Simulated Welfare Effects of the EU's EV Transition

Ben Roble

October 2026

Abstract

The European Union aims to have nearly all new cars be electric by 2035 utilizing a series of progressively stricter automaker emission targets. Car manufacturers either can comply with the standard or pay a fine commensurate to their emissions deviation from the target. The scale of the EU's program is unprecedented, and until now EV transitions have been largely theoretical. This paper simulates the short-run effects of the future EU emission targets on total welfare and average car emissions. I first estimate automobile demand using 2023-2024 data on new car registrations and car list prices in France, Germany, and Italy. I then construct a Nash-Bertrand equilibrium model of profit maximization with respect to the EU emission standard. Automakers face heterogeneous compliance costs, but the model does not strictly bind them to the target. Firms are free to choose to over-pollute and pay the corresponding fine. Implementing the EU's future emission targets in 2024, I find the EU could reach the 2025 target with only a 1.18% decrease in welfare. The stricter targets beginning in 2030 are unachievable and result in significant drops in consumer surplus and profits. Manufacturers' ability to lower emissions by lowering EV prices is constrained by contemporaneous EV demand. Average car emissions in the EU reach a basin of 78 g/km with the 2030 and 2035 targets while EV shares plateau at 28%. This suggests the need for policymakers to establish programs to concurrently boost consumer EV demand.

JEL Codes: C60, L11, L98

1. Introduction

In the auto industry, governments have historically intervened to induce change on subjects deemed of national importance. The 1970s oil crises for example spurred the introduction of fuel economy standards in the Western world. The prototypical example of these interventions is the CAFE program in the United States, where car companies were subject to penalties on every new car sold if their average fuel economy fell below the standard. The justification frequently employed is that, but for government intervention, automakers would never make the necessary adjustments. In light of increasing concerns about climate change, the EU resolved to have net-zero greenhouse gas emissions by 2050. To this end, the EU established a program beginning in 2020 where car manufacturers are evaluated against a series of progressively tighter emission targets. The impetus to transition their car sales to electric vehicles is provided by a large fine on their EU-wide sales, made proportional to the degree of violation. The onus is on automakers alone to make the EV transition a reality.

My research question is: How would the European Union's implementation of different emission targets in 2024 impact aggregate EV shares, prices, and total welfare? Since electric car sales in the European Union have fallen well short of expectations, it is vital to conduct analysis of EV policies for the short-term. My paper aims to fill this gap by evaluating the welfare effects of future emission targets as if they were implemented in 2024. I first estimate car demand in the EU using a 2023-2024 dataset on national car registrations and car list prices. I then construct a equilibrium model that incorporates the EU emissions target program and use the observed prices and the estimated price derivatives to impute implied marginal costs. The model is used to analyze counterfactuals which separately implement the 2025, 2030, 2035, and 0 g/km emission targets in 2024 and a scenario in which the emissions target program is abolished. The simulated equilibrium outcomes are used to construct welfare estimates that I then analyze.

My paper follows several strands of literature. Two papers examine how auto manufacturers reacted to earlier EU mandates to reduce emissions. In relation to the 2009 EU program, Reynaert (2021) found that the emissions mandate reduced total welfare regardless of which method of compliance automakers used (selling more light-emissions vehicles, adopting green technology, or gaming the emissions test). Yilmaz (2021) studied the welfare impacts of the first target in 2020 under the current system. He incorporates the EU fine directly into the profit maximization problem with the caveat being that all pools are assumed to over-pollute. Similar to Reynaert (2021), he found the EU's emissions reduction program in 2020 to be

welfare-reducing with automakers shouldering more of the loss.

Other works study how firms adapted to CAFE fuel economy standards in the United States. Goldberg (1998) is an important paper which analyzed the effects of CAFE on prices and average fuel economy. Importantly, her equilibrium model accounted for the CAFE penalty differentially affecting firm compliance costs. Jacobsen (2013) studied the effects of a 1 mpg increase in the CAFE fuel economy standard and found total welfare would decrease, with producer surplus losses falling mostly on American car manufacturers. His methodology expanded on Goldberg's framework by allowing for the standard to be exactly binding for some automakers. In contrast to Jacobsen (2013), Anderson & Sallee (2011) found firms' marginal costs of compliance were much lower when the law's flexible-fuel provision is taken into account. Durrmeyer & Samano (2018) compared CAFE and the French feebate system on average fuel efficiency and welfare, and they found that the feebate system resulted in smaller welfare losses vis-a-vis CAFE.

My paper also intersects the literature on EV pricing and EV demand. Cruz et al. (2026) examined the EU market from 2014-2023 to determine factors that drive EV adoption; one of the most important adoption factors they found was EVs having a lower price relative to petroleum-fueled cars. The European Commission's 2025 Impact Assessment Report also identified higher relative EV prices as a strong determinant of uptake being slower than expected. Consumers appear to be myopic when setting expectations of future fuel prices (Grigolon et al., 2018; Sallee, 2011). In contrast, present-day fuel prices influence the decision to buy a fuel-efficient or electric vehicle (Alberini et al., 2019; Cincotta & Thomassen, 2024; Fei et al., 2025).

The contributions of my paper are twofold. Mine is the first contemporaneous study to examine how future EU emission targets would affect car prices, EV shares, and welfare in the short-run. Second, like in Jacobsen (2013) and Reynaert (2021), I allow for heterogeneity in compliance costs but expand their frameworks to cover both individual firms and emission pools. The latter is especially important considering that most automakers that sell in the EU belong to emission pools that the EU evaluates collectively.

I find that the EU could have achieved its 2025 emissions target of 93.6 g/km with only a 1.18% decline in total welfare from observed levels. For the 2030, 2035, and 0 g/km targets, the EU would fail to reach its goal in each scenario. With these 3 targets, aggregate EV shares plateau at 28.2% and EU average emissions hover around 79 g/km. These upper bounds reveal a structural weakness of the EU's approach: the assumption that EV demand will sufficiently evolve on its own. Given contemporaneous demand for EVs and PHEVs, automakers can only lower their emissions by so much through reducing the relative prices

of EVs/PHEVs to petrol cars. While it is not shocking that EV demand was too low to have successfully implemented the 2030 or 2035 targets in 2024, the EU has made no provisions to ensure EV demand will be high enough in 2030 or 2035. The welfare losses are mitigated by reductions in CO2 emissions, but both consumer and producer surplus would face significant declines (in the range of 7-15%). The abolishment of the emissions target program would result in a minuscule welfare increase of 0.049%, and average car emissions would be almost identical to those observed in 2024 (105.9 vs. 106.85). The latter is evidence that nearly all emission pools were truly unconstrained by their 2024 emission targets.

The structure of the paper is as follows. Section 2 lays out the details of the EU emissions target program and its history. Section 3 details my dataset's construction and its format. Section 4 describes the demand estimation techniques, price instruments, and estimation results. Section 5 derives the equilibrium model. Section 6 discusses the counterfactual policies and the welfare results Section 7 concludes with a discussion of my findings, some limitations of this study, and future work.

2. EU Policy

The EU's emissions target program started in 2020 with the enactment of Regulation (EU) 2019/631. In line with the Paris Agreement and the European Green Deal, the EU set to substantially reduce automobile emissions by the 2030s. The European Union's approach to this decarbonization is to incentivize automakers to primarily sell electric vehicles in the future via an emissions target program. The average emissions of the new cars sold by a car manufacturer is evaluated against the EU's specific emissions target that year[1]. Car manufacturers whose average emissions are below or at the target are left undisturbed. Conversely, manufacturers which have average emissions greater than their targets are fined by the EU. These manufacturers pay a lump-sum equal to 95 Euros times the positive difference between its average emissions and its target multiplied by the number of new cars sold in EU territory. This penalty structure applies to the majority of auto manufacturers, but there are a few manufacturers that are granted derogations or altogether exempted from the emissions target on the basis of their very low sales figures[2].

To alleviate some automakers' concerns that they would be unable to individually meet their targets, a provision in Regulation (EU) 2019/631 allows for the formation of emission pools. Emission pools are

[1]Emissions are measured in grams of CO2 per km. This statistic is generated during a car model's WLTP test, and it is published by the manufacturer.

[2]I exclude these firms from my analysis on account of their tiny market shares and the difficulty of applying the byzantine derogation rules in my counterfactuals

voluntary associations of automakers for whom the EU evaluates their combined average emissions. Like with individual firms, emission pools are left undisturbed if their average emissions lie below the EU target and are fined otherwise. The formula for a pool's emissions penalty is analogous to the single-firm version with all firm-specific terms being replaced by their pool equivalents. The EU requires this fine to be paid in one lump-sum by the designated pool manager. How the pool collects the requisite funds from its members is left to the discretion of the pool. There are two types of emission pools: open and closed. Open pools are defined by the lack of a common owner across members. Closed pools are only comprised of subsidiaries all belonging to the same corporate owner. The European Union imposes information sharing restrictions on open pools to prevent anti-competitive collaboration by competing firms. Members can only share the following information with each other: a) their average specific emissions of CO_2, b) their specific emissions target, and c) the total number of vehicles that they registered that year. This means that, unlike automakers in closed pools, car manufacturers in open pools cannot coordinate on price or sales strategy to ensure group compliance. The predominant internal mechanism for open pools is a transfer payment scheme. The details of pools' payment schemes and the magnitude of intra-pool payments are not publicly observable, but their widespread existence has been documented by several news sources (Automotive World, 2026; Campbell & McGee, 2019; Dahil, 2025).

In practice, a firm/pool's specific emissions target is not known in advance. The specific emissions target is equal to the sum of the EU fleet-wide target and a secondary term that accounts for a firm's weight mix. The EU fleet-wide target is a threshold that every manufacturer must meet and is something adjusted every 5 years. If the average mass of the new cars sold by a firm or pool exceeds the established standard that year, its emission target is tightened. If the converse occurs, the EU sets a lower emissions target for that firm. These deviations are made proportional by a coefficient which is altered every 5 years. The exact formulations can be found in Table 1 in the Appendix.

European policymakers recognized the impracticality of an immediate transition to electric vehicles, and so they set up a future schedule of declining emissions targets. The emissions reduction goals set in the 2019 version of the regulation were relatively modest in comparison to today. At that time, the EU was not planning on a large-scale EV transition. The EU aimed to have firms pollute at least 15% less than the EU's 2021 CO_2 target between 2025-2029 and 37.5% less than the 2021 CO_2 target beginning in 2030. The 2021 target itself was nebulous in character back then, as it was to be based off of the average emissions of new cars sold in 2020. Nevertheless, the schedule of yearly emissions targets was designed to encourage firms to switch more and more of their sales to electric vehicles.

On May 15, 2023, the EU announced a paradigm shift in the program's goal. It was announced that, starting in 2035, car manufacturers were required to have their new cars' average emissions be 0 g/km of CO_2. In other words, the EU mandated that, post-2034, all new cars had to be electric. Existing combustion-powered cars would be allowed to remain on the road and their secondary sales would be unimpeded. Car buyers would thus not be forced to substitute to electric vehicles in the short-run. But since all automobiles eventually become inoperable, the long-run consequence of this policy is a complete transition to electric vehicles.

Although every emissions pool met their targets in 2023 and 2024, there were concerns raised by many firms when 2025 started. The 2020-2024 EU fleet-wide target was 115.1 g/km, and the 2025-2029 target was around 22 points lower at 93.6 g/km. Even though they had already known about the 2025-2029 fleet-wide target since 2023 (Commission Implementing Decision (EU) 2023/1623, Annex II), European car makers argued that they could not meet the target that year. In July 2025, policymakers relented and a one time exception was made wherein the EU would evaluate every firm/pool's performance over the period 2025-2027. Yearly emission targets will be implemented once again starting in 2028. This was intended as a short-term relief, but the underlying issues led to a comprehensive reform later that year.

On December 16, 2025, the EU announced the "Automotive Package". It is a tentative proposal by the EU Commission which alters several aspects of the existing program. The impetus behind this proposal was that contemporaneous EV demand had been well short of projections. EVs only had a 16.5% aggregate market share in the EU in the first ten months of 2025; the car manufacturer association ACEA estimated that the aggregate market share had to be at least 25% for automakers to meet their targets (European Commission, 2025a, p. 9). Several prominent manufacturers like BMW and Volkswagen also published profit warnings in 2025, citing sluggish growth in EV sales. The "Automotive Package" eliminates the 2035 EV mandate and replaces it with a new 2035 target of 11.5 g/km of CO_2: a 90% reduction in emissions relative to 2021. The new target still incentivizes firms to transition most of their EU sales to EVs by 2035, but it grants more flexibility by allowing for plug-in electric vehicle sales. PHEVs and range extender electric cars (together denoted OVC-HEVs) will be treated as emitting 0 g/km so long as they can travel 150 km or more on just battery power. However, only a limited number of OVC-HEVs are granted this special treatment. If OVC-HEVs make up 10% or less of an automaker's sales, every OVC-HEV is considered as being CO_2-free. However, if OVC-HEVs make up more than 10% of an automaker's sales, the EU assigns 0 g/km to the OVC-HEVs that comprise the first 10% of units sold. Further compliance flexibility will be

granted through the extension of ZEV credits[3] and a carbon credit mechanism[4].

Despite this, contemporary reports indicate emission pools are well on their way to meet the 2025-2027 emissions targets (Jones, 2026; The International Council on Clean Transportation, 2026). Their specific emissions targets are still unknown as they are set *ex post*, but an article by the ICCT[5] lists each pool's likely target being very close to the 93.6 g/km 2025-2029 EU fleet-wide target (Monteforte & Diaz de Aguilar, 2026). Average car emissions in the European Economic Area dropped by 10 g/km between 2024 and 2025; the aggregate EV share increased from 14.2% to 18.9% in the same period (European Environmental Agency, 2026).

3. Data

2023-2024 EU car sales data comes from the "Monitoring of CO2 emissions from passenger cars Regulation (EU) 2019/631" database published by the European Environmental Agency (EEA). All new car registrations must be annually reported by member states to the EEA with the following information: manufacturer name, TVV (Type-Variant-Version) code, model name, mass in running order, wheelbase, total new registrations, fuel type, fuel mode, engine power (kW), engine capacity, CO_2 emissions in g/km (both under WLTP and NEDC), fuel consumption (liters per 100 km), electricity consumption (kWh per 100 km), electric range (km), and eco-innovations[6]. Observations are at the car registration level. A car's national sales for a given year are computed by counting the rows which it appears in that country-year (as each row contains exactly 1 registration).

The above dataset does not provide the transaction price or list price of newly registered cars. I use other sources to obtain car list prices[7] of cars. For Italy and France, I used the Quattroruote and AutoPlus car magazines respectively. I obtained German price data by purchasing a web-scraped dataset on every car listing on ADAC[8] from a private data mining firm. The ADAC dataset also included a rich collection of car characteristics and amenities by submodel. I assume that the list price, adjusted for national taxes and subsidies, is what consumers pay. I add France's CO_2 and weight registration tax, known as malus, to the

[3]Zero-emission vehicle
[4]Cars that run on sustainable renewable fuels (SRFs) will be counted as 0 g/km. The emission targets will be positively adjusted by the proportion of these car sales, up to a limit.
[5]The International Council on Clean Transportation
[6]These are defined in Regulation (EU) 2019/631.
[7]Transaction prices are ideal, but this type of data is not made publicly available
[8]Allgemeiner Deutscher Automobileclub

prices of French cars. EV and PHEV purchase subsidy programs were in existence in Germany, Italy, and France in 2023 and France and Italy in 2024. I calculated the relevant subsidy for each eligible EV or PHEV and subtracted that from their observed list price.

1:1 matches between car submodels found in the sales data and cars found in the pricing data were not possible in most cases. The EEA dataset does not provide the car's submodel name and the sparse set of car characteristics provided were not enough to differentiate between the submodels of that car. These submodels can share the same observables but differ on key unobservables like drivetrain, geartrain, number of seats, body type, et cetera. A similar issue occurred with the pricing dataset where AutoPlus and Quattroruote did not provide the full set of specifications like ADAC did. To get around this problem, I create the "median" version of each car model following Grieco et al. (2024). In both the sales and pricing datasets, I define a car model to be the unique combination of make, model name, fuel type, fuel mode[9], and body type. I then take the median, across a car's submodels, of price, weight, engine power, CO_2 emissions, and other car characteristics. Variation in price and car characteristics is maintained because a model's "median" version will look differently on a per-country basis. Manufacturers sell different portfolios of cars in different markets, and so a model's median price, weight, etc will vary accordingly. There are, however, several cars whose submodels vary in body type (e.g., Toyota Corolla has a station wagon, hatchback, and SUV version). I assume these models' "median" body type is the body type with the highest number of associated submodels.

Monthly national fuel prices are obtained from the World Bank's Global Fuel Prices Database. Each price is in the form of Euros per liter of fuel. I compute, by country, average yearly fuel prices for petrol and diesel. For the cost of electricity, I collected from Eurostat bi-annual, national household electricity prices for households who use between 2500 and 4999 kWh a year[10]. These are incorporated with fuel and electricity consumption statistics to create the energy cost variable "Cost of energy consumption per 10 km". This is a unified measure that applies to all cars: mono-fueled, hybrid, and electric. The fuel consumption measure (liters per 10 km) is converted to a fuel consumption cost measure (euros per 10 km) by multiplying it by the fuel price (euros/liter). Electric vehicles purely run on electricity, so an analogous process occurs to compute electricity consumption cost. I take the electricity consumption measure (kWH / 10 km) and multiply it by the average EU household electricity price (euros / kWh). Plug-in electric vehicles (PHEVs) are capable of running on two modes: engine mode and battery mode. I calculate the number of kilometers each PHEV

[9]Fuel mode indicates how an engine processes fuel in generating energy. Examples included mono-fuel (M), a mild or full hybrid EV (H), purely electric (E), plug-in hybrid (P), etc.

[10]Eurostat defines the average household as falling into this range: https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Electricity_price_statistics

can travel on engine mode and then the number of kilometers that can be driven on battery mode. Assuming that a PHEV's total range is equal to the sum of these two numbers, I calculate the proportion of total range that a PHEV can travel solely using electricity r_e. I weigh the electricity consumption cost measure by r_e and the fuel consumption cost measure by $(1 - r_e)$ and then add them together. Thus, a single variable is created that measures the energy cost of driving.

I obtain data on national GDP per capita, Gini coefficents of income, and the number of households in each country from Eurostat. GDP per capita serves as a proxy for consumer income in the logit specification. The Gini coefficients are used to back out each country's standard deviation of income when assuming a lognormal income distribution. I define the market size to be the number of households in a country following previous literature (BLP, 1995; Nevo, 2000). I also construct lagged real exchange rates (RXR) between a car's country of production and the country of sale. The country of manufacture for each car model is identified from car production data provided by MarkLines. Currencies' nominal exchange rates to the Euro are collected from Eurostat, and PPP exchange rates are collected from the World Bank's World Development Indicators dataset.

The EU publishes yearly performance metrics of automakers in relation to that year's targets. These performance metrics are generated for individual manufacturers as well as for pools in aggregate. The information includes: manufacturer/pool name, indicator for a derogation or exemption, total number of registrations, average mass in running order (kg), average specific emissions of CO_2 (g/km), specific emissions target (g/km), distance to target, CO_2 savings from eco-innovations and error margin. Data for 2020-2023 has been published, but the official EU reports for the years 2024 and 2025 have not been released yet. The pooling structure in 2024 and 2025 was obtained from a report in the EU's CIRABC database.

3.1. Summary Statistics

Table 4 reports the summary statistics. Inclusive of national VAT tax and any relevant subsidies, the mean car price is 57,581 Euros. It is dragged upwards by the presence of high-end luxury cars in the dataset. The median car price of 42,655 Euros is more reflective of typical car prices in France, Germany, and Italy. In terms of GDP per capita, the average new car has a price 34% above it. The cars in the bottom 50th percentile are priced under than GDP per capita and the cars in the top 50th percentiles are priced over. This motivates my later decision to use the price/income approximation when estimating the

price coefficient (Van Biesebroeck & Verboven, 2026, p. 991). The average emissions of new cars in the EU is 107.74 g/km. This is below the EU's 2024 standard of 115.1 g/km but above any of the EU's future emission targets. The vast majority of car models have internal combustion engines; only 17% of car-market observations concern electric vehicles. SUVs are by far the most common body type[11] at 55%, followed by hatchbacks with 20% and station wagons in third place at 11%. The mean number of doors is 4.81 and has a small standard deviation of 0.6. Taken together, they imply most new cars have between 4 and 5 doors with the latter being significantly more common.

4. Demand Estimation

4.1. Multinomial Logit

I use a representative utility function to model logit discrete-choice demand that takes the form:

$$u_{ijn} = x_j\beta - \alpha\frac{P_{jn}}{y_n} + \xi_j + \epsilon_{ijn} \tag{1}$$

The indices i, j, n denote individual, car, and market respectively. I define a market to be a country-year (e.g., Germany 2024). x_j are observable car characteristics, P_{jn} is the consumer price[12], y_n is GDP per capita (proxy for consumer income), ξ_j denotes utility from unobserved market demand characteristics, and ϵ_{ijn} is the utility from unobserved individual demand characteristics. ϵ_{ijn} is assumed to be uncorrelated with x, independent and identically distributed among consumers, and distributed according to a Type-1 Extreme Value Distribution. Utility u_{ijn} can be written as the sum of mean utility δ_{jn} and an individual level term ϵ_{ijn}. The share of car j in market n (s_{jn}) and that of the outside option (s_{0n}) can be expressed as functions of mean utility:

$$s_{jn}(\delta_n) = \frac{exp(\delta_{jn})}{1 + \sum_{k=1}^{J} exp(\delta_{kn})} \tag{2}$$

[11] It should be noted that the definition of an SUV encompasses both traditional SUVs and crossovers, the latter being much more numerous.

[12] Consumer price is the car price inclusive of taxes and subsidies. I will later introduce the pre-tax price p_{jn} in my equilibrium model, and so I follow my model's notation in denoting the price the consumer receives as P_{jn}.

$$s_{0n}(\delta_n) = \frac{1}{1 + \sum_{k=1}^{J} exp(\delta_{kn})} \tag{3}$$

Mean utility can be expressed as a function of market shares through an inversion process (Berry, 1994). Equation (4) is the classic multinomial logit form of mean utility:

$$\delta_{jn} = ln(s_{jn}) - ln(s_{0n}) = x_j\beta - \alpha\frac{P_{jn}}{y_n} + \xi_j \tag{4}$$

Mean utility of car j in market n is the log difference between its market share and the market share of the outside good. An IV regression of mean utility on product characteristics x_j and instruments for price P_{jn} produces estimates for taste parameters β and the marginal disutility of price relative to income α. α is used to construct the own-price derivative $\frac{\partial s_{jn}}{\partial P_{jn}} = -\alpha s_{jn}(1 - s_{jn})$ and the cross-price derivative $\frac{\partial s_{jn}}{\partial P_{kn}} = \alpha s_{jn} s_{kn}$.

The earlier i.i.d. assumption on ϵ_{ijn} and the lack of any consumer heterogeneity in (α, β) implies that δ is the only factor which distinguishes one car from another (Berry, 1994). Two cars j and l which share the same mean utility have identical market shares. This means that the market shares for j and l react in the same way to a change in price of any other car k. This holds even if one car happens to be a better substitute for k (e.g., an SUV if car k is also an SUV). The unrealistic substitution patterns become even more apparent when looking at the elasticity formulae:

Own-price elasticity

$$\frac{\partial s_{jn}}{\partial P_{jn}} * \frac{P_{jn}}{s_{jn}} = -\alpha P_{jn}(1 - s_{jn}) \tag{5}$$

Cross-price elasticity

$$\frac{\partial s_{jn}}{\partial P_{kn}} * \frac{P_{kn}}{s_{jn}} = \alpha P_{kn} s_{kn} \tag{6}$$

With respect to a given car, every other car share identical cross-price elasticities. The ease of estimation and elasticity construction is largely offset though by a rigid structure of product substitution. This feature motivates my decision to estimate a random coefficients model of demand which allows for consumer heterogeneity in income and tastes. The logit estimates provide a benchmark to compare these new estimates.

4.2. Random Coefficients

The utility function in a random coefficients demand model takes a different functional form than in multinomial logit. The utility equation is listed below:

$$u_{ijn} = x_j\beta - \alpha\frac{P_{jn}}{y_i} + \sum_k \sigma_k \nu_{ik} x_k + \xi_j + \epsilon_{ijn} \tag{7}$$

Where:

i = Individual

j = Car

k = Characteristic

n = Nation

x_j = Observable vector of characteristics for car j

P_{jn} = Price of car j in nation n

y_i = Income of individual i, drawn from a lognormal distribution

σ_k = Standard deviation of preferences for attribute k

ν_{ik} = Random taste shock for consumer i for car characteristic k

ξ_j = Utility of unobserved product characteristics for car j

The indices i, j, and n denote individual, car, and market respectively. A market is defined to be a country-year, and there are 6 markets in my dataset: France in 2023, France in 2024, Germany in 2023, Germany in 2024, Italy in 2023, and Italy in 2024. I assume that a country's income distribution in any year is lognormal with a mean equal to its GDP per capita that year and a standard deviation derived from its Gini coefficient. In the estimation procedure, I populate each market with 200 representative individuals. Individuals' incomes y_i are drawn randomly from their market's lognormal income distribution. Individuals' random taste shocks for car characteristic k, ν_{ik}, are drawn from a standard normal distribution. Both y_i and ν_{ik} are then held fixed during demand estimation.

An individual's random coefficient for characteristic k is $\beta_{ik} = \beta_k + \sigma_k \nu_{ik}$. The first term is the mean taste for k across consumers, and the second term is how individual i's preference differs from the mean. Specifically, his heterogeneity is modelled as being the standard deviation in taste for k across consumers weighted by his random standard normal shock for k. I estimate random coefficients for all measurable car

characteristics: engine power, mass, car base (length times width), number of doors, number of seats, and the energy consumption cost per 10 kilometers driven. The marginal disutility of price relative to income, α, is not a random coefficient as it is identical across consumers. However, individuals effectively differ in price sensitivity since they have different incomes. The random draws of consumer income are what require α to be estimated alongside σ, even though α itself does not vary by individual.

Market shares are calculated through the following process. The computer chooses a vector $[\alpha, \sigma]$ as a guess to satisfy the GMM equation. The mean utilities for each product j are initialized to be equal to their logit forms: $\delta_{jn} = ln(s_{jn}) - ln(s_{0n})$. For each consumer i in market n, their utility for car j, u_{ijn}, is calculated by summing the mean utility δ_{jn} and their individual deviation from that mean $\mu_{ijn} = -\alpha \frac{p_{jn}}{y_i} + \sum_k \sigma_k \nu_{ik} x_k$. The simulated market share for car j in country n, $\tilde{s_{jn}}$, is created by taking the average of the individual choice probabilities ($P_{ij} = \frac{e^{\delta_{jn} + \mu_{ijn}}}{1 + \sum_{j=1}^{J_n} e^{\delta_{jn} + \mu_{ijn}}}$) across all consumers. The simulated market shares are then used to construct a new guess for mean utilities that correspond to observed market shares. This is done according to the formula:

$$\delta_j^{t+1} = \delta_j^{(t)} + ln(\tilde{s}_{jn}) - ln(s_{jn}(\delta, \sigma)) \tag{8}$$

Where

t = Iteration index

If the difference between the new market share iteration and the old market share iteration is greater than the tolerance band 0.00000001, then the consumer utilities and simulated market shares are recomputed. The end result is one where the computed mean utilities correspond to observed market shares. The "squarem" Pyblp method is used to accelerate this contraction mapping.

The next step is to back out the unobserved[13] product utility ξ_{jn}. ξ is ultimately a function of (α, σ) and is assumed to be mean-independent of Z: $E(\xi|Z) = 0$. It is calculated by subtracting $x_j\beta$ from the product's mean utility δ_{jn}. To compute it, β must first be estimated:

$$\hat{\beta} = (X^T Z W Z^T X)^{-1} (X^T Z W Z^T \xi) \tag{9}$$

[13]To the econometrician

Where:

X = Matrix of car characteristics

Z = Matrix of demand instruments

W = Weighting matrix[14]

With the weighting matrix $W1$, demand instrument matrix Z, and the unobserved product utility vector ξ vector, the GMM objective function $g(\theta) = \tilde{\xi}(\theta)' Z W Z' \tilde{\xi}(\theta)'$ is evaluated. The value produced is unique to the computer's guessed parameter vector $\theta = (\alpha, \sigma)$, and the goal is to minimize $g(\theta)$ with respect to θ. Each guess of the parameter vector involves running a contraction mapping for δ_{jn}, obtaining values for ξ_{jn}, and plugging in ξ into $g(\theta)$. Once the computer finds the θ vector that minimizes the GMM function, the weighting matrix W2 is constructed from the ξ vector associated with θ^*. Optimization of the GMM function is done with the L-BFGS-B algorithm in Pyblp.

The second step of GMM is identical to the first step of GMM except W2 is used as the weighting matrix and the mean utilities from GMM-1 are used as the initial guess for δ_{jn}. Once $g(\theta)$ is once again minimized with respect to $\theta = (\alpha, \sigma)$, the coefficient estimates for β, α, and σ are obtained with corresponding heteroskedastic-robust standard errors.

4.3. Instruments

Price is an endogenous variable when estimating demand; thus, instruments for it are required. My primary instruments for price are the lagged real exchange rate (RXR) between the country of production and the country of sale, BLP instruments[15], and the car's own characteristics x_{jn}. The lagged real exchange rate instrument comes from Remmy (2026) and Grieco et al. (2024), and it instruments for price through two channels. The real exchange rate equals the PPP exchange rate between the country of production m and the country of sale n divided by the nominal exchange rate between the two currencies:

$$RXR_{m,n} = \frac{PPP\ Exchange\ Rate_{m,n}}{Nominal\ Exchange\ Rate_{m,n}} \tag{10}$$

[14] In the first step of GMM, the 2SLS weighting matrix $W1 = (Z^T Z)^{-1} = (\frac{1}{N} \sum_{jn} Z_{jn}^T Z_{jn})^{-1}$ is used. For the second step of GMM, the optimal weighting matrix $W2 = (\frac{1}{N} \sum_{jn} g_{jn} g_{jn}^T)^{-1} = (\frac{1}{N} \sum_{jn} Z_{jn}^T \xi_{jn} \xi_{jn}^T Z_{jn})^{-1}$ is used.

[15] From Berry, Levinsohn, & Pakes (1993)

Even if the car was made in the Eurozone, the RXR varies through the PPP exchange rate component (unless $m = n$). It is reflective of local labor costs shifts in a car's country of production as that directly affects the PPP exchange rate and through it the RXR. Secondly, if the country in which a car is made has its currency depreciate relative to the Euro, automakers will have an incentive to lower prices in the Eurozone (Grieco et al., 2024, p. 1208). It is lagged by one year to account for the fact that firms plan ahead using the most recent data obtainable. I modify the RXR instrument to account for the fact that some car models are exported to Europe from multiple countries. My MarkLines data allows me to see where every car model was manufactured and how many units were produced. I also obtained information from auto manufacturers' websites about which factories supply the car models sold in France, Germany, and Italy although the exact inflows are unknown. Since there are many new cars sold in Europe that were sourced from several countries, a unique RXR for them cannot be determined. In these instances, I create a simple average of all possible real exchange rates. For example, if units of car A sold in Germany are produced in either country 1 or country 2, I take a simple average of the RXR between country 1 and Germany & the RXR between country 2 and Germany.

BLP instruments which are commonly used in automobile demand estimation serve as my other main price instruments. For each product characteristic other than the electric car dummy, I construct the own-firm and rival-firm BLP instrument. The own-firm BLP instrument sums up the product characteristic across all other cars made by the same manufacturer; the rival-firm BLP instrument does the same thing but across all cars made by rival manufacturers. The last instruments for price are the product characteristics x_j themselves.

4.4. Estimates

The logit and random coefficient results are reported in Table 4. Column 1 contains the estimates for the logit specification, and column 2 contains the estimates for the random coefficients specification. The marginal disutility of price relative to income, α, jumps from -0.993 in the logit specification to -3.257 in the random coefficients specification. The logit price coefficient leads to a median own-price elasticity of around -0.99 and roughly 51% of cars having inelastic demand. No Nash-Bertrand equilibrium exists where firms set car prices such that their demands are inelastic. It is always more profitable to raise price when demand is inelastic: the increase in marginal revenue offsets any decrease in sales. The unrealistic elasticity estimates are clear evidence that the restrictive substitution patterns under logit do not accurately model car demand.

Conversely, the estimate for α (-3.257) from the random coefficients specification is within reason, albeit it resides on the lower end of the spectrum in the European auto market literature[16]. The median own-price elasticity is -3.03, and there are no cars with inelastic demands. I use this estimate for α to construct the own-price and cross-price derivatives utilized in my equilibrium model.

The logit estimates of the mean coefficients β are noisy but somewhat informative. Engine power, mass, and the base of a car (length times width) all come out insignificant. Like with α, this is reflective of the limitations of logit demand in my setting. Consumers value cars with fewer doors (correlated with small cars), and they incur disutility from electric cars and from higher energy consumption costs.

The estimates of β in the random coefficients model are, with one exception, nailed down with greater precision. The coefficient on engine power changes from the insignificant 1.126 estimate under logit demand to the significant -1.317 estimate under random coefficients demand. This is counterintuitive because it is well-established that consumers value engine power. The β estimate in the logit specification reflects this although it is insignificant. However, the β estimate for engine power in the random coefficients specification is negative and significant. I interpret this to be an anomalous estimate rather than a novel discovery. The random coefficients model provides much better estimates in general, but it is not surprising that there is an outlier given the small size of the dataset and the reduced variation from cars being median "composites"[17]. European drivers on average prefer heavier, more compact (i.e., a lower surface area) cars. The mean coefficient on the number of doors increased significantly in magnitude. I believe this to be the result of a high correlation with car size instead of a deep dislike for more doors. The mean taste for the number of seats is not statistically different from zero (neither was the logit estimate at a 95% significance level). I conjecture that the combination of body type fixed effects and the inclusion of car doors in the estimation caused its insignificance. The disutility from electric cars is even stronger than in the logit specification. The negative β for electric cars comes out as expected; the reverse sign would imply higher EV shares than were observed in 2023 and 2024. Drivers have even stronger mean preferences towards cars with lower energy costs.

The σ estimates are listed below the mean coefficients in column 2. The σ estimate for engine power is statistically insignificant despite numerous car papers finding consumers have heterogeneous tastes regarding engine power. This lends support to the hypothesis that the estimated mean coefficient for engine

[16]EU: -6.709 (Reynaert, 2021), -5.33 (Grigolon et al., 2018). Norway: -3.6835 (Springel, 2021). Germany: -7.112 (Remmy, 2026)

[17]As mentioned earlier, cars were able to be matched between the pricing and sales data by creating their median "trims".

power is an anomaly. The standard deviation for car mass collapsed to zero during the two-step GMM estimation. This might have been caused in part by the inclusion of body type fixed effects as they already account for general differences in weight between car types (e.g., SUVs weigh more than sedans). I would argue that the null estimate for σ_{mass} is more reflective of a limitation of using random coefficients in my small dataset than a novel discovery that preference heterogeneity for car weight does not exist. The large, negative β estimate for number of doors (-8.001) is mollified by the σ estimate for the same characteristic. The standard deviation estimate is slightly greater in magnitude (8.487) and has the opposite sign. This means that if any consumer has a positive shock from the standard normal distribution, his random coefficient will be skewed towards zero. I.e., there will be consumers that are either relatively indifferent about the number of doors or have strong dislikes towards cars with more doors (which has a strong correlation to the car being large). Heterogeneity in car size preference is also captured by the significant σ for a car's base (surface area). The estimated preference heterogeneity for number of seats is logically consistent: those who prefer more doors also prefer more seats and vice versa. Consumers have heterogeneous valuations of the energy cost of driving as evidenced by the σ estimate.

5. Equilibrium Model

I create a Nash-Bertrand model of profit maximization under the EU emissions target program. Individual firms and emission pools decide whether or not to meet their emission target according to how costly compliance is relative to the €95 fine per excess gram of CO_2 assessed on all their cars. Some firms easily meet their EU emission targets and are not constrained by them. Even if emissions targets did not exist, these automakers would still have average emissions below them. I refer to such firms as being unconstrained. A good example is Tesla as it only sells EVs. Other firms incur costs in complying with their target, but they still find compliance less costly than paying an emissions violation fine. I refer to these firms as being constrained. Still other firms find it more profitable to emit above their targets and pay the associated fine instead of complying. I refer to these companies as "fine-optimal".

Each firm has a corresponding group index g which indicates emissions pool membership. Firms in a given emissions pool all share the same group index; firms that do not belong to an emissions pool instead receive a unique group index. For example, say there are 3 firms indexed by f = 1, 2, 3. Firms 1 and 2 are in the same emissions pool whilst firm 3 chooses to individually meet the EU's emissions target. The group index $g = \{1, 2\}$ is created where the two {f, g} pairings are $\{(1, 2), 1\}$ and $\{3, 2\}$. It is important

to note that a firm f is defined to be the set of car manufacturers that operate under one common parent company. Automakers under common ownership are assumed to coordinate pricing internally even though these subdivisions are marketed as being distinct (e.g., Audi and Volkswagen). I classify any firm that is not in an emissions pool in its own eponymous, closed emissions pool. Note that the second series of summations disappears for closed emission pools. I denote the group g that firm f belongs to as $g(f)$ going forward.

All auto manufacturers belonging to an emission pool share the same constraint status since the decision to meet or violate the EU emissions target is made at the pool level. This heterogeneity in compliance costs has important implications when solving for equilibrium prices and market shares. Every carmaker or emissions pool cannot be assumed to shoot for the EU emissions limit when profit maximizing. I allow for heterogeneity in compliance by solving each firm's profit maximization problem subject to its emissions pool's mixed complementarity condition. Equation (11) is the firm's profit function, equation (12) denotes the pool's mixed complementarity condition, and equation (13) details the consumer price that market shares are a function of.

$$\pi_f = \sum_{n=1}^{N} \sum_{j \in \mathcal{F}_{fn}} (p_{jn} - c_j - \lambda_{g(f)}(co2_j - T_{g(f)}))s_{jn}(P_n)M_n \tag{11}$$

$$0 \leq \lambda_g \leq 95 \quad \perp \quad -\sum_{n=1}^{N} \sum_{j \in G_{gn}} (co2_j - T_g)s_{jn}(P_n)M_n \tag{12}$$

$$P_{jn} = \tau_n p_{jn} + m_{jn} - b_{jn} - \rho_{jn} \tag{13}$$

Where:

f = Firm

g = Group (emissions pool)

j = Car

n = Country

$g(f)$ = Index of the group (pool) that firm f belongs to

$\mathcal{F}_{fn}$ = Set of cars that firm f sells in country n

$G_{g,n}$ = Set of cars that pool g sells in country n

VAT_n = VAT rate

$\tau_n = 1 + VAT_n$

p_{jn} = Pre-tax price

c_j = (Emissions-unadjusted) marginal cost

λ_g = Shadow cost of emissions

$co2_j$ = CO_2 emissions in g/km

T_g = EU emissions pool target

s_{jn} = Market share

M_n = Market size

$a(j)$ = Fuel mode of car j

$$m_{jn} = \begin{cases} co2_malus_{jn} + weight_malus_{jn} & \text{if } co2_malus_j + weight_malus_j \leq 70000 \text{ and } n = \text{France} \\ 70000 & \text{if } co2_malus_j + weight_malus_j > 70000 \text{ and } n = \text{France} \\ 0 & \text{otherwise} \end{cases}$$

$$b_{jn} = \begin{cases} 0.27\tau_n p_{jn} + 1431.33 & \text{if } n = \text{France } \& \ mass_j < 2400 \ \& \ co2_j = 0 \ \& \ \tau_n p_{jn} \leq 18518.52 \\ 6431.33 & \text{if } n = \text{France } \& \ mass_j < 2400 \ \& \ co2_j = 0 \ \& \ (18518.52 < \tau_n p_{jn} \leq 47000) \\ 0 & \text{otherwise} \end{cases}$$

$$\rho_{jn} = \begin{cases} 6484.95 & \text{if } n = \text{Italy } \& \ a(j) = \{\text{EV, PHEV}\} \ \& \ co2_j \leq 20 \ \& \ p_{jn} \leq 35000 \\ 4323.30 & \text{if } n = \text{Italy } \& \ a(j) = \text{PHEV } \& \ (20 < co2_j \leq 60) \ \& \ p_{jn} \leq 45000 \\ 0 & \text{otherwise} \end{cases}$$

In equation (11), each firm f sets pre-tax prices p_{jn} to maximize profit given its cars' emissions-adjusted marginal costs[18] $c_j + \lambda_{g(f)}(co2_j - T_{g(f)})$. Firms take their pool's shadow cost of compliance $\lambda_{g(f)}$ as given when setting pre-tax prices. It is important to note that market share is a function of more than just pre-tax prices. The consumer faces a final price (P_{jn}) that is the sum of four elements: the VAT-inclusive price $\tau_n p_{jn}$, the malus registration tax[19] m_{jn}, the "écologique bonus" b_{jn} for eligible EVs sold in France, and a purchase subsidy for eligible EV/PHEVs in Italy ρ_{jn}. The two subsidies, b_{jn} and ρ_{jn}, are price-dependent; I will discuss in the next subsection how I tackle these when solving for optimal pre-tax prices p^*.

The emissions adjustment to marginal cost, $\lambda_{g(f)}(co2_j - T_{g(f)})$, reflects the cost of emission target

[18] Hereafter denoted $\tilde{c_j}$

[19] Only applicable for French cars

compliance associated with car j. Each firm f pays it on each unit of car j sold. Sales of cars which emit less than the target reduce its average emissions, and sales of cars that emit more than the target raise them. The scale of a car's deviation from the pool target, $co2_j - T_{g(f)}$, determines how costly selling that car is in terms of emission target compliance. The net emissions deviation for a firm and for a pool are given in equations (14) and (15) respectively:

$$E_f(P) = \sum_{n=1}^{N} \sum_{j \in \mathcal{F}_{fn}} (co2_j - T_{g(f)}) s_{jn}(P_n) M_n \tag{14}$$

$$E_g(P) = \sum_{n=1}^{N} \sum_{j \in G_{gn}} (co2_j - T_g) s_{jn}(P_n) M_n \tag{15}$$

Firm f's net payment is equal to $\lambda_{g(f)} * E_f(P)$ and pool g's net payment is equal to $\lambda_g * E_g(P)$. This emissions-based component of marginal cost only occurs for firms belonging to pools with positive shadow costs.

Intra-pool transfer payments only occur in open emission pools with $\lambda_g > 0$. These transfer payments work in the following manner. A pool member f that has a positive net emissions deviation ($E_f(P) > 0$) sends the pool manager a transfer payment equal to $\lambda_g * E_f(P)$. From this fund, the pool manager disperses transfer payments to pool members with negative net emissions deviations[20]. Members that exactly meet the pool target receive nothing as they neither raised nor lowered the pool's average emissions. This mechanism rewards the members who lowered pool average emissions and punishes the members who raised them. Intra-pool payments net to zero when the pool meets its target ($E_g(P) = 0$), and they net to the EU fine when the pool exceeds its target ($E_g(P) > 0$). I illustrate with numerical examples how the pooling payment scheme works in Table 2 of the Appendix.

In the real world, emission pools do not know their exact specific emission target *ex ante*. However, in my model, pools know their pool-specific target T_g in advance. I made this modeling decision for the following reason. Specific emission targets are not known in advance because pools cannot perfectly predict their sales and average weight of the cars they sell. If they could, they would know their exact target. I would need to make car demand itself an uncertainty for firms to make their contemporaneous emission targets unknown. The first-order conditions of the model require knowledge of $\frac{\partial s_{jn}(P_n)}{\partial P_{jn}}$, and demand estimates

[20]These transfer payments follow the same formula: $\lambda_g * E_f(P)$. The difference here is that $E_f(P)$ is negative.

are the only source for these partial derivatives. They have to be held constant in order for the vector of first-order conditions to be solved. Therefore, the incorporation of *ex post* EU emissions targets is not possible. The pool-specific targets T_g are set in the following way. When calibrating the model with 2024 data, I use the specific emissions targets published by European automotive organizations (Díaz et al., 2025). The published 2024 pool targets are better suited to calibrating the model because it prevents pools from being misclassified. For example, Volkswagen has average emissions of 121 g/km in my data and its specific emissions target was 122 g/km. If I had instead used 115.1 g/km as Volkswagen's target, then it would be mistakenly classified as fine-optimal instead of unconstrained. For any counterfactual involving a different emissions target in 2024, I set every pool's target T_g equal to the alternate EU fleet-wide target T. Future years' specific emission targets cannot be calculated because they rely on the average test mass of cars sold in the preceding two years. For example, the 2030 specific emission target requires knowledge of the average test mass of cars in 2028 and 2029. Since future data is unavailable in the present, I use the EU fleet-wide target as the next best alternative. And owing to the impossibility of calculating specific pool targets, ZEV credits[21] as well as the proposed 2035 carbon credit mechanism and OVC-HEV credits are not modeled.

The shadow cost $\lambda_{g(f)}$ in equation (11) which firms take as given is the solution to the pool's mixed-complementarity problem in (12). A mixed complementarity problem is a type of optimization problem where one minimizes a function $F(x)$ given bounds on the values of the input x: $[l, u]$. If the optimal value of x is l, $F(x)$ is greater than or equal to zero. If the optimal value of x lies between l and u, $F(x)$ equals 0. Finally, if the optimal value of x is u, $F(x)$ is less than or equal to zero. In other words, this optimization method allows for the constraint to be slack, binding, or broken depending on the value of x. Since emission pools can choose to comply or violate their EU emission targets, I solve for each pool's shadow cost λ_g using a mixed complementarity approach.

The left expression denotes the effective range that any pool's shadow cost can take. The lower limit is 0 as shadow costs can never be negative. The upper limit to λ_g is 95 in this setting since pools are given the choice to break the EU emissions "constraint" (average emissions being at or below target) altogether. Pools would rather emit above the threshold and pay a fine if their shadow cost of compliance was 95 or greater. Profit maximization entails cost minimization, and it is not cost minimizing to incur more than 95 Euros per car to comply with standards when one could pay 95 Euros to the EU and over-pollute.

The right hand expression in (12) is the function which the optimal shadow cost λ_g^* solves. Here, it

[21] Used to calculate pool-specific emission targets

is the negative of a pool's net emissions deviation ($E_g(P)$). A pool's shadow cost does not directly appear in $-E_g(P)$, but firms' choices of prices are influenced by λ via equation (11). The optimal pool shadow cost vector $\lambda = [\lambda_1, \lambda_2, ..., \lambda_G]$ results in the vector of chosen car prices minimizing $E_g(P)$ for every pool under mixed complementarity conditions. I.e., if a pool has a shadow cost of zero, it pollutes under the threshold. If a pool has a shadow cost greater than 0 but less than 95, it aligns its average emissions exactly to the EU target. And if a pool has a shadow cost of 95, it pollutes over the target and pays a fine to the EU. It is important to note that pools do not "select" their shadow costs; the shadow costs represent the per-car cost to the pool of marginally deviating one g/km of CO_2 above the threshold.

An unconstrained emissions pool is defined by having average emissions below the threshold. $-E_g(P) \geq 0$ only occurs as a solution if the pool's optimal shadow cost is 0. I therefore classify pools with $\lambda_g^* = 0$ as "unconstrained". A constrained emissions pool aligns its average emissions to its target T_g but faces costs in doing so. $-E_g(P) = 0$ only occurs when λ_g* falls between 0 and 95. I classify pools with $0 < \lambda_g^* < 95$ as "constrained". Finally, a fine-optimal pool pollutes over the EU target and pays a fine. $-E_g(P) \leq 0$ (meaning the pool has a positive net emissions deviation) only when $\lambda_g^* = 95$. Pools with λ_g^* are classified as "fine-optimal". The method by which the mixed complementarity problem is solved is described later in the next subsection.

Given the implied marginal costs c and equilibrium shadow costs λ^*, the equilibrium price vector p^* is the one that satisfies the stacked first-order conditions of each firm's profit function:

$$p - \Omega^{-1}(P)(s(P) \oslash \gamma) = \tilde{c}(\lambda, T) \tag{16}$$

Where:

$J = \sum_{n=1}^{N} J_n$ = country-car observations

N = number of markets

G = number of pools

$p = (J \times 1)$ vector of pre-tax prices

$\Omega(P) = (J \times J)$ ownership matrix

$P = (J \times 1)$ vector of final prices

$s(P) = (J \times 1)$ vector of market shares

$\gamma = (J \times 1)$ vector of $\frac{\partial P_{jn}}{\partial p_{jn}}$

$\tilde{c}(\lambda, T) = (J \times 1)$ vector of emissions-adjusted marginal costs.

$\chi = (J \times G)$ matrix with elements $[\chi_n]_{jg} = \begin{cases} 1 & \text{if product } j \text{ is made by a firm in group } g \\ 0 & \text{otherwise} \end{cases}$

$co2 = (J \times 1)$ vector of CO_2 emission ratings

$T = (G \times 1)$ vector of pool emission targets

The implied marginal costs c are backed out by plugging in the observed price vector in 2024. However, c_j and λ_g cannot be simultaneously recovered. Emissions-adjusted marginal cost $\tilde{c}_j(\lambda, T)$ is calculated using the expression on the left and is itself equal to $c + \chi\lambda \odot (co2 - \chi T)$. The issue is that for any value of $\tilde{c}_j$, there are an infinite number of combinations (c, λ) which would satisfy it. An increase in λ would lead to a proportional decrease in c and vice versa. Papers in the CAFE literature have worked around this issue in several ways. Goldberg (1998) takes advantage of the fact that no firm exactly complied with CAFE to bind the shadow costs to a specific range. In the context of the European Union, no automaker exactly met their target. I apply the same approach to fix the values of λ_g in 2024. Any pool with $E_g(P) < 0$ are unconstrained and have null shadow costs. Emissions-unadjusted marginal cost c_j is equivalent to the emissions-adjusted marginal cost $\tilde{c}_j$. Fine-optimal pools have $E_g(P) > 0$, and their shadow costs are accordingly fixed at 95. c_j is backed out by subtracting $95(co2_j - T_g)$ from $\tilde{c}_j$. I obtain every car's implied marginal cost c_j using this method and hold them fixed across all counterfactuals.

5.1. Using the Model for Counterfactuals

The simulation of a counterfactual equilibrium requires two steps. First, given a guess for the pool shadow cost vector $\lambda = [\lambda_1, \lambda_2, ..., \lambda_G]$ and the implied marginal costs c, the price vector p is calculated[22]. Through p the market share vector $s(P)$ and pool emission deviations $E_g(P)$ are obtained. Next, the computer evaluates the guess for λ using a residual. If the residual does not equal zero for all pools, then a new guess for λ is made[23]. The shadow costs which satisfies each pool's mixed complementarity problem compose the equilibrium shadow cost vector λ^*; the price vector associated with it is the equilibrium price vector p^*. I first will describe how p is calculated, and then I will explain the residual minimization process.

p is the price vector which solves equation (16. Given a guess for λ and the implied marginal costs c, the computer can easily back out p. However, this process is complicated by the fact that subsidy-eligible cars

[22] I use the Morrow-Skerlos (2011) fixed-point method to decrease the solve time for p.

[23] I use SciPy's *hybr* method to run iterations for λ

have kinked first-order conditions. When a French EV eligible for the 2024 "ecologique bonus" program[24] has a VAT-inclusive price between 0 and 18518.52 Euros, the consumer gets a purchase subsidy equal to 27% of this price. But when the same eligible EV is priced between 18518.52 and 47000 Euros, this purchase subsidy is fixed at €5000. The Italian EV/PHEV purchase subsidy program in 2024 also had threshold kinks. If a EV or PHEV emits between 0-20 g/km of CO_2 and has a pre-tax price below or at €35000, the buyer is eligible for a subsidy with an expected value of €6484.95. If that car is priced any higher, the consumer no longer gets the subsidy. And if a PHEV emits between 20 and 60 grams of CO_2 per km and has a pre-tax price below €45000, the buyer gets a subsidy with an expected value of €4323.30. Thus, depending on what price the manufacturer sets, the first-order condition for these cars will change and thus the price vector. I use the following procedure to handle these kinks. I conduct an initial price solve where every potentially eligible car[25] in France and Italy gets the subsidy regardless of whether the manufacturer abides by the pricing requirement. I compare each car's optimal price in relation to their respective subsidy threshold. If the car was priced under the subsidy threshold when it was assigned a subsidy, then the automaker will always price the car to be subsidy-eligible. However, if the car was priced above the threshold when it got the subsidy (a "switcher"), then that price is not applicable. I then solve the pricing FOCs again but now remove the subsidies from the prices of switchers. I examine the new price vector to see if any car was priced contrary to their subsidy assignment (i.e., unsubsidized cars being priced below the threshold or subsidized cars being priced above the threshold). For any car which is priced contrary to the subsidy assignment, I mark for the next round whether to add or remove the subsidy for that car. This process repeats until there are no cars which switch subsidy status.

With this final subsidy assignment, I flag any car that is priced above the subsidy threshold despite receiving the subsidy in the first round. They are flagged because one of two possibilities exist for each of them. The first possibility is that the price lying above the threshold is optimal. The second possibility is that pinning the price of the flagged car at the threshold is more profitable. The FOC at the threshold is non-differentiable so using the price derived from the first-order condition is not sufficient to determine optimality. Holding all other subsidy assignments fixed in the firm's portfolio, I compute the profits when the car is priced above the threshold versus when the car is priced at the threshold. The comparison of the firm's profits in both scenarios determines that car's equilibrium price. The above-threshold price is chosen if the firm's profits are higher than if it had priced that car at the subsidy threshold. After sequentially working through each "switcher", a list of cars pinned at the subsidy thresholds is generated. With this

[24] i.e., weighing less than 2400 kg

[25] All EVs under 2400 kg in France and all EV/PHEVs in Italy with emissions less than or equal to 60 g/km.

list, I fix these cars' prices to be at their respective subsidy thresholds and solve for the pricing first-order conditions.

With the equilibrium price vector having been determined, the computer evaluates the corresponding vector of pool shadow costs $\lambda = [\lambda_1, \lambda_2, ..., \lambda_G]$ with the following residual:

$$\lambda_g - \min\{\max\{0, \lambda_g + E_g(P)\}, 95\} \tag{17}$$

The computer searches for the unique λ that makes the above residual equal zero for all pools. With each guess of λ, the computer must solve for the price vector p. The prices then determine the market share of each car and thereby pool emission deviations $E_g(P)$.

The presence of the minimum function in the residual effectively binds all pool shadow costs to lie in the range [0, 95] and restricts the range of possible emission deviations $E_g(P)$ for a given shadow cost. The maximum function nested inside the minimum function tests whether $\lambda_g + E_g(P)$ is positive. If $\lambda_g + E_g(P)$ is less than zero, the function returns zero. This corresponds to an unconstrained pool which by definition has average emissions less than the target. The maximum function returns $\lambda_g + E_g(P)$ though if this sum is positive. With the output having been determined, the minimum function is evaluated.

If pool g is unconstrained, meaning the maximum function returned zero, the minimum function outputs zero. The entire residual then equals λ_g. As the residual's value for pool g is set equal to zero, it must be the case that $\lambda_g = 0$.

If pool g is not unconstrained, the maximum function returns the positive value $\lambda_g + E_g(P)$. The minimum of $\{\lambda_g + E_g(P), 95\}$ is then returned. If the left element is returned ($\lambda_g + E_g(P)$, then the residual evaluates to $\lambda_g - (\lambda_g + E_g(P)) = E_g(P)$. The computer's guess for λ_g is correct if the residual equals zero, and this can only happen if $E_g(P) = 0$. This scenario corresponds to pool g being constrained as the pool is not constrained by virtue of $\lambda_g + E_g(P) > 0$ and the shadow cost λ_g is greater than zero.

If the right element of the minimum function is returned (95), the residual evaluates to $\lambda_g - 95$. The only way this expression equals zero is if λ_g itself equals 95 and $E_g(P) \geq 0$. If the pool's average emissions were below the the target ($E_g(P) < 0$), the residual would evaluate to $95 - (95 + E_g(P)) = -E_g(P)$ and therefore would fail to equal zero. Pools that satisfy their residual in this manner are fine-optimal.

The above rules allow for cases where a pool that exactly meets its target has a shadow cost of 0 or 95. This would be problematic for my identification strategy for cars' emissions-unadjusted marginal costs c if at least one pool exactly meet its target. Every shadow cost must have a 1:1 mapping with a constraint status, but it would be impossible to map the "true" shadow cost to an emissions pool that hit its exact target. Fortunately no such case existed and so this identification problem never applied. The computer runs SciPy's *hybr* iterative procedure to find λ given an initial guess for the shadow costs λ_0. If an iteration's λ produces prices such that at least one pool's residual is non-zero, then the computer runs a new iteration. When λ satisfies each pool's mixed complementarity problem, the equilibrium price vector p^* is obtained.

5.2. Pooling Structure

Emission pools are formed endogenously in reaction to a year's emission targets. The choice to join or stay apart from an emissions pool depends on two factors: individual incentive and the perceived likelihood of meeting the emissions target individually versus collectively. Individual incentive refers to whether membership is more profitable than not for a given emissions target T_g. A firm that is individually unconstrained finds it strictly more profitable to belong to an emissions pool regardless of the pool's emission performance that year. Although unconstrained pools have null shadow costs in the model, one can assume in actuality they have a infinitesimally small intra-pool transfer rate. It is effectively zero, but it rationalizes membership for firms that otherwise would be individually unconstrained. A firm that is individually constrained prefers to be in an emissions pool g' i.f.f. $\lambda_{g'} < \lambda_g < 95$ where g is the group index that corresponds to the firm being by itself. If the pool penalty rate $\lambda_{g'}$ exceeded its individual shadow cost of compliance λ_g, the firm would rather join another pool where the reverse is true. Failing that, it would individually adhere to the EU's goal. Fine-optimal firms would join an emissions pool if and only if $\lambda_{g'} < 95$.

The perceived likelihood of meeting the emissions target is also a crucial factor in pool formation. It affects both the willingness to belong to an emissions pool and the willingness to accept a particular automaker into the pool. A firm that sells primarily high-polluting cars[26] will not want to join an emissions pool if it is convinced that the pool will fail to keep its average emissions below the target. Emission pool members will also be less willing to accept a low-emissions firm if they think that the pool will already meet the target. This is the most likely explanation as to why Tesla did not belong to any emissions pool in 2023 and 2024. In 2025, the same year many car manufacturers publicly voiced concerns that they could not meet

[26] i.e., fine-optimal firms

the 93.6 g/km target, emission pool membership had dramatically changed. Table 3 in the Appendix details the pooling structure in the years 2024 and 2025.

It is important to keep in mind that a firm's shadow cost of compliance is relative both to the EU target T_g and the pool that it belongs to. An automaker that merely finds itself constrained as an emissions pool member might have found it optimal to over-pollute and pay an EU fine if it was alone. Conversely, a solo firm that finds the EU fine less costly than emissions reduction might instead reduced its emissions if it belonged to a pool. All pools are not created equal either: the member composition matters. Pools with more high-emissions firms will *ceteris paribus* have correspondingly higher shadow costs and vice versa. In theory, the optimal grouping of automakers for a counterfactual emissions target T_g is simple. For a given grouping of automakers, one would simulate the equilibrium and calculate producer surplus. The unique assortment of pools that yields the highest producer surplus is the optimal pooling structure. The difficulty lies in the time cost necessary to test every possible combination of pools. With 23 car companies, the number of unique combinations is $\sum_{r=1}^{23} \binom{23}{r}$ = 8,388,607. With each equilibrium taking several hours to solve, the required time would be astronomical. The time cost would still be enormous if the program was ran in parallel. This is why I hold the pools fixed when simulating the implementation of the 2030 or 2035 EU emission targets in 2024. I use the 2025 pools since they are the last observable grouping and were formed in reaction to the EU's lowest fleet-wide target yet (93.6 g/km).

5.3. Welfare

Total welfare (W) is defined to be the sum of four components: consumer surplus (CS), producer surplus (PS), government net receipts on autos (G), and the social cost of carbon ($CO2$). The mean consumer surplus in market n is given by:

$$CS_n = M_n \sum_{i \in I_n} \frac{1}{I_n} * \frac{y_i}{-\alpha} log(1 + \sum_{j \in J_n} exp(V_{ijn})) \tag{18}$$

Where:

I_n = Number of consumers in market n

$V_{ijn} = x_j\beta - \alpha \frac{p_{jn}}{y_i} + \sum_k \sigma_k \nu_{ik} x_k$

I compute the consumer surplus in each market n and then take the sum of these to create my aggregate consumer surplus measure. Producer surplus is calculated by taking the summation of gross profits[27] across all firms and then subtracting the total amount of EU fines paid. Government net receipts are equal to total government revenue from cars minus expenditure on auto purchase subsidies. The sources of revenue that governments collect are value-added tax, malus (specific to France), and EU emission fines. Subsidies are the only source of spending in this model and come in the form of "écologique bonus" (France) and EV/PHEV purchase subsidy (Italy). I separately compute government net receipts by country because there are country-specific taxes and revenues (e.g., malus). EU-wide [28] net receipts are calculated by adding each nation's net receipts and the EU emission target fines[29]. I conduct a back-of-the-envelope calculation for the social cost of carbon following earlier studies on EU emissions standards (Parry, Walls, & Harrington, 2017; Reynaert, 2021; Yilmaz, 2021). I adopt the following assumptions used in these papers to calculate this cost. Each new car is driven 14000 km yearly with a total lifetime of 15 years, and the discount rate is assumed to be 0.9. The externality cost of car ownership (e.g., congestion and accidents) is €0.12/km (Parry, Walls, & Harrington, 2017). With these, I utilize the formula from Yilmaz (2021, p. 31) to calculate the social cost of carbon[30]

$$CO2 = -\sum_{n=1}^{N}\sum_{j\in J_n}(0.12 + 77\frac{co2_j}{1000000})s_{jn}(P)M_n * \sum_{t=0}^{14}14000(0.9)^t \tag{19}$$

Where:

t = Year

$\frac{co2_j}{1000000}$ = Metric tons of CO_2 emitted per km

6. Counterfactuals

I simulate how firms would contemporaneously react to alternate EU emissions targets in 2024. The counterfactual is constructed by altering T_g in the constraint and then solving for the new equilibrium prices and shares. The following counterfactual targets are examined: 93.6 g/km (2025-2029 target), 49.5 g/km

[27]Gross profits are the firm's profits gross of any intra-pool transfer payments or EU fines and ignoring fixed costs which I do not observe or estimate.

[28]In this model, I define this to be the 3 countries available in my sample: France, Germany, and Italy.

[29]Note that EU emission target fines cancel out in the welfare calculation. The reduction in producer surplus from EU fines equals the gain in net government revenue from these fines

[30]As the social cost of carbon contributes negatively to total welfare, I multiply the estimate by -1 to make it negative. The more negative the estimate is, the higher the social cost of carbon is and vice versa.

(2030-2034), 11.51 g/km (current post-2034 target), and 0 g/km (original post-2034 target). I also evaluate what would happen if the emissions target program did not exist in 2024. Each firm would set prices to maximize its profits according to the equation:

$$\pi_f = \sum_{n=1}^{N} \sum_{j \in \mathcal{F}_{fn}} (p_{jn} - c_j) s_{jn}(P_n) M_n \tag{20}$$

The emissions penalty term $\lambda_g(co2_j - T_g)$ disappears altogether since emission pools do not exist in this scenario. Automakers set prices only with respect to emissions-unadjusted marginal cost (c_j) and car demand. The equilibrium prices are the ones which satisfy the vector of first-order conditions:

$$p - \Omega^{-1}(P)(s(P) \oslash \gamma) = c \tag{21}$$

These pricing first-order conditions are nearly identical to those in my main model. The only difference is that the right-hand side contains c instead of $\tilde{c}$. The simulated prices reflect what automakers would have set had they never needed to worry about their average emissions.

I report the welfare estimates in Table 6. Total welfare drops monotonically as the counterfactual EU emissions target shrinks. The largest welfare drop, 16 billion Euros in magnitude, occurs if the EU implemented its original 2035 EV mandate in 2024. This translates to an almost 5% decline in welfare from the observed 2024 equilibrium. The implementation of the current 2035 emissions target would result in a similar welfare decrease of 4.36%. The reduction in the CO_2 externality does much to flatten these welfare losses. Without taking this into consideration, they would be considerably higher. Interestingly, welfare would only improve by 0.049% if the EU scrapped its emissions target problem. EU average emissions from new cars would be slightly lower than the 106.88 g/km level observed in 2024. This result aligns with an ICCT research brief which reported that automakers did not feel constrained by the emission targets imposed in 2024 (Diaz et al., 2025).

Consumer surplus is the dominating factor in driving these changes; it is the largest component in every counterfactual. Its decline with the 93.6 g/km target is marginal at -0.64%, but it quickly shoots up to around -7.1% for the 49.5 g/km and around 15% for the 11.5 and 0 g/km targets. Ignoring possible positive effects on utility from decreased CO_2 emissions, consumers lose out in aggregate from tighter emission targets.

This is not to say that every individual consumer is worse off in these scenarios. Consumers with preferences for electric vehicles will *ceteris paribus* have greater utility as a result of decreased EV prices. Conversely, consumers that have preferences towards cars running on petroleum will have decreased consumer surplus from rising petrol car prices.

Unsurprisingly, producer surplus decreases when the EU imposes stricter emission targets. A pool's cost of compliance is non-decreasing in nature. The pool finds it equally, if not more costly, to meet a lower emissions target. As such, most pools will incur greater costs when the EU emissions target is lower. This in turn directly affects firm profits. Table 7 displays each automaker's percent change in profits from the observed 2024 equilibrium. In general, the only firms which had higher profits under lower EU emission targets are the car manufacturers which solely make electric vehicles: BYD, Leapmotor, Smart, and Tesla. Automakers that already had EVs or PHEVs as major components of their 2024 portfolio also tended to fare better (e.g., Mercedes-Benz or Volvo). Legacy car manufacturers like Volkswagen or Nissan found it especially difficult to adapt to tighter EU emission targets. The worst decrease in profits in the 2025 target counterfactual was Nissan's 16.34% decline. The implementation of the 2030 target in 2024 causes almost every automaker to lose profits aside from EV manufacturers and manufacturers with proportionally higher EV sales like Mercedes-Benz. Multiple automakers lose over 20% of their profits with Subaru in particular losing over a third of its profits. With the 11.5 and 0 g/km counterfactuals, almost every automaker has at least 10% fewer profits than in 2024. Automakers with large market shares like Volkswagen or Stellantis have more than 20% fewer profits. Whilst most of these car companies conduct business globally, these profit decreases would nevertheless be devastating. Contrary to expectations, abolishing the emissions target program slightly reduces producer surplus (-0.553%). Most carmakers have slightly less profits in this scenario than their observed 2024 profits. An outlier is Jaguar which especially benefits from the program's removal: its profits increase by almost 20%. Jaguar paid the EU fines in 2024 and in every counterfactual scenario, so the program's removal would be a welcome relief.

The difficulty of many automakers to comply with future EU targets implemented in 2024 is largely due to the insufficient level of observed EV demand. As seen in Figure 1, aggregate EV share never rises above 29%. The EV share is 19.2% (7.4 percentage point increase from observed) in the 93.6 g/km target counterfactual. In the counterfactuals with the 2030 target (49.5 g/km), the current 2035 target (11 g/km), and the old 2035 target (0 g/km), EV share plateaus at around 28%. Average EU emissions also plateau with stricter targets. Table 8 reports the average emissions in the EU under the counterfactual 2024 emission targets. Average emissions would drop to 93.51 g/km if the 2025 emissions target (93.6 g/km) was imple-

mented. However, the EU would not be able to successfully achieve its 2030 target (49.5 g/km), the current 2035 target (11.5 g/km), or the defunct 2035 EV mandate (0 g/km) in 2024. What is readily apparent here is the constraining influence of contemporaneous EV demand. In all three of these scenarios, EU average emissions hover around 79 g/km. Knowing this, EU policymakers whose sole goal is to minimize short-run emissions might want to set a 79 g/km target instead. EU average emissions would decline to a level very close to this target, and the welfare loss would almost assuredly be minimized. The ceilings which EV shares and average EU emissions hit indicate short-run EV demand is simply insufficient for automakers to meet the EU's long-run goals. Lowering the prices of electric cars relative to cars with ICEs only can accomplish so much. This suggests the need to establish government programs to boost EV demand among consumers.

This conclusion is specific however to the short-run. In the future, improvements in EV technology are likely to positively shift consumer EV demand. Automakers will also alter their car lineups in line with evolving market conditions and advances in auto technology. My model does not include these endogenous factors as its focus is on the counterfactual equilibrium in 2024, and the they are unknowable *ex ante*. It is possible that, with time, EV demand will sufficiently increase to the point where EVs become the dominant type of car sold in the EU. This is what the European Union assumes will happen with its exclusively supply-side approach. Time will tell whether this was a good bet or a strategic error.

Table 9 shows how well each emission pool performed relative to the counterfactual emissions targets. For the 2025 target, nine out of eleven pools have average emissions less than or equal to the target. Jaguar and Volkswagen are fine-optimal in this scenario with Jaguar being over 35 g/km above target. BMW, Mercedes-Benz, and SAIC are unconstrained with the 93.6 g/km target. This is not surprising considering EVs made up a higher proportion of these pool members' 2024 sales, and they had average emissions well below target. With the 2030 target, the standard is roughly halved at 49.5 g/km. The results are now reversed where nine out of eleven pools have average emissions exceeding the target. Only Mercedes-Benz and SAIC manage to hit the EU's emission target exactly. And with both the old 2035 EV mandate target and the current 2035 target, no pool manages to comply. This is to be expected with targets that are intended to be applied 11 years later. However, the 2035 target and 0 g/km counterfactuals still serve well as proofs-of-concept.

It is useful to decompose how government spending changes in these counterfactuals. Table 10 splits up the change in government expenditure by each constituent component. EU emission target fines drastically increase if the EU had imposed tighter emission targets in 2024. The EU collects an additional €350 million

fines with the 2025 target, and the increases are in the billions with the 2030 and 2035 targets as well as the defunct 0 g/km target. It is evident that the magnitude of these fines had a significant negative impact on producer surplus. The total amount of value-added tax (VAT) collected by Germany and Italy increased as the targets got tighter. Curiously, in the 0 g/km target counterfactual, French VAT tax revenue decreased by 1% relative to France's 2024 VAT revenue. Outlays on the French ecologique bonus and Italian EV/PHEV subsidy program skyrocket when the EU emissions target tightens. France spends over a billion euros more on écologique bonus subsidies in every scenario save the "no target" counterfactual. Italian outlays on its EV/PHEV subsidy increase between 300 and 500 percent. The concomitant rise in EV and PHEV sales are what triggers the increase in subsidy outlays. Whether or not the French and Italian governments would be willing to ramp up spending on these subsidy programs is another matter. These programs were designed with specific budgets in mind. For example, France allocated €1.5 billion for the écologique bonus in 2024 but slashed €500 million from this program for 2025 (Brice, 2024). Italy allocated €793 million for car purchase subsidies in 2024 with only €390 million of that sum going towards EVs and PHEVs (Randall, 2024). I included the programs in the equilibrium model as-is because I am investigating the hypothetical effect of the EU implementing a different emissions target in 2024 *ceteris paribus*. Nevertheless, the limited budgets these programs had in 2024 should be kept in mind when evaluating the model's results.

7. Conclusion

The EU's implementation of tighter emission targets in 2024 would result in a limited reduction in auto emissions at the cost of welfare decreases across the board. The EU could achieve its 2025 target of 93.6 g/km with around a 1% decrease in total welfare. The future, stricter emission targets (like the 2030 49.5 g/km objective) could not be met in 2024 though. Average emissions hover around 79 g/km in the 2030, 2035, and "no target" counterfactuals. Consumers bear the brunt of the welfare losses if the EU tightened its emission standard. Consumer surplus falls by 2 billion Euros[31] with the 93.6 g/km target, but the decline increases elevenfold with the 49.5 g/km target and over twentyfold with the 11.5 or 0 g/km targets. Producer surplus declines with every smaller emissions target, but the effects on firm profits are heterogeneous. Firms which had a EV-heavy lineup in 2024 tend to have higher profits whereas automakers reliant on ICE car sales face enormous decreases in profits. The welfare decreases in these counterfactuals are noticeably dampened by decreases in long-term CO_2 emissions. In the worst scenario for total welfare

[31] Less than 1 % relative to 2024

(the 0 g/km target), the decline is less than 5%. If the emissions target program did not exist in 2024, CO_2 emissions would stay around their observed levels. Total welfare would marginally improve (0.049%), but there would be no evidence that EV shares would naturally increase in later years. The existence of basins for EV shares and average EU emissions in the 2030, 2035, and 0 g/km counterfactuals indicate the limits of the EU's supply-side approach. Automakers can only boost EV sales so much by lowering their relative prices to petrol cars. This finding suggests the need to investigate methods to boost EV demand, especially since electric car sales in the European Union have consistently fallen below projections.

There are some limitations that should be considered when evaluating the results. First, the small number of markets in my dataset poses some difficulty in estimating car demand. Year or market-year fixed effects were unable to be estimated since there were only two years of data. There was too little market variation to allow for the random coefficient on mass to be properly estimated. And although I think the price coefficient of -3.257 to be reasonable in the context of the literature's estimates, it is on the smaller side. Given that α has been known to be biased when the number of markets is small (Conlon & Gortmaker, 2020, pp. 45-46), it is possible that it is an underestimate. This would imply that my welfare estimates constitute a lower bound instead. The role of local charging station availability on EV adoption is also not addressed in this paper. Despite recent advances in the maximum distances electric cars can travel on a full battery, range anxiety is still widely recognized as being an inhibiting factor towards mass EV adoption (Cruz et al., 2026; Springel, 2021). I was not able to address this in my demand estimation or counterfactual analysis as the EEA's national car registration data is too granular. Finally, I fixed the pooling structure in all counterfactuals (save the "no target" scenario) to consist of the 2025 emission pools. I did this because it was too computationally intensive to find realistic counterfactual pooling structures.

Future research on this subject can be expanded in a number of ways. One important direction to consider is modelling endogenous pool formation. Firms choose whether to pool or not, and with whom, based on individual incentive and the perceived likelihood of meeting the target. Every emission pool which has ever existed formed in reaction to the EU emissions standard and the market idiosyncrasies that year. There would be a different pooling structure in 2024 if the EU had set a different emissions target[32]. The value of my equilibrium model would be further enhanced by discovering a feasible way to incorporate endogenous pool formation. Firm R&D investments and product entry, which this study abstracts from, are another interesting avenue to explore. The EU emissions target program was designed to be staggered in part because future advancements in EV technology would make production of electric cars less expensive

[32] An exception could arise if the change was minuscule (e.g., 0.0000001).

and thereby compliance. Automobile manufacturers will have other emission reduction options available to them than just adjusting the adjustment of EV/PHEV relative prices. Some examples include the relocation of gasoline car sales to other markets, minor emission modifications to existing cars, research into battery technology, et cetera. Product design and technology can be assumed to be fixed in the short-run, but they should be taken into account by researchers interested in simulating the long-term equilibrium effects of EV policies.

8. References

Alberini, A., Di Cosmo, V., & Bigano, A. (2019). How are fuel efficient cars priced? Evidence from eight EU countries. *Energy Policy*, *134*, 110978. https://www.sciencedirect.com/science/article/pii/S0301421519305658

Anderson, S., & Sallee, J. (2011). Using loopholes to reveal the marginal cost of regulation: The case of fuel-economy standards. *American Economic Review*, *101*(4), 1375-1409. https://www.aeaweb.org/articles?id=10.1257/aer.101.4.1375

Automotive World (March 4, 2026). *Toyota and Stellantis walk out of Tesla's CO2 credit pool.* https://www.automotiveworld.com/news/toyota-and-stellantis-walk-out-of-teslas-eu-co2-credit-pool/

Berry, S., Levinsohn, J., & Pakes, A. (1993). *Automobile prices in market equilibrium: Part I and II* (NBER Working Paper No. w4264). https://www.nber.org/papers/w4264

Brice, M. (October 10, 2024). *France plans to cut assistance for EV purchases by third, toughen penalties on some vehicles.* Reuters. https://www.reuters.com/business/autos-transportation/france-plans-cut-assistance-ev-purchases-by-third-toughen-penalties-some-2024-10-10/

Campbell, P., & McGee, P. (April 7, 2019). *Fiat Chrysler pools fleet with Tesla to avoid EU emission fines.* Financial Times. https://www.ft.com/content/7a3c8d9a-57bb-11e9-a3db-1fe89bedc16e?syn-25a6b1a6=1

Cincotta, C., & Thomassen, Ø. (2025). Evaluating Norway's electric vehicle incentives. *Energy Economics*, *146*, 108490. https://www.sciencedirect.com/science/article/pii/S0140988325003147?ref=cra_js_challenge&fr=RR-1

Commission Implementing Decision (EU) 2023/1623 of 3 August 2023 specifying the values relating to the performance of manufacturers and pools of manufacturers of new passenger cars and new light commercial vehicles for the calendar year 2021 and the values to be used for the calculation of the specific emission targets from 2025 onwards, pursuant to Regulation (EU) 2019/631 of the European Parliament and of the Council and correcting Implementing Decision (EU) 2022/2087, 2023 O.J. (L 200) 5.

Conlon, C., & Gortmaker, J. (2020). Best practices for differentiated products demand estimation with PyBLP. *RAND Journal of Economics*, *51*(4), 1108-1161. https://jeffgortmaker.com/files/Best_Practices_for_Differentiated_Products_Demand_Estimation_with_PyBLP.pdf

Cruz, A., Fuinhas, J., & Belucio, M. (2026). Battery electric vehicle adoption in the European Union: Analysis of the determining factors. *The Electricity Journal*, *39*, 107531. https://www.sciencedirect.com/science/article/pii/S1040619026000060

Dahil, J. (May 9, 2025). *How carmakers fought EU rules and won.* Politico. https://www.politico.eu/article/carmakers-co2-rules-auto-eu-renault-cars/

Díaz, S., Tietge, U., Dornoff, J., & Mock, P. (2025). *CO2 emissions from new passenger cars in Europe: Car manufacturers' performance in 2024.* The International Council on Clean Transportation. https://theicct.org/publication/co2-emissions-from-new-passenger-cars-in-europe-car-manufacturers-performance-in-2024-dec25/

Durrmeyer, I., & Samano, M. (2018). To rebate or not to rebate: Fuel economy standards versus feebates. *The Economic Journal*, *128*(616), 3076-3116. https://academic.oup.com/ej/article/128/616/3076/5251687

European Commission (2025). *Commission Staff Working Document: Impact assessment report accompanying the document Proposal for a regulation of the European Parliament and of the Council amending Regulation (EU) 2019/631 as regards CO2 emission performance standards for new light duty vehicles and vehicle labelling and repealing Directive 1999/94/EC.* https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52025SC1058

European Commission. (2025, December 16). *Commission takes action for clean and competitive automotive sector* [Press release]. https://ec.europa.eu/commission/presscorner/detail/en/ip_25_3051

European Environmental Agency (2026, June 25). *Average CO2 emissions from new cars and vans signifi-*

cantly decreased in 2025 [Press release]. https://www.eea.europa.eu/en/newsroom/news/average-co2-emissions-from-new-cars-and-vans-significantly-decreased-in-2025

Goldberg, P. (1998). The effects of the Corporate Average Fuel Efficiency standards in the US. *The Journal of Industrial Economics*, *46*(1), 1-33. https://onlinelibrary.wiley.com/doi/full/10.1111/1467-6451.00059

Grieco, P., Murry, C., & Yurukoglu, A. (2024). Evolution of market power in the U.S. automobile industry. *The Quarterly Journal of Economics*, *139*(2), 1201-1253. https://academic.oup.com/qje/article/139/2/1201/7276495

Grigolon, L., Reynaert, M., & Verboven, F. (2018). Consumer valuation of fuel costs and tax policy: Evidence from the European car market. *American Economic Journal: Economic Policy*, *10*(3), 193-225. https://www.aeaweb.org/articles?id=10.1257/pol.20160078

Jacobsen, M. (2013). Evaluating US fuel economy standards in a model with producer and household heterogeneity. *American Economic Journal: Economic Policy*, *5*(2), 148-187. https://www.aeaweb.org/articles?id=10.1257/pol.5.2.148

Jones, L. (2026, March 12). Tougher CO2 targets pushed EV prices down in Europe after 5 years of increases, report finds. *Automobile News Europe*. https://www.autonews.com/ev/ane-ev-prices-te-europe-0312/

Monteforte, M., & Diaz de Aguilar, S. (2026). *European car market monitor: February 2026*. The International Council on Clean Transportation. https://theicct.org/publication/european-car-market-monitor-feb-2026/

Nevo, A. (2000). Mergers with differentiated products: The case of the ready-to-eat cereal industry. *The RAND Journal of Economics*, *31*(3), 395-421. https://www.jstor.org/stable/2600994?seq=1

Randall, C. (February 2, 2024). *Italy sets aside nearly one billion euros for cleaner transportation*. Electrive. https://www.electrive.com/2024/02/02/italy-sets-aside-nearly-one-billion-euros-for-cleaner-transportation/

Regulation (EU) 2019/631 of the European Parliament and of the Council of 17 April 2019 setting CO2 emission performance standards for new passenger cars and for new light commercial vehicles, and repealing Regulations (EC) No 443/2009 and (EU) No 510/2011, 2019 O.J. (L 111) 13.

Remmy, K. (2026). Adjustable product attributes, indirect network effects, and subsidy design: The case of electric vehicles. *American Economic Journal: Economic Policy*, *18*(2), 107-140. https://doi.org/10.1257/pol.20230294

Reynaert, M. (2021). Abatement strategies and the cost of environmental regulation: Emission standards on the European car market. *The Review of Economic Studies*, *88*(1), 454-488. https://academic.oup.com/restud/article/88/1/454/5917028

Sallee, J. (2014). Rational inattention and energy efficiency. *The Journal of Law and Economics*, *57*(3), 781-820. https://www.journals.uchicago.edu/doi/full/10.1086/676964

Springel, K. (2021). Network externality and subsidy structure in two-sided markets: Evidence from electric vehicle incentives. *American Economic Journal: Economic Policy*, *13*(4), 393-432. https://www.aeaweb.org/articles?id=10.1257/pol.20190131

The International Council on Clean Transportation. (2026, January 25). *Europe's battery electric car market closes 2025 at a 19% average, up 4 percentage points from 2024* [Press release]. https://theicct.org/pr-europe-battery-electric-market-closes-2025-19-average/

Van Biesebroeck, J., & Verboven, F. (2026). Demand, competition, and public policy in the automobile industry. *Journal of Economic Literature*, *64*(3), 984-1033. https://doi.org/10.1257/jel.20251792

Yilmaz, Y. (2021). *Microeconomic analyses of the effects of regulation and elections across Europe and Turkey* [Doctoral dissertation, Northeastern University]. Northeastern University Digital Repository Service. https://repository.library.northeastern.edu/files/neu:4f16j584s/fulltext.pdf

9. Appendix

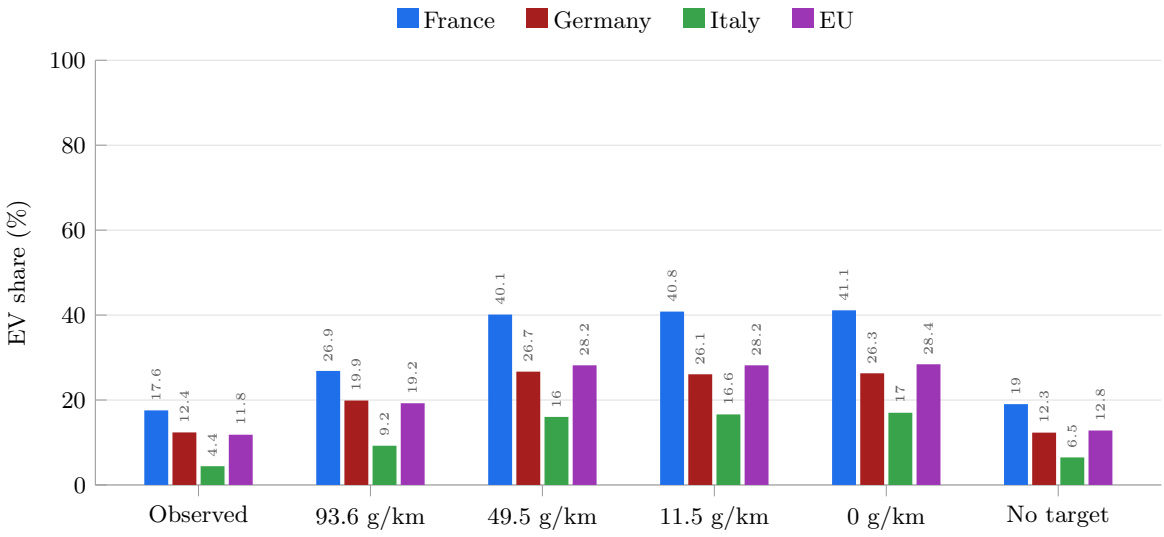


Figure 1: EV share by market under alternative CO_2 targets

Table 1: EU CO_2 emission targets

Years	**EU fleet-wide target**	**Specific emissions reference target**	**Specific emissions target**
2021	115.1	WLTP Reference Target$_f$ $= WLTP_CO2_{f,2020} \times \frac{NEDC_{2020\text{target}}}{NEDC_CO2_{f,2020}}$	WLTP Reference Target$_f$ $+ 0.0333\,(M_{ft} - M_{f,2020})$
2022–2024	115.1	WLTP Reference Target$_f$ $= WLTP_CO2_{f,2020} \times \frac{NEDC_{2020\text{target}}}{NEDC_CO2_{f,2020}}$	WLTP Reference Target$_f$ $+ 0.0333\,\big((M_{ft} - 1398.50) - (M_{f,2020} - 1379.88)\big)$
2025–2026	93.6	$93.6 - 0.0144\,(TM_{ft} - 1650.15)$	Specific emissions reference target$_{ft}$ $\times$ ZLEV factor$_{ft}$
2027–2029	93.6	$93.6 - 0.0144\,(TM_{ft} - TM_0)$	Specific emissions reference target$_{ft}$ $\times$ ZLEV factor$_{ft}$
2030–2034	49.5	$49.5 - 0.0076\,(TM_{ft} - TM_0)$	Specific emissions reference target$_{ft}$ $\times$ ZLEV factor$_{ft}$
2035-onwards	11.5	$11.5 - 0.0018\,(TM_{ft} - TM_0)$	Specific emissions reference target$_{ft}$ $\times$ ZLEV factor$_{ft}$

Where:

$$x = \begin{cases} 0.15 & \text{if } 2025 \leq t < 2030 \\ 0.35 & \text{if } 2030 \leq t \end{cases}$$

y = Share of low-emitting or zero-emissions vehicles in a manufacturer's fleet of cars in year t

$$ZLEV factor_{ft} = \begin{cases} 1 & \text{if } (1 + y - x) < 1 \\ 1 + y - x & \text{if } 1 < (1 + y - x) < 1.05 \\ 1.5 & \text{if } 1 + y - x > 1.05 \end{cases}$$

TM_{ft} = Average test mass of all cars made by firm f in year t

TM_0 = Average test mass of all cars registered in the 2 preceding years

M_{ft} = Average mass in running order for firm f in year t

$M_{f,2020}$ = Average mass in running order for firm f in 2020

Table 2: Pool Transfer Payments (T_g = 115 g/km of CO_2)

Case 1: $E_g(P) > 0, \lambda_g = 95$				
Firm	λ_g	**Sales-weighted emissions deviation**	**# of cars sold**	**Payment**
A	95	-100	10	−95000
B	95	10	50	47500
C	95	50	40	190000
Group	95	15	100	142500
Case 2: $E_g(P) < 0$				
Firm	λ_g	**Sales-weighted emissions deviation**	**# of cars sold**	**Payment**
A	0	-100	20	0
B	0	10	50	0
C	0	30	30	0
Group	0	-6	100	0
Case 3: $E_g(P) = 0$				
Firm	λ_g	**Sales-weighted emissions deviation**	**# of cars sold**	**Payment**
A	50	-100	20	−100000
B	50	10	40	20000
C	50	40	40	80000
Group	50	0	100	0

Table 3: EU Emission Pools by Year[33]

Year	**Pool**	**Makes**	**Closed**
2024	BMW	BMW, Mini, Rolls-Royce	1
2024	BYD	BYD	1
2024	FORD	Ford	1
2024	HONDA	Honda	1
2024	HYUNDAI MOTOR EUROPE	Hyundai	1
2024	JAGUAR	Jaguar, Land Rover	1
2024	KIA	Kia	1
2024	LEAPMOTOR	Leapmotor	1
2024	MERCEDES-BENZ AG	Mercedes-Benz, Smart	1
2024	RENAULT-NISSAN-MITSUBISHI	Alpine, Dacia, Mitsubishi, Nissan, Renault	0
2024	SAIC MOTOR CORPORATION LTD	MG	1
2024	STELLANTIS	Alfa Romeo, Citroen, DS Automobiles, Fiat, Jeep, Lancia, Peugeot, Opel	1
2024	SUBARU-MAZDA-TOYOTA	Lexus, Mazda, Subaru, Toyota	0
2024	TESLA	Tesla	1
2024	VOLKSWAGEN	Audi, Bentley, Cupra, Lamborghini, Porsche, Seat, Skoda, Volkswagen	1
2024	VOLVO CARS POLESTAR SUZUKI	Polestar, Suzuki, Volvo	0

[33]Firms that nominally do not belong to emission pools are given their own eponymous emission pool

Year	Pool	Makes	Closed
2025	BMW	BMW, Mini, Rolls-Royce	1
2025	HYUNDAI MOTOR EUROPE	Hyundai	1
2025	JAGUAR	Jaguar, Land Rover	1
2025	KIA	Kia	1
2025	MERCEDES-BENZ	Mercedes-Benz, Polestar, Smart, Volvo	0
2025	MITSUBISHI	Mitsubishi	1
2025	NISSAN-BYD	Nissan, BYD	0
2025	RENAULT	Alpine, Dacia, Renault	1
2025	SAIC MOTOR CORPORATION LTD	MG	1
2025	TESLA	Alfa Romeo, Citroen, DS Automobiles, Fiat, Ford, Honda, Jeep, Leapmotor, Lancia, Lexus, Mazda, Opel, Peugeot, Subaru, Suzuki, Tesla, Toyota	0
2025	VOLKSWAGEN	Audi, Bentley, Cupra, Lamborghini, Porsche, Seat, Skoda, Volkswagen	1

Table 4: Summary statistics

Variable	N	Mean	SD	Min	Median	Max
Market share	2,694	1.10×10^{-4}	2.24×10^{-4}	2.41×10^{-8}	3.40×10^{-5}	3.67×10^{-3}
Price (EUR)	2,694	57,581	79,395	12,415	42,655	3,272,500
Price / GDP per capita	2,694	1.34	1.66	0.29	0.99	64.38
Engine power (kW)	2,694	156.61	91.56	33.00	124.00	782.00
Mass (kg)	2,694	1,687.45	372.00	947.50	1,623.00	2,975.00
Base of car (m^2)	2,694	8.33	0.89	4.48	8.33	12.07
Number of doors	2,694	4.81	0.60	2.00	5.00	5.00
Number of seats	2,694	4.93	0.51	2.00	5.00	7.00
Cost of energy consumption per 10km (EUR)	2,694	1.07	0.42	0.34	1.02	3.66
CO_2 emissions (g/km)	2,694	107.74	74.80	0.00	125.00	456.00
Lagged average RXR	2,694	0.91	0.19	0.25	0.92	1.56
Electric car	2,694	0.17	0.37	—	—	—
Body types						
Hatchback	2,694	0.20	0.40	—	—	—
Station wagon	2,694	0.11	0.31	—	—	—
SUV	2,694	0.55	0.50	—	—	—
Other	2,694	0.14	0.35	—	—	—

Notes: N refers to the number of car-market observations (where a market is a country-year) and SD stands for standard deviation. Market share is reported in scientific notation for ease of reading. I omit the minimum, median, and maximum for the electric car and body type dummies. The body type 'Other' encompasses convertibles, coupes, mini vans, off-road vehicles, roadsters, and sedans.

Table 5: Demand estimates

	Logit	Random coefficients
Price coefficient (α)		
Price / income	-0.993^{***}	-3.257^{***}
	(0.153)	(0.668)
Mean coefficients (β)		
Engine power (KW / 100)	1.166	-1.317^{*}
	(0.750)	(0.794)
Mass (kg / 2000)	-3.691	11.033^{***}
	(3.528)	(3.362)
Length $\times$ width (m^2)	-0.115	-1.049^{**}
	(0.360)	(0.427)
Number of doors / 5	-2.101^{*}	-8.001^{***}
	(1.220)	(2.947)
Number of seats / 5	-2.155	1.598
	(1.408)	(1.778)
Electric car	-0.632^{*}	-2.966^{***}
	(0.343)	(0.724)
Cost of energy consumption (€) per 10km	-0.533^{***}	-4.458^{*}
	(0.188)	(2.329)
Standard deviations (σ)		
Engine power (KW / 100)		0.261
		(0.364)
Mass (kg / 2000)		0.000
Length $\times$ width (m^2)		0.268^{***}
		(0.046)
Number of doors / 5		8.487^{***}
		(1.773)
Number of seats / 5		2.095^{***}
		(0.429)
Cost of energy consumption (€) per 10km		3.256^{***}
		(1.250)
Observations	2694	2694

Notes: Observations are at the car-country-year level. The dependent variable is mean utility δ_{jn}, which is set to the log difference between the car share and the share of the outside good in the logit specification. Mean utility is solved through the BLP contraction mapping in the random coefficients specification. The price/income term represents price divided by GDP per capita in the logit specification and price divided by agent income in the random coefficients specification. Instruments for price include the own-firm and rival-firm BLP instruments, own product characteristics, and the lagged average RXR. Brand and body type fixed effects are not reported for brevity. Standard errors are heteroskedastic-robust. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 6: Welfare effects of alternate 2024 emission targets

Welfare measure	**Difference from observed**				
	$T_{2025} = 93.6$	$T_{2030} = 49.5$	$T_{2035} = 11.5$	$T_0 = 0$	No target
CS	-2.00 (-0.64%)	-22.13 (-7.12%)	-41.61 (-13.39%)	-47.00 (-15.12%)	1.53 (+0.49%)
PS	-1.41 (-2.21%)	-6.12 (-9.57%)	-9.96 (-15.57%)	-11.14 (-17.41%)	-0.35 (-0.55%)
G	-1.24 (-4.44%)	9.84 (+35.09%)	23.19 (+82.72%)	26.53 (+94.64%)	-0.73 (-2.59%)
CO2	0.69 (-1.03%)	7.42 (-11.02%)	13.76 (-20.43%)	15.41 (-22.87%)	-0.28 (+0.42%)
W	-3.96 (-1.18%)	-10.99 (-3.28%)	-14.63 (-4.36%)	-16.20 (-4.83%)	0.17 (+0.05%)

Notes: CS = Consumer surplus, PS = Producer surplus, G = government net expenditure concerning automobiles, CO2 = Social cost of CO_2, and W = Total welfare. Each cell reports the level difference (in billions EUR) with the percentage change in parentheses below. All differences are relative to the observed 2024 equilibrium.

Table 7: Percent change in firm profits relative to the observed 2024 equilibrium

Firm	$T_{2025} = 93.6$	$T_{2030} = 49.5$	$T_{2035} = 11.5$	$T_0 = 0$	No target
BMW	2.79	6.70	0.39	-1.54	-0.12
BYD	111.93	96.79	77.77	71.76	7.80
Ford	2.22	-6.22	-10.09	-11.38	-1.35
Honda	1.21	-10.95	-15.18	-16.55	-0.39
Hyundai	1.51	-7.26	-12.41	-14.00	-1.15
Jaguar	-5.45	-9.37	-11.65	-12.45	19.60
Kia	-2.70	-12.99	-18.17	-19.84	-0.25
Leapmotor	33.74	132.66	93.56	83.00	-2.65
Mazda	0.61	-15.32	-18.67	-19.67	-1.18
Mercedes-Benz	3.31	10.52	5.32	4.38	-1.10
Mitsubishi	2.60	-7.32	-13.63	-15.42	-0.75
Nissan	-16.34	-25.00	-28.67	-29.88	-1.11
Polestar	-5.72	37.08	47.43	47.33	-0.69
Renault	-4.41	-21.46	-30.52	-33.07	-1.82
SAIC	2.34	11.85	5.06	2.88	-1.03
Smart	1.03	64.20	74.12	71.79	4.11
Stellantis	0.43	-11.57	-20.53	-23.29	-1.54
Subaru	-6.80	-34.60	-35.24	-35.65	-1.25
Suzuki	-2.43	-25.71	-31.94	-33.90	-1.65
Tesla	25.68	70.61	58.68	53.89	7.18
Toyota	-1.06	-16.61	-22.94	-24.88	-1.83
Volkswagen	-7.67	-16.16	-20.23	-21.44	-0.64
Volvo	5.19	27.95	26.53	25.73	-1.05

Notes: Profits are computed using emission-adjusted marginal costs $\tilde{c}$ and thus are net of any EU fines or intra-pool transfer payments. Each entry is $100 \times (\pi_f^{cf} - \pi_f^{obs})/\pi_f^{obs}$ for firm f.

Table 8: EU average emissions under alternate 2024 targets

Alternate 2024 emissions target	EU average emissions	Difference from target
93.6	93.51	-0.09
49.5	79.25	+29.75
11.5	79.20	+67.70
0	78.92	+78.92
N/A	105.91	—

Notes: EU average emissions reflects the sales-weighted average number of grams of CO2 emitted per kilometer by new cars sold in France, Germany, and Italy.

Table 9: Pool average emissions and distance from target, by counterfactual

Pool	$T = 93.6$ g/km		$T = 49.5$ g/km		$T = 11.5$ g/km		$T = 0$ g/km	
	Average emissions	Difference from target	Average emissions	Difference from target	Average emissions	Difference from target	Average emissions	Difference from target
BMW	90.1	-3.5	51.9	2.4	54.9	43.4	55.0	55.0
HYUNDAI MOTOR EUROPE	93.6	0.0	78.4	28.9	76.6	65.1	76.0	76.0
JAGUAR	128.7	35.1	136.0	86.5	140.4	128.9	141.7	141.7
KIA	93.6	0.0	88.4	38.9	86.4	74.9	85.6	85.6
MERCEDES-BENZ	76.5	-17.1	49.5	0.0	43.7	32.2	43.7	43.7
MITSUBISHI	93.6	0.0	90.5	41.0	89.7	78.2	89.5	89.5
NISSAN-BYD	93.6	0.0	92.0	42.5	94.0	82.5	94.1	94.1
RENAULT	93.6	0.0	81.2	31.7	87.8	76.3	87.9	87.9
SAIC	92.1	-1.5	49.5	0.0	46.7	35.2	46.6	46.6
TESLA	93.6	0.0	74.5	25.0	73.8	62.3	73.5	73.5
VOLKSWAGEN	96.6	3.0	98.2	48.7	98.0	86.5	97.5	97.5

Notes: Pool average emissions are in g/km of CO_2. The observed 2025 pooling structure is used for every emission target counterfactual (save the "no target" counterfactual).

Table 10: Change in government expenditure relative to the observed 2024 equilibrium

Government	**Expenditure**	**Difference from observed**				
		2025 target	2030 target	2035 target	0 g/km target	No target
France	VAT	195.8 (+2.3%)	249.1 (+2.9%)	9.7 (+0.1%)	-89.6 (-1.1%)	17.4 (+0.2%)
	Malus	-290.6 (-27.3%)	-512.6 (-48.2%)	-596.2 (-56.1%)	-619.3 (-58.2%)	-5.4 (-0.5%)
	Ecobonus	-1,073.5 (+74.2%)	-1,992.8 (+137.8%)	-1,602.3 (+110.8%)	-1,495.4 (+103.4%)	-319.6 (+22.1%)
Germany	VAT	105.4 (+0.8%)	644.3 (+5.0%)	1,134.1 (+8.7%)	1,274.7 (+9.8%)	19.9 (+0.2%)
Italy	VAT	232.9 (+3.3%)	728.3 (+10.2%)	821.8 (+11.5%)	830.6 (+11.7%)	101.4 (+1.4%)
	EV/PHEV subsidy	-759.5 (+309.7%)	-1,251.1 (+510.1%)	-1,060.8 (+432.5%)	-1,054.7 (+430.0%)	-425.8 (+173.6%)
EU	Emission target fines	345.2 (+304.7%)	11,970.4 (+10567.3%)	24,479.0 (+21609.8%)	27,681.2 (+24436.7%)	-113.3 (-100.0%)
Total		-1,244.4 (-4.4%)	9,835.7 (+35.1%)	23,185.3 (+82.7%)	26,527.6 (+94.6%)	-725.2 (-2.6%)

Notes: Each cell reports the level difference between the counterfactual and observed (in millions EUR) and the percentage change in parentheses below. Negative entries for the ecological bonus and the EV/PHEV subsidy denote larger outlays. Total denotes the aggregate change in government revenues and expenditures across all authorities (Germany, France, Italy, and the EU).