Reconciling Environmental Goals and Progressivity: The Income Cap Design for U.S. Electric Vehicle Subsidies [*]

Eva Wang[†]

This version: November 6, 2025
Click here for the latest version

Abstract

This paper examines the U.S. federal EV subsidy's income cap, introduced under the 2022 Inflation Reduction Act, focusing on its environmental and redistributive effects. Using variation in exposure to the income-cap across ZIP-codes, I first estimate that a 22-percentage-point drop in subsidy eligibility reduces EV registrations by 21 percent, implying a price elasticity of 2.9 for the marginal group. I then develop an optimal-tax framework for income-capped subsidies and derive cap-conditional sufficient statistics for the jointly optimal per-unit subsidy and income cap. The optimal subsidy rate is larger when higher marginal external benefits coincide with greater price responsiveness, and when subsidy take-up is inversely related to income; the cap itself determines which households enter these moments. Calibrating a homothetic CES demand block, I construct an equity–environment frontier: as inequality aversion rises, the optimal cap moves down and the optimal subsidy falls; total environmental benefits decline, yet net fiscal efficiency can increase because inframarginal transfers are concentrated among high-income, high-adoption bins. The framework generalizes to other income-capped clean-energy programs and yields implementable rules for jointly choosing generosity and eligibility.

---

[*] I am grateful to my advisors Joel Slemrod, Catherine Hausman, James Hines Jr., and Damian Vergara for their invaluable guidance on this project. This paper has greatly benefited from conversations with Owen Kay, Wojciech Kopczuk, Jing Li, Gerardo Sanz-Maldonado, Candice Wang, Mazhar Waseem, Niaoniao You, as well as many participants in the Public Finance Seminar of the University of Michigan, UM/MSU Environmental and Energy Economics Workshop, National Taxation Association Conference, and Michigan Tax Invitational. All errors are my own.

[††] Department of Economics, University of Michigan. Email: yiyunw@umich.edu.

1 Introduction

Electric-vehicle (EV) subsidies are a cornerstone of U.S. climate policy, but have been criticized for favoring high-income households. Because early adopters are disproportionately wealthy, broad subsidies risk being both fiscally costly and distributionally regressive. The Inflation Reduction Act (IRA) of 2022 addressed these concerns by introducing income-based eligibility limits: individuals with adjusted gross income above $150,000 (or $300,000 for joint filers) can no longer claim the federal EV credit. While this reform makes the program more progressive, its efficiency and environmental implications are not *a priori* clear. If high-income households are the most price-responsive, restricting their eligibility could reduce adoption among marginal buyers and thereby lower the policy's environmental impact. If, however, their purchases would have occurred irrespective of the subsidy, the income cap could improve both equity and fiscal efficiency by reducing fiscal transfers to inframarginal high-income buyers.

To illustrate the distributional concern that motivated the reform, Table 1 reports the distribution of federal EV tax credits across income groups in Tax Year 2020, using data from the IRS Statistics of Income. More than half of all credits were claimed by households with adjusted gross income above $200,000, although these households accounted for less than 10 percent of all taxpayers and only about one-third of total income reported. This concentration of benefits among high-income taxpayers raised concerns that the policy no longer aligned with its original objectives. For instance, one Los Angeles Times op-ed (2024) argued, "the federal tax credit for electric-vehicle purchases has far outlived its purpose and now stands as a glaring example of government overreach and economic inequity."

Table 1: Tax Credits by Size of Adjusted Gross Income, Tax Year 2020

AGI Range	**# of Returns**	**Thousands of $**	**Cumulative %**
Under $15,000	5,005	2,147	1
$15,000–$40,000	4,155	7,160	3
$40,000–$75,000	9,302	35,931	14
$75,000–$100,000	8,042	46,201	29
$100,000–$200,000	12,851	66,006	50
$200,000–$500,000	15,708	111,314	85
$500,000 and above	6,729	44,359	100

Note: Table reports the distribution of tax credits by adjusted gross income brackets for tax year 2020. *Source:* IRS Statistics of Income (SOI).

These facts underscore the tension that motivates this paper: EV subsidies designed as uniform incentives are simple and transparent, but in practice they direct a disproportionate share of fiscal resources to wealthy early adopters. The income cap introduced under the IRA represents a deliberate shift toward redistributive targeting within a corrective policy instrument, providing a natural experiment for studying how equity and efficiency consequences interact.

This paper quantifies these trade-offs and develops a welfare framework for optimal subsidy design under income-based eligibility. Empirically, I exploit variation in income composition across ZIP codes to estimate how EV adoption responded to the introduction of the IRA income cap. The analysis combines new-vehicle registration data from 2017 to 2024 with vehicle characteristics and local income distributions across several states. Using a Poisson pseudo-maximum-likelihood specification with ZIP and time fixed effects, I compare areas with higher and lower shares of high-income households before and after the cap took effect in January 2023. The estimates indicate that ZIP codes with larger high-income shares experienced a 21 percent decline in EV registrations relative to lower-income areas, corresponding to a price elasticity of roughly 2.9 among affected buyers. This sizable response shows that upper-income consumers remain sensitive to price incentives.

Motivated by these findings, I develop a welfare framework that formalizes the design of income-capped subsidies when externalities and responsiveness vary across the income distribution. Standard Pigouvian logic prescribes a uniform subsidy equal to the marginal external benefit, assuming that homogeneous responsiveness and redistributive concerns are handled entirely through the income tax. In practice, governments often embed equity objectives directly into commodity subsidies, making eligibility thresholds an active policy lever. The framework introduced here treats the income cap as an additional margin of design: it determines which households receive the subsidy and, therefore, which groups enter the key welfare-relevant moments. The model delivers implementable sufficient-statistics formulas for the jointly optimal per-unit subsidy and income threshold, clarifying how corrective and redistributive forces interact in shaping the optimal policy.

To put empirical flesh on the theory, I calibrate a homothetic CES demand block using empirically observed adoption rates, externalities inferred from household travel data, and elasticity estimates consistent with the reduced-form results. The simulations reveal how welfare varies with the income threshold and subsidy generosity. When social weights are neutral and external benefits rise with income, the welfare-maximizing policy is broad and generous: expanding eligibility improves efficiency because high-income adopters drive more and respond strongly to prices. When social weights tilt toward lower-income households, the optimal policy becomes more targeted: both the subsidy and the cap fall, reducing total adoption but improving aggregate welfare by lowering fiscal costs and focusing transfers on higher-weight households.

These simulations clarify why the observed policy debate cannot be reduced to whether

EV subsidies are progressive or regressive in a purely distributional sense. The welfare effect depends on the joint distribution of responsiveness, external benefits, and social weights, all of which vary systematically with income. The model provides a tractable structure to interpret these empirical objects and to evaluate how much efficiency loss, if any, accompanies more progressive designs.

Although the analysis focuses on EVs, the logic extends to other clean-energy programs that use income-based eligibility, such as rooftop-solar rebates, home-efficiency upgrades, or electric-heat-pump incentives, where early adopters are typically the wealthy. In each case, policymakers must balance broad generosity against targeted efficiency within a single instrument. By linking observable behavioral heterogeneity to a transparent welfare accounting, this paper provides a unified framework for understanding and designing such second-best environmental policies.

This paper makes a few contributions to the literature. First, theoretically, I provide an optimal subsidy formula in the context of preference and externality heterogeneity. Imposing an income cap on a purchase subsidy departs from the canonical Atkinson and Stiglitz (1976) result that, with an optimal nonlinear income tax in place, commodity taxes should be uniform across goods and not used as instruments for redistribution. However, commodity subsidies with income caps are particularly prevalent in clean energy policies, where early adopters are often the wealthy. The introduction of income caps reflects policymakers' redistributive motive within the commodity subsidies, rather than through external income taxation, making it essential to address the redistributive gains within the framework.

Second, my paper contributes to the topic of environmental justice and the distributional consequences of clean energy policies (Muehlegger and Rapson, 2022; Boomhower and Davis, 2014; Borenstein and Davis, 2024; Holland et al., 2019). While existing papers primarily examine the heterogeneous effects of a uniform subsidy, my paper instead focuses on a policy with an income-based eligibility threshold, introducing a subsidy differential across income groups. By studying it in an optimal taxation framework, this approach identifies key empirical moments that are crucial in different empirical settings for informing the optimal design of subsidy thresholds and levels.

Third, this paper contributes to the literature by examining the trade-off between environmental outcomes and progressivity. While means-tested transfer programs with eligibility thresholds such as the Supplemental Nutrition Assistance Program (SNAP) are extensively studied in public finance, they typically do not involve goods with heterogeneous externalities (Chetty, Friedman, and Saez (2013)). Conversely, the literature on optimal corrective taxation and environmental subsidies, such as Pigouvian taxes, primarily focuses on efficiency and externality correction, with limited attention to redistribution within the policy. This paper bridges these two strands of literature by studying an income-capped subsidy for EV.

Notably, Allcott, Lockwood, and Taubinsky (2019) study the optimal design of sin taxes, incorporating both internalities and externalities. This paper more directly addresses the design of income caps, a growingly common policy instrument in clean energy programs, rather than focusing on a general, uniform tax setting where an external non-linear income tax is also available. Moreover, this paper explicitly incorporates heterogeneous externalities across income groups, highlighting how this heterogeneity affects the optimal design of income-based subsidy thresholds. This paper is closest to Xing, Leard, and Li (2021), which simulates the environmental impacts of restricting subsidies to low-income households, taking into account the non-random replacement of gasoline vehicles. While their analysis focuses on counterfactual emission outcomes and compares it with a universal subsidy, this paper models the optimal income eligibility threshold and subsidy level within an optimal taxation framework and provides insights into the key determinants of the subsidy level. Thus, this model can be generalized to welfare analyses in other contexts where income caps are applied, particularly in cases where there is reason to suspect that regressivity may generate positive externality impacts.

Finally, the estimates in this paper add to the growing body of literature on consumer behavioral responses to EV subsidies and the evaluation of the subsidy programs (Sallee, 2011; Gallagher and Muehlegger, 2011; Jenn, Springel, and Gopal, 2018; Sheldon and Dua, 2019; Xing, Leard, and Li, 2021; Muehlegger and Rapson, 2022; Lohawala, 2023). While much of the existing work focuses on the overall effectiveness of EV subsidies in driving adoption and reducing emissions, my paper provides new evidence on how consumer responses vary with income-based eligibility criteria. By examining both the intensive and extensive margins of EV adoption, it sheds light on how eligibility thresholds affect not only aggregate adoption rates but also the composition of EV buyers, with implications for both environmental outcomes and distributional equity.

The rest of this paper is organized as follows. Section 2 discusses the policy background of the US electric vehicle subsidies. Section 3 describes the data. Section 4 examines the effect on new EV adoption of introducing income eligibility, using a reduced-form continuous difference-in-difference model. Section 5 examines the heterogeneous externality across income groups. Section 6 introduces the theoretical welfare framework and derives the optimal subsidy formula. Section 7 calibrates the theoretical model and discusses the trade-off. Section 8 concludes.

2 Policy Background and Variations

The history of EV subsidies in the United States reflects significant changes over time, driven by evolving federal and state-level policies. At the federal level, the Energy Improvement and Extension Act of 2008 introduced a tax credit program for plug-in hybrid electric vehicles (PHEVs) and battery electric vehicles (BEVs). Effective for purchases made after December 31, 2009, the program offered nonrefundable credits of $2,500 per vehicle, with an additional $417 per kilowatt-hour of battery capacity exceeding 4 kWh, capped at $7,500 (IRS 2009). This structure generated variation in the subsidy amount based on vehicle characteristics, with BEVs generally qualifying for higher credits due to their larger battery capacities.

The federal EV tax credit program featured a manufacturer-specific phaseout mechanism, which introduced additional variation over time. This phaseout was activated once a manufacturer reached 200,000 qualifying vehicle sales for U.S. use. After reaching this threshold, the full credit remained available for the quarter in which the 200,000th vehicle was sold and the following quarter. It then phased down to 50% of the original amount for the next two quarters, followed by a reduction to 25% for the subsequent two quarters, after which the credit fully expired. For instance, Tesla reached the 200,000-vehicle threshold in July 2018, resulting in a phaseout that reduced the credit to $3,750 by January 2019, $1,875 by July 2019, and $0 by the end of 2019 (IRS 2018). General Motors reached the threshold in November 2018, with its credits expiring entirely by April 2020 (IRS 2019). Other manufacturers, whose sales volumes did not reach the threshold, continued to offer the full credit, contributing to temporal differences in subsidy availability across brands.

The Inflation Reduction Act (IRA) of 2022 introduced significant changes to the federal EV subsidy structure, altering the timeline and eligibility criteria for the tax credit program. Announced in August 2022 and signed into law on August 16, 2022, the IRA replaced the manufacturer-specific phaseout with a uniform marketwide deadline of December 31, 2032. Under the new framework, all manufacturers, including those that had previously exhausted their credits (e.g., Tesla and General Motors), regained eligibility, provided their vehicles meet additional criteria.

The updated credit, known as the Clean Vehicle Credit, introduced several new requirements effective at different times. As of August 16, 2022, to qualify for the credit, vehicles had to be assembled in North America, a provision that immediately eliminated the eligibility of some models.[1] Starting January 1, 2023, income caps were introduced, limiting eligibility

---

[1] For analysis of industrial policy aspect of the IRA 2022, see Allcott, Reigner-Kane, et al., 2024. The authors find that the "Buy American" policy goal decreases global carbon emissions but uses profit shifting to decrease foreign producer surplus. My analysis focuses on the income-based eligibility aspect of the IRA.

to individual taxpayers with modified adjusted gross incomes below \$150,000, heads of household below \$225,000, and joint filers below \$300,000. Additionally, vehicle price caps were set, restricting eligibility to cars priced under \$55,000 and vans, SUVs, and trucks priced under \$80,000. Further provisions, including battery component and critical mineral sourcing requirements, were scheduled to phase in starting in mid-2023, adding more layers to vehicle eligibility criteria.

State-level incentives also contributed to temporal and model-specific variation in subsidies. These programs, which include rebates, tax credits, and exemptions from fees or taxes, frequently undergo policy revisions. Changes in eligibility criteria, subsidy amounts, and expiration timelines are common, reflecting adjustments to state budgets, fiscal considerations, and local policy priorities. This dynamic policy environment introduces significant heterogeneity in subsidy availability over time, across geographic regions, and for specific vehicle models. For example, New York introduced its Drive Clean Rebate program in 2017, offering up to \$2,000 for the purchase of eligible EVs. The rebate amount varies by vehicle type and electric range, with higher rebates for fully electric vehicles compared to plug-in hybrids. Over time, adjustments to the program's budget and vehicle eligibility criteria have influenced the number of vehicles qualifying for the subsidy. Similarly, California's Clean Vehicle Rebate Project (CVRP) offers rebates of up to \$7,500, and the program has undergone multiple revisions to incorporate income caps, prioritize low- and middle-income applicants, and adjust rebate amounts based on vehicle type and price. Other states, such as Colorado, have shifted their policies over time as well. Colorado initially provided flat-rate EV tax credits but transitioned to a tiered system in 2021, linking credit amounts to vehicle battery size and price.

The combination of federal and state-level policies generates variation in subsidy exposure over time and across vehicles, driven by factors such as policy expiration, eligibility criteria, and vehicle characteristics. I will explore these variations in my empirical analyses for demand estimations.

3 Data

3.1 Data Sources

The data for this paper come from multiple sources.

The vehicle sales data consist of new vehicle registrations from 2018 to 2024 across seven states: Colorado, Connecticut, Maine, Minnesota, New Mexico, New York, and Oregon. These are the states that publicly release disaggregated registration records at the model level with consistent coverage over time, enabling construction of a panel suitable for analysis. The reporting frequency varies by state, ranging from daily to quarterly. Each record represents a unique combination of model year, make, model, and fuel type. New EV registrations serve as a proxy for vehicle sales, as registration typically occurs within one to two weeks of purchase under state requirements. While these seven states together account for less than 20 percent of the national EV market, California alone represents about 35 percent. Given California's outsized influence on nationwide adoption, this dataset still captures a substantial share of the market outside California and provides a representative view of adoption patterns across other states. Although registrations are available as early as 2010, the main analysis focuses on data from 2018 onward, when the EV market became more developed and a wider range of models became available.

Car Specifications: I obtain model-year level characteristics from the Wards Intelligence Data Center. Although I observe vehicle characteristics at the trim level, registrations are only available at the make-model-fuel level. I use the characteristics and prices of the **base model** to match to the registrations to avoid bias from certain models having a higher number of high-end trim versions. Vehicle characteristics I consider are Manufacturer Suggested Retail Price (MSRP), SUV dummy, and battery range.

Prices: While average transaction prices would be ideal for estimating precise pass-through effects and tax incidence, such data are not readily available. To approximate the prices faced by consumers, I combine Manufacturer's Suggested Retail Price (MSRP) data with manufacturer-specific retail discounts from Automotive News and federal and state-level subsidies. Retail discounts vary by week and model, while subsidies differ across models and over time. For direct-to-consumer manufacturers like Tesla, I manually collect historical price changes from archived versions of Tesla's website and contemporaneous news reports. Assuming no systematic differences in consumers' bargaining power, federal and state subsidies provide valid instruments for consumer prices in my demand estimation. All prices are deflated using the Consumer Price Index from the Bureau of Labor Statistics.

Externalities/Carbon Emission: I calculate carbon emissions/externalities using data from the National Household Travel Survey (NHTS) 2017 and 2022. The survey provides detailed information on vehicle usage patterns, including fuel efficiency of the vehicle owned by the household, miles driven by each vehicle, as well as household demographics. By leveraging these data, I estimate the variation in environmental costs associated with vehicle emissions across different income groups. This allows me to assess how the environmental benefits of EV adoption differ depending on consumer incomes.

3.2 Summary Statistics

Table 2: Summary Statistics of EV Sales and Vehicle Characteristics, 2018–2023

	2018		2019		2021		2022		2023	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Dealer Incentive	219	686	265	836	508	1505	173	619	320	902
MSRP	47,855	23,163	49,021	24,433	48,328	13,637	55,205	16,806	55,880	16,593
Driving Range	207.6	48.1	216.4	54.8	238.3	40.2	258.1	45.3	267.3	40.8
Is SUV	0.37	0.48	0.28	0.45	0.23	0.42	0.54	0.50	0.62	0.49
Buyer Net Price	45,308	18,861	44,679	16,573	43,437	12,406	50,407	17,203	51,309	15,631
Sales	165,737		309,333		825,933		1,288,627		2,101,962	

Notes: This table reports sales-weighted averages of EV characteristics across seven states (Colorado, Connecticut, Maine, Minnesota, New Mexico, New York, and Oregon) from 2018 to 2023. EV sales are measured as new vehicle registrations. "Buyer Net Price" reflects MSRP minus dealer incentives. "SUV Dummy" equals one for sport-utility vehicles. Dealer incentives and prices are expressed in real 2023 dollars.

Table 2 summarizes the EV sales and the sales-weighted car characteristics from 2018 to 2023 in my sample. Overall, EV sales increased tremendously over this period. The average driving range increased from 207.6 miles in 2018 to 267.3 miles in 2023, which indicates in part an improvement of the battery technology. There is also a shift in consumer preferences toward larger vehicles. The proportion of SUVs rose from 0.37 to 0.62 over the same period.

Buyer (discount-inclusive) prices remained relatively stable from 2018 to 2021, averaging around $44,000, before rising to over $51,000 in 2023, reflecting both changes in product mix and cost structures. Averaged dealer incentives peaked at $508 in 2021, but with large variation over the years.

4 Impact of the Income Cap on EV Sales

This section investigates how the income cap introduced by the Inflation Reduction Act (IRA) of 2022 has affected EV adoption. To identify the policy's impact, I exploit variation in income distribution across zip codes, focusing on how areas with different income compositions responded to the policy change. The identification strategy relies on the assumption that, despite differences in baseline EV adoption levels, zip codes with varying income composition exhibited similar underlying trends prior to 2023.

Figure 1 shows the sales trajectories of electric vehicles across zip codes with varying income levels. The first vertical line marks the announcement date of August 16, 2022, and the second vertical line represents the effective date of the income cap on January 1, 2023. Despite differences in income composition, the overall trajectory of EV sales is similar across different zip codes. Broad macroeconomic factors, such as supply chain disruptions from the COVID-19 pandemic and the global chip shortage, appear to have affected all zip codes similarly. Given this evidence, I apply a two-way fixed effects (TWFE) model, using the lower-income zip codes as the control group. Importantly, the parallel trends assumption holds under various income definitions, including the percentage of households above the $200,000 AGI threshold and the mean of the zip code incomes (both by quantiles), and remain robust under alternative measures such as median income and the $150,000 threshold (see Appendix A.1). This robustness across different income measures strengthens the credibility of the empirical strategy, and indicates that the results are not sensitive to how income is defined. [2]

Figure 1: Parallel Trend Across Zip Codes with Varying Income

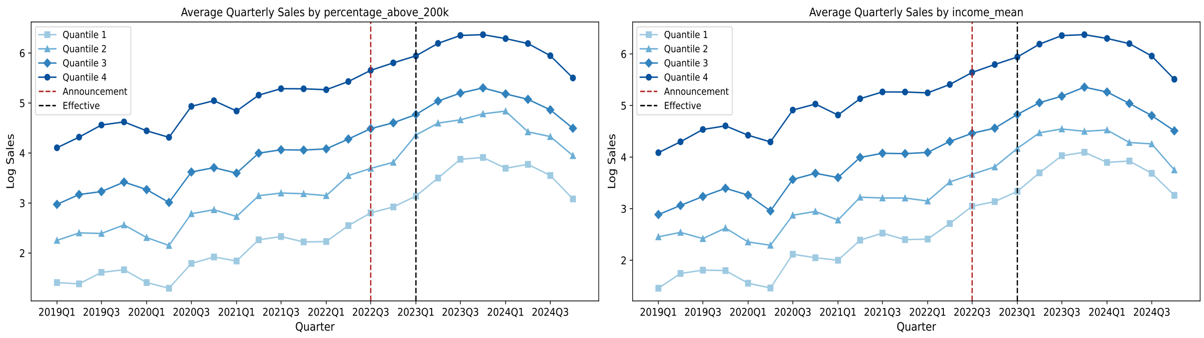

(a) Sales in zip code, by percentage of households above the $200,000 AGI cap

(b) Sales in zip code, by zipcode mean income

4.1 A Poisson Pseudo Maximum Likelihood Design

The dependent variable in the TWFE specification is the number of new EV registrations in each ZIP code and period. These data are counts, with a large share of zeros in rural areas and early periods of the sample. Standard log-linear OLS regressions require transforming the outcome as $\log(Q_{it} + c)$ for some constant c (or using related transformations such as $\text{asinh}(Q_{it})$), but this approach can give biased and scale-dependent results (Chen and Roth,

---

[2] There is an "EV lease loophole" that allows high-income consumers to access the federal EV credit because leased vehicles are treated as commercial sales and thus exempt from income and vehicle price limits. The credit is claimed by the lessor, who passes partial or full credits on through lower lease prices or payments. The registration data this paper uses includes leased vehicles. Recent evidence Allcott, Reigner-Kane, et al., 2024 shows high pass-through to consumers and a shift from purchases to leasing.

2024). Intuitively, adding a constant shifts small counts much more than large counts, so the transformation disproportionately affects low-adoption ZIPs and produces estimates that depend arbitrarily on the chosen constant. To address this issue, I follow Chen and Roth (2024) in estimating a Poisson Pseudo Maximum Likelihood (PPML) model. PPML naturally accommodates zeros, yields coefficients interpretable as semi-elasticities of expected counts, and remains consistent under general forms of heteroskedasticity. The PPML specification is as follows:

$$\mathbb{E}\big[Q_{it} \mid X_{it}\big] = \exp\Big(\alpha + \delta_i + \gamma_t + \beta_{\text{PPML}}\,\text{Treat}_i + \zeta\, X_{it}\Big)$$

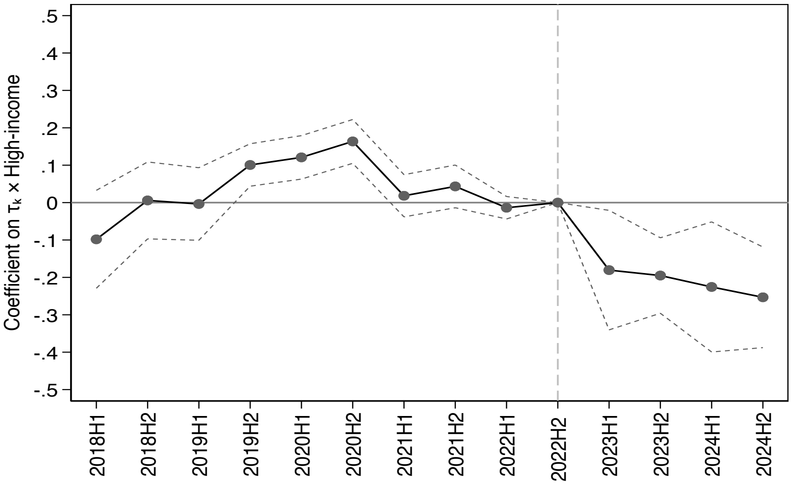

Figure 2: PPML Binary Treatment

Notes: Each point represents a coefficient from the PPML event-study regression of new EV registrations on event-time indicators interacted with a high-income dummy (ZIP codes above the median share of households with AGI >$200,000). The dependent variable is the number of new registrations per ZIP code and half-year period. 95% confidence intervals are shown.

where Q_{it} is number of new EV registration in ZIP i in time t, δ_i and γ_t are ZIP and time fixed effects. To allow for departures from a strictly linear response in ZIP-level income composition, I use a binary treatment indicator in the preferred PPML specification. The indicator equals one for ZIP codes whose pre-treatment share of households with AGI above $200,000 exceeds the sample median, and zero otherwise.[3] The continuous-share specification, which exploits variation in the intensity of exposure, is presented later as a robustness check. I recover β_{PPML} by maximizing the pseudo–log-likelihood. The binary treatment event study result is shown in figure 2. The corresponding regression estimates for all pre and post periods are presented in Table A.2. Column (1) reports coefficients from the PPML event-study regressions with ZIP- and half-year fixed effects. Column (2) allows for state-specific linear time trends, capturing

---

[3] ZIP code–level income data are top-coded at $200,000, so the $200,000 threshold reflects the highest observable bin in the ACS. While the federal income cap for EV credit eligibility is $150,000 for single filers and $300,000 for joint filers, over 70% of EV subsidy claimants are married couples, making the $200,000 cutoff a reasonable proxy for identifying ZIP codes likely affected by the policy.

gradual differences in within-state EV adoption unrelated to the policy change. The inclusion of state trends modestly improves model fit without altering the magnitude and pattern of post-policy effects. Overall, the point estimates range from roughly –0.18 to –0.25 across post periods, and the average post-IRA effect is about –0.21 (s.e. = 0.07).

The estimated β_{PPML} implies an approximately 21 percent larger decline in registrations in high-income ZIPs relative to low-income ZIPs (controls) following the income-cap implementation.[4]An event-study in the PPML context confirms no pre-policy differential trends. [5]

Transforming the Aggregate Coefficient into a Micro-level Price Elasticity of Demand: The estimated coefficient of –0.21 reflects the effect of a 22 percentage point difference in exposure, comparing ZIP codes above versus below the sample median share of \$200,000+ households. On average, 28 percent of households in treated ZIP codes (above the median) lost eligibility due to the income cap, compared to just 6 percent in control ZIP codes (below the median). Because this is an aggregate regression, the coefficient does not represent the effect of moving from full eligibility to complete ineligibility at the individual level. Rather, it captures the semi-elasticity of EV registrations with respect to the average change in the share of ineligible households across ZIP codes.

Importantly, this -0.21 effect is defined relative to total registrations in the ZIP (Q_{total}), even though only a subset of buyers (high income) actually faced a price change. Simply dividing by the household treatment share (0.28-0.06=0.22) would implicitly assume that EV purchases are proportional to household counts, which is unlikely because higher-income households are more likely to purchase EVs. To correct for this, let

$$s = \frac{Q_{\text{treated}}}{Q_{\text{total}}}$$

denote the sales share attributable to treated households. Then,

$$\frac{\Delta Q_{\text{treated}}}{Q_{\text{treated}}} = \frac{\Delta Q_{\text{total}}}{Q_{\text{total}}} \Big/ s,$$

and the implied own-price elasticity is

$$\varepsilon = \frac{\Delta Q_{\text{treated}}/Q_{\text{treated}}}{\Delta p/p} = \frac{-0.21}{s \cdot (\Delta p/p)}.$$

---

[4] The estimated β_{PPML} is –0.21, which translates into an approximately $e^{-0.21}-1 \approx -19\%$ decline in expected registrations. For ease of exposition, I refer to this effect as a 21 percent decline, noting that the difference between the exact and approximate interpretation is negligible.

[5] A modest uptick appears during the COVID-related chip shortage, likely reflecting temporary supply prioritization of high-income ZIP codes. Industry evidence confirms that automakers prioritized chip-limited vehicles with higher profitability (e.g., Hyundai allocated semiconductors preferentially to high-margin models, see Shih, 2022). Because this was a transitory supply shock that had faded before the IRA income-cap took effect, it does not threaten the identification strategy.

because $\beta_{PPML} = -0.21$ describes $\frac{\Delta Q_{\text{total}}}{Q_{\text{total}}}$, and $\frac{\Delta p}{p} \approx 0.15$ (the average price increase as a result of the removal of subsidy is approximately 15%).[6] Depending on the assumption for s, the elasticity lies in a wide but interpretable range. If one assumes that all EV buyers are in treated households ($s = 1$), the implied elasticity is $-0.21/0.15 \approx -1.4$. This is a lower bound, because it attributes the aggregate decline to the broadest possible base of buyers. At the other extreme, if one assumes that EV adoption is proportional to household counts so that $s = 0.22$, the elasticity rises to $-0.21/(0.22 \times 0.15) \approx -6.3$. This is an upper bound, since it attributes the same aggregate decline to a much smaller base of buyers. Using evidence from the NHTS survey that higher-income households are disproportionately represented among EV adopters (so that treated households account for roughly 47% of baseline EV sales, $s = 0.47$), the preferred estimate lies between these extremes at

$$\varepsilon \approx \frac{-0.21}{0.47 \times 0.15} \approx -2.9.$$

This elasticity estimate pertains specifically to higher-income households and indicates that their EV adoption remains highly responsive to the subsidy removal. Despite being less subsidy-constrained, these households reduce purchases in response to losing eligibility. The magnitude of the estimated elasticity is broadly consistent with existing evidence: for example, Xing, Leard, and Li (2021) estimate model-specific own-price elasticities of EV demand ranging from –2 to –4 across vehicle models, while Springel (2021) reports elasticities between –1.5 and –2.1 in structural simulations.

4.2 Robustness: Addressing Model-Level Eligibility Shifts and Low-Income Entry

In the above preferred specification, the analysis aggregates *all* EV models, despite variation in model-specific subsidy eligibility. Around the policy period, Tesla and GM cars regained eligibility while cars that went through the final assembly process outside of the US (mostly EVs with foreign makers) lost eligibility. Given the complexity of model-level eligibility changes and the unanticipated nature of IRA, it is plausible that many consumers failed to fully account for these variations when making purchasing decisions. Hence, aggregating across models captures the policy's aggregate effect on EV adoption without explicitly modeling each vehicle

---

[6] Before the Inflation Reduction Act, the federal EV tax credit provided up to \$7,500 per vehicle. Nearly all BEVs qualified for the full amount, while plug-in hybrids with smaller batteries typically received partial credits of around \$3,000–\$5,000. Because PHEVs accounted for roughly 20–25 percent of plug-in sales during this period, the sales-weighted average federal subsidy across new EVs was about \$6,500 per vehicle, which is roughly 15% of the full price.

separately. To interpret the estimated coefficient as a well-identified price elasticity of demand of the high-income population, the underlying assumption made is that individuals meeting the income eligibility threshold can substitute an ineligible EV model with an eligible one.

The main threat to identification is that the observed post-policy drop in EV registrations in high-income ZIP codes may not reflect a behavioral response to subsidy loss by high-income households, but instead results from increased purchases by low-income households. This could occur if low-income buyers began entering the market in response to newly-subsidized models such as Teslas. In that case, the estimated decline could be driven by compositional changes in model eligibility rather than by substitution behavior among high-income groups. To address this concern, I re-estimate the TWFE model using a restricted sample that includes only EV models that *lost* eligibility after the income cap was introduced. In this restricted sample, all models faced the same or reduced subsidy generosity following the policy change. Low-income households may still receive partial subsidies, while high-income households lose eligibility entirely. As a result, any remaining differential trend in this sample would likely understate the true response of high-income buyers, because the comparison group (low-income ZIPs) is also partially treated. Indeed, I continue to find a negative and significant effect ($\beta = -0.24$, col 3 in A.2) (detailed event study plot in A.3). This supports the interpretation that the observed decline in high-income ZIP registrations reflects a genuine behavioral response to subsidy removal, rather than a mechanical result of shifting model composition.

4.3 Robustness Checks: Continuous Treatment, Log-Linear OLS, Frequency Aggregation, and COVID period exclusions

Table 3 presents a series of additional robustness checks. First, I replace the binary treatment indicator with a continuous measure of treatment intensity, defined as the pre-policy share of households earning above $200,000 in each ZIP code (spline functions in Appendix A.5). This "dose–response" specification allows for a more flexible estimation of treatment effects under continuous exposure. Following the framework of Callaway, Goodman-Bacon, and Sant'Anna (2024), I interpret the estimated coefficient as an average causal response: specifically, the effect of excluding an additional one percent of the ZIP population from a subsidy worth approximately $6,500. Identification requires the absence of unobserved treatment effect heterogeneity. The estimates are consistent in sign and magnitude with the binary specification, reinforcing the interpretation that ineligibility reduces EV adoption even at the upper end of the income distribution. A more detailed explanation about the interpretation of the coefficient and its limitation is in Appendix A.4.

Second, I test sensitivity to alternative outcome transformations by estimating a log-linear version of the model with $\log(\text{count} + c)$ as the dependent variable. Although this specification permits level-based interpretations, it performs poorly in the presence of zero-inflated and highly skewed outcomes. Indeed, I find that the estimated treatment effect is highly sensitive to the choice of constant c, and in some cases reverses sign. In contrast, the preferred PPML estimator handles zero outcomes naturally and yields consistent results. I therefore treat the log-linear estimates as suggestive and maintain PPML as the baseline specification.

Third, I examine robustness to temporal aggregation in col 3 of table 3. While some states report EV registration data at finer temporal resolutions (e.g., monthly or daily), I aggregate to the quarterly or half-year level to smooth idiosyncratic variation caused by delivery cycles and reporting noise. Re-estimating the model at the quarterly level yields comparable results, suggesting that the core findings are not driven by transitory fluctuations in registration timing.

Finally, I address the possibility that pandemic-related disruptions confound the estimates. I re-estimate the model excluding the 2019H1–2020H1 period, which captures the onset of COVID-19 and associated supply-chain and demand shocks. The results are stable across this restriction.

Together, these checks confirm that the observed decline in EV adoption among high-income ZIP codes reflects a robust behavioral response to subsidy ineligibility. The results are not driven by model composition, temporal noise, functional form assumptions, or COVID-related disruptions.

Table 3: Robustness to Estimation Method, Treatment Definition, and Panel Frequency

	PPML Binary	PPML Ex-C19	PPML Qtrly	PPML CTS	Log OLS
Post × High-INC Binary	-0.244***	-0.234***	-0.220***		-0.702***
	(0.089)	(0.089)	(0.078)		(0.050)
Post × Share > 200k				-1.181***	
				(0.310)	
Observations	58,566	45,574	153,320	58,220	58,678

Notes: Each column reports estimates from a separate regression on post-policy indicators interacted with ZIP-level income (treatment) measures. Column (1) is the preferred baseline PPML specification using a binary treatment indicator equal to one for ZIP codes with an above-median share of households earning over $200k. Column (2) replicates Column (1) but excludes all observations from 2019H1 to 2020H1, the period most affected by COVID-related disruptions. Column (3) uses quarterly data instead of half-yearly. Column (4) replaces the binary treatment with a continuous measure: the ZIP-level share of households earning over $200k, interacted with the post-policy period. Column (5) estimates a log-linear OLS model with log(count + 0.01) as the dependent variable. All specifications include ZIP-code and time fixed effects. Standard errors are clustered at the ZIP × time level.

5 Heterogeneity in EV Adoption and Externality

This section documents empirical patterns from the 2017 National Household Travel Survey (NHTS) that motivate the heterogeneous structure of the model. The NHTS provides detailed information on household vehicles, including annual miles traveled, vehicle type, ownership, and income, which together allow for an assessment of how environmental externalities vary across income groups.

Figure 3 highlights several dimensions of heterogeneity. Panel (a) shows that EV adoption rises steeply with income, consistent with income-based constraints or preferences shaping adoption decisions. Panel (b) reports average miles driven per person, which increase with income before flattening, likely reflecting differences in household size, ride-sharing and rural-urban discrepancies. Panel (c) shows average household CO_2 emissions,[7] which remain higher among richer households despite their ownership of newer, more fuel-efficient vehicles. Panel (d) shows homeownership rates increasing with income, consistent with lower charging frictions for higher-income households. Together, these facts suggest that richer households both drive more and emit more, while also facing lower fixed costs of EV adoption.

Additional evidence, reported in Appendix B, reinforces the central role of income. A logistic regression using the 2017 and 2022 surveys shows that urban households and those in single-family homes are more likely to own an EV, but income is by far the strongest predictor: the probability of EV ownership rises monotonically across income groups. Moreover, EV-owning households tend to have slightly larger vehicle fleets. Appendix Table 6 shows that, conditional on income, EV adoption is associated with about 0.18 additional vehicles per household on average (col. 1). This indicates that EVs are often acquired as additions rather than one-for-one replacements of internal combustion engine vehicles. However, when I allow the EV effect to vary by income (col. 2), the interaction terms are small and imprecise, suggesting that this "second-car" margin is not concentrated exclusively among high-income households.

Taken together, these results underscore the importance of income-based heterogeneity in both adoption and externalities. Higher-income households not only adopt EVs at greater rates, but also contribute disproportionately to vehicle miles traveled and CO_2 emissions. These empirical patterns motivate the model's heterogeneous setup, with a positive slope of household-level externalities E(z) and income-linked differences in adoption costs.

---

[7] Emissions are calculated by multiplying each vehicle's reported annual miles by fuel-specific emission factors (e.g., gasoline $\approx$ 444 g/mi, diesel $\approx$ 463 g/mi), then aggregating to the household level and averaging across income bins using survey weights.

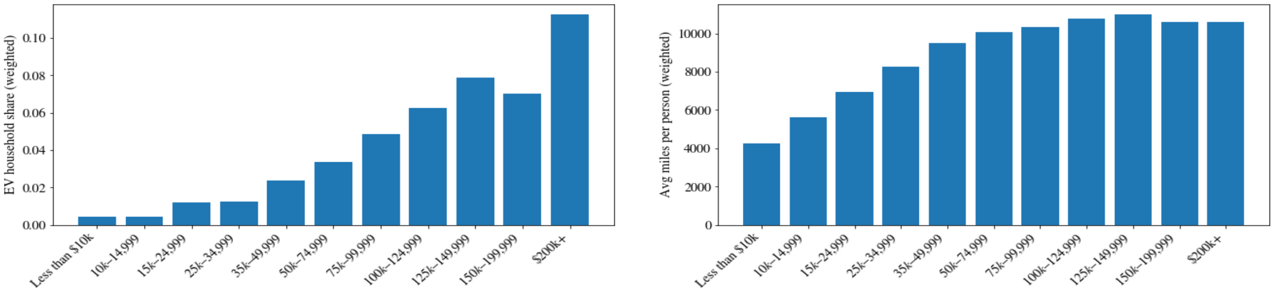

(a) Average EV Adoption Rate

(b) Average Person Mileage Driven

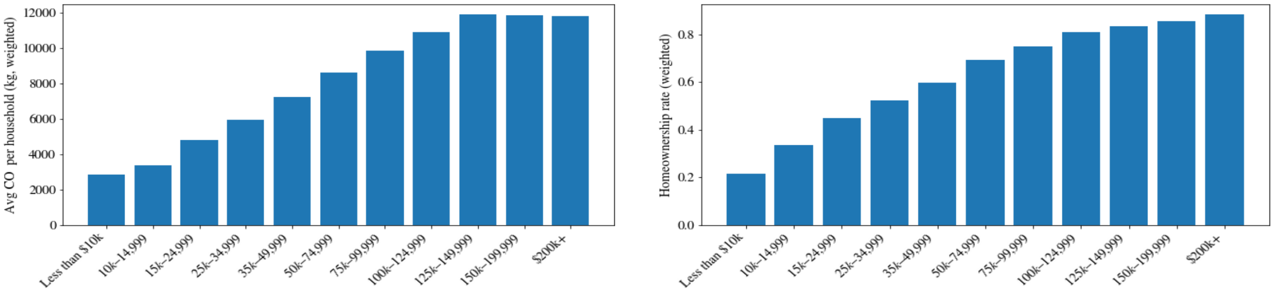

(c) Average Emission per HH in kg of CO2

(d) Percentage of Home-owners

Figure 3: Heterogeneity Stats Across Income Groups

6 Welfare Model

Motivated by two empirical facts, (i) a nontrivial behavioral response to the EV subsidy at the current income threshold, and (ii) sizable heterogeneity in environmental externalities across income groups (descriptive evidence from the NHTS), I develop an optimal tax framework for income-capped commodity subsidies. The framework nests income–dependent externalities and private costs and delivers empirically implementable sufficient statistics for the optimal pair $(t^\star, \bar{z}^\star)$ (per–unit subsidy and income cap). I first clarify why modeling the cap explicitly is essential, and then characterize the dynamics between the two policy instruments.

6.1 Differential Commodity Taxation

Pigouvian policy is a form of differential commodity taxation used for *correction*: prices are adjusted so that private marginal rates of substitution equal social marginal costs inclusive of external damages (Pigou 1920). This corrective role is conceptually distinct from using commodity taxes for *redistribution*. The Atkinson–Stiglitz theorem, and extensions such as Saez (2002) show that, under weak separability between consumption and labor and homogeneous preferences, and with an optimally set non–linear income tax, there is no additional welfare role for non–Pigouvian differential commodity taxation to achieve redistribution. In other words, once Pigouvian adjustments are made to internalize externalities, further differentiation of commodity taxes or subsidies should not be used for equity when the income tax can do that job efficiently (Atkinson and Stiglitz 1976). Intuitively, commodity taxation distorts consumption choices by altering relative prices while failing to improve labor market outcomes.

Income caps on subsidies break this benchmark *in practice*. Policymakers routinely impose caps in clean–energy programs—where early adopters tend to be the wealthy (e.g., EVs, residential solar, home-efficiency upgrades)—to steer resources away from inframarginal high-income adopters and toward lower-income households. In such settings, the subsidy program is effectively asked to shoulder *residual distributional objectives* that are not (or cannot be) accomplished through the income tax. When there is a redistributive motive, and when caps are present and external benefits and private adoption response to subsidy vary with income, a purely Pigouvian rule is incomplete. The income cap and the subsidy rate together govern the welfare calculus, trading off corrective efficiency, fiscal costs, and distributional priorities *within* the commodity instrument. The framework below formalizes these departures and derives implementable sufficient statistics for the optimal cap and subsidy.

In the following framework, I interpret z as *post–tax* income under the existing non–linear income tax, which is held fixed. Because the income tax is not re–optimized here, the planner may retain residual distributional concerns across post–tax incomes. These are captured by Pareto weights $\alpha(z)$ in the social objective; they represent normative priorities over gains rather than Mirrleesian redistribution that would be undone by re–optimizing the income tax. [8]

Notation. A policy is a pair $(t, \bar{z})$, where $\bar{z}$ is the income cutoff (cap) and eligibility is the set $\{z \leq \bar{z}\}$. For each fixed cap $\bar{z}$, let $t^\star(\bar{z})$ denote the subsidy that maximizes welfare holding $\bar{z}$ fixed (the "ridge"). The global optimum is $(t^\star, \bar{z}^\star) \in \arg\max_{t,\bar{z}} W(t, \bar{z})$, with $t^\star = t^\star(\bar{z}^\star)$. Expectations and covariances "among eligibles" use $\mathbb{E}[\cdot \mid z \leq \bar{z}]$ and $\mathrm{Cov}(\cdot, \cdot \mid z \leq \bar{z})$. Let $S'(z) \equiv \partial s(P; z)/\partial t$ for $z \leq \bar{z}$ and $\overline{S'}_{\text{elig}} = \mathbb{E}[S'(z) \mid z \leq \bar{z}]$.

---

[8] I acknowledge the "income–tax–does–all" benchmark that, if the income tax can be freely adjusted, non–tax instruments (including corrective policy) should ignore distribution and leave equity entirely to the income tax (Kaplow 2006b; Kaplow 2006a; Kaplow and Shavell 2000). However, my analysis takes second–best perspective more suited to the reality that income taxes are not optimized this way.

6.2 Individuals

There is a continuum of agents indexed by post–tax income $z \in \mathcal{Z}$ with CDF μ (density $f > 0$ and $\int_{\mathcal{Z}} d\mu = 1$). Each agent allocates disposable resources to a numéraire good c (price 1) and an EV good s. The pre–subsidy EV price is p. The per–unit subsidy is

$$t(z) = \begin{cases} t, & z \leq \bar{z}, \\ 0, & z > \bar{z}, \end{cases} \qquad P(z) = p - t(z),$$

Preferences are given by a monotone, concave utility index $u(c, s)$; in the calibration we use a homothetic CES form, but the theory below requires only standard regularity.[9] A lump–sum rebate R returns monetized climate damages avoided by EV adoption; R is financed within the government's resource constraint and is taken as given by individuals. The problem is

$$\max_{c,s} \; u(c, s) \qquad \text{s.t.} \quad c + P(z)\, s \leq z + R.$$

At an interior solution,

$$u_c(c^*, s^*)\, P(z) = u_s(c^*, s^*), \tag{1}$$

with the usual Kuhn–Tucker conditions if $s = 0$. For reference, write $(c^E(z), s^E(z))$ for the choices when $z \leq \bar{z}$ (eligible, price $p - t$) and $(c^N(z), s^N(z))$ when $z > \bar{z}$ (ineligible, price p). An adopter at income z generates a monetized external benefit $e(z) > 0$ (avoided damages). I do not claim a causal link between income and emissions intensity; $e(z)$ summarizes heterogeneity that the instrument cannot otherwise tag.

6.3 Government

The planner chooses $(t, \bar{z})$ to maximize a social welfare functional with income–specific Pareto weights $\alpha(z)$ and shadow cost of public funds $\lambda > 0$:

$$W(t, \bar{z}) \;=\; \int_{\mathcal{Z}} \alpha(z)\, U\big(P(z); z\big)\, d\mu(z) \;+\; \int_{\mathcal{Z}} e(z)\, s\big(P(z); z\big)\, d\mu(z) \;-\; \lambda\, t \int_{z \leq \bar{z}} s\big(P(z); z\big)\, d\mu(z), \tag{2}$$

where $U(P; z)$ is indirect utility and $s(P; z)$ the EV demand at price P. The second term records environmental benefits; the last term is fiscal cost. The rebate R is a wash in (2). This formulation treats the externality as entering the planner's welfare additively, a representation that is equivalent to the canonical Pigouvian model under the assumption of a common marginal utility of environmental quality. [10]; see Appendix C.3 for more details.

---

[9] At the bin level, s may be interpreted as an adoption share generated by within–bin taste dispersion.

[10] In the canonical Pigouvian formulation, individuals have utility $u(c, s, E)$ and the aggregate environmental state E enters preferences directly.

FOC in t (fixed $\bar{z}$). Using the money–metric envelope $dU/dt = s(P;z)$ and writing $E_z = e(z)$, one obtains the decomposition

$$t^{\star}(\bar{z}) = \underbrace{\bar{E}_{\mathrm{elig}} + \frac{\mathrm{Cov}_{\mathrm{elig}}\big(E_z, S'(z)\big)}{\overline{S'}_{\mathrm{elig}}}}_{\text{corrective}} + \underbrace{\frac{\mathrm{Cov}_{\mathrm{elig}}\big(\Delta g(z), s(z)\big) + \mathbb{E}_{\mathrm{elig}}[\Delta g]\mathbb{E}_{\mathrm{elig}}[s]}{\overline{S'}_{\mathrm{elig}}}}_{\text{redistributive}}, \tag{3}$$

with $\bar{E}_{\mathrm{elig}} = \mathbb{E}[E_z \mid z \leq \bar{z}]$, $\mathrm{Cov}_{\mathrm{elig}}(\cdot,\cdot) = \mathrm{Cov}(\cdot,\cdot \mid z \leq \bar{z})$, and $\Delta g(z) = \alpha(z)u_c/\lambda - 1$, where u_c is the marginal utility of the numéraire evaluated at eligible prices, λ is the marginal value of the public fund, $\Delta g(z)$ is the deviation from the baseline social marginal weight. These weights are endogenous to the tax system, but are useful for characterizing the necessary conditions that must hold at the optimum (Allcott, Lockwood, and Taubinsky (2019)). The detailed proof can be found in Appendix C.1.

Discussion. This representation highlights three ideas. First, the "corrective" component captures the $S'(z)$–weighted average of external benefits among eligible consumers. It combines the standard Pigouvian term, $\bar{E}_{\mathrm{elig}}$, with a targeting correction term, $\mathrm{Cov}_{\mathrm{elig}}\big(E_z, S'(z)\big)$, which measures whether the individuals who generate larger externalities (e.g., higher pollution per vehicle) are also those who respond more strongly to the subsidy (higher adoptions per \$ of subsidy). When this alignment is positive, that is, when the most polluting adopters are also the most price–responsive, the weighted average external benefit exceeds the simple mean, increasing the optimal subsidy. This logic is consistent wtih the findings in Griffith, O'Connell, and Smith (2019).

Second, with a fixed income cap $\bar{z}$, all statistics in the corrective term are compute over the eligible subpopulation; the cap therefore governs the reference population for the relevant averages and covariances. Therefore, empirical estimates of heterogeneity in external benefits and price responsiveness across income groups are essential for the corrective calculus.

Third, the redistributive bracket captures equity gains when the commodity instrument must shoulder residual redistribution: it aggregates social marginal welfare differences *within* the eligible set in exactly the way the program delivers transfers—through who already adopts, $s(z)$, relative to who is induced at the margin via $\overline{S'}_{\mathrm{elig}}$. Put it differently, it answers: how much should we raise or lower the per–unit subsidy so that the marginal fiscal cost on newly induced adoption exactly offsets the equity value of topping up current adopters? Its sign and magnitude depend on the incidence of inframarginal transfers: if adoption is concentrated among rich households, the bracket decreases $t^{\star}(\bar{z})$; under pro–poor weights and homothetic demand (so s rises with income), it is typically negative and lowers $t^{\star}(\bar{z})$. Lowering the cap (i.e. excluding high income population) makes this covariance less negative when high–income households buy more EVs, attenuating the downward pressure on $t^\star(\bar{z})$. [11]

FOC in $\bar{z}$ (optimal cap). A marginal expansion of eligibility at $\bar{z}$ changes welfare by the difference in social surplus when the threshold type is treated as eligible rather than ineligible. At the global optimum $(t^\star, \bar{z}^\star)$,

$$\Big\{\alpha(\bar{z}^\star)\big[U^E(\bar{z}^\star) - U^N(\bar{z}^\star)\big] + e(\bar{z}^\star)\big[s^E(\bar{z}^\star) - s^N(\bar{z}^\star)\big]\Big\} f(\bar{z}^\star) = \lambda\, t^\star\, s^E(\bar{z}^\star)\, f(\bar{z}^\star), \tag{4}$$

where U^E and U^N denote indirect utility at $P = p - t^\star$ and $P = p$, respectively, and s^E, s^N are the associated EV choices. The detailed proof can be found in Appendix C.2.

Discussion. Condition (4) is a local rule at the cutoff: expand the cap if and only if the marginal social gain from making the threshold type eligible (left-hand side) exceeds the marginal fiscal cost (right-hand side). Two cases are instructive. (i) Pure inframarginality ($s^E = s^N$): expanding eligibility confers only a transfer, $t\, s^E$, with no environmental gain. Locally, $U^E - U^N \approx u_c\, s^E\, t$, so (4) reduces to $\alpha(\bar{z}^\star) u_c\, s^E\, t = \lambda\, t\, s^E$, i.e. expand if $g(z) \equiv \alpha u_c/\lambda > 1$ and contract if $g(z) < 1$. Intuitively, when the marginal type would buy regardless, the cap is set by whether a dollar in their hands is socially more valuable than a dollar in the budget. (ii) Behavioral margin ($s^E > s^N$): eligibility also induces adoption, adding $e(\bar{z}^\star)\,[s^E - s^N]$ to the benefit side. Larger responsiveness at the cutoff or higher $e(\bar{z}^\star)$ tilts toward a higher cap; conversely, low responsiveness or low external benefits favors a tighter cap.

Consistent with this logic, the difference-in-differences estimates document a discrete increase in EV adoption at the current threshold—formally, $s^E - s^N > 0$—so the "rich people would have bought anyway" critique does not apply at the current margin. At the same time, the level of s^N around the cutoff is nontrivial, indicating a sizable inframarginal mass and making the redistributive bracket in (9) quantitatively relevant for policy design.

---

[11] Under the Atkinson–Stiglitz conditions with a fully flexible, re–optimized non–linear income tax and separable, homogeneous preferences, the planner sets $g(z) \equiv 1$, so $\Delta g(z) = 0$ and only the corrective component remains. If income is a good tag for EV adoption and external harms, there is an additional efficiency reason to raise the size of the subsidy.

7 Calibration Results: Who to Subsidize and By How Much

This section implements and quantifies the framework in Section 6 by imposing a tractable homothetic CES demand and assigning empirically disciplined values for key moments, tracking how they vary with eligibility. The model imposes structure on preferences and demand responses to deliver a full money-metric welfare accounting—private surplus (via compensating variation), fiscal costs, and environmental benefits. Rather than re-deriving optimality conditions, I use them to: (i) compute the welfare surface and report the empirically grounded optimal subsidy and cap; (ii) run comparative statics (e.g., neutral vs. pro-poor social weights), highlighting which empirical moments shift the balance between corrective and redistributive forces; and (iii) trace an equity–environment frontier as inequality aversion varies. The goal is to deliver realistic orders of magnitude and policy-relevant comparisons while keeping the mapping from assumptions to results transparent.

I use a homothetic CES block for two reasons with important implications. First, it generates EV adoption that is increasing in income via standard income effects, in line with observed empirical patterns.[12] Second, homotheticity implies the marginal utility of income is constant across income at given prices, so any redistributive motive comes only from explicit social weights, not from curvature of utility. [13] This keeps the efficiency–equity comparison transparent. See more details in Appendix D.1 for the features of Homothetic CES.

Calibration discipline (magnitudes and rationale). I discipline substitution and levels using observables, keeping assumptions minimal:

1. **Substitution.** Curvature ρ (elasticity $\sigma = 1/(1-\rho)$) is chosen to match price elasticities around the observed p. The baseline uses $\rho = 0.15$ ($\sigma \approx 1.18$), consistent with modest substitution; robustness exercises consider higher elasticities (e.g., $\sigma \in [1.3, 1.8]$) to reflect recent estimates.[14]
2. **Levels.** The share parameter θ is set to reproduce a baseline EV adoption at (p, y); $\theta = 0.80$ matches a low initial EV share consistent with current data.

---

[12] I do not separately identify income effects from preference heterogeneity across income. If heterogeneity were perfectly taggable, a more efficient policy could target it; I abstract from tagging here.

[13] The curvature still exists between EV and the numarie good due to complementarity of the two goods.

[14] Results are monotone in σ: higher σ raises $S'(z)$ proportionally and steepens the ridge, but does not reverse the comparative statics on targeting.

3. **Income distribution and weights.** Households are grouped into five income bins with midpoints $z = (25, 57, 100, 160, 250)$ (k\$) and population weights $w = (0.20, 0.25, 0.25, 0.20, 0.10)$. $M(z^*) = \sum_{i:\, z_i \le z^*} w_i$. Targeting results are robust to finer binning because, under homothetic CES, s_i and s_i' scale with y_i.

4. **Externalities.** I calibrate the external benefit vector $e(z)$ to increase with income, reflecting higher avoided carbon damages when higher-income households drive more. The calibration is based on a social cost of carbon (SCC) of $150 per metric ton of CO_2 (Rennert et al., 2022) and the observed gradient in annual driving mileage by income from NHTS, together with the per-mile emissions gap between internal combustion and electric vehicles (404g versus 150g CO_2 per mile, U.S. Environmental Protection Agency, 2024). Details of the conversion from emissions to monetary values are provided in Appendix D.3.

5. **Fiscal environment.** The policy instrument is a per-unit subsidy t and an income cap $\bar{z}$. The shadow cost of funds is set to $\lambda = 1$ in the baseline (interpreted as an already-optimized income tax system); changing λ scales t^* proportionally without altering the ranking of caps.

6. **Welfare metric.** Private benefits are valued in *money* via compensating variation (CV). For large policies CV yields slightly smaller t^* than consumer surplus (CS) and larger than equivalent variation (EV); the ridge shape and optimal caps are qualitatively unchanged.

Welfare forces. Two forces drive the results, as described in the derived FOCs in section 6. A *corrective* force rewards subsidizing adopters who deliver larger external benefits and respond more to price. A *redistributive* force tilts resources toward lower-income households when social weights are pro-poor. In homothetic CES, income only matters for adoption and responsiveness; equity enters solely through the social weights.

Benchmarks (Appendix D.2). Benchmark A sets externalities to zero and assigns neutral social weights; the optimal policy is a zero subsidy for any cap. Benchmark B assigns a positive, flat externality with neutral weights; the solution lines up with the classic Pigouvian logic: under consumer surplus the optimal subsidy equals e/λ and the preferred cap is broad; with compensating variation the optimal t is slightly smaller for large policies but the targeting pattern is the same. With those references in hand, I now turn to the two cases of interest.

Case 1: External benefits rise with income, neutral social weights

With equal social weights and an externality profile that increases with income, the welfare-maximizing policy is broad and generous. As the eligibility cap widens, the subsidy that maximizes welfare at that cap rises. Intuitively, higher-income households both generate larger external benefits per adoption and are more responsive in level terms; expanding eligibility directs marginal dollars toward adopters with the highest corrective return. In the calibration, the global maximum occurs at the universal cap (top bin included) with a sizable per-unit subsidy, $t^* \approx \$4{,}340$.[15]

Policy translation. When the objective is purely corrective and external benefits scale with income, means testing is costly: it removes precisely the adopters with the highest per-dollar environmental payoff. Expanding eligibility yields larger gains than pushing the subsidy rate within a narrow cap. As shown in Figure 4, the ridge of welfare-maximizing points is upward sloping, illustrating the comparative statics between the cap and the subsidy size: because empirically higher-income households both pollute more and respond more to prices, including them would pushing up the optimal subsidy size.

Case 2: External benefits rise with income, pro-poor social weights

Introducing pro-poor social weights reverses the slope of the ridge. With the illustrative Pareto weights gradient $(1.4, 1.2, 1, 0.6, 0.4)$. As the cap widens to admit richer households, the welfare-maximizing subsidy falls. The corrective force still points toward broader eligibility, but the redistributive force pulls the other way: high-income adopters have low social weight precisely where take-up is strongest, so broad programs spend heavily where the social value of transfers is lowest. The optimum shifts to a larger subsidy with a higher cap: $z^* \approx \$100{,}000$ and $t^* \approx \$6{,}960$.

Policy translation. With neutral weights, keep eligibility broad and size of subsidy to external benefits. With very progressive social weights, prioritize targeting via a cap or smooth phase-out, and size the subsidy by how e(z) varies with income: if e(z) is concave or only mildly rising, equity and correction reinforce each other—set a higher subsidy than under neutral weights; if e(z) rises sharply with income, lean more on targeting and keep the subsidy smaller than the neutral benchmark.

---

[15] Figures reference: the upward-sloping ridge in the welfare heat map makes the pattern visually transparent.

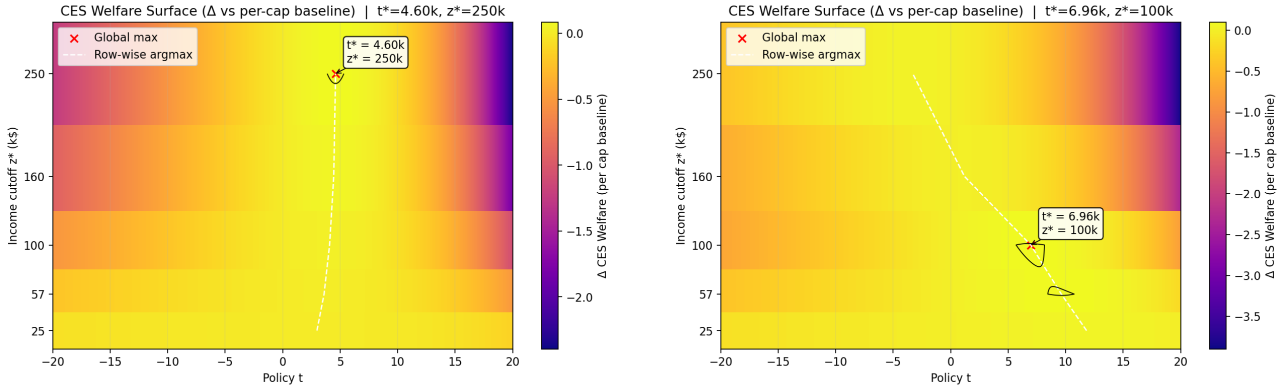

Case 1: Rising externalities, neutral weights. Global max at a broad cap with a high subsidy ($t^* \approx \$4{,}340$, $z^* \approx \$250\text{k}$).

Case 2: Rising externalities, pro-poor weights. Global max at a tighter cap with a smaller subsidy ($t^* \approx \$6{,}960$, $z^* \approx \$100\text{k}$).

Figure 4: Reading guide. Each panel plots welfare relative to the per-cap baseline as a function of the subsidy t (horizontal axis) and the income cap z^* (vertical axis). Lighter colors indicate higher welfare. The dashed white line shows, for each cap, the subsidy that maximizes welfare holding the cap fixed (the "ridge"); the red cross marks the global maximum. An upward-sloping ridge (left) means that broader eligibility calls for a higher t; a downward-sloping ridge (right) means that broadening the cap reduces the optimal t.

Design implications. When social weights are neutral, pursue broad eligibility and set a high subsidy guided by external benefits. With pro-poor weights, prioritize targeting (the cap or a phase-out) and use a moderate subsidy inside the eligible set. If external benefits truly rise with income (e.g., more miles displaced), usage-linked credits can recover part of the corrective gain without sacrificing equity.

Equity–Environment Frontier

To summarize the policy trade-off, I trace an equity–environment frontier. Figure 5 plots outcomes under the globally optimal policy $(t^*, \bar{z}^*)$ for each degree of progressivity κ, which governs inequality aversion in the social welfare function. Higher κ places less weight on high-income households through Pareto weights $\alpha_\kappa(z) = K(\kappa)\, z^{-\kappa}$, with $K(\kappa)$ chosen so that $\mathbb{E}_w[\alpha_\kappa] = 1$. When $\kappa = 0$, the planner is utilitarian and values all households equally; as κ increases, welfare weights decline more steeply with income.

The left panel of Figure 5 reports total environmental benefits $\sum_i e_i\, s_i$, while the right panel reports net benefits $\sum_i e_i\, s_i - \lambda\, t^* \sum_i s_i$ (environmental benefits net of government cost), both evaluated at each κ-optimal policy. Marker color indicates the optimal income cap z^*

(lighter shades denote broader eligibility), and marker size reflects the optimal per-unit subsidy t^*. The connected points therefore trace how the optimal cap and subsidy evolve as the planner becomes more progressive. I do not compare SWF levels across points; the frontier summarizes realized outcomes under each κ-optimal policy.

Overall, the figure shows that greater inequality aversion tightens the optimal cap and lowers the subsidy, reducing total environmental benefits but, given the observed incidence, often increasing or stabilizing net benefits by avoiding inframarginal transfers. The frontier thus summarizes how progressive preferences shift the balance between environmental ambition and fiscal efficiency.

When $\kappa = 0$, the utilitarian planner implements a nearly universal and moderately sized subsidy that maximizes total welfare given the externality calibration. The resulting policy generates large environmental benefits but also high fiscal outlays. As κ increases, the planner discounts high-income households' welfare gains. The optimal response is first to reduce the subsidy size t^* and then to tighten the eligibility threshold $\bar{z}^*$. Both t^* and $\bar{z}^*$ decline as the planner trims inframarginal transfers to rich adopters whose high baseline adoption rates make subsidies fiscally costly. Over this range, total environmental benefits fall as the policy reaches fewer high-income, high-mileage households, but net benefits typically rise or remain stable because the reduction in fiscal cost outweighs the loss in avoided emissions. The efficiency gain arises because inframarginal high-income adopters receive fewer subsidies while the induced adoption share among remaining bins remains high.

At very high levels of inequality aversion, eligibility becomes narrow and the marginal bin dropped contributes relatively more to avoided emissions per subsidy dollar. Beyond this point, further increases in κ reduce both external and net benefits, as the planner begins to sacrifice substantial environmental gains for redistribution. In this region, optimal progressivity is achieved not by further tightening eligibility but by increasing the generosity of the subsidy to remaining lower-income adopters—a shift from broad efficiency to targeted generosity, which comes at greater fiscal cost.[16]

The increase in net benefits with higher progressivity in the calibration arises from the interaction between the policy structure—a per–unit subsidy—and the empirical inputs, namely how $e(z)$ and price responsiveness scale with income. A local accounting around a cap change makes this mechanism transparent. If the planner tightens eligibility and excludes bin i while

---

[16] This inversion—higher subsidies for a narrower eligible group—appears only at the upper end of the progressivity range. It reflects the planner's preference for redistribution within the eligible population once the policy becomes highly targeted.

holding the subsidy at t, the change in net benefits is

$$\Delta\mathrm{NB}_i = \underbrace{e_i\, s_i(p)}_{\text{baseline external benefit}} - \underbrace{\big[e_i\, s_i(p-t) - \lambda t\, s_i(p-t)\big]}_{\text{benefit} - \text{fiscal cost when eligible}} = -\, e_i \big[s_i(p-t) - s_i(p)\big] + \lambda t\, s_i(p-t).$$

Equivalently,

$$\Delta\mathrm{NB}_i > 0 \quad\Longleftrightarrow\quad \lambda t \;>\; e_i \cdot \delta_i, \qquad \delta_i \equiv \frac{s_i(p-t) - s_i(p)}{s_i(p-t)} \in (0,1),$$

where δ_i is the *incremental share* of subsidized units—the fraction of $s_i(p-t)$ that is induced rather than inframarginal.

Two features of the calibration make this inequality hold for the higher–income bins that are excluded as κ rises. First, under homothetic demand, both $s_i(\cdot)$ and its slope scale with income, so $s_i(p-t)$ is large in rich bins while the induced share δ_i is small—subsidies pay for many inframarginal units. Second, the externality gradient $e(z)$ rises with income but not steeply enough to offset the resulting fiscal cost. As eligibility tightens and $t^*(\bar{z}^*)$ falls with higher κ (through the redistributive term in the first-order condition), the program sheds precisely the expensive mass—high $s_i(p-t)$ with low δ_i—so $\lambda t \sum s$ declines faster than $\sum e\, s$. Net benefits therefore rise, even as total environmental benefits fall.

If $e(z)$ rose sharply with income, if low–income bins were much more elastic (high δ_i), if λ were larger, or if the optimal t^* increased rapidly (for instance, under equivalent variation rather than compensating variation), then $e_i\delta_i$ could exceed λt and net benefits would fall. In the figure, the upward slope of net benefits is thus an empirical implication of the calibration: per–unit subsidies with sizable inframarginal payments concentrated at the top, rather than a generic property of the CES structure itself.

Overall, the equity–environment frontier quantifies how increasing progressivity reshapes the composition of subsidy recipients and the efficiency of public spending. The policy moves from broad, fiscally intensive support of high-income adopters toward more targeted subsidies for middle- and lower-income households. While total environmental benefits inevitably decline, fiscal efficiency improves and aggregate net benefits remain broadly stable over a wide range of inequality-aversion parameters.

Figure 5: Equity-Environment Frontier Across Progressivity

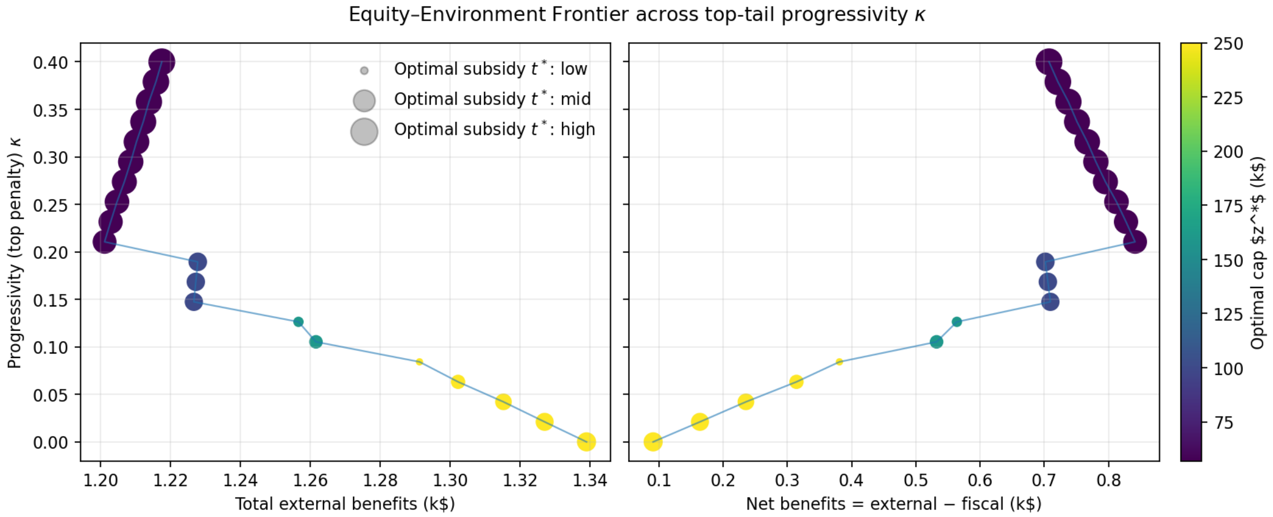

8 Conclusions

This paper provides an integrated empirical and theoretical analysis of the income-capped EV subsidy introduced under the 2022 Inflation Reduction Act. Using variation in income composition across ZIP codes, I estimate that the removal of subsidy eligibility for high-income households reduced EV registrations by roughly 21 percent, implying an elasticity of about 2.9 among affected buyers. The magnitude of this behavioral response, despite occurring at the upper end of the income distribution, demonstrates that even relatively unconstrained consumers remain highly sensitive to policy incentives. This finding motivates the development of a comprehensive welfare framework that jointly evaluates corrective and redistributive goals. If high-income households are unresponsive to subsidy removal, tightening eligibility can improve progressivity and fiscal efficiency without sacrificing environmental outcomes; if they are responsive, the trade-off becomes explicit and quantifiable within such a framework.

The welfare model embeds heterogeneity in both externalities and adoption responses across the income distribution. It delivers closed-form sufficient statistics for the jointly optimal subsidy and income cap, clarifying how corrective and redistributive forces interact. The corrective term rewards targeting adopters who generate high environmental benefits and respond strongly to subsidies, while the redistributive term discounts transfers to inframarginal, high-income adopters when social weights decline with income. The income cap determines which households enter these empirical moments and therefore governs the balance between fiscal cost and corrective efficiency.

I calibrate a homothetic CES demand system using empirically disciplined parameters on adoption, externalities, and price elasticity. The simulations yield several insights. When external benefits rise with income and social weights are neutral, the optimal policy is broad and generous: expanding eligibility raises welfare because higher-income adopters both drive more and respond strongly to prices. When social weights tilt toward lower-income households, the welfare-maximizing cap tightens and the optimal subsidy falls. The pattern reflects the interaction of two forces: tightening the cap removes high-cost fiscal transfers to rich adopters, while lowering the subsidy mitigates redistributive inefficiency among those remaining eligible. These simulations help quantify the trade-offs implied by the model even without fully tracing a global efficiency–equity frontier.

Although the application here focuses on EVs, the same logic extends to many other clean-energy programs where subsidies promote goods with positive externalities and unequal participation across income groups, such as rooftop solar, home-efficiency upgrades, or electric heat pumps, where early adopters are typically the wealthy. In these settings, policymakers face the same trade-off between broad generosity and targeted efficiency, and the framework provides a transparent way to evaluate that balance.

Future work can extend this analysis in several directions. Individual-level registration data would allow more precise identification of behavioral responses and better estimation of how externalities vary with income. Linking registration and usage information would also clarify who benefits from subsidies and who drives the resulting emissions reductions. Incorporating dynamic adoption or learning effects could further illuminate how optimal caps evolve as markets mature and technologies diffuse.

References

Allcott, Hunt, Benjamin B. Lockwood, and Dmitry Taubinsky (2019). “Regressive Sin Taxes, with an Application to the Optimal Soda Tax”. In: *Quarterly Journal of Economics* 134.3, pp. 1557–1626.

Allcott, Hunt, Thomas Reigner-Kane, et al. (Oct. 2024). *The Effects of “Buy American”: Electric Vehicles and the Inflation Reduction Act*. NBER Working Paper 33032. Revision: December 2024. National Bureau of Economic Research.

Atkinson, A.B. and J.E. Stiglitz (1976). “The design of tax structure: Direct versus indirect taxation”. In: *Journal of Public Economics* 6.1, pp. 55–75. ISSN: 0047-2727.

Boomhower, Judd and Lucas W. Davis (2014). “A Credible Approach for Measuring Inframarginal Participation in Energy Efficiency Programs”. In: *Journal of Public Economics* 113, pp. 67–79.

Borenstein, Severin and Lucas W. Davis (2024). *The Distributional Effects of U.S. Tax Credits for Heat Pumps, Solar Panels, and Electric Vehicles.* NBER Working Paper 32688. National Bureau of Economic Research.

Chen, Jiafeng and Jonathan Roth (2024). "Logs with zeros? Some problems and solutions". In: *The Quarterly Journal of Economics* 139.2, pp. 891–936.

Chetty, Raj, John N. Friedman, and Emmanuel Saez (2013). "Using Differences in Knowledge across Neighborhoods to Uncover the Impacts of the EITC on Earnings". In: *American Economic Review* 103.7, pp. 2683–2721.

Gallagher, Kelly Sims and Erich Muehlegger (2011). "Giving Green to Get Green? Incentives and Consumer Adoption of Hybrid Vehicle Technology". In: *Journal of Environmental Economics and Management* 61.1, pp. 1–15.

Griffith, Rachel, Martin O'Connell, and Kate Smith (2019). "Tax design in the alcohol market". In: *Journal of Public Economics* 172, pp. 20–35.

Holland, Stephen P. et al. (2019). "Distributional Effects of Air Pollution from Electric Vehicle Adoption". In: *Journal of the Association of Environmental and Resource Economists* 6.S1, pp. 65–94.

Jenn, Alan, Katalin Springel, and Anand R. Gopal (2018). "Effectiveness of Electric Vehicle Incentives in the United States". In: *Energy Policy* 119, pp. 349–356.

Kaplow, Louis (2006a). "On the Undesirability of Commodity Taxation Even When Income Taxation Is Not Optimal". In: *Journal of Public Economics* 90.6–7, pp. 1235–1250.

— (2006b). "Public Goods and the Distribution of Income". In: *European Economic Review* 50.7, pp. 1627–1660.

Kaplow, Louis and Steven Shavell (2000). "Should Legal Rules Favor the Poor? Clarifying the Role of Legal Rules and the Income Tax in Redistributing Income". In: *The Journal of Legal Studies* 29.2, pp. 821–835.

Lohawala, Nafisa (2023). *Roadblock or Accelerator? The Effect of Electric Vehicle Subsidy Elimination.* RFF Working Paper 23-13. Resources for the Future.

Muehlegger, Erich and David S. Rapson (2022). "Subsidizing Low- and Middle-Income Adoption of Electric Vehicles: Quasi-Experimental Evidence from California". In: *Journal of Public Economics* 216, p. 104757.

Pigou, Arthur C. (1920). *The Economics of Welfare.* London: Macmillan and Co.

Rennert, Kevin et al. (2022). "Comprehensive evidence on the social cost of carbon". In: *Nature* 610.7933, pp. 687–692.

Saez, Emmanuel (2002). "The Desirability of Commodity Taxation Under Non-Linear Income Taxation and Heterogeneous Tastes". In: *Journal of Public Economics* 83.2, pp. 217–230.

Sallee, James M. (2011). "The Surprising Incidence of Tax Credits for the Toyota Prius". In: *American Economic Journal: Economic Policy* 3.2, pp. 189–219.

Sheldon, Tamara L. and Rubal Dua (2019). “Measuring the cost-effectiveness of electric vehicle subsidies”. In: *Energy Economics* 84, p. 104545.

Shih, W. C. (2022). “Global Supply Chains in a Post-Pandemic World”. In: *Management and Organization Review* 18.3, pp. 446–452.

Springel, Katalin (Nov. 2021). “Network Externality and Subsidy Structure in Two-Sided Markets: Evidence from Electric Vehicle Incentives”. In: *American Economic Journal: Economic Policy* 13.4, pp. 393–432.

U.S. Environmental Protection Agency (2024). *Greenhouse Gas Emissions from a Typical Passenger Vehicle*. `https://www.epa.gov/greenvehicles/greenhouse-gas-emissions-typical-passenger-vehicle`.

Xing, Jianwei, Benjamin Leard, and Shanjun Li (2021). “What Does an Electric Vehicle Replace?” In: *Journal of Environmental Economics and Management* 107, p. 102432.

Appendix

A Event Studies

A.1 Parallel Trend

The parallel trends assumption remains robust under alternative measures such as the median and the $150,000 threshold.

Figure 6: Parallel Trend Across Zip Codes with Varying Income

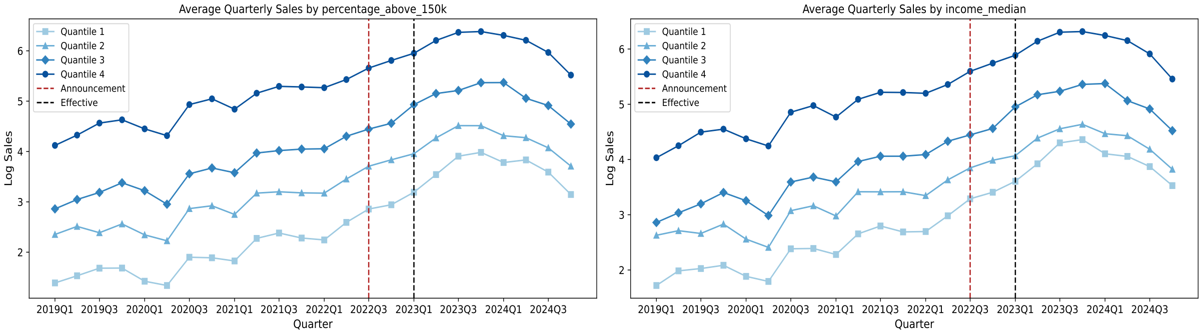

(a) Sales in zip code, by percentage above 150k (b) Sales in zip code, by zipcode median income

Table 4: Event-Study: $\tau_k\times$ High-income (reference τ_0 = 2022H2)

	(1) No trends b/se	(2) State linear trends b/se	(3) Restricted sample b/se
τ_−9 (× High-income)	-0.098	-0.094	-0.110
	(0.067)	(0.054)	(0.056)
τ_−8 (× High-income)	0.006	0.010	-0.252***
	(0.052)	(0.036)	(0.055)
τ_−7 (× High-income)	-0.004	-0.001	-0.111
	(0.049)	(0.040)	(0.066)
τ_−6 (× High-income)	0.101***	0.103***	-0.023
	(0.029)	(0.024)	(0.073)
τ_−5 (× High-income)	0.121***	0.123***	0.124
	(0.030)	(0.027)	(0.075)
τ_−4 (× High-income)	0.164***	0.165***	-0.048
	(0.030)	(0.023)	(0.075)
τ_−3 (× High-income)	0.018	0.019	0.008
	(0.029)	(0.023)	(0.054)
τ_−2 (× High-income)	0.043	0.044	0.009
	(0.029)	(0.024)	(0.053)
τ_−1 (× High-income)	-0.014	-0.013	-0.082
	(0.015)	(0.016)	(0.052)
τ_1 (× High-income)	-0.181*	-0.181*	-0.367*
	(0.082)	(0.087)	(0.154)
τ_2 (× High-income)	-0.195***	-0.196***	-0.280**
	(0.051)	(0.052)	(0.100)
τ_3 (× High-income)	-0.226*	-0.226*	-0.114*
	(0.089)	(0.089)	(0.054)
τ_4 (× High-income)	-0.253***	-0.254***	-0.216
	(0.069)	(0.074)	(0.122)
Constant	6.862***	6.895***	5.520***
	(0.011)	(0.260)	(0.017)
State trends	No	No	Yes
Observations	58566	58566	40040
Avg. post effect	-0.214	-0.214	-0.244
SE (avg. post)	0.071	0.074	0.097
Post periods (K)	4	4	4
Pseudo R^2	0.938	0.950	0.893
Log pseudolikelihood	-1132561	-919337	-327412

Notes: Robust standard errors in parentheses, two-way clustered by ZIP code and half-year. Stars indicate significance at the 10%, 5%, and 1% levels. Columns (2) and (3) include *state-specific linear time trends*, allowing for differential pre-trends across states. Reported pseudo-R^2 values are based on log pseudo-likelihoods and are comparable only across models estimated on the same sample.

A.2 Main Specification: Full Table

A.3 Robustness Check: Limiting the Sample to EV Models Whose Subsidy Either Remained Constant or Decreased

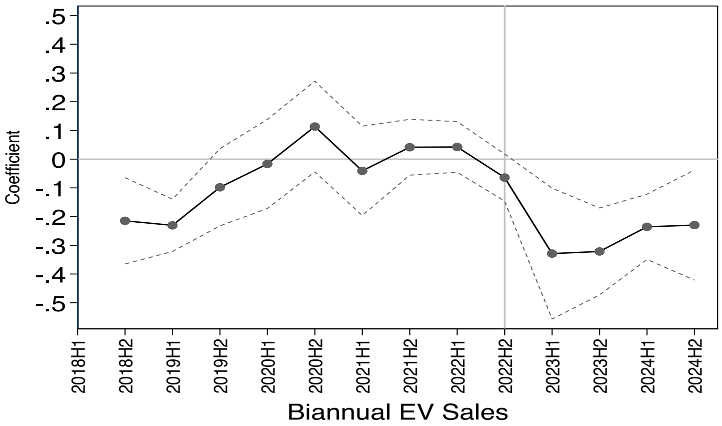

Figure 7: PPML Sample with non-increasing subsidies

A.4 Continuous Treatment

The result of the "dose-response" PPML-DID is shown in figure 8.

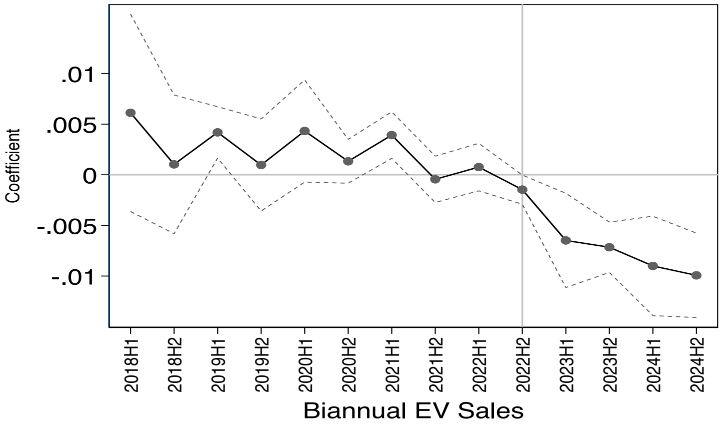

Figure 8: PPML Continuous Treatment

Interpretation of the coefficient of the continuous $Treat_{it}$: I adopt the "average causal response" interpretation, following the framework of Callaway, Goodman-Bacon, and Sant'Anna (2024), to interpret the TWFE coefficient estimates. Specifically, β^{TWFE} captures the effect of excluding an additional 1% of the population from eligibility for the approximately $4,500 (on average) EV subsidy on sales in a typical zip code. In this context, the non-binary treatment reflects continuous variation in the intensity of the policy exposure, driven by differences in the share of ineligible populations across zip codes.

To interpret the coefficient as an average causal response, the key assumption is the absence of unobserved treatment effect heterogeneity, also known as heterogeneity bias. While the parallel trends assumption ensures that untreated outcomes for different groups would have evolved similarly in the absence of the policy change, it does not rule out potential differences in how these groups respond to varying treatment intensities. A heterogeneity bias would arise in this context if high-income households living in the rich zip codes systematically differ from those living in the low-income zip codes, due to factors beyond those already controlled for. To mitigate this concern, I control for EV charging infrastructure, political ideology, and urban-rural characteristics.

Another important limitation of the reduced-form analysis is that increasing the ineligibility effectively raises the income cap, meaning that the analysis extrapolates the effects to a slightly different population. This marginal change may not capture the full range of behavioral responses that would occur under larger shifts in the income cap. As such, while the estimates are informative for small, incremental adjustments to the policy, they may not generalize well to broader eligibility changes. This signals a need for a more structural model as well as a theoretical welfare framework.

A.5 Spline Functions

To allow for potential heterogeneity in treatment effects across income levels, I estimate a piecewise linear spline in the ZIP-level share of households earning over $200,000, with knots at 10%, 15%, and 20%. This flexible specification serves as a middle ground between the binary (above/below median) and fully continuous approaches, allowing treatment effects to vary across the high-income distribution without imposing strict linearity. The estimated effects are smallest in ZIP codes with around 15% high-income households. Effects are larger in ZIPs with low and high density of high-income population.

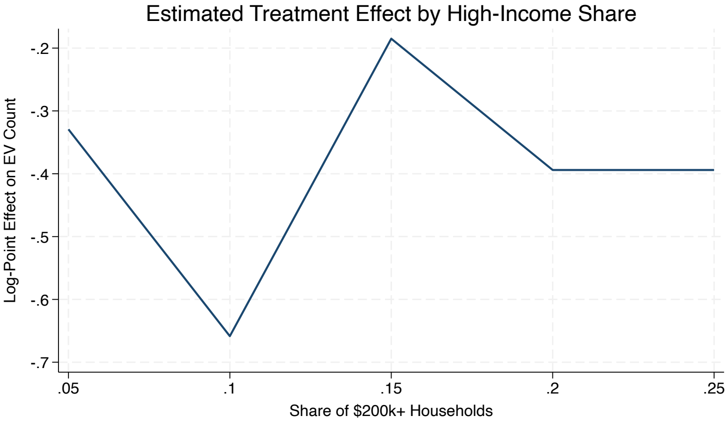

Figure 9: Hetereogenous Treatment Effects by Treatment Intensity

.

B Heterogeneity

I examine whether households that own at least one EV tend to have larger vehicle fleets than non-EV households, controlling for income and conditional on having a car. Table 6 reports the results of an OLS regression of the number of vehicles owned by a household on an indicator for EV ownership and a full set of income group fixed effects. The coefficient on EV ownership is positive and statistically significant at the 5% level, indicating that, conditional on income, EV-owning households have approximately 0.18 more vehicles on average than non-EV households. This finding suggests that EVs are often purchased as additions to existing vehicle fleets rather than as replacements for ICE vehicles. However, such effect is not particularly unique among higher-income groups where EV adoption is more prevalent.

Income is also a strong predictor of vehicle ownership. Relative to the lowest income group, households in income deciles 5 through 11 own significantly more vehicles, with the difference increasing monotonically and reaching nearly one additional vehicle ($0.91; p \leq 0.001$) in the top income group. Together, these results underscore that income is the primary driver of differences in vehicle fleet size. Nevertheless, the modest but significant conditional association between EV ownership and vehicle count provides evidence that at least some EV adoption represents fleet expansion rather than fleet turnover, which has implications for the environmental effectiveness of untargeted EV subsidies.

Table 5: Probability of Owning an EV

	2017	2022
Own EV		
One Family House		0.273
		(1.64)
Own Home	0.00575	0.142
	(0.14)	(0.95)
Urban	0.510***	0.429***
	(14.85)	(3.42)
10 to 15	-0.202	-0.0233
	(-0.99)	(-0.04)
15 to 25	0.251	-0.347
	(1.53)	(-0.65)
25 to 35	0.449**	0.0218
	(2.85)	(0.05)
35 to 50	0.643***	0.216
	(4.23)	(0.48)
50 to 75	0.928***	0.548
	(6.28)	(1.27)
75 to 100	1.203***	0.755
	(8.14)	(1.74)
100 to 125	1.386***	0.970*
	(9.37)	(2.23)
125 to 149	1.568***	1.173**
	(10.47)	(2.67)
150 to 199	1.613***	1.326**
	(10.79)	(3.03)
200 or more	1.900***	1.626***
	(12.83)	(3.78)
Constant	-5.213***	-3.997***
	(-35.24)	(-9.09)
Observations	248006	7321

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 6: EV Ownership and Number of Vehicles in Household

	Dependent variable: Number of Vehicles	
	(1)	(2)
EV owner (1 = yes)	0.180**	0.037
	(0.079)	(0.713)
Income group 2	–0.062	–0.071
	(0.105)	(0.106)
Income group 3	–0.000	–0.002
	(0.088)	(0.088)
Income group 4	0.077	0.074
	(0.085)	(0.085)
Income group 5	0.280***	0.278***
	(0.081)	(0.081)
Income group 6	0.417***	0.417***
	(0.077)	(0.078)
Income group 7	0.582***	0.580***
	(0.078)	(0.079)
Income group 8	0.665***	0.659***
	(0.080)	(0.081)
Income group 9	0.859***	0.857***
	(0.083)	(0.084)
Income group 10	0.885***	0.889***
	(0.084)	(0.084)
Income group 11	0.909***	0.910***
	(0.081)	(0.082)
EV owner × Income group 2		1.571
		(1.234)
EV owner × Income group 3		0.000
		(0.000)
EV owner × Income group 4		0.426
		(1.007)
EV owner × Income group 5		0.222
		(0.844)
EV owner × Income group 6		–0.008
		(0.775)
EV owner × Income group 7		0.170
		(0.757)
EV owner × Income group 8		0.260
		(0.737)
EV owner × Income group 9		0.189
		(0.776)
EV owner × Income group 10		0.044
		(0.738)
EV owner × Income group 11		0.106
		(0.726)
Observations	7,417	7,417
R-squared	0.085	0.086
Adjusted R-squared	0.083	0.083

Notes: Each column reports OLS estimates of the number of vehicles in a household as a function of EV ownership and income group. The omitted income category is group 1 (lowest). Column (2) includes interaction terms between EV ownership and income group. Standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

C Proofs

C.1 Proof of the FOC for the optimal subsidy t

The planner's objective is

$$W(t,\bar{z}) = \int_{\mathcal{Z}} \alpha(z)\, U\big(P(z);z\big)\, d\mu(z) + \int_{\mathcal{Z}} e(z)\, s\big(P(z);z\big)\, d\mu(z) \\ - \lambda\, t \int_{z\le\bar{z}} s\big(P(z);z\big)\, d\mu(z), \tag{5}$$

where $P(z) = p - t$ for $z \le \bar{z}$ (eligible) and $P(z) = p$ otherwise. Let u_c denote the marginal utility of the numéraire, and define $S'(z) \equiv \partial s(P(z);z)/\partial t$ as the behavioral response of adoption with respect to the subsidy.

For eligibles ($z \le \bar{z}$), the envelope theorem implies

$$\frac{\partial U(P(z);z)}{\partial t} = u_c\, s(P(z);z), \qquad \frac{\partial s(P(z);z)}{\partial t} = S'(z), \tag{6}$$

while for ineligibles ($z > \bar{z}$), both derivatives are zero since $P(z)$ does not depend on t.

Differentiating (5) with respect to t yields

$$\frac{\partial W}{\partial t} = \underbrace{\int_{z\le\bar{z}} \alpha(z)\, u_c\, s(P(z);z)\, d\mu(z)}_{\text{(a) utility gain for current eligibles}} + \underbrace{\int_{z\le\bar{z}} e(z)\, S'(z)\, d\mu(z)}_{\text{(b) external benefit from newly induced adopters}} \\ - \lambda \int_{z\le\bar{z}} \Big[\underbrace{s\big(P(z);z\big)}_{\text{(c) fiscal cost for existing eligibles}} + \underbrace{t\, S'(z)}_{\text{(d) fiscal cost for newly induced adopters}} \Big]\, d\mu(z). \tag{7}$$

Setting $\partial W/\partial t = 0$ and rearranging terms gives

$$\int_{z\le\bar{z}} \Big[(e(z) - \lambda t)\, S'(z)\Big] d\mu(z) + \int_{z\le\bar{z}} \big(\alpha(z)u_c - \lambda\big) s(P(z);z)\, d\mu(z) = 0. \tag{8}$$

Rewriting (8) in terms of expectations over the eligible population and normalizing by the mean behavioral response $\bar{S}'_{\text{elig}} \equiv \mathbb{E}_{\text{elig}}[S'(z)]$ yields

$$t^*(\bar{z}) = \bar{E}_{\text{elig}} + \frac{\text{Cov}_{\text{elig}}\big(E_z, S'(z)\big)}{\bar{S}'_{\text{elig}}} + \frac{\text{Cov}_{\text{elig}}\big(\Delta g(z), s(z)\big) + \mathbb{E}_{\text{elig}}[\Delta g]\mathbb{E}_{\text{elig}}[s]}{\bar{S}'_{\text{elig}}}, \tag{9}$$

where $\bar{E}_{\text{elig}} = \mathbb{E}[E_z \mid z \le \bar{z}]$ and $\Delta g(z) \equiv \alpha(z)u_c/\lambda - 1$ measures the deviation of a household's social marginal welfare weight from its fiscal value.

Equation (9) corresponds to the ridge expression in the main text.

C.2 Proof of the first–order condition for the optimal cap $\bar{z}$

Write the objective with an eligibility indicator:

$$\begin{aligned} W(t,\bar{z}) = &\int_{\mathcal{Z}} \alpha(z)\, U\big(P(z;\bar{z}); z\big)\, d\mu(z) + \int_{\mathcal{Z}} e(z)\, s\big(P(z;\bar{z}); z\big)\, d\mu(z) \\ &- \lambda\, t \int_{\mathcal{Z}} \mathbb{1}\{z \le \bar{z}\}\, s\big(P(z;\bar{z}); z\big)\, d\mu(z), \end{aligned} \tag{10}$$

where $P(z;\bar{z}) = p - t\,\mathbb{1}\{z \le \bar{z}\}$.

Only the *boundary* type at $z = \bar{z}$ changes status when $\bar{z}$ moves. Let $s_E(z) \equiv s(p-t;z)$, $s_N(z) \equiv s(p;z)$, and $U_E(z) \equiv U(p-t;z)$, $U_N(z) \equiv U(p;z)$. By Leibniz' rule,

$$\begin{aligned} \frac{\partial W}{\partial \bar{z}} = \Big\{ &\alpha(\bar{z})\big[U_E(\bar{z}) - U_N(\bar{z})\big] + e(\bar{z})\big[s_E(\bar{z}) - s_N(\bar{z})\big] \\ &- \lambda\, t\, s_E(\bar{z}) \Big\}\, f(\bar{z}), \end{aligned} \tag{11}$$

where $f(\bar{z})$ is the density of z at the cutoff.

The first two terms in braces are the marginal social gain from treating the cutoff type as eligible rather than ineligible (private gain plus external benefit from induced adoption); the last term is the marginal fiscal cost at the cutoff. At an interior optimum, $\partial W/\partial \bar{z} = 0$, hence

$$\alpha(\bar{z})\big[U_E(\bar{z}) - U_N(\bar{z})\big] + e(\bar{z})\big[s_E(\bar{z}) - s_N(\bar{z})\big] = \lambda\, t\, s_E(\bar{z}), \tag{12}$$

which is the cap condition in the main text.

C.3 Interpreting the Externality Term

The planner's welfare is written as the sum of private utility and an additive externality term,

$$W = \int \alpha(z)\, U(P(z); z)\, d\mu(z) + \int e(z)\, s(P(z); z)\, d\mu(z) - \lambda t \int_{z \le \bar{z}} s(P(z); z)\, d\mu(z).$$

This specification is one step removed from the canonical formulation in which individuals have utility $u(c, s, E)$ and the environmental state $E = \int \tilde{e}(\tilde{z}) s(\tilde{z})\, d\mu(\tilde{z})$ enters directly into preferences. Under the assumption that the marginal utility of environmental quality, $u_E \equiv \partial u/\partial E$, is approximately constant across individuals and around the observed level of E, the planner's problem can be linearized as

$$W \simeq \int \alpha(z)\, U(P(z); z)\, d\mu(z) + \underbrace{u_E}_{\text{common value of } E} \int \tilde{e}(z)\, s(P(z); z)\, d\mu(z).$$

Defining $e(z) \equiv u_E\,\tilde{e}(z)$ yields the separable form above. Thus, the model is isomorphic to the canonical framework under a common marginal value of environmental quality, preserving the same corrective logic while simplifying the welfare representation.

Following Allcott, Lockwood, and Taubinsky (2019), the same structure can equivalently be written with the externality embedded in the resource constraint rather than in utility:

$$\max_{t,\bar{z}} \int \alpha(z)\, U(P(z);z)\, d\mu(z) \quad \text{s.t.} \quad \lambda\,\mathcal{R} = \lambda t \int_{z\le\bar{z}} s(P(z);z)\, d\mu(z) - \int e(z)\, s(P(z);z)\, d\mu(z),$$

where $\mathcal{R}$ denotes net fiscal outlays. Here, $e(z)$ acts as an implicit credit per adopter, offsetting the subsidy's fiscal cost by the monetized value of avoided emissions. When environmental damages are expressed in monetary units, it is more consistent with the concept of the social cost of carbon: $e(z)$ naturally belongs in the budget rather than the utility term. In a later calibration section, this interpretation allows direct mapping between estimated emission reductions, SCC values, and the planner's resource constraint.

D Calibration

D.1 Homothetic CES Utility Functional Forms

Demand and welfare primitives. With effective price $P \equiv p - t$ and $\sigma = 1/(1-\rho)$, Marshallian demand for group i with income y_i is

$$s_i(t) = y_i \cdot \frac{(1-\theta)P^{-\sigma}}{\theta^{\sigma} + (1-\theta)P^{-\sigma}}.$$

The CES price index is

$$\Pi(P) = \big(\theta^{\sigma} + (1-\theta)P^{1-\sigma}\big)^{\frac{1}{1-\sigma}},$$

so indirect utility is $V(P,y) = y/\Pi(P)$ and the shadow value of income equals the marginal utility of the numeraire:

$$U_c(c_i^*, s_i^*) \;=\; \frac{\partial V(P,y_i)}{\partial y} \;=\; \frac{1}{\Pi(P)}.$$

With constant α_i and λ, the welfare weight $g_i = \alpha_i U_c/\lambda$ is invariant across income; redistribution enters only through α_i.

D.2 Benchmark Cases

Case A: Neutrality (no externality, neutral weights) Set $e_i \equiv 0$, $\alpha_i \equiv 1$, and $\lambda = 1$. For every income cap z^* the optimal policy is $t = 0$, and (under the per-cap normalization used

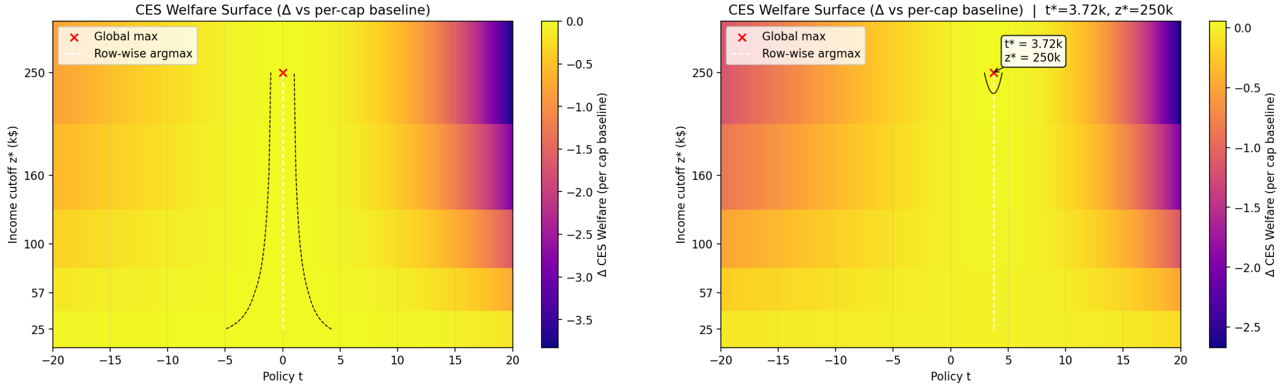

Case A (no externality, neutral weights). The ridge is *flat at* $t = 0$ for every cap z^*; the global maximum is attained at $t^* = 0$ (the planner is indifferent over z^* under the per-cap normalization). Any nonzero tax/subsidy creates only Harberger losses.

Case B (flat externality $e > 0$, neutral weights; CS). The ridge is *vertical at* $t = e/\lambda$ for all caps (here $t^* = \$3{,}720$). Because e and α are uniform, widening eligibility weakly raises welfare, so the global maximum occurs at the broadest cap (top bin included).

Figure 10: Benchmarks. Each panel shows welfare relative to the per-cap baseline as a function of the subsidy t (horizontal) and the income cap z^* (vertical). Lighter colors indicate higher welfare. The dashed line traces, for each cap, the subsidy that maximizes welfare holding the cap fixed (the “ridge”); the red cross marks the global maximum.

in the figures) welfare at $t = 0$ is identical across caps. Intuition: with no external benefits and neutral funds, any nonzero commodity tax or subsidy creates a Harberger-type second-order loss. The ridge is flat at $t = 0$ for all z^*.

Case B: Homogeneous Positive externality, neutral weights Set $e_i \equiv e = \$4,000 > 0$, $\alpha_i \equiv 1$, and $\lambda = 1$ (baseline demand parameters as in the main text). Under consumer surplus, the envelope $d(\Delta CS)/dt = s$ delivers the Pigouvian solution $t^* = e/\lambda$, and the preferred cap is broad because e and α are uniform. Using compensating variation instead of CS lowers the optimal t slightly for large policies (by the standard $r(t) = \Pi(P)/\Pi(p) < 1$ scaling) but leaves the upward-sloping ridge and cap ordering unchanged. In the calibration, $t^*_{\text{CV}} \approx \$3{,}720$ with z^* at the top bin; under CS, $t^* = e/\lambda = \$4{,}000$.

D.3 Externality Calibration

To quantify the environmental externality from EV adoption, I monetize avoided carbon emissions using a social cost of carbon (SCC) of \$150 per metric ton of CO_2, consistent with recent integrated assessment updates in Rennert et al., 2022. A conventional internal combustion vehicle emits about 404g CO_2 per mile, whereas the average EV under the current U.S. grid mix emits roughly 150 g CO_2 per mile U.S. Environmental Protection Agency, 2024, implying an emissions gap of approximately 254g CO_2 per mile. Combining this gap with the observed gradient in driving intensity, from roughly 7,000 miles per year at \$25,000 income to 12,000 miles per year at \$250,000, yields annual avoided emissions of 1.8 to 3 tons CO_2. Over a 12 year vehicle lifetime, these translate into cumulative climate benefits of approximately \$3,200 to \$5,500 per household, increasing with income as higher income households drive more and thus offset more emissions. These magnitudes define the external benefit vector

$$e_i = [3.2,\ 4.1,\ 4.8,\ 5.3,\ 5.5],$$

expressed in thousands of dollars (k\$) per EV over twelve years, which I incorporate into the welfare analysis to capture the income related heterogeneity in environmental benefits.

E Estimating Key Parameters for the Optimal EV Subsidy with Income Cap

In this section, I estimate the empirical parameters derived above to calibrate the optimal subsidy level and the income cap.

E.1 Demand Elasticity and Incidence Using Reduced-form Approach

In this section, I estimate the price and income elasticities of demand, ζ and ξ, and how they vary by income. To address the simultaneity bias from potential correlation between price and unobserved demand shifters, I use federal and local electric vehicle subsidies to instrument for consumer prices. Prior to 2023, several manufacturers (including Tesla, General Motor) reached the threshold and experience a cut of subsidy over the course of a year. Additionally, in 2022, the North America Final Assembly requirement makes many foreign makes ineligible. At the state level, programs such as New York's EV rebate, which ranges from \$2,000 to \$1,000 depending on the vehicle's battery range, provide additional variation in consumer prices. In sum, my identification is obtained from within-model variation over time. Graph 11 shows the change in subsidy level over time by model.

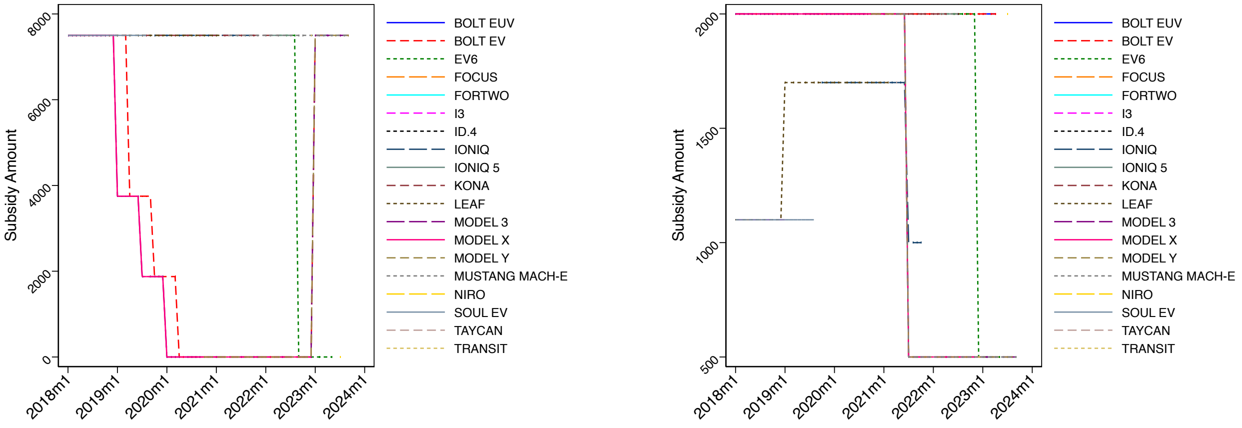

(a) Federal Subsidy Variation Over Time

(b) Local Subsidy by Model

Figure 11: Subsidy Variation Over Time

Let q_{ijt} denote the sales of a particular car model j at a given zipcode i in time t. Let p_{it} denote the average price of the car model at time t, and let $feature_{jt}$ denote the vector of feature variables of the car. Z_{it} is the median income reported by the Census Bureau for zipcode i for the five year estimates, γ_t is a vector of month of sample indicators, $\rho_{city-quarter}$ is a vector of city-quarter fixed effect. My base regression specification to estimate uniform elasticities that do not vary by income is

$$\ln(q_{ijt}) = \zeta \ln(p_{ijt}) + \nu feature_{jt} + \gamma_t + \alpha_i + \delta_j + \rho_{city-quarter} + \epsilon_{ijt} \tag{13}$$

with standard errors clustered by county and manufacturer with $ln(p_{it})$ instrumented in a manner I describe below. To allow elasticities to vary by income, I add linear interaction terms.

$$ln(q_{ijt}) = \zeta ln(\hat{p}_{ijt}) + \xi ln(Z_{it}) + ln(\hat{p}_{ijt}) * Z_{it} + \nu feature_{jt} + \gamma_t + \alpha_i + \delta_j + \rho_{city-quarter} + \epsilon_{ijt} \tag{14}$$

The first-stage regression, where I use federal and state EV subsidies as instruments for price, is specified as follows:

$$ln(p_{ijt}) = ln(subsidy)_{jt} + \nu feature_{jt} + \gamma_t + \alpha_i + \delta_j + \rho_{city-quarter} + \epsilon_{ijt} \tag{15}$$

E.1.1 Estimation Results

Table 7: Price Elasticity Estimation - Log(sales)

	(1) Log	(2) Log	(3) Log
ln(Buy price)	-1.447*	-1.359**	-2.540***
	(0.655)	(0.554)	(0.622)
SUV	0.631***	0.620***	0.633***
	(0.184)	(0.173)	(0.186)
Driving Range	-0.001	-0.001	-0.001
	(0.002)	(0.002)	(0.002)
Income ($100k) × ln(Buy price)			1.284***
			(0.226)
Month, Zipcode, and Model FE	Yes	Yes	Yes
City-quarter FE	No	Yes	No
Kleibergen-Paap first-stage F stat	63.9	63.1	8.0
N	399,896	399,896	399,896

Table A2 in the appendix shows the first-stage result. Table 7 presents the results of the demand estimation. The unit of observation is at the month–zip code–model level. All regressions include fixed effects for month, model, and zip code to control for time-invariant and location-specific factors. Robust standard errors, clustered at the city and manufacturer levels, are reported in parentheses.

Since there are a substantial number of zero values at the model-month-zip code level in the balanced panel dataset, I modified the log regressions by adding one to the original dependent variable. Adding one to the dependent variable in log regressions can potentially distort the proportionality and scale of the data, leading to biased estimates and incorrect interpretations of the relationships between variables. Therefore, I also applied the Inverse Hyperbolic Sine (IHS) transformation on both the dependent and independent variable as a robustness check.

Columns (1) and (2) report IV estimates of the baseline specification (Equation 13), using federal and state EV subsidies as instruments for price. The estimated price elasticity of demand is approximately -1.4. This estimate is somewhat less elastic compared to prior literature, which often finds elasticities around -2 (e.g., Gallagher and Muehlegger, 2011; Li et al., 2017). The difference may reflect variations in market maturity, consumer awareness, or differences in the policy environment.

To examine heterogeneity in price elasticities, Column (3) extends the baseline model by interacting price with median zip code income (Equation 14). The results suggest that lower-income households are more price elastic: the estimated price elasticity is -1.90 at an income level of \$50,000, compared to -1.26 at \$100,000. This finding aligns with theoretical expectations and existing literature, which suggest that lower-income consumers are more sensitive to price changes due to tighter budget constraints (Hausman and Newey, 1995; Chetty et al., 2009).

To address the issue of a large number of zeroes in the dataset, I applied the Inverse Hyperbolic Sine (IHS) transformation. This approach is effective because it allows for the inclusion of zero and negative values while retaining the interpretability of elasticity measures. Table 8 shows the results. The implied price elasticities of demand were -1.8 and -1.7, respectively. These elasticity estimates are relatively consistent with those obtained from the logarithmic transformation.

Table 8: Price Elasticity Estimation - IHS Transformation

	(1) IHS	(2) IHS
asinh(Buy Price)	-1.710*	-1.606**
	(0.776)	(0.657)
SUV	0.726***	0.713***
	(0.216)	(0.202)
Driving Range	-0.001	-0.001
	(0.002)	(0.002)
Month, Zipcode, and Model FE	Month	Month
City-quarter FE	Yes	No
Kleibergen-Paap first-stage F stat	63.7	62.9
Elasticity	-1.8	-1.7
N	399,896	399,896

Table 9: Price Elasticity Estimation - First Stage

	(1) ln(Price)	(2) ln(Price)	(3) ln(Price)	(4) ln(Price) × Income	(5) asinh(Price))	(6) asinh(Price))
ln(Subsidy)	-0.089***	-0.089***	-0.089***	0.101**		
	(0.011)	(0.011)	(0.011)	(0.049)		
asinh(Subsidy)					-0.089***	-0.089***
					(0.011)	(0.011)
SUV	0.083	0.083	0.083	0.069	0.083	0.083
	(0.070)	(0.070)	(0.070)	(0.058)	(0.070)	(0.070)
Driving Range	0.001***	0.001***	0.001***	0.001***	0.001***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Income ($100k) × ln(Subsidy)			-0.000***	-0.211***		
			(0.000)	(0.053)		
Month, Zipcode, and Model FE	Yes	Yes	Yes			
City-quarter FE	No	Yes	No			
Kleibergen-Paap first-stage F stat	63.9237.0	63.1320.4	86.0			
N	399,896 57,818	399,89 53,556	399,896 57,818			