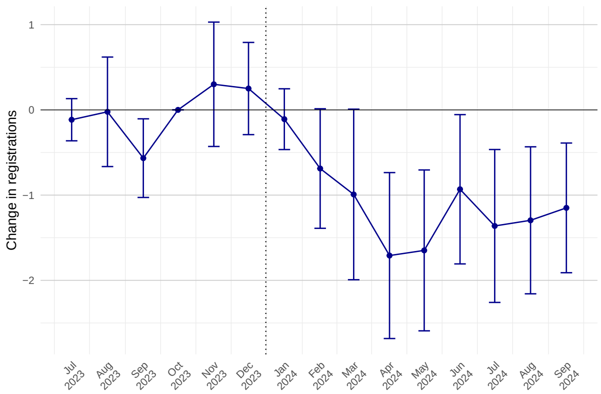

Vehicle + month fixed effects

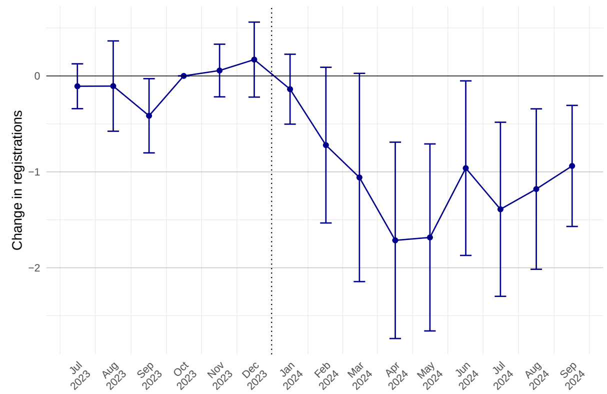

Vehicle + segment x month fixed effects

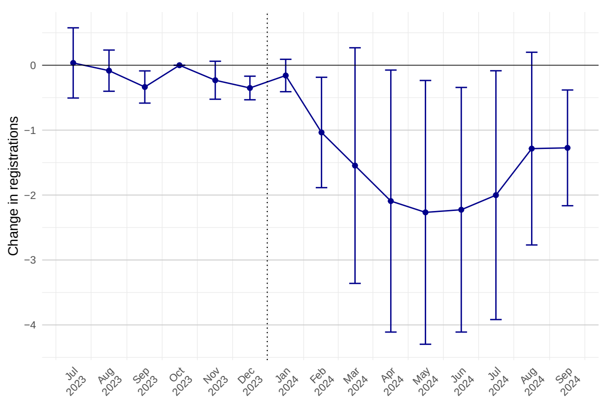

vehicle + firm x segment x month fixed effects

Notes: This figure presents results of a dynamic difference-in-differences regression of number of new vehicle registrations on loss of eligibility according to equation **??**. *Vertical bars indicate a 95% confidence interval.* ***Monthly observations are at the Brand x Model x PF level****. The main specification can be found in figure 8.*

Table B.2.1: Static regression of number of new vehicle registrations on loss of eligibility in 1000 of euros (cnit level)

Dependent Variable:	Market Share		
Model:	(1)	(2)	(3)
Variables			
Treatment (TP)	-0.169***	-0.152***	-0.236***
	(0.047)	(0.052)	(0.066)
Fixed-effects			
Month	Yes		
Carline	Yes	Yes	Yes
Month × Segment		Yes	
Month × Segment × Firm			Yes
Fit statistics			
Observations	9,490	9,479	9,345
Pseudo R^2	0.262	0.265	0.277
Implied own-price elasticity	-4.75	-4.27	-6.64

Notes: This table presents results of a static difference-in-differences regression of number of new vehicle registrations on loss of eligibility according to equation **??**. Coefficients capture the effect of a 1000 euro price increase. Observations are grouped at the cnit level. The same regression with observations grouped at the Brand x Model x Power level can be found in table 3. Clustered (carline) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

B.2.2 No aggregation (cnit level analysis)

Figure B.2.8: Event study of loss of eligibility among all EVs at the CNIT level

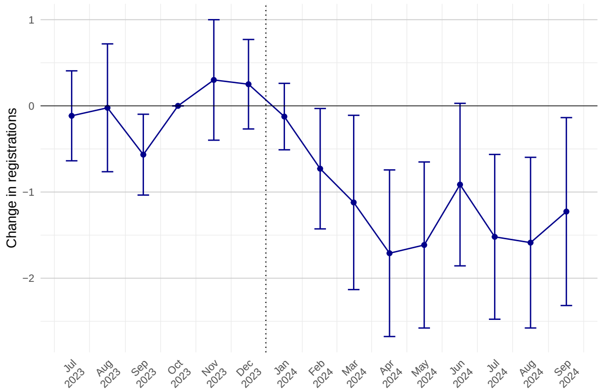

Vehicle + month fixed effects

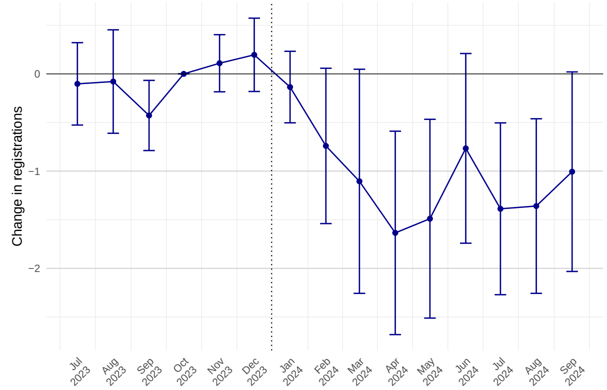

Vehicle + segment x month fixed effects

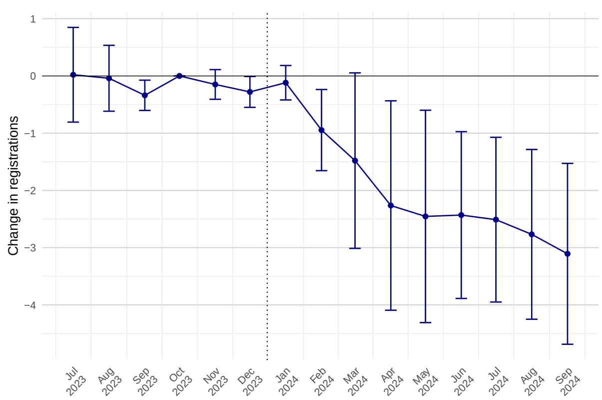

Vehicle + firm x segment x month fixed effects

Notes: This figure presents results of a dynamic difference-in-differences regression of number of new vehicle registrations on loss of eligibility according to equation **??**. *Vertical bars indicate a 95% confidence interval.* ***Monthly observations are at the CNIT level***. *The main spceification can be found in 8.*

B.3 Malus: Robustness and Further Event Studies

B.3.1 Malus: Further Static Regressions

Table B.3.1: Controlling for model-implied changes in markups in the main DiD analyses

	Baseline ICE	Iterated est. w/ ctrl for mkups	Baseline EV	Iterated est. w/ ctrl for mkups
Model:	(1)	(2)	(3)	(4)
Variables				
Malus (per €1k)	-0.2276***	-0.2242***		
	(0.0474)	(0.0466)		
Bonus Loss (per €1k)			-0.1546***	-0.1537***
			(0.0503)	(0.0503)
Fixed-effects				
Month × Segment × Motorization	Yes	Yes	Yes	Yes
Vehicle	Yes	Yes	Yes	Yes
Fit statistics				
Observations	44,627	44,627	1,965	1,965
Pseudo R^2	0.19658	0.23013	0.17480	0.19719

Notes: This table presents results of a static difference-in-differences regression of market share on continuous malus increase according to equation **??** in column (1) and (2). Column (1) is equivalent to column (2) in table 5. Column (2) controls for estimated markups. Column (3) and (4) are the EV counterpart to columns (1) and (2). The time period considered covers six pre- and six post-reform months. Observations are grouped at the brand × model × Co2 g/km level for ICEs and brand × model × power leve for EVs.

Table B.3.2: Static regression among ICE of market share on binary indicator for malus increase bigger than 500 euro

Dependent Variable:	Market Share		
Model:	(1)	(2)	(3)
Variables			
$\mathbb{1}\{\Delta\tau_j > 500\} \times Post$	-0.425***	-0.422***	-0.343**
	(0.115)	(0.138)	(0.144)
Fixed-effects			
Month x Motorisation	Yes		
Vehicle	Yes	Yes	Yes
Month × Segment x Motorisation		Yes	
Month × Segment × Firm x Motorisation			Yes
Fit statistics			
Observations	25,696	25,685	24,908
Pseudo R^2	0.22616	0.22642	0.22652

Note: This table presents results of a static difference-in-differences regression of market share on a binary indicator for malus increase larger than 500 euros according to equation **??**. The time period considered covers six pre- and six post-reform months. Observations are grouped at the brand × model × Co2 g/km level. Clustered (carline) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

B.3.2 Malus: Event Studies

Figure B.3.4: Event study of increase in malus among ICE, 2022 to 2023 increase

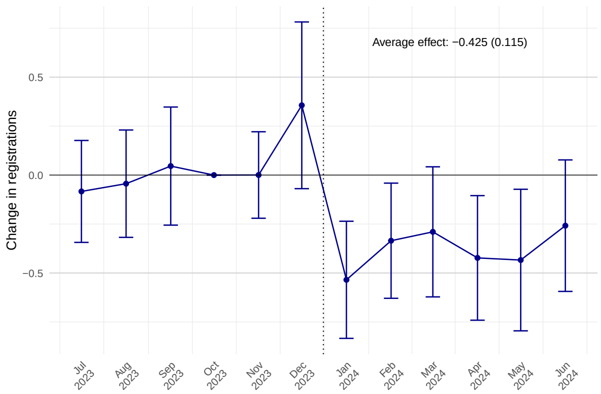

Vehicle + month x motorization fixed effects

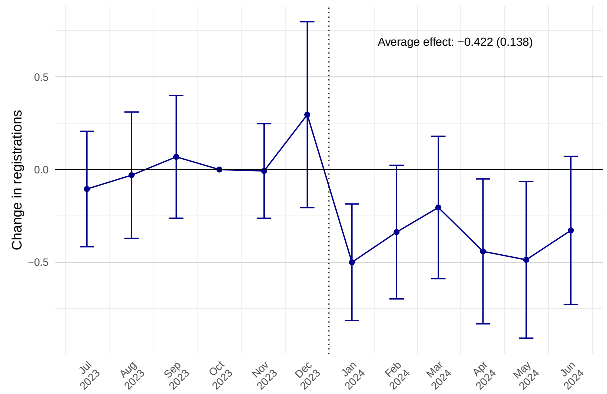

Vehicle + segment x motorization x month fixed effects

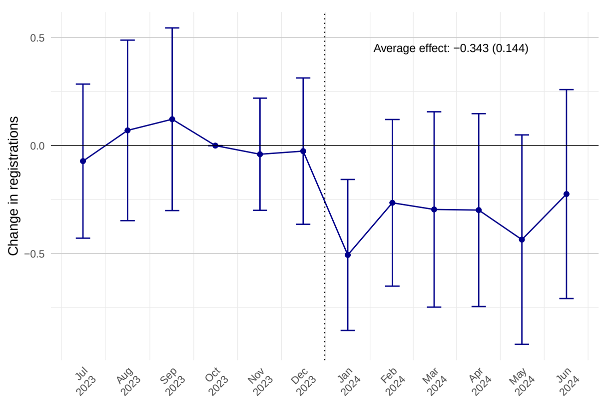

Vehicle + firm x segment x motorization x month fixed effects

Notes: This figure presents results of a dynamic difference-in-differences regression of number of new vehicle registrations on a binary indicator for malus increase larger than 500 euros according to equation **??**. *Vertical bars indicate a 95% confidence interval. The main specification can be found in figure 9.*

C Model appendix

C.1 Markups in the nested logit model

C.1.1 Analytical formula for markups

Profit-maximising prices in a differentiated-product market. Each firm f owns a set of models $\mathcal{M}_f$ and chooses the price vector $p_f = \{p_j\}_{j \in \mathcal{M}_f}$ to maximise joint profits

$$\Pi_f(p) = \sum_{j \in \mathcal{M}_f} (p_j - c_j)\, s_j(p), \qquad s_j(p): \text{market share of model } j. \tag{C.16}$$

The first-order condition for every product $i \in \mathcal{M}_f$ is

$$\sum_{j \in \mathcal{M}_f} \frac{\partial s_j}{\partial p_i} (p_j - c_j) + s_i = 0. \tag{C.17}$$

Stacking derivatives in the $J \times J$ matrix $\Delta(p) = \partial s / \partial p'$ and defining the ownership matrix Θ_f with entries $\Theta_{jk} = 1$ if $j, k \in \mathcal{M}_f$ and 0 otherwise, the system can be written compactly as

$$\big(\Theta_f \odot \Delta(p)\big)\,(p - c) = -\, s(p), \tag{C.18}$$

where $\odot$ denotes the Hadamard product. The optimal mark-up vector is therefore

$$m \;=\; p - c \;=\; -\big(\Theta_f \odot \Delta(p)\big)^{-1} s(p). \tag{C.19}$$

For a single-product firm this collapses to the familiar formula $m_i = -s_i / \big(\partial s_i / \partial p_i\big)$.

Closed-form mark-ups under a two-level (energy-segment) nest logit. With a common upper-nest scale μ_3 and group-specific lower-nest scales $\{\mu_{4,g}\}$, the mark-up on product i of firm f can be expressed entirely in terms of observed shares and structural parameters:

$$m_i = \frac{1}{\alpha(1 - S^f)} \; \frac{1}{\dfrac{1}{\mu_{4,g}} - \Big(\dfrac{1}{\mu_{4,g}} - \dfrac{1}{\mu_3}\Big) s_{f|s}} \left[1 + \Big(\tfrac{1}{\mu_3} - 1\Big) \frac{Z_g^f}{s_g^{\circ}\big[1 - (\frac{1}{\mu_3} - 1) Z_g^f / s_g^{\circ}\big]}\right], \tag{C.20}$$

where: s_g° is the share of engine-type g in total sales; $s_{f,s}$ the share of firm f in segment s; $s_{f|s} = s_{f,s}/s_s^{\circ}$ the firm–segment share; $D_s^f = 1/\mu_{4,g} + (\frac{1}{\mu_3} - \frac{1}{\mu_{4,g}}) s_{f|s}$; $Z_g^f = \sum_s s_{f,s}/D_s^f$; $S^f = \sum_g \dfrac{Z_g^f}{1 + (1 - \frac{1}{\mu_3}) Z_g^f / s_g^{\circ}}$.

Equation (C.20) shows that mark-ups are a function of market shares and the parameters $\theta = \{\alpha, \mu_3, \mu_{4,\text{EV}}, \mu_{4,\text{ICE}}\}$. In particular, all products of a given firm share the same mark-up within a segment–engine nest (s, g), a property exploited in the estimation procedure. This property has been shown by Hottman et al. (2016) to hold for nested CES demand system.

C.1.2 Iterative procedure

Here, we describe the iterative process that can be used to recover $\alpha/\mu_{4,g}$ by controlling for markups as given by equation C.20, in the spirit of Head and Mayer (2021) and Breinlich et al. (2024).

We start with a guess $b^{(n)} = \{\alpha^{(n)}, \mu_{4,\texttt{ICE}}^{(n)}\}$ and estimates for $\{\mu_3, \mu_{4,\texttt{EV}}\}$.

1. Use equation C.20 to compute markups $m_{jt}^{(n)}$.
2. Construct a new dependent variable $y_{jt}^{(n)} = \ln s_{j|g,t} - \alpha^{(n)}/\mu_{4,g}^{(n)} m_{jt}^{(n)}$, where $g = \texttt{ICE}, \texttt{EV}$.
3. Run regression using $y_{jt}^{(n)}$ as dependent variables and recover estimates from which we can get $b^{(n+1)}$.
4. Start from step (1) until $\text{dist}(b^{(n)}, b^{(n+1)}) < \texttt{tol}$.

C.2 Parameters used for the cost benefit analysis

We detail here the main input to the cost-benefit analysis.

C.2.1 Parameters

Social cost of carbon is set at €200 per ton. Tank to wheel emissions are retrieved from the WLTP number. Following findings from the EU commission (European Commission, 2024) they are inflated by a factor of 1.2 for gasoline and diesel engines including non-plugin hybrids and by a factor of 3.5 for hybrids. We assume vehicles drive 180 000 km during their lifetime.

C.2.2 Method used to compute manufacturing–phase emissions

For every vehicle we estimate the embodied CO_2 of the **manufacturing stage only**. Three contributions are added:

$$\text{CO}_{2,\text{mfg}} = \text{CO}_{2,\text{assembly}} + \text{CO}_{2,\text{battery}} + \text{CO}_{2,\text{transport}}.$$

Inputs.

- Curb mass m (kg) and engine type
- Final-assembly country (share when multiple plants)
- Battery sourcing region (share, when relevant)

1. Assembly & materials

1. **Materials**: curb mass is multiplied by $f_{\text{mat}} = 3.0$ kg CO_2/kg vehicle.
2. **Assembly energy**: the same mass is multiplied by the country factor $f_{\text{ass,c}}$:

France	0.58
Europe (avg.)	0.76
Mainland China	1.60
Asia w/o China	1.56
Rest of World	1.03

Result: $CO_{2,\text{assembly}} = m(f_{\text{mat}} + f_{\text{ass,c}})$.

2. Battery production (EV and PHEV)

1. Pack capacity E_{bat} (kWh) is inferred from curb mass:

$m \leq 1200$ kg	30 kWh
1200–1600 kg	50 kWh
1600–2000 kg	75 kWh
$m > 2000$ kg	90 kWh

Regional factors $g_{\text{bat,r}}$ (kg CO_2/kWh):

Europe	53
USA	55
Asia w/o China	63
Mainland China	68

Result: $CO_{2,\text{battery}} = E_{\text{bat}}\, g_{\text{bat,r}}$.

3. Transport to France

1. Sea freight (if applicable): $e_{\text{sea}} = 0.101$ kg $CO_2/(\text{t km})$ with distances $d_{\text{sea}} = 20\,000$ km (Asia) or 9 000 km (N. America).
2. Truck legs: $e_{\text{EU}} = 0.256$, $d_{\text{EU}} = 1\,000$ km; $e_{\text{FR}} = 0.208$, $d_{\text{FR}} = 300$ km (all units kg CO_2/t·km).

3. With payload $m_t = m/1000$:

$$CO_{2,\text{transport}} = m_t \left(d_{\text{sea}} e_{\text{sea}} + d_{\text{EU}} e_{\text{EU}} + d_{\text{FR}} e_{\text{FR}}\right).$$