Electric vehicle demand and electricity prices shocks*

Mattias Vesterberg

Department of Economics
Umeå School of Business, Economics and Statistics,
Umeå University, Sweden.
mattias.vesterberg@umu.se

March 2025

Abstract

This study estimates the causal impact of the recent years' high electricity prices on electric vehicle (EV) adoption. Utilizing Swedish registry data and leveraging regional discontinuities in electricity prices, I demonstrate that higher electricity prices reduce EV demand, but also the demand for combustion-engine vehicles. Additionally, the response to electricity prices varies across different types of EVs and socio-economic groups. Based on these findings, I explore a counterfactual policy that reduces electricity prices for EV buyers, and show that under plausible assumptions, this policy is less cost-effective in boosting EV demand compared to subsidies for EV purchases or charging infrastructure.

Keywords: Electrification, Transportation, Regression discontinuity
JEL Classification: D12, Q41, R41

1 Introduction

The transition to electric vehicles (EVs) represents an important shift in the global effort to reduce greenhouse gas emissions and combat climate change. As governments and policymakers worldwide implement measures to promote EV adoption, understanding the factors that influence consumer demand for these vehicles is increasingly critical. Among these factors, electricity prices may play an important role, impacting both the operating costs of EVs and the overall attractiveness of switching from traditional

*The paper has benefited from comments from Alejandro Del Valle Suarez, Mathilda Eriksson, Lassi Ahlvik, Sef Meens Eriksson, David Granlund, Hanna Lindström and Gauthier Lanot, and from participants at the USBE internal seminar. Financial support from Formas (grant number 2023-01082) is gratefully acknowledged.

internal combustion engine vehicles. High electricity prices also reduces the purchasing power of households, since they face higher expenditure for electricity use at home.

While numerous studies have examined the impact of purchase incentives (e.g., Muehlegger and Rapson, 2022; Singh et al., 2020), access to charging infrastructure (e.g., Springel, 2021) and environmental concerns on EV adoption (e.g., Carley et al., 2019), the specific influence of electricity prices on consumer decisions has not been thoroughly investigated. This research gap is particularly concerning given the recent reports of a slowdown in EV uptake across Europe[1], more or less coinciding with a significant rise in electricity prices in most European countries during the last five years.

In this paper, I leverage comprehensive Swedish registry panel data to analyze the responsiveness of EV demand to the recent years' electricity price shocks, and I analyze heterogeneity in price responsiveness in several dimensions. This rich dataset provides detailed information on vehicle registrations and car characteristics, electricity prices and socio-economic characteristics of consumers, allowing for a precise examination of consumer behavior in response to varying electricity costs. Notably, Sweden leads in EV adoption, supported by progressive policies, environmental consciousness, and substantial charging infrastructure investments. As shown in Figure 2, EVs (including plug-in hybrids) became the most common fuel type among new cars in 2021. Given Sweden's similarities with other developed nations, this high EV penetration provides valuable insights for other countries aiming to increase EV adoption.

The identification strategy I use exploits the fact that the Swedish electricity market since 2011 is divided into four distinct price areas (see Section A in the Appendix for details). This unique market structure, found only in Sweden, Norway, Denmark, and Italy, creates price discontinuities at the borders of these areas. These spatial discontinuities results from regional variations in demand and generation mix, together with constraints on transmission capacity, creating natural experiments that facilitate the identification of causal relationships.

Importantly, these discontinuities were particularly pronounced in 2021 and 2022 (see Figure 1) due to the combined effects of the war in Ukraine and the shutdown of nuclear power plants in Sweden. Furthermore, the price area borders do not generally follow other administrative borders, such as the border between municipalities or counties. To illustrate, 81% of the individuals in the data used in this paper lives in a municipality where the price area border cuts through the municipality.

By comparing households that are geographically close but on opposite side of the borders, I can isolate the effect of electricity prices on EV demand. In particular, these households face identical car supply and access to charging infrastructure, but substan-

[1] See, for example, https://www.reuters.com/business/autos-transportation/eu-electric-car-sales-drop-may-german-demand-slumps-industry-says-2024-06-20/ and https://swedenherald.se/article/electric-car-sales-continue-to-fall

tially different electricity prices. This allows me to attribute differences in EV demand to local variation in electricity prices, thereby constructing a credible counterfactual scenario.[2]

In the first part of the empirical analysis, I use a regression discontinuity (RD) design (Calonico et al., 2014, 2017) applied on cross-sectional data aggregated to 1000×1000 meter plots to show that there is a causal effect of electricity prices on EV demand at the border and year with the largest price discontinuity (a price differential of approximately 0.022€/kWh). The estimated treatment effect amounts to ten fewer EVs per 1000 households in the treatment group, compared to the control group that faced substantially lower electricity prices. This can be compared to the mean outcome in the control group, which is 23 EVs per 1000 households. Thus, the estimated treatment effect corresponds to approximately 44% of the mean outcome. This result is robust to alternative specifications, and I also address possible concerns relating to sorting and spillover of treatment effects.

Furthermore, I show that this effect is predominantly driven by the demand for battery electric vehicles (BEVs), whereas there is no effect of electricity price on the demand for plug-in hybrids (PHEVs), and that the demand for leasing EVs is less responsive to electricity prices than purchased EVs. I also show that increasing electricity prices not only reduce the demand for EVs, but also reduces the demand for combustion-engine vehicles (CVs). This latter result suggests that the income effect of an increase in the electricity price is relatively large, and that policies to promote EV demand needs to take such effects into account.

The RD model only measures the effect of electricity prices on EV demand close to the price area border (i.e., the cutoff) and, since it is a model for cross-sectional data, at one border and year. Furthermore, as I detail in Section 3.1 and Section D in the Appendix, there is only one out of three border regions that has a high enough population density to ensure a high enough statistical power when using the RD model.

To understand if the treatment effect estimated from the RD is generalizable to other price areas and years, the second part of the empirical analysis consist of estimating a Poisson pseduo-maximum likelihood (PPML) model applied to plot-level data covering several years and price area borders. To account for unobserved heterogeneity in the PPML model, I compare plots located within 20 km on opposite side of each price area border by including price area border dummy variables in the spirit of Black (1999). Conceptually, this methodology is equivalent to calculating differences in mean EV demand per 1000 households on opposite sides of price area borders (controlling for household characteristics at the plot level) and relating this to differences in electricity

[2]As I detail in Section B in the Appendix, while there has been some changes to policies pertaining to the electricity market and the market for EVs, these policy changes do not affect my empirical results since they do not change at the price-area borders, or at any other administrative borders.

prices within price area border regions. This approach allows me to pool the data across price areas and years, while still accounting for unobserved heterogeneity in space in a similar spirit to the RD approach.

The estimated effect of electricity price on EV demand using the PPML model corroborates the results from the RD model in that there is a sizable effect of electricity prices on EV demand, with estimated elasticities ranging from -0.568 to -2.132, depending on specification. This, combined with the insights from the RD model, leads to the conclusion that EV demand is sensitive to electricity prices—more so than what previous studies have shown (e.g., Bushnell et al., 2022; Mauritzen, 2025).

The third part of the empirical analysis is concerned with heterogeneity in price responsiveness. A growing literature have shown that EVs are predominantly purchased by high-income households (see, for Borenstein and Davis, 2024; Haan et al., 2024), which raises concerns about equity regarding the electrification of transportation. Specifically, the literature argues that while the benefits of transportation electrification are significant for low-income households (e.g., because they spend a larger share of their income on transportation costs), a disproportionately low number of EVs have been sold in these communities. The usual explanation is that low-income households are less likely to have home charging facilities or afford the installation of such infrastructure, have smaller budgets for vehicle purchases, and generally own fewer vehicles.

To understand if heterogeneity in the response to electricity prices contributes to differences in EV demand across consumer types, I estimate a logit specification applied to household-level data to explore heterogeneity in price responsiveness across a broad range of household characteristics, including income, age, number of kids and commuting distance. Similar to the PPML model, I include price area border dummy variables to facilitate the comparison between households located on opposite sides of each price area border. As far as I am aware, this is the first such analysis in the literature, and Swedish registry data provides a unique opportunity to study such heterogeneity, given the high level of detailed socio-economic data available. In brief, I show that high-income households, households with long commuting distances, and households with more than one existing car, are relatively more responsive to prices. The price responsiveness is, however, surprisingly homogeneous, with small and in most cases statistically insignificant differences across consumer types.

Finally, I discuss policy implications of estimated causal effect of electricity prices on EV demand, and consider a hypothetical removal of the Swedish electricity tax for EV buyers as a counterfactual policy. Given some strong but plausible assumptions, back-of-the-envelope calculations show that removing the tax for EV buyers may increase EV uptake with approximately 21%, using the number of EVs bought in 2022 as baseline. Next, I compare this effect with those from previous studies that have examined subsidies on public charging infrastructure and EV purchases. This comparison reveals that a subsidy on electricity is less effective in boosting EV demand than a same-sized subsidies to charging infrastructure or cars. Finally, since EVs are predominantly bought by high-income households, and since the heterogeneity in price responsiveness is small, the policy is likely regressive, mostly benefiting high-income households.

The paper makes several contributions to the existing literature on EV demand. In particular, the sensitivity to prices is not very well studied, and as far as I am aware, the only two papers that so far have measured the sensitivity of EV demand to electricity prices are Bushnell et al. (2022) and Mauritzen (2025).[3] The current paper expands the analysis in these studies in several directions.

First, this paper provides comprehensive evidence that the demand for EVs is sensitive to electricity prices - more so than previously reported in earlier studies. This finding highlights the importance of price as a determinant of EV adoption. The analysis draws on a comprehensive dataset that captures not only car purchases, but also a wide range of consumer characteristics, allowing for a more nuanced understanding of how electricity prices impact consumer decisions than in the existing literature. For example, Mauritzen (2025) only observe income and population density, and only at the county level. Bushnell et al. (2022) do not observe consumer characteristics at all. In comparison, the data used in the current paper includes detailed individual-level information about income, age, family size, commuting distance, existing cars and precise geographic location.

This detailed data not only enhances the accuracy of the findings but also allows for an exploration of heterogeneity in price responsiveness among different consumer groups. This aspect of heterogeneity in price sensitivity has not been explored in prior research, making this a novel and significant contribution to the literature on heterogeneity in EV demand (e.g., Bigler and Radulescu, 2022; Borenstein and Davis, 2024; Haan et al., 2024; Hardman et al., 2016, 2021).

Second, the research extends the scope of analysis to include the impact of electricity prices on the demand for CVs. This aspect of the relationship between electricity prices and vehicle demand has not been estimated in previous literature, making this an important contribution. The results highlight the broader economic implications of electricity price changes, beyond their direct impact on EV adoption.

Lastly, the paper contributes to the growing body of literature on policy measures aimed at incentivizing EV adoption (e.g., Clinton and Steinberg, 2019; Haan et al., 2024; Halse et al., 2025; Jenn et al., 2018; Muehlegger and Rapson, 2022; Münzel et

[3]While the literature on the effect of electricity prices on EV demand (the extensive margin) is sparse, there is a somewhat larger body of work on the effects of electricity prices on charging behavior (the intensive margin). This literature has shown that consumers are willing to pay a premium for fast charging Dorsey et al., 2022; Wolbertus et al., 2018, that the charging of EVs at home is inelastic (Nehiba, 2024), and that automation can increase the response to electricity prices (Burkhardt et al., 2023).

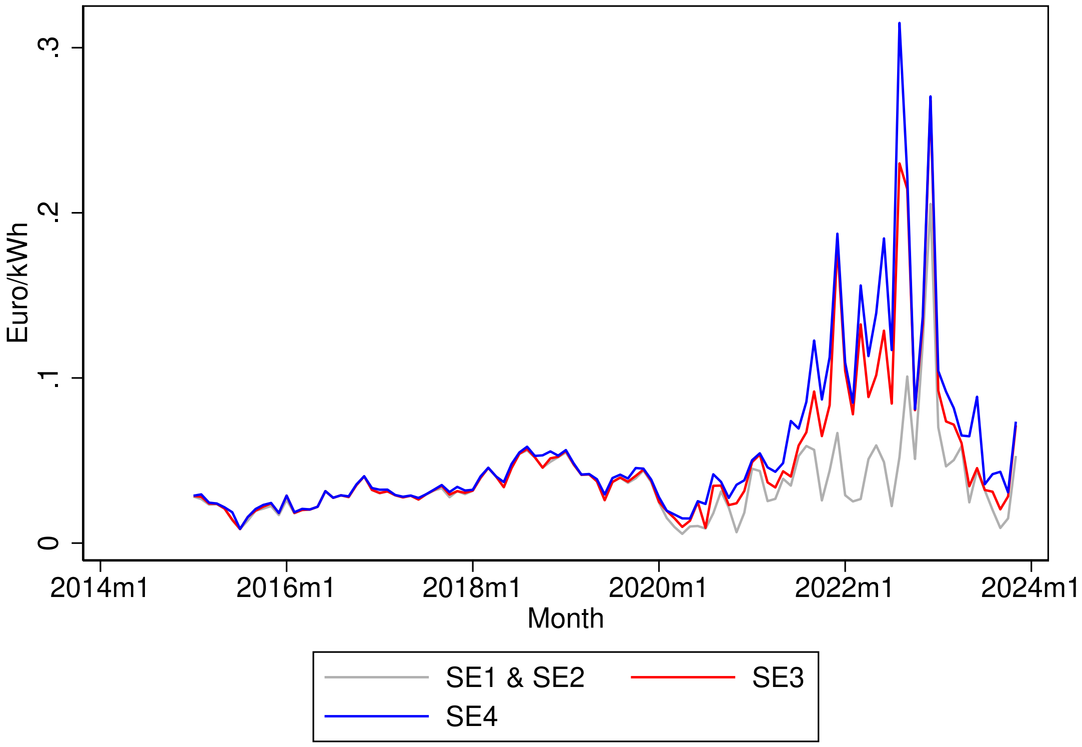

Figure 1: Monthly average wholesale electricity price from 2015 until 2024
Note: SE1 and SE2 faced more or less identical prices during the period 2015 to 2022, and are shown as one line

al., 2019; Springel, 2021). It explores the potential of subsidizing electricity prices as an alternative policy measure to the more commonly studied subsidies on vehicles or charging infrastructure. This is the first study to make this comparison, providing a fresh perspective on policy alternatives for encouraging the uptake of EVs.

The rest of the paper is structured as follows: In Section 2, I describe the data used, detailing sources, variables collected, and preprocessing steps. Section 3 outlines the empirical methods for analyzing the data and estimating the effects of electricity prices. Section 3.1 explains the RD approach, including the rationale, identification strategy, and implementation. Section 3.2 examines heterogeneity in price responsiveness across BEVs, PHEVs, and leasing EVs, while Section 3.3 analyzes the impact of electricity prices on the demand for CVs.

In Section 3.4, I address sorting and spillover effects, discussing their potential impact on the RD design and mitigation steps. Section 3.5 covers the robustness and sensitivity of the RD results, including robustness checks and sensitivity analyses. Section 3.6 presents the PPML model, explaining its motivation, specification, and estimation, and Section 3.7 explores heterogeneity in price responsiveness across household characteristics. Next, I discuss policy implications of my findings in Section 4, and also

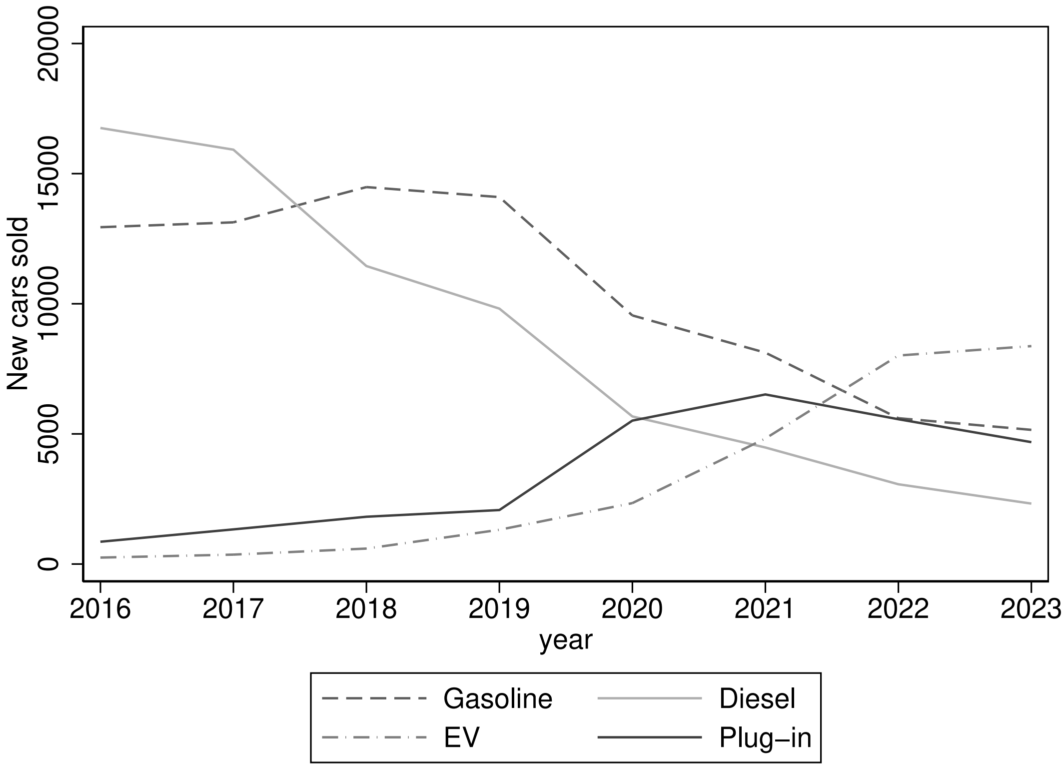

Figure 2: New cars sold in Sweden, by fuel type.
Note: The lines represent car sales for the whole of Sweden and includes leasing cars. Source: https://www.scb.se/hitta-statistik/statistik-efter-amne/transporter-och-kommunikationer/vagtrafik/fordon/pong/statistiknyhet/fordonsstatistik-juni-2018/

provide back-of-the-envelope calculations of a policy counterfactual where the Swedish electricity tax is removed for EV buyers, evaluating its impact on EV demand and CO_2 emissions. Finally, Section 5 concludes.

2 Data

The data used in this paper originates from Statistics Sweden (https://www.scb.se/en/services/ordering-data-and-statistics/microdata/). Specifically, I combine data from the longitudinal integrated database for health insurance and labor market studies ("LISA") and the geographic database ("Geografidatabasen") with data from the Swedish vehicle register ("Fordonsregistret"). The merged data measures annual car purchases and leasing per individual living in villas[4] located within 20 kilometers of an electricity price area border. The data covers the years 2019-2022. Before 2019, both EV demand and electricity prices were relatively low and stable, with minimal variation across price areas.

For each individual, I observe the fuel type, make, model, and vintage of both purchased and leased cars during the sample period, as well as existing cars.[5] The data is rich in individual characteristics, and includes individual-level information about which electricity price area they live in, the closest neighboring electricity price area, and the distance to the nearest electricity price area border as the crow flies (all these variables are provided by Statistics Sweden). Additionally, I observe income, education level, age, number of children under the age of 18, and commuting distance. The data allows me to identify individuals living together, enabling aggregation of individual variables to the household level. In the remaining text, the unit of study is the household with the assumption that for individuals living together, the decision to purchase an EV is made at the household level.

I merge this data with public data on wholesale electricity prices (in €/kWh) from Nordpool (https://www.nordpoolgroup.com/). These prices are measured at the price area level and are annual averages. While I do not observe households' retail electricity prices or their specific tariff choices, most Swedish households have tariffs that vary over time. The Swedish electricity market is deregulated, allowing households to choose from approximately 200 retailers, each offering various tariff options. About 65% of households have tariffs with monthly price fluctuations, and an additional 10% have real-time pricing contracts with hourly price changes. In contrast, less than 20% of

[4] Individuals living in flats often have electricity costs included in their rent, making their response to electricity prices less clear. Additionally, individuals in flats may not be able to charge their EVs at home, and fewer flat-dwellers own cars compared to those living in villas (e.g., Chaisemartin et al., 2022).

[5] For older cars, I observe milage, but this is only recorded during annual inspections, from which newly produced cars are exempt. Consequently, I do not observe mileage for most EVs, as they are predominantly new.

households have opted for tariffs with prices fixed for a year or longer. Thus, variations in wholesale electricity prices are reflected in the retail prices paid b households. Monthly changes in wholesale prices translate into similar variations in retail prices. Households with fixed-price contracts are expected to base their investment decisions on price expectations, influenced by wholesale price trends.

Summary statistics of key variables, by price area, are presented in Table 1. Given that price areas 1 and 2 experienced almost identical electricity prices during the sample period, I aggregate them into a single price area for this table and the subsequent analysis. All variables are household-level aggregates, except for age which is the age of the oldest person in the household.

Starting with the outcome variable, the data reveals that the number of EVs per household purchased varies significantly, both within and across different price areas. Specifically, households living in price areas 3 and 4 purchase more EVs, averaging between 0.014 and 0.015 EVs per household per year. In contrast, households in price areas 1 and 2 purchase fewer EVs, with an average of 0.008 EVs per household per year.[6] This lower uptake of EVs in the northernmost price areas does not appear to be influenced by electricity prices, which are consistently lower in these regions. Instead, a plausible explanation for this phenomenon is the sparse population in northern Sweden, longer commuting distances, and the cold winters that negatively impact battery capacity (e.g., Song et al., 2021).

Figure 3 provides a visual representation of the average number of EVs per price area over the years. One notable observation is that throughout the sample period, EV demand is higher in price area 3 than in price area 4. Interestingly, while the demand for EVs has increased throughout the sample period, this growth appears to slow down in 2022 for price areas 3 and 4, with a more pronounced deceleration in price area 4 compared to price area 3. Price area 1 and 2, that faced only small increases in the electricity price during the sample period, have seen a rather constant increase in EV demand.

Electricity prices vary both spatially and temporally, with the highest and most variable prices observed in price area 4, followed by price area 3, and then price areas 1 and 2. This variation is also illustrated in Figure 1 and in Table A.1 in the Appendix. In addition to the wholesale electricity price, the retail price per kWh includes an energy tax (approximately €0.02 per kWh for consumers in price areas 1 and 2 and €0.03 per kWh for consumers in price areas 3 and 4) and a sales tax (25% on the total cost for electricity).

To put the price differentials across price areas into perspective, an average Swedish household consumes approximately $15,000$ kWh per year (not including charging of

[6] A similar pattern is observed for EVs purchased prior to 2019, with higher numbers in the southern price areas compared to the northern ones.

	Mean	Std. dev.	Min	Max
Price areas 1 & 2				
EV per household	0.008	0.091	0.000	1.000
Elec. price (Euro/kWh)	0.037	0.016	0.014	0.060
No. of cars	1.309	0.548	1.000	3.000
Commuting dist. (km)	3.460	8.083	0.000	115.657
Distance to border (km)	11.611	5.109	0.134	19.999
Household inc. (1000 Euro)	0.232	0.196	0.000	2.551
No. of children	0.457	0.888	0.000	4.000
Age	55.888	16.393	18.000	99.000
Vintage	2010.201	7.075	1946.000	2022.000
Obs.	57,945			
No. of hh.	17,850			
Price area 3				
EV per household	0.015	0.123	0.000	1.000
Elec. price (Euro/kWh)	0.061	0.040	0.020	0.125
No. of cars	1.297	0.537	1.000	3.000
Commuting dist. (km)	2.465	5.739	0.000	129.393
Distance to border (km)	11.652	5.591	0.001	20.000
Household inc. (1000 Euro)	0.286	0.238	0.000	10.171
No. of children	0.576	0.944	0.000	4.000
Age	54.404	16.152	17.000	102.000
Vintage	2012.633	6.075	1923.000	2022.000
Obs.	295,946			
No. of hh.	90,628			
Price area 4				
EV per household	0.014	0.118	0.000	1.000
Elec. price (Euro/kWh)	0.071	0.047	0.024	0.146
No. of cars	1.317	0.553	1.000	3.000
Commuting dist. (km)	2.332	5.356	0.000	131.565
Distance to border (km)	10.356	4.749	0.003	19.999
Household inc. (1000 Euro)	0.268	0.216	0.000	4.367
No. of children	0.543	0.931	0.000	4.000
Age	55.259	16.216	18.000	103.000
Vintage	2012.540	6.076	1927.000	2022.000
Obs.	183,340			
No. of hh.	55,747			

Table 1: Summary statistics

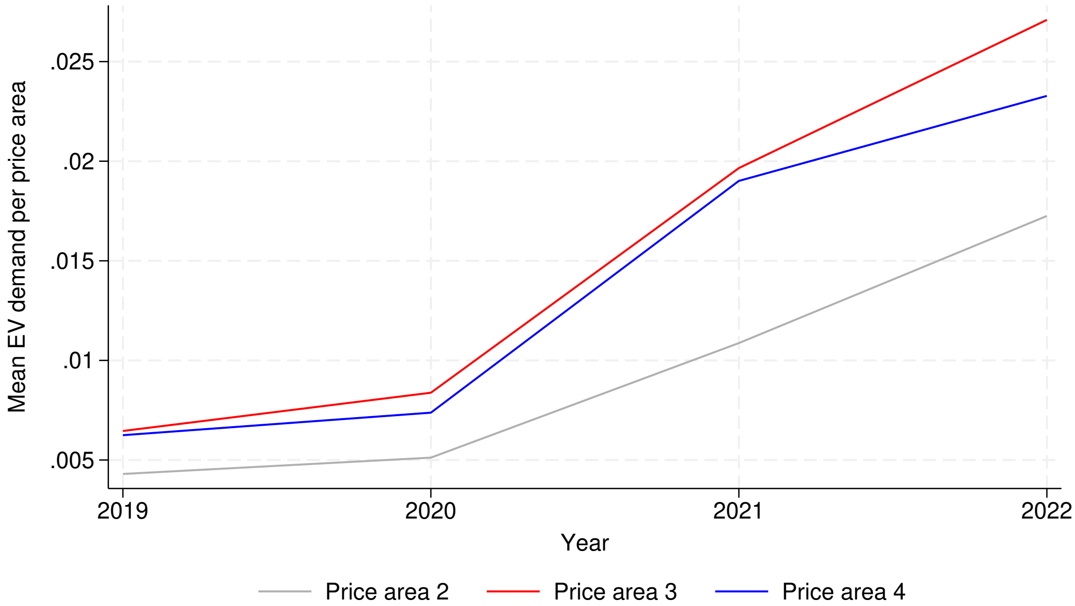

Figure 3: EVs per price area and year

Note: The lines represents average EV demand per price area and year, and includes both battery electric vehicles, plug-in hybrids and leasing cars. SE1 is excluded from the figure and the main analysis, which focuses on price areas 2 to 4. This is because there was no discontinuity in the electricity price at the border between SE1 and SE2.

EVs), see, for example, Lanot and Vesterberg (2019, 2021) and Vesterberg (2018). In 2022, when the price differential was the largest between price areas 3 and 4 (at about €0.022 per kWh), a household in price area 4 would pay approximately €400 (including the 25% sales tax) more in electricity expenditure than a similar household in price area 3, on average. In terms of charging costs, a household that charges an EV at home and drives 15,000 km per year, consuming 0.2 kWh per km, would incur about €65 (again including the 25% sales tax) more in annual charging costs if located in price area 3, compared to price area 4.

Although the differences in expenditure may seem minor, it's crucial to recognize that households in price areas 3 and 4, and to a lesser extent those in price areas 1 and 2, faced significant price increases during 2021 and 2022. For instance, a household in price area 4 consuming 15,000 kWh saw its electricity expenditure in 2022 rise to more than four times the amount spent in 2019 (from approximately €712 to €2737, including sales tax). Similarly, a household in price area 3 experienced an increase of slightly more than three times.[7]

Income levels vary slightly across price areas, with higher income in price areas 3 and 4, compared to price areas 1 and 2, and with substantial variation within price areas.

[7] This price variation is much larger than in Bushnell et al. (2022), and similar in magnitude to the price variation in Mauritzen (2025).

Since the previous literature repeatedly has emphasized income as a key determinant of EV demand, I further explore the role of income by plotting Engel curves for EV demand by year in Figure 4 presents Engel curves for EV demand by year. The sample is divided into ten equally-sized income groups, and the share of households purchasing an EV each year is plotted for each income group. For this figure, I have dropped households with zero reported income.

A first thing to notice is that EVs are owned predominantly by high-income households (above the median income). Second, the relationship between income and EV demand is non-linear. The share of EVs is relatively homogeneous across lower income groups (i.e., below an annual income of approximately €35,000), where the income elasticity is close to zero), whereas EV demand increases rapidly among higher income groups, indicating a relatively high income elasticity in this segment. Third, the share of EVs per income group has increased relatively more among the low-income groups, compared to the high-income groups. For income group 5, the share of EVs is eight times higher in 2022 compared to 2019. For income group 10, the share of EVs in 2022 is seven times larger than it was in 2019. Fourth, the relative difference in the share of EVs across income groups has decreased somewhat during the sample period. In 2019, the share of EVs was 3.5 times larger in income group 10 compared to income group 5, whereas it was only 2.75 times larger in 2022. In the coming sections, I detail how I account for this income feature in the empirical approaches used in this paper.

It is noteworthy that other variables, such as commuting distance and the number of children, exhibit substantial variation, primarily within price areas and to a lesser extent across price areas. This is evident when comparing the standard deviation to the mean of these variables. For instance, mileage and commuting distance show large differences across households, reflecting diverse travel patterns and transportation needs. Similarly, the number of children per household varies widely, which can influence household transportation choices and vehicle ownership. Age, which here is defined as the age of the person registered to the car the household owns[8], is relatively high, with a mean of 54 to 55. However, this is explained by the fact that the sample only includes people living in villas, which are generally older than people living in flats.

The number of cars per household, and the vintage of the newest of these cars shows less variation compared to the other variables, both across price areas and within price areas. This variation is primarily observed across different households rather than over time within the same household or plot. In other words, while households may experience some changes in this variable over the years, the most significant differences are seen when comparing different households.

The differences in household characteristics across price areas decrease when fo-

[8] In the case of several cars registered on different family members, I use the average over those persons

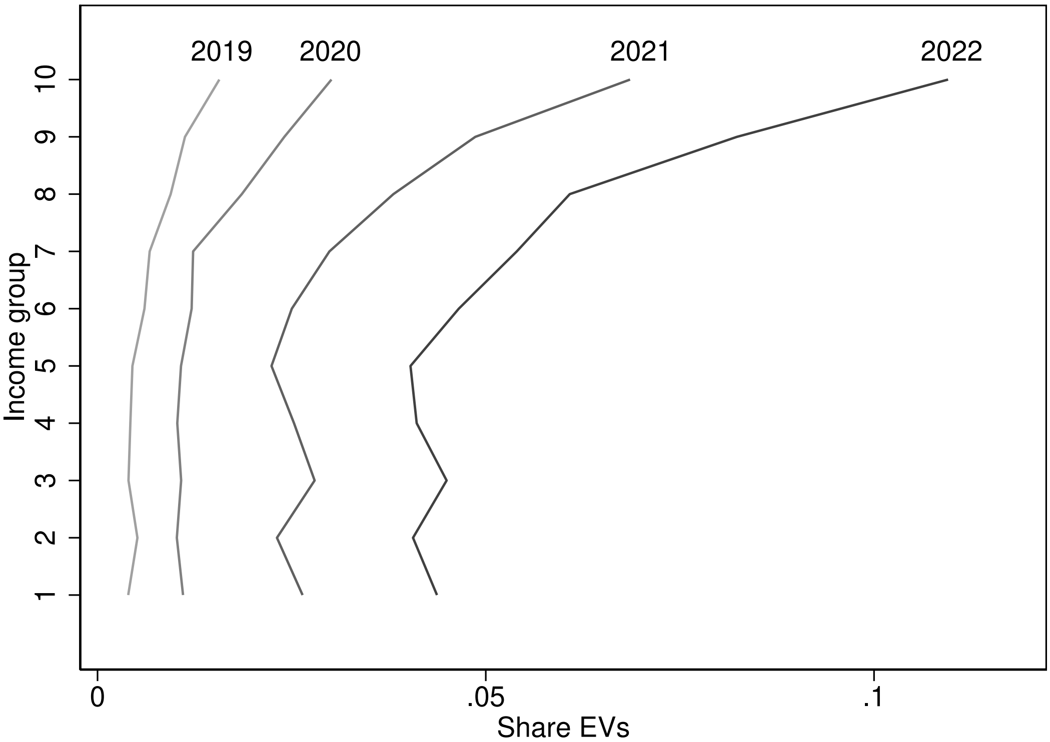

Figure 4: Engel curves for EV demand, by year

Note: i) The figure is constructed by dividing the sample into ten equally-sized income groups, and then computing the share of households purchasing an EV each year for each of these income groups. ii) To simplify the understanding of the figure, I drop households with zero income before splitting the sample into income groups. iii) The levels of income across deciles are presented in Table C.2 in the Appendix cusing on households living close to the price area borders. To illustrate, I focus on households living in price areas 3 or 4 and regress a price area dummy variable on, for example, income and commuting distance, respectively. I estimate separate regressions for bandwidths equal to 15, 10 and 5 km of distance to the price area border. As the bandwidth shrinks, the estimated effect of the dummy variable decreases and very close to zero for small bandwidths. This is reassuring, since it suggest that a design that relies on the discontinuous jump in prices and balanced covariates at the price area borders is valid. I further explore whether covariates jump at the border in more detail in Section 3.1.

Finally, for most new EVs, I observe the car prices. In Table C.1 in the Appendix, I list the five most common EV makes and models along with their purchase prices, using figures from 2022 (the summary statistics are similar for the other sample years). These figures are provided for the entire sample and broken down by price area. Notably, the five most common EVs in terms of make and model are mostly consistent across different price areas. For example, Kia Niro is the most common EV in price areas 1, 2 and 3, and the fourth most common EV in price area 4. Furthermore, the prices of the most popular EVs are very similar across price areas. This is explained by the fact that car dealers typically offer free delivery of new cars to the purchaser's location, allowing consumers to buy from the dealer with the lowest price. The reason car prices vary somewhat between price regions is likely because consumers choose different options and specifications for their cars. The similarity in EV make and model preferences and prices is reassuring for the current paper, as it suggests that spatial differences in EV demand are not influenced by car prices or preferences for specific makes and models of EVs.

3 Empirical analysis

In line with the previous literature on vehicle demand (e.g., Allcott and Wozny, 2014; Bushnell et al., 2022; Busse et al., 2013), it is assumed that consumers base their purchase decisions on the price of the car and the expected (discounted) operational costs over the car ownership period (in my data, the average replacement rate for relatively new cars is approximately five years). These operational costs are influenced by several factors, including electricity prices, the car's efficiency (i.e., kWh per kilometer driven), and the extent of the car's usage (i.e., kilometers driven).

When electricity prices increase, the expected operational costs of owning an EV also rise. This increase in operational costs lead to a decrease in the demand for EVs, as consumers may find them less economically attractive. Additionally, higher electricity prices can have an income effect on consumers (e.g., Mauritzen, 2025). As electricity prices rise, the cost of baseline electricity consumption becomes more expensive. This increase in baseline consumption costs reduces consumers' overall purchasing power, making it more difficult for them to afford new vehicle purchases, including EVs.

In the subsequent sections, I will outline the empirical methods used to identify the causal effect of electricity prices on the demand for EVs. These methods isolate the impact of electricity prices from other factors that may influence consumer behavior and vehicle demand.

3.1 Regression discontinuity

To measure the causal effect of electricity prices on EV demand, I begin by employing a RD design (e.g., Imbens and Lemieux, 2008; Lee and Lemieux, 2010 and more recently Cattaneo and Titiunik, 2022).[9] This approach relies on the assumption that households located near the price area borders are similar in all respects except for the electricity prices they face. Specifically, these households experience the same interest rates, inflation, car supply, access to charging infrastructure, and proximity to labor markets, but they face different electricity prices.

Formally, I estimate the local average treatment effect τ based on the following model:

$$\tau = \mu_{+} - \mu_{-} \tag{1}$$

where

$$\mu_{+} = \lim_{x \to 0^{+}} \mu(x), \quad \mu_{-} = \lim_{x \to 0^{-}} \mu(x), \quad \text{and} \quad \mu(x) \equiv \mathbb{E}[Y_i | X_i = c] \tag{2}$$

The parameters μ_{+} and μ_{-} represent the limit of the expectation of the outcome variable Y given the running variable X as it approaches the cutoff threshold $X = \overline{x}$ (in my case, the price area border) from above and below, respectively. Thus, τ measure the magnitude of the jump in the outcome variable at the point of the discontinuity.

To estimate the treatment effect, I follow the approach outlined in Calonico et al. (2015, 2017) to obtain bias-corrected point estimates using local linear functions, data-driven optimal bandwidths and confidence intervals that are robust to large bandwidths choices (hereinafter referred to as robust confidence intervals). This approach aims to balance bias (a smaller bandwidth reduces bias) with variance (a larger bandwidth reduces variance) by minimizing the mean squared error of the estimator. Furthermore, I use a first-order local polynomial to construct the point estimator (i.e., a local linear regression) following the arguments provided in, for example, Gelman and Imbens (2019). In the main specification, I use a triangle kernel but provide estimation results for the uniform and the Epanechikov kernels in Section F in the Appendix.

I aggregate the household data to plot level using 1000×1000 meter plots in sparsely populated areas and 250×250 meter plots in densely populated areas. The plots are

[9] Spatial discontinuities have been leveraged as identification strategy in, for example, Black (1999), Butts (2023), Grembi et al. (2016), and Keele and Titiunik (2015).

pre-defined by Statistics Sweden.[10]. The outcome variable is defined as number of purchased EVs per 1000 households, and the treatment effect of interest is the exposure to high electricity prices. Distance to the price area border is the forcing variable, and the localization of the border is the cutoff. Since I do not observe the coordinates of the price area borders, but only the distance to the border as the crow flies, I am unable to estimate a RD model with two forcing variables (e.g., longitude and latitude) as in, e.g., Keele and Titiunik (2015).

I focus the analysis on the most recent year in the data, which is 2022. This is the year in the data with the largest price differential (i.e., the largest discontinuity in price). I will also estimate the model using data for the years prior to 2022, which in the case of 2019 and 2020, can be viewed as placebo tests. Since there was no discontinuity in price during these years or the period before this, we should expect the treatment effect to be zero for these years. In 2021, there was a non-zero price discontinuity that was smaller in size than in 2022, and we should then expect the estimated treatment effect to be negative but smaller than the treatment effect in 2022. These are indeed the results I find, see Section 3.4 and Section F in the Appendix.

It is crucial to understand that any differences in outcomes measured in 2022 may be influenced by both recent price differentials during that year and a potentially long history of price variations. Prices began to diverge in 2021, and this divergence continued up until 2022. Therefore, the effects I measure reflect the impact of the price discontinuity at the border, without making assumptions about whether these differences in outcome are due to recent price differentials or the accumulation of price differentials over a period of time.

Furthermore, I focus on the border between price areas 3 and 4 since this is the most densely populated price border area, whereas the other price area borders are more sparsely populated. As I detail in Section D in the Appendix, this difference in population density translates to large differences in statistical power (i.e., the probability that the test will correctly reject a false null hypothesis), with much higher power at the price area border between price areas 3 and 4, compared to the other border regions.

I estimate three RD specifications: First, I estimate the model using EV sales per 1000 households in price area 3 and 4 as the outcome variable, distance to the border between price area 3 and 4 as the forcing variable, and without any additional covariates. Second, I estimate the same specification but this time with covariates, following the method outlined in Calonico et al. (2019).[11] These covariates are included to re-

[10] I do not observe the location (longitude and latitude) of individual households. However, Statistics Sweden has divided Sweden into plots, and information about which of these plots each household is located in. I observe the coordinates (longitude and latitude in sweref99) of each plot, and the plots never overlap price area borders (but may overlap other administrative borders, such as municipalities and counties).

[11] An alternative is to residualize the outcome variable and then estimate the RD on the residuals of the outcome. See, for example, Greenstone et al. (2022) and Meng and Yu (2023). This can be done

duce the residual variance in the outcome variable, leading to more precise estimates of the treatment effect of interest. Specifically, I control for the number of existing cars and the vintage of these cars, commuting distance, income group (using the same income bins as in Figure 4), age and number of kids per household. All these variables are included as plot-level averages. Including more covariates, such as family situation (married, living together or single household), capital income, squared terms (e.g., income and commuting distance squared) or interaction terms (e.g., between income and age or commuting distance) does not lead to better precision, and does not change the estimated treatment effect.

Because two plots located at the same distance to the price area border may still be separated in longitude, I also account for longitudinal differences. Specifically, I divide the longitude into 20 bins (the results are robust to a different number of bins), and include these as covariates linearly. This is a common approach in spatial RD (see, e.g., Keele and Titiunik, 2015; Lehner, 2021) and ensures that I compare only units in close proximity (within the boundary segment) and rule out comparisons of units that are close to the border but geographically far away from each other.

For all covariates, I use the value of these covariates prior to the treatment (i.e., in 2019), and all the covariates are "balanced" (meaning that they have equal conditional expectations at the cutoff). I test this by estimating the RD model with the covariates as outcome variables (see Calonico et al., 2019; Lee, 2008), and results are presented in Table E.1 in the Appendix. Based on these empirical results, I find that all covariates have an RD treatment effect indistinguishable from zero at a 95% significance level. In other words, I cannot reject the null hypothesis of equal conditional expectations at the price area border.

I also estimate a third specification that is identical to the second one except for that I winsorize the data for all observation where the outcome variable (EVs per 1000 households) is above the 99th percentile. This accounts for the fact that the RD is sensitive to outliers, especially if these outliers happens to be located close to the cutoff.[12]

Estimation results from these three specifications are presented in Table 2 and the

using both parametric and non-parametric techniques, as well as machine learning techniques (Kreiss and Rothe, 2023; Noack et al., 2021). However, these approaches does not improve the precision of the estimated treatment effect in my analysis. Estimation results using these approaches are available upon request.

[12]An alternative approach to account for the influence of outliers close to the border is to estimate a so-called donut RD (e.g., Barreca et al., 2011), where observations close to the border are dropped. The donut RD also addresses the concern that households might manipulate the running variable to fall just above or below the cutoff to avoid the treatment. This manipulation can distort the estimated treatment effect. On the other hand, in Section 3.4, I use the panel structure of the data to show that sorting along the borders is not very likely in my context. In any case, Figure E.2 in the Appendix, I present results from the donut RD, where I vary the donut hole from 100 meters to 500 meters. While this also reduces the estimated treatment effect, compared to the results in the second column in Table 2, it does so to a lesser extent than winsorizing. The winsorized results may therefore to be viewed as conservative estimates.

treatment effect from the two specifications with covariates are presented in Figure 5.[13] A first thing to note is that for all specifications, the estimated treatment effect is negative, meaning fewer EV purchases in price area 4 compared to price area 3. However, the estimate for the specification without covariates is relatively noiser, as indicated by the relatively large robust bias-corrected confidence interval. For the specifications where I include covariates, the treatment effect is estimated more precisely. For specification 2, the treatment effect is approximately -10, and the corresponding estimate for specification 3 is -5.6.[14]

To give a sense of magnitude, the average outcome for the control group near the discontinuity (within the MSE-optimal bandwidth) is approximately 23, with a standard deviation of 94. Thus, the treatment effect for specification 2 corresponds to approximately 44 percent of the mean outcome, or approximately 11 percent of the standard deviation of the outcome variable for the control group. The electricity price in price area 4 in 2022 was 0.146€/kWh, which is approximately 18% higher than in price area 3, which faced a price of 0.124 €/kWh.

This relatively large estimated treatment effect suggests that the price shock in 2022 has influenced consumers' expectations about future prices, and that consumers expect the shock to be relatively persistent. In particular, if consumers would expect the price shock in 2022 to be temporary and the electricity price to revert back to its mean after that, it seems likely that the response would be much smaller, since such a temporary shock would have limited effects on the operational costs of an EV.

Comparing with estimates from the existing literature, Mauritzen (2025) finds that higher electricity prices reduce the proportion of electric car purchases by 2-5%, depending on the specification. This reduction appears small given the relatively large price differential between the control and treatment groups, with average prices of 0.029 Euro/kWh in low-price areas and 0.140 Euro/kWh in high-price areas. In Bushnell et al. (2022), the estimated marginal effect ranges from approximately -0.2 to -0.8. With a mean price of 30 (cents per kWh) and a mean quantity of 10 (EVs per 10,000 households), this implies an elasticity of approximately -0.07 to -0.27.

3.2 Battery EVs, Plug-in hybrids and leasing EVs

I have so far aggregated battery electric vehicles (BEV) and plug-in hybrids (PHEV), but there might be differences across these two types of EVs. In particular, since the latter type relies to a lesser extent on electricity, it seems plausible that the response to the electricity price should be smaller for PHEVs than for BEVs. To address this, I estimate the local treatment effect separately for BEVs and PHEVs using the same

[13] Furthermore, I also provide a power plot for the main specification in Section E in the Appendix.

[14] This relatively large difference in estimated treatment effects suggest that outliers have a relatively sizeable influence on the estimates.

	(1) EVs per 1000 hh.	(2) EVs per 1000 hh.	(3) EVs per 1000 hh.
RD_Estimate	-8.201	-10.332***	-5.618**
Robust 95% CI	[-24.986 ; 5.532]	[-20.18 ; -1.542]	[-11.454 ; .32]
Kernel type	Triangular	Triangular	Triangular
BW type	mserd	mserd	mserd
Bandwidth	8769	7976	6140
Obs.	13074	9599	9599
Robust p-value	0.212	0.022	0.064
Order Loc. Poly. (p)	1.000	1.000	1.000
Order bias (q)	2.000	2.000	2.000
Price areas	3 & 4	3 & 4	3 & 4
Year	2022	2022	2022
Control variables	No	Yes	Yes
Winsorizing	No	No	Yes
Mean outcome, treatment	18.123	18.141	14.712
Mean outcome, control	23.121	23.534	17.106

Table 2: RD estimation results, EVs per 1000 households

Note: i) Treatment effects are estimated using the robust RD approach suggested by Calonico et al. (2017) and with a triangular kernel. ii) Bandwidth is chosen using the data-driven MSE-optimal approach in Calonico et al. (2014). iii) The order bias specifies the order of the local polynomial used to construct the bias correction. This bias correction helps to adjust the RD point estimator to account for the bias introduced by the polynomial approximation. iv) The mean outcome for the treatment and control groups are estimated within the MSE-optimal bandwidth and reflect the mean outcome near the discontinuity.

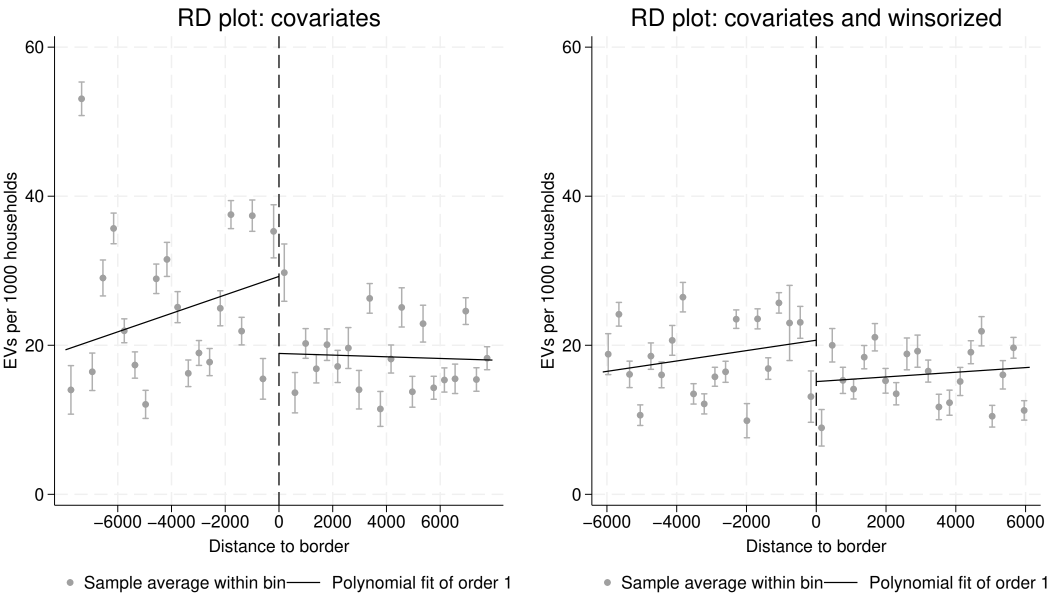

Figure 5: RD treatment plots, EVs per 1000 households

Note: i) The solid lines represents local linear polynomial regression curve estimates for control (to the left of the dashed vertical line) and treatment (to the right of the dashed vertical line) units separately. ii) The bins illustrate local sample averages of the outcome variable within evenly spaced bins of the running variable (distance to the border), and the number of bins are chosen using the data-driven approach outlined in Calonico et al. (2014).

	(1) BEVs per 1000 hh.	(2) PHEVs per 1000 hh.	(3) Leas. EV per 1000 hh.
RD_Estimate	-7.304*	-2.008	-2.092**
Robust 95% CI	[-17.428 ; 1.079]	[-7.018 ; 4.162]	[-4.809 ; .368]
Kernel type	Triangular	Triangular	Triangular
BW type	mserd	mserd	mserd
Bandwidth	8470	6605	6199
Obs.	9599	9599	9599
Robust p-value	0.083	0.617	0.093
Order Loc. Poly. (p)	1.000	1.000	1.000
Order bias (q)	2.000	2.000	2.000
Price areas	3 & 4	3 & 4	3 & 4
Year	2022	2022	2022
Control variables	Yes	Yes	Yes
Winsorizing	No	No	No
Mean outcome, treatment	12.995	8.602	11.028
Mean outcome, control	14.937	9.748	11.181

Table 3: RD estimation results, BEVs, PHEVs and leasing cars per 1000 households
Note: i) In the first column, the outcome variable is number of battery EVs per 1000 households. In the second column, the outcome variable is the number of plug-in hybrids per 1000 households, and in the third column, the outcome variable is the number of leasing EVs (both BEV and PHEV) per 1000 households. ii) Treatment effects are estimated using the robust RD approach suggested by Calonico et al. (2017) and with a triangular kernel. iii) Bandwidth is chosen using the data-driven MSE-optimal approach in Calonico et al. (2014). iv) The order bias specifies the order of the local polynomial used to construct the bias correction. This bias correction helps to adjust the RD point estimator to account for the bias introduced by the polynomial approximation. v) The mean outcome for the treatment and control groups are estimated within the MSE-optimal bandwidth and reflect the mean outcome near the discontinuity.

approach as above

Furthermore, there may be differences in price responsiveness between leasing and purchasing EVs. Recent literature emphasizes the growing importance of the EV leasing market and suggests that consumers leasing EVs may differ from those purchasing them, particularly in terms of financial constraints (see Gonzalez-Salazar et al., 2023; Li et al., 2021). To understand if leasing and purchased EVs differ in their sensitivity to electricity prices, I estimate a specification that includes only leasing EVs.

Estimated treatment effects for these three specifications are presented in Table 3, and as expected, there is a negative treatment effect for BEVs that is equal to -7.3, whereas the treatment for PHEVs is -2 and statistically insignificant. This difference in price responsiveness is in line with the results presented in Mauritzen (2025). Furthermore, the estimated treatment effect for leasing EVs is smaller than the previous estimates, suggesting that the demand for leasing EVs is less elastic.

3.3 Combustion engine vehicles

There are two economic effects of higher electricity prices on car demand: an income effect and a substitution effect. The income effect occurs because higher electricity prices reduce overall purchasing power, leading to a decrease in the total demand for cars. Households have less disposable income to spend on vehicles after accounting for their higher electricity bills. The substitution effect, on the other hand, arises because the relative cost of operating an EV compared to a CV changes due to the increase in electricity prices. As a result, households may find it more economical to opt for CVs over EVs, leading to an increase in the demand for CVs relative to EVs.

To investigate this, I estimate the RD model previously described but using CVs as the dependent variable. These are predominantly vintage CVs: there are approximately 180 vintage CVs bought per 1000 households, and approximately 10 new CVs (the vintage of the car is the same as the year of purchase).[15] Apart from this change of the dependent variable, the models remain identical to those presented earlier.

The results from this specification are presented in Table 4, and show that high electricity prices not only reduce the demand for EVs but also for CVs, with an estimated treatment effect of -27. This corresponds to approximately 14 percent of the mean or 11% of the standard deviation of the outcome, which means that the effect is relatively smaller than the effect on EVs. A plausible explanation to this result is that higher electricity prices reduces households' purchasing power since they face a higher cost for their baseline electricity consumption (electricity used for heating, cooking, lighting, etc), therefore leading to lower demand for all cars. In particular, if there was no income effect and only a substitution effect, or if the income effect was smaller than the substitution effect, we would expect the demand for CVs to increase. I discuss policy implications of this finding in Section 4.

3.4 Placebo tests, sorting and spillover of treatments

To assess the robustness of the RD results, I conduct placebo tests and discuss sorting and spillover of treatments. To begin with, I estimate the treatment effect for the other sample years. The price differential between price areas 3 and 4 was the highest in 2022 (€0.022/kWh). In comparison, the price differential was smaller in 2021 (approximately €0.014/kWh), and even smaller in 2019 and 2020 (see Table A.1 in the Appendix). We should thus expect the treatment effect to be smaller in 2021, and zero in 2019 and 2020. Estimating the RD for these two latter years can then serve as a placebo treatment, where a zero treatment effect is indicative of support for the identifying assumption in the RD model. As is illustrated in Table F.1 in the Appendix, this is precisely what I

[15] I have estimated the RD model separately for old and new CVs, and the estimated treatment for old CVs is approximately -33 and zero for new CVs.

	(1) CVs per 1000 hh.
RD_Estimate	-27.992**
Robust 95% CI	[-64.746 ; -.515]
Kernel type	Triangular
BW type	mserd
Bandwidth	6391
Obs.	9588
Conventional p-value	0.057
Robust p-value	0.046
Order Loc. Poly. (p)	1.000
Order bias (q)	2.000
Price areas	3 & 4
Year	2022
Control variables	Yes
Winsorizing	No
Mean outcome, treatment	189.739
Mean outcome, control	194.785

Table 4: RD estimation results, new and old CVs per 1000 households

Note: i) Treatment effects are estimated using the robust RD approach suggested by Calonico et al. (2017) and with a triangular kernel. ii) Bandwidth is chosen using the data-driven MSE-optimal approach in Calonico et al. (2014). iii) The order bias specifies the order of the local polynomial used to construct the bias correction. This bias correction helps to adjust the RD point estimator to account for the bias introduced by the polynomial approximation. iv) The mean outcome for the treatment and control groups are estimated within the MSE-optimal bandwidth and reflect the mean outcome near the discontinuity.

find when estimating the RD for these years.[16] In particular, the treatment effects are close to zero for 2019 and 2020 (albeit negative), and the estimated treatment effect in 2021 is approximately half of the estimated treatment effect for 2022.

Next, I conduct placebo tests by estimating the treatment effects at points $6,000$, $4,000$ and $2,000$ meters north and south of the electricity price area border between price areas 3 and 4. These placebo tests serve to check whether any discontinuities in the outcome variable are present at locations where no actual treatment occurs.

The results of these placebo tests, presented in Figure F.1 in the Appendix, indicate that the estimated treatment effects at all placebo points are statistically indistinguishable from zero. This finding supports the assumption that any observed discontinuity at the actual border is attributable to the treatment effect rather than other confounding factors. Specifically, the absence of significant treatment effects at the placebo points suggests that the running variable (distance to the border) does not inherently influence the outcome variable in the absence of the treatment.

Furthermore, I address possible concerns about households sorting around borders. For example, households sensitive to electricity prices might move from expensive to cheap price areas. I investigate this by using the panel structure of the data to measure how many households move to a different price area per year, the extent to which they move from expensive to cheap price areas, and how these quantities differ across years.

If households sort along borders based on electricity prices, we should expect more households to move from an expensive price area to a cheap one, especially when electricity price differentials are larger. However, this exercise reveals that few households move across price areas each year (less than one percent per year), while many more households move within price areas (approximately four percent per year). See Table F.2 in the Appendix for details. Additionally, there are as many households moving to cheaper price areas (i.e., moving north) as there are households moving to more expensive price areas (i.e., moving south). Finally, the number of households moving across price areas is similar across years. For example, regressing yearly dummy variables on the first difference of price area (which is then zero if a household do not move to a different price area, negative if a household move north to a cheaper price area and positive if they move south to a more expensive price area) shows no statistically significant differences across years.

Spillover of treatment effects through equilibrium prices (e.g., as discussed in Stiglitz and Kosenko, 2024) is unlikely. At the beginning of 2021, there were approximately $340,000$ EVs in Sweden, of which approximately $150,000$ were BEVs and $190,000$ were PHEVs. An EV consume on average 0.2 kWh per km. The average driving distance per

[16]Since I do not observe covariates prior to 2019, I am unable to include pre-treatment covariates for this year. Results are from a specification where covariates are from 2019, and this should be kept in mind when interpreting the results

car in Sweden in 2021 was $11,260$km per year (and let's assume that PHEVs drive half of that on electricity).[17] This amounts to $150,000 * 0.2 * 11,260 + 190,000 * 0.2 * 5,630 = 787,500,000$ kWh or approximately 0.79 TWh in electricity charging. This, in turn, corresponds to approximately 0.4% of Sweden's total electricity use. This is assumed to be too small to have any substantial effect on the electricity price, which is therefore assumed exogenous.

Furthermore, it is equally unlikely that consumers drive to other side of the border to charge, since it is always more expensive to charge at charging station than at home, irrespective of where the consumer live. To illustrate, the highest spot price I observe in my data is 0.16 €/kWh, and the prices at public charging stations starts at around $0.3-0.5$ €/kWh, and with even higher prices for fast chargers (see, for example, https://www.ionity.eu/network/access-and-payments which is one of the largest charging networks in Sweden).

3.5 Sensitivity to estimation options

To check the sensitivity of the RD results to estimation options, I provide a series of sensitivity analyses. Results from these analyses are presented in Section F in the Appendix and are briefly discussed here. First, I estimate the main RD specification for different bandwidths. Specifically, the optimal bandwidth is approximately 5000, and I provide point estimates and confidence intervals for bandwidths from 3500 to 7500, in increments of 500. Results from these estimations are presented in Figure F.2. Evidently, the point estimates are very similar across different bandwidths, but the confidence intervals decrease as the bandwidth increases. Second, I estimate the main RD specification for different polynomials (linear and quadratic) and for different kernels (uniform and Epanechnikov). Again, while point estimates are very similar across specifications, the precision varies somewhat, with the uniform kernel providing the lowest precision.

Finally, since the plots differ in population density, they contain different amount of information. In particular, a densely populated plot may carry more information than a sparsely populated plot. On the other hand, it should be noted that population density in my sample is relatively homogeneous across plots, and furthermore, there is i) no jump in population density at the border, and ii) there is no statistically significant difference in population density at opposite sides of the border. Nevertheless, I estimate the model excluding plots below the 10th percentile and above the 90th percentile of population density (i.e., number of households per plot).[18] Although this reduces the

[17] See https://www.scb.se/hitta-statistik/redaktionellt/tredubbling-av-elbilar-pa-tva-ar2/ for statistics on the number of BEVs and PHEVs, and https://www.trafa.se/globalassets/statistik/vagtrafik/korstrackor/2021/korstrackor-2021.pdf for driving distance.

[18] An alternative would be to weight observations by their population density. However, it is not obvious how to include such weights in the RD model, and I therefore use the simpler approach of

sample size and, therefore, the precision, the estimated treatment effect is very similar to the main specification (approximately -8.5), suggesting that the population density per plot does not influence the results to any greater extent.

3.6 Poisson pseudo-maximum likelihood

The RD approach outlined above is a robust quasi-experimental method, but it does come with limitations. Most importantly, RD estimates the treatment effect precisely at the cutoff point where the assignment changes. This specificity means that the results are highly localized and may not be applicable to other points in the distribution or to the broader population. This localized nature of RD estimates can thus limit the generalizability of the findings. Furthermore, I have so far only used data around the border between price areas 3 and 4, and for the year 2022. Since the price discontinuity at the border vary in size over time and space, any analysis that includes more than one year and border needs to include electricity price as a continuous treatment variable. This also allows me to express the response to prices as an elasticity.

To validate the findings obtained from the RD approach, and to estimate the response to electricity prices for all data, I estimate a Poisson Pseudo-Maximum Likelihood (PPML) model (Silva and Tenreyro, 2006) using aggregate plot-level data covering all sample years and several price area borders.[19] As in the RD analysis, EV demand per 1000 households is the outcome variable, and the log of electricity price (€/kWh) is the independent variable of interest. I control for the same set of covariates as in the RD model.

To account for unobserved heterogeneity in the PPML model, I include price area border dummy variables in the spirit of Black (1999). Specifically, I include dummy variables for each price area border. For example, the dummy variable for the border between price area 3 and 4 takes the value one if plot i is located within 20 km from this border and zero otherwise, and similar for the remaining borders. Price areas 1 and 2 faced identical prices during the sample period. This means that there is no price variation within that border region, and therefore this border region is excluded from the estimation.

In principle, this methodology is equivalent to calculating differences in mean EV demand per 1000 households on opposite sides of price area borders (controlling for household characteristics at the plot level) and relating this to differences in electricity prices. Importantly, this accounts for unobserved heterogeneity that is specific to any given price area border region, such as car supply, charging infrastructure or climate.

excluding outliers from the estimation.

[19]This model is particularly well-suited for this analysis as it effectively handles count data and is robust to heteroskedasticity, ensuring that the estimates remain consistent and unbiased even when the variance of the errors is not constant. Moreover, this approach allows me to directly estimate the price elasticity of EV demand.

In other words, by comparing households that are geographically close but on opposite side of the borders, I can isolate the effect of electricity prices on EV demand. This approach allows me to pool the data across price areas and years, while still accounting for unobserved heterogeneity in space in a similar spirit to the RD approach (Busse et al. (2006) refers to this as a regression discontinuity approach with a continuous treatment.)

Formally, the regression equation is

$$Q_{jt} = \exp(\beta \ln P_{jt} + \boldsymbol{Z}_{jt}\theta + \gamma_j + \gamma_t) + \epsilon_{jt} \tag{3}$$

where Q is the number of EVs bought per 1000 households in year t and plot j, $\ln P_{jt}$ is the log electricity price, $\boldsymbol{Z}$ is a vector of household characteristics at the plot level (income, number of cars, commuting distance, number of kids, age and vintage of existing cars), γ_j is the price area border fixed effect and γ_t is a year fixed effect. Note that year fixed effects account for the price of cars and the price of petrol and diesel, since these variables only vary across time, not space. The key parameter of interest is β, the price elasticity of EV demand with respect to electricity prices.

Estimation results from the PPML model are presented in Table 5 for four different specifications. In the first column, I include all price areas (2, 3 and 4) and years but no spatial fixed effects (price area border and longitude fixed effects) and no household characteristics, only year fixed effects. In the second column, I present estimation results for a similar specification but where I add price area border and longitude fixed effects. In the third column, I also add household characteristics as control variables. The fourth column present results for a similar specification as in the third column, but this time only for price areas 3 and 4. Thus, the data used in the fourth column corresponds approximately to the data used to estimate the RD, except for that I here include all years.

For the first specification with no spatial fixed effects, the estimated effect of electricity price on EV demand is positive. However, once I add spatial fixed effects, the effect of the electricity price on EV demand is negative. Adding more control variables improves the model fit (as indicated by, for example, the psedo-R2 in the bottom panel of Equation (3)) and also improve the precision of the estimated price elasticity slightly. The estimated elasticity in specification 3 is −0.568, meaning that as EV demand is inelastic. Finally, in specification 4, the estimated elasticity is −2.132, which is very similar to the RD estimates presented previously, which implies an elasticity of approximately −2. In addition, I provide estimation results for a specification identical to column 3 but for different bandwidths in Table G.1 in the Appendix. Reducing the bandwidth increases the estimated elasticity, but leads to lower precision.

The fact that the RD and PPML estimates are similar is an important finding. Remember that the RD estimates capture the response to contemporaneous and historical

	(1) EVs per 1000 hh.	(2) EVs per 1000 hh.	(3) EVs per 1000 hh.	(4) EVs per 1000 hh.
ln(€/kWh)	0.290**	-0.670***	-0.568***	-2.132***
	(0.133)	(0.216)	(0.207)	(0.336)
Household inc.			0.109***	0.116***
			(0.011)	(0.013)
No. of cars			0.280***	0.286***
			(0.063)	(0.069)
Commuting dist.			0.000*	0.000
			(0.000)	(0.000)
No. of kids			0.231***	0.195***
			(0.041)	(0.045)
Age			0.008***	0.009***
			(0.003)	(0.003)
Vintage			0.028***	0.022***
			(0.007)	(0.008)
Constant	3.476***	0.944	-55.660***	-48.150***
	(0.355)	(0.576)	(15.075)	(16.323)
Spatial Fe	No	Yes	Yes	Yes
Price areas	2-4	2-4	2-4	3-4
Observations	55325	55325	49390	39315
ll	-1974250.475	-1956376.028	-1572775.112	-1290787.686
aic	3948504.949	3912756.057	3145566.224	2581591.372
bic	3948522.791	3912773.899	3145636.684	2581660.007
pr2	0.045	0.054	0.078	0.073

Standard errors in parentheses

* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Table 5: Poisson Psedo-Maximum Likelihood, EVs per 1000 households

price differentials, allowing for past variations in prices to influence behavior over time. On the other hand, the PPML model focuses on the response to current, or contemporaneous, prices. If households (or households' price expectations) respond more to historical price differentials, we would expect the RD estimates to be large in absolute value, compared to the PPML estimates. However, the fact that they are very similar suggests that the demand for EVs is primarily driven by current prices rather than historical price trends.

The difference between columns 3 and 4 in Table 5 suggests that consumers located at the border between price areas 3 and 4 are much more responsive to the electricity price than consumers located at the border between price areas 2 and 3. There are several possible explanations to this difference. Importantly, it is important to remember that EV adoption is much more pronounced in the south of Sweden, compared to the north (e.g., price areas 1 and 2), and that the electricity increased much more in the south of Sweden than in the north. This may affect the response to prices. For example, consumers in the south of Sweden may be more concerned about price increases in the future. Furthermore, climate and temperature may be important determinants in the north of Sweden (because of cold and long winters), and that this might make consumers less responsive to prices.

In addition, there are differences in average income levels and socioeconomic status across regions, which may influence price sensitivity and not only the level of demand. For example, in the south of Sweden, income is higher and commuting distances are shorter than in the north. In particular, the south of Sweden has more urban areas compared to the north, and urban residents might have different driving patterns and energy consumption habits, making them more sensitive to electricity prices. In the next section (Section 3.7), I explore heterogeneity in price responsiveness in more detail.

3.7 Heterogeneity in price responsiveness

The final part of my empirical analysis explores how sensitivity to electricity prices varies across different socio-economic characteristics, including age, income, number of cars, and commuting distance. For this analysis, I utilize disaggregated household-level data. I estimate a logistic regression with EV demand per household as the outcome variable, log of electricity price as the key independent variable of interest, and I control for covariates at the household level rather than at the plot level. To allow for non-linear effects of household characteristics, I split each variable (income, age, commuting distance and vintage of existing cars) and include dummy variables for each decile. For the variables number of children and number of cars, I include dummy variables for each level of these variables (1-3 cars[20] and 0-4 children). I then evaluate the marginal effect

[20] The number of households purchasing an EV with zero existing cars are virtually non-existent in my sample

of the log electricity price across the household characteristics deciles to understand how the price sensitivity vary across consumer types.

Formally, the regression equation is

$$\Pr[q_{it} = 1] = \frac{1}{1 + \exp\left[-(\beta_0 + \beta_1 \ln(p)_{it} + \sum_{k=1}^{K} \sum_{j=1}^{J_k} \beta_{kj} D_{kj,i} + \gamma_j + \gamma_t)\right]} \tag{4}$$

where q_{it} takes the value one if household i purchase an EV in year t, p_{it} is the electricity price and $D_{kj,i}$ is a dummy variable for the j-th category of the k-th categorical control variable, taking the value 1 if the k-th control variable equals j and 0 otherwise. K is the number of categorical control variables (in my case six: income, age, commuting distance, vintage, number of children and number of cars) and J_k is the number of categories for the k-th control variable (10 categories for income, age, commuting distance, vintage, 5 categories for number of children and 3 categories for number of cars). As before, γ_i is the price area border fixed effect, and γ_t is the year fixed effect, which, in a similar spirit to before account for the supply of cars and the price of petrol and diesel. I use robust standard errors.

Predicted probabilities across household characteristics are presented in Figures H.1 to H.3 in the Appendix and briefly summarized here. In general, there is substantial heterogeneity in EV demand across income and commuting deciles and across the number of children and cars, and to a lesser extent across age and vintage deciles. Specifically, the probability of purchasing an EV is much higher for high-income households and households with long commuting distances. Furthermore, young and old households are more likely to purchase an EV, compared to middle-aged households, and the probability increases with the number of cars. Households with two children are more likely to purchase an EV than households with zero or four children.

Next, I turn to the marginal effect of a 1% increase in the electricity price on the probability of purchasing an EV, conditional on the income decile a household belongs to. Given the specification in Equation (4), the marginal effect describes the percentage points change in the probability of purchasing an EV from a 1% change in the electricity price. The levels of the household characteristics across deciles are presented in Table C.2 in the Appendix.

The estimated price responsiveness across deciles for each of the regressions described above are presented in Figures 6 to 8. The estimated effect of price on the outcome is on average approximately 0.015, meaning that a 1% increase in price is a decrease in the probability of purchasing an EV of approximately -0.015% points. The mean outcome of approximately 0.015 (1.5% of households purchase an EV). This means that a 0.015% points reduction in the probability is equal to a 1.3% decrease in probability. The magnitude of this effect is relatively close to the effect estimated using the PPML model (in which a 1% increase in price reduces demand by between 0.568%-2.132%, depending on specification).

To compare the estimates in the figures to those estimated previously using the PPML model, the average marginal effect of a 1% increase in price is a decrease in the probability of purchasing an EV of approximately -0.015% points according to the results in Figures 6 to 8. The mean outcome in the data is approximately 0.015, (i.e., 1.5% of households purchase an EV). A 0.02% points reduction in the probability is then equal to a 1.3% decrease in probability. If we interpret the outcome as EVs per household, so that a 1% increase in price reduces the number of EVs per household by 1.3%, the estimated effect here is relatively close to the PPMP estimates (in which a 1% increase in price reduces EVs per 1000 households by between 0.568%-2.132%, depending on specification).

Turning to heterogeneity across consumer types, and beginning with the effect of household income on price responsiveness as presented in the left panel of Figure 6, price sensitivity is highest among income deciles 1-5 and then decreases for higher income deciles. For the top income decile, price sensitivity is indistinguishable from zero. The right panel of the same figure shows price sensitivity across age groups, revealing that both young and old households are slightly more sensitive to electricity prices than middle-aged households. However, for both household income and age, the differences in price responsiveness are small, as indicated by the overlapping confidence intervals. This suggests that consumers are relatively homogeneous in terms of price response across both income and age.

Next, Figure 7 shows that price sensitivity decreases with both commuting distance (left panel) and vintage of existing cars (right panel). However, the differences are small. For the top decile of commuting distance, the differences in price responsiveness across deciles are not statistically significant.

Finally, the left panel of Figure 6 illustrates that households with either no children or four children are relatively more responsive to electricity prices than households with 1-3 children. Additionally, as is shown in the right panel, consumers with no previous car are more responsive to electricity prices than those who already have one or more cars. However, the differences in price responsiveness between households owning one, two, or three cars are statistically insignificant.

To summarize, while there is some heterogeneity in price responsiveness across consumer types, the differences are in general small. This is an important result: first, it shows that heterogeneity in price responsiveness does not appear to contribute to heterogeneity in EV demand to any larger extent. Secondly, the results also suggest that it may be difficult to target subsidies to electricity prices to specific consumer groups, since most consumers would respond relatively similar to such policies.

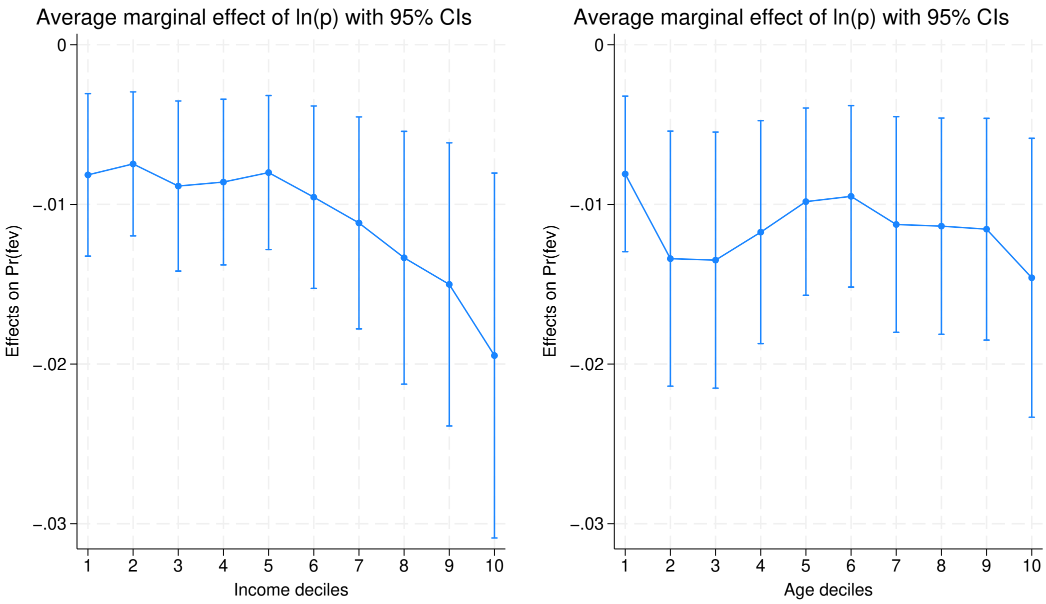

Figure 6: Price responsiveness across income and age

Note: i) The figure presents the marginal effect of a 1% increase in electricity price on EV demand across income and age deciles, as described in Equation (4). ii) The levels of the household characteristics across deciles are presented in Table C.2 in the Appendix

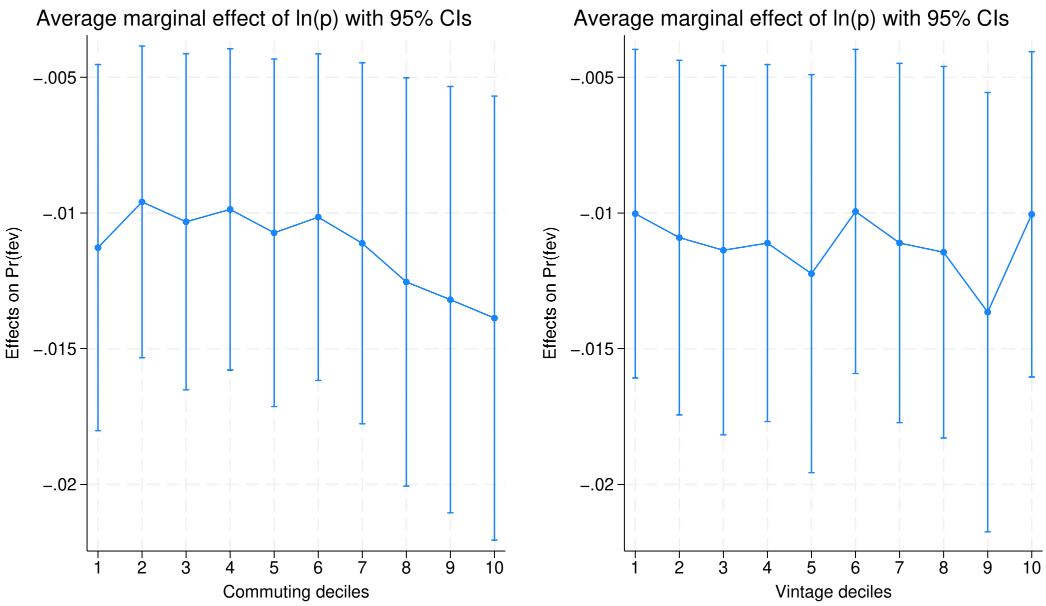

Figure 7: Price responsiveness across commuting distance and vintage

Note: i) The figure presents the marginal effect of a 1% increase in electricity price on EV demand across commuting distance and vintage deciles, as described in Equation (4). ii) The levels of the household characteristics across deciles are presented in Table C.2 in the Appendix

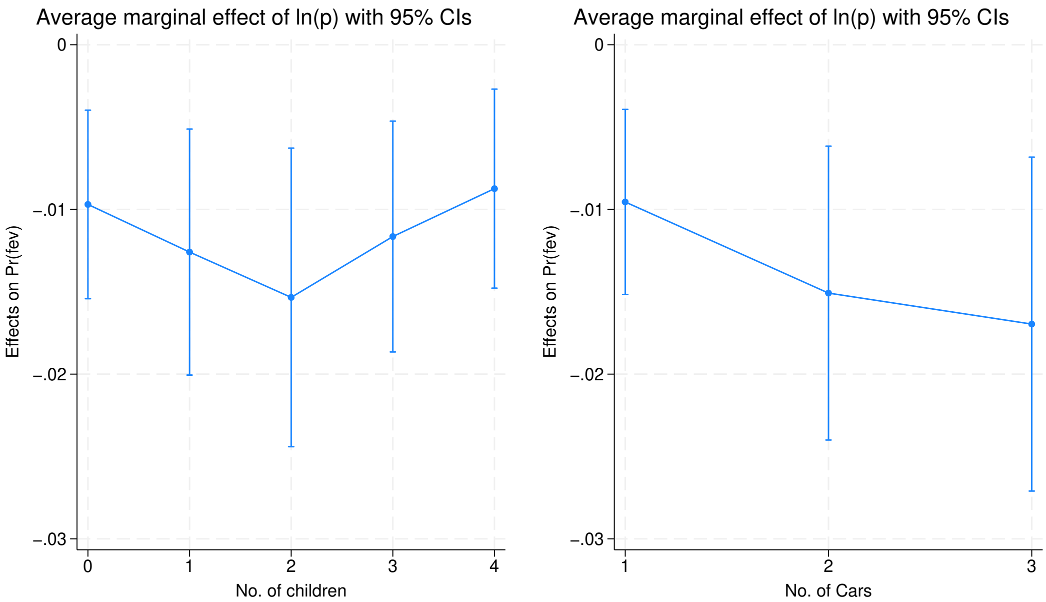

Figure 8: Price responsiveness across number of children and cars

Note: i) The figure presents the marginal effect of a 1% increase in electricity price on EV demand across commuting distance and vintage deciles, as described in Equation (4). ii) The number of households purchasing an EV with zero existing cars are virtually non-existent in my sample.

4 Policy implications

I have so far demonstrated that high electricity prices lead to a reduction in the demand for EVs. This reduction is influenced by both income effects, where higher electricity costs reduce overall purchasing power, and substitution effects, where consumers opt for CVs instead of EVs due to the increased operational costs. Furthermore, EVs are predominantly purchased by high-income households, who exhibit a greater sensitivity to electricity prices compared to low-income households. However, the differences in sensitivity between income groups are relatively small.

These findings have several important policy implications. First, they suggest that subsidizing electricity could effectively increase the demand for EVs. This type of policy could serve as an alternative to direct subsidies for car prices or investments in charging infrastructure. While subsidies for car prices or infrastructure represent a one-time cost for policymakers, an electricity subsidy would likely need to be maintained over several years. This is because consumers tend to base their purchase decisions on the expected operational costs over the period they own the car, which, according to the data, averages around five years.

Second, if there are both income and substitution effects from electricity price shocks, it may be necessary to subsidizing both the electricity used specifically for charging EVs to make them more cost-competitive relative to CVs, as well as subsidizing baseline electricity consumption to enhance overall household purchasing power. In any case, it would be challenging for policymakers to distinguish between different types of electricity use, making the implementation of a policy that targets only charging very complex.

Third, unless policymakers can precisely target the subsidy to those who would not have purchased an EV without it, many households that benefit from the policy would have bought an EV regardless. This makes it difficult to ensure that the subsidy is effectively encouraging additional EV purchases rather than simply providing financial benefits to existing buyers (this is the case for other policies as well; see, for example, Haan et al., 2024; Springel, 2021).

Fourth, a uniform electricity subsidy is likely to be regressive. This is because EVs are mainly bought by high-income households, which are more sensitive to changes in electricity prices. Consequently, a subsidy targeting low-income households would need to be larger to have a meaningful impact, as these households are less responsive to electricity price changes.

To accurately assess the efficiency of an electricity subsidy in terms of generating additional EV purchases, and to compare it with alternative policies like subsidies for car prices and charging infrastructure, I present a policy counterfactual that provides a rough estimate of the impact of an electricity subsidy based on the empirical findings in this paper and some strong but plausible assumptions. Specifically, I consider the effects of a counterfactual removal of the Swedish electricity tax. The electricity tax in Sweden is €0.045 per kWh for households in southern Sweden and €0.035 per kWh for households in northern Sweden.[21] In my counterfactual scenario, the electricity tax is set to the EU minimum level (0.0005€/kWh)[22] for EV buyers during the lifespan of the EV. In my data, the average replacement rate of cars is approximately five years. To account for both income and substitution effects of electricity prices, the policy considered here is a subsidy on all electricity consumption, and not just the electricity used for charging the EV.

I assume that the incidence of the tax completely falls on consumers. Thus, a tax cut leads to a same-size reduction in consumers' electricity prices. This appears reasonable, given that electricity demand is close to completely inelastic in the short run (e.g., Lanot and Vesterberg, 2021) and that supply is considerably more elastic (Lundin and

[21] The motivation for this tax was originally to raise fiscal revenues for the government, but it is also used to provide monetary incentives for promoting energy conservation and efficiency (for a detailed discussion of the electricity tax, see Brännlund and Kriström, 2020) A tax on electricity is problematic for several reasons. Firstly, it is unclear what market failure this tax is intended to address, given that the Swedish electricity generation mix is clean in terms of CO2 emissions (see Section A in the Appendix). Secondly, it is particularly problematic to tax electricity when electrification is crucial for achieving climate goals. Taxing electricity in this context can be counterproductive, as it may discourage the adoption of EVs and other electric technologies essential for reducing greenhouse gas emissions. Indeed, this is precisely what the results presented in this paper show.

[22] See https://taxation-customs.ec.europa.eu/taxation/excise-taxes/excise-duties-energy_en

Tangerås, 2020). See also Fabra and Reguant (2014). I also assume that the increased demand for EVs due to the tax cut does not lead to higher vehicle prices.[23]

Using the estimated price elasticity from column 3 in Table 5 in Section 3.6 as a measure of the response to this policy change, this policy would result in approximately 22,924 additional EVs in Sweden (20,812 in the south and 2112 in the north). Using 2022 as the baseline scenario, this corresponds to an increase of EV demand by approximately 21%, which is a sizable increase.

To compare this policy to subsidies on car prices and charging infrastructure, I need to relate the increase in EV demand to resources spent by the government (in this case, foregone tax revenues). I assume that electricity consumption is on average 15,000 kWh per household per year. This corresponds approximately to the figures presented in recent studies of residential electricity consumption in Sweden; see, for example, Lanot and Vesterberg (2019, 2021) and Vesterberg (2018). Previous literature have repeatedly shown that residential electricity consumption is close to completely inelastic in the short run, and this baseline demand is therefore assumed fixed in the short run. Further, I assume that EVs replace conventional cars one-for-one, and that cars on average are driven 12,000 km per year (this corresponds approximately to the average mileage in the data), irrespective of whether it is an EV or CV. Bushnell et al. (2022) makes similar assumptions in their policy analysis. The average kWh/km is 0.2 in my data, meaning that cars consume 0.2 kWh per km driven. Given these assumptions, the policy results in a reduction in tax revenues of €86,326,194 per year.[24]

The increase in EVs and the associated foregone tax revenue can be compared to the estimated cost of subsidies for either cars or charging stations provided in the previous literature. For example, Springel (2021), using Norwegian household-level data, estimate that every €100,000 spent on charging station subsidies results in 13.5 additional EV purchases, while the same amount spent on purchase price subsidies lead to 4.8 additional EVs being purchased. Haan et al. (2024) finds that EV subsidies in Germany increase EV purchases by approximately 7 EVs per €100,000, and Clinton and Steinberg (2019) and Muehlegger and Rapson (2022) finds similar effects of EV subsidies in the US.

In comparison, removing the electricity tax leads to approximately 5.3 more EVs per €100,000 spent, which suggests that this policy is less efficient per spent tax money than subsidies to charging stations, and similar in size to the effect of subsidies to car prices.

[23]This assumption can, at least to some extent, be motivated by the recent years' slowdown in EV demand, suggesting that there is an excess supply of EVs. In particular, many automakers have overestimated the demand for EVs, leading to higher production levels than the market can currently absorb. An increased demand for EVs (due to the tax cut) should therefore not lead to higher car prices.

[24]I arrive at this number by multiplying the tax reduction by the number of EVs bought, the electricity use (baseline plus charging), the electricity use per km and the km's driven.

5 Conclusions

In this paper, I estimate the causal impact of electricity prices on the demand for EVs by leveraging unique and highly detailed registry data from Sweden. This analysis takes advantage of the discontinuities in electricity prices across Sweden's four distinct price areas. I employ two empirical strategies to exploit this feature: a non-parametric RD design, and a PPML regression where price area border fixed effect facilitates the comparison between observations located close to each other but at opposite side of the price area borders. Both methodologies consistently indicate that electricity prices have a substantial influence on the demand for EVs. Furthermore, I show that this effect is mainly driven by BEVs, and that electricity prices reduces the demand not only for EVs, but for CVs as well.

Next, I use a logit regression on household-level data to explore heterogeneity in price responsiveness across a broad range of socio-economic household characteristics. I show that high-income households and households with many children, and, to a lesser extent, households with many cars and longer commuting distances, are relatively less responsive to prices, and that older households respond stronger to electricity prices.

The finding that high electricity prices reduces the demand for EVs is particularly concerning given the broader context of ongoing electrification efforts. As industries and transportation sectors continue to electrify, the demand for electricity will increase, potentially driving prices even higher. This creates a paradoxical situation: the more society moves towards electrification, the higher the electricity prices become, which in turn makes further electrification less appealing. This catch-22 underscores the need for careful consideration in policy design to balance the benefits of electrification with the economic realities of electricity pricing. Policy makers may respond to this either by subsidizing (or in other ways enabling) an expansion of electricity supply, or through subsidizing electricity prices for EV buyers.

The implications of these findings are important for policy-making. The transition to EVs is a critical component of global strategies aimed at reducing greenhouse gas emissions and mitigating climate change. However, recent years have seen a deceleration in the growth of EV demand. This paper posits that elevated electricity prices may be a key contributing factor to this slowdown, alongside other economic pressures such as rising interest rates and inflation. Nevertheless, subsidizing electricity prices may be less efficient in boosting EV demand than same-sized subsidies to charging infrastructure. To illustrate this, I use the estimated effects to consider a counterfactual policy scenario in which the Swedish electricity tax is removed for households purchasing an EV. I show that this policy is likely less cost-efficient in terms of increasing EV demand than alternative policies such as, for example, subsidies to either EVs or charging infrastructure.

References

Allcott, H. and N. Wozny. “Gasoline prices, fuel economy, and the energy paradox”. In: Review of Economics and Statistics 96.5 (2014), pp. 779–795.

Barreca, A. I. et al. “Saving babies? Revisiting the effect of very low birth weight classification”. In: The Quarterly Journal of Economics 126.4 (2011), pp. 2117–2123.

Bigler, P. and D. M. Radulescu. “Environmental, redistributive and revenue effects of policies promoting fuel efficient and electric vehicles”. In: CESifo Working Paper (2022).

Black, S. E. “Do better schools matter? Parental valuation of elementary education”. In: The Quarterly Journal of Economics 114.2 (1999), pp. 577–599.

Borenstein, S. and L. W. Davis. “The distributional effects of US tax credits for heat pumps, solar panels, and electric vehicles”. In: National Bureau of Economic Research (2024).

Brännlund, R. and B. Kriström. “Svensk energi och miljöbeskattning–ett reformförslag”. In: (2020).

Burkhardt, J., K. T. Gillingham, and P. K. Kopalle. “Field experimental evidence on the effect of pricing on residential electricity conservation”. In: Management Science 69.12 (2023), pp. 7784–7798.

Bushnell, J. B., E. Muehlegger, and D. S. Rapson. Energy prices and electric vehicle adoption. Tech. rep. National Bureau of Economic Research, 2022.

Busse, M., J. Silva-Risso, and F. Zettelmeyer. “$1,000 cashback : Thepass - throughofautomanufacturerpromotions$”. In: American Economic Review 96.4 (2006), pp. 1253–1270.

Busse, M. R., C. R. Knittel, and F. Zettelmeyer. “Are consumers myopic? Evidence from new and used car purchases”. In: American Economic Review 103.1 (2013), pp. 220–256.

Butts, K. “Geographic difference-in-discontinuities”. In: Applied Economics Letters 30.5 (2023), pp. 615–619.

Calonico, S., M. D. Cattaneo, and R. Titiunik. “Robust nonparametric confidence intervals for regression-discontinuity designs”. In: Econometrica 82.6 (2014), pp. 2295–2326.

— “Optimal data-driven regression discontinuity plots”. In: Journal of the American Statistical Association 110.512 (2015), pp. 1753–1769.

Calonico, S. et al. “rdrobust: Software for regression-discontinuity designs”. In: The Stata Journal 17.2 (2017), pp. 372–404.

— “Regression discontinuity designs using covariates”. In: Review of Economics and Statistics 101.3 (2019), pp. 442–451.

Carley, S., S. Siddiki, and S. Nicholson-Crotty. “Evolution of plug-in electric vehicle demand: Assessing consumer perceptions and intent to purchase over time”. In: Transportation Research Part D: Transport and Environment 70 (2019), pp. 94–111.

Cattaneo, M. D. and R. Titiunik. “Regression discontinuity designs”. In: Annual Review of Economics 14.1 (2022), pp. 821–851.

Cattaneo, M. D., R. Titiunik, and G. Vazquez-Bare. “Power calculations for regression-discontinuity designs”. In: The Stata Journal 19.1 (2019), pp. 210–245.

Chaisemartin, C. de et al. “Difference-in-differences estimators for treatments continuously distributed at every period”. In: arXiv preprint arXiv:2201.06898 (2022).

Clinton, B. C. and D. C. Steinberg. “Providing the Spark: Impact of financial incentives on battery electric vehicle adoption”. In: Journal of Environmental Economics and Management 98 (2019), p. 102255.

Dorsey, J., A. Langer, and S. McRae. “Fueling alternatives: Gas station choice and the implications for electric charging”. In: National Bureau of Economic Research (2022).

Fabra, N. and M. Reguant. “Pass-through of emissions costs in electricity markets”. In: American Economic Review 104.9 (2014), pp. 2872–2899.

Gelman, A. and G. Imbens. “Why high-order polynomials should not be used in regression discontinuity designs”. In: Journal of Business & Economic Statistics 37.3 (2019), pp. 447–456.

Gonzalez-Salazar, M., G. Kormazos, and V. Jienwatcharamongkhol. “Assessing the economic and environmental impacts of battery leasing and selling models for electric vehicle fleets: A study on customer and company implications”. In: Journal of Cleaner Production 422 (2023), p. 138356.

Greenstone, M. et al. “Can technology solve the principal-agent problem? Evidence from China’s war on air pollution”. In: American Economic Review: Insights 4.1 (2022), pp. 54–70.

Grembi, V., T. Nannicini, and U. Troiano. “Do fiscal rules matter?” In: American Economic Journal: Applied Economics (2016), pp. 1–30.

Haan, P., A. Santonja, and A. Zaklan. “Effectiveness and heterogeneous effects of purchase grants for electric vehicles”. In: Environmental and Resource Economics (2024), pp. 1–39.

Halse, A. H. et al. “Local incentives and electric vehicle adoption”. In: Journal of the Association of Environmental and Resource Economists 12.1 (2025), pp. 145–180.

Hardman, S., E. Shiu, and R. Steinberger-Wilckens. “Comparing high-end and low-end early adopters of battery electric vehicles”. In: Transportation Research Part A: Policy and Practice 88 (2016), pp. 40–57.

Hardman, S. et al. “A perspective on equity in the transition to electric vehicle”. In: MIT Sci. Policy Rev 2 (2021), pp. 46–54.

Holmberg, P. and T. P. Tangerås. “The Swedish electricity market–today and in the future”. In: Sveriges Riksbank. (2023).

Imbens, G. W. and T. Lemieux. “Regression discontinuity designs: A guide to practice”. In: Journal of Econometrics 142.2 (2008), pp. 615–635.

Jenn, A., K. Springel, and A. R. Gopal. “Effectiveness of electric vehicle incentives in the United States”. In: Energy policy 119 (2018), pp. 349–356.

Keele, L. J. and R. Titiunik. “Geographic boundaries as regression discontinuities”. In: Political Analysis 23.1 (2015), pp. 127–155.

Kreiss, A. and C. Rothe. “Inference in regression discontinuity designs with high-dimensional covariates”. In: The Econometrics Journal 26.2 (2023), pp. 105–123.

Lanot, G. and M. Vesterberg. “An empirical model of the decision to switch between electricity price contracts”. In: Journal of Business Analytics 2.1 (2019), pp. 24–46.

— “The price elasticity of electricity demand when marginal incentives are very large”. In: Energy Economics 104 (2021), p. 105604.

Lee, D. S. “Randomized experiments from non-random selection in US House elections”. In: Journal of Econometrics 142.2 (2008), pp. 675–697.

Lee, D. S. and T. Lemieux. “Regression discontinuity designs in economics”. In: Journal of Economic Literature 48.2 (2010), pp. 281–355.

Lehner, A. “A note on spatial regression discontinuity designs”. In: (2021).

Li, K., T. Zhou, and B. Liu. “The comparison between selling and leasing for new and remanufactured products with quality level in the electric vehicle industry.” In: Journal of Industrial & Management Optimization 17.3 (2021).

Lundin, E. and T. P. Tangerås. “Cournot competition in wholesale electricity markets: The Nordic power exchange, Nord Pool”. In: International Journal of Industrial Organization 68 (2020), p. 102536.

Mauritzen, J. “With great power (prices) comes great tail pipe emissions? A natural experiment of electricity prices and electric car adoption”. In: Transportation Research Part A: Policy and Practice 192 (2025), p. 104356.

Meng, X. and Y. Yu. “Does the Russia-Ukraine conflict affect gasoline prices?” In: Energy Economics 128 (2023), p. 107113.

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

Münzel, C. et al. “How large is the effect of financial incentives on electric vehicle sales?–A global review and European analysis”. In: Energy Economics 84 (2019), p. 104493.

Nehiba, C. “Electric vehicle usage, pollution damages, and the electricity price elasticity of driving”. In: Journal of Environmental Economics and Management 124 (2024), p. 102915.

Noack, C., T. Olma, and C. Rothe. “Flexible covariate adjustments in regression discontinuity designs”. In: arXiv preprint arXiv:2107.07942 (2021).

Silva, J. S. and S. Tenreyro. “The log of gravity”. In: The Review of Economics and Statistics (2006), pp. 641–658.

Singh, V., V. Singh, and S. Vaibhav. “A review and simple meta-analysis of factors influencing adoption of electric vehicles”. In: Transportation Research Part D: Transport and Environment 86 (2020), p. 102436.

Song, Z. et al. “Effects of temperature on the performance of fuel cell hybrid electric vehicles: A review”. In: Applied Energy 302 (2021), p. 117572.

Springel, K. “Network externality and subsidy structure in two-sided markets: Evidence from electric vehicle incentives”. In: American Economic Journal: Economic Policy 13.4 (2021), pp. 393–432.

Stiglitz, J. E. and A. Kosenko. “The economics of information in a world of disinformation: A survey part 2: Direct communication”. In: (2024).

Vesterberg, M. “The effect of price on electricity contract choice”. In: Energy Economics 69 (2018), pp. 59–70.

Wolbertus, R. et al. “Fully charged: An empirical study into the factors that influence connection times at EV-charging stations”. In: Energy Policy 123 (2018), pp. 1–7.

Appendix A Electricity price areas in Sweden

In 2011, Sweden’s electricity market was restructured to address regional differences in electricity supply and demand, as well as to manage transmission constraints more effectively. This restructuring resulted in the division of the country into four distinct electricity pricing areas, also known as bidding areas. The division into these four areas was implemented to better reflect the regional variations in electricity supply and demand, and to manage the transmission network more efficiently. By having distinct pricing areas, the electricity market can reflect local market conditions and encourage efficient use of the grid by highlighting areas with transmission constraints, thereby encouraging investments in grid infrastructure and local generation capacity where it is most needed.

The northernmost areas, price areas 1 and 2, are characterized by a surplus of electricity production, primarily from hydroelectric power plants. The low population density and significant hydro resources contribute to lower electricity prices compared to the other price areas. Price area 3 is located in central Sweden (and includes the capital Stockholm). This is the most densely populated area, and includes major urban centers like Stockholm and Uppsala. The high demand for electricity in this region, combined with limited local production, often results in higher electricity prices. This area relies more on electricity imports from other regions and countries. Finally, price area 4 is located in the southernmost part of Sweden. This price area experiences the highest electricity prices among the four areas. This is due to a combination of high demand, limited local production capacity, and transmission constraints. The region often imports electricity from other areas and neighboring countries. The four price areas are illustrated in Figure A.1.

One of the significant challenges facing Sweden's electricity market is the constrained transmission capacity between different regions, particularly from the north to the south. The transmission grid, managed by Svenska Kraftnät (see https://www.svk.se/), consists of about $16,000$ kilometers of high-voltage lines. Despite this extensive network, the capacity to transfer electricity from the north to the south is limited. This bottleneck means that the electricity generated in the north cannot always be efficiently transported to meet the demand in the south, leading to regional price disparities and potential supply constraints.

In more detail, the Swedish main grid is an alternating current grid with transmission in the north-south direction. The grid is built to transfer energy produced in the north to consumers in the southern parts of the country. The need for transfer varies with demand and the hydrological situation. A bottleneck is a section in the transmission network that is often at risk of being congested. Congestion risks occurring when the market demand to transmit electricity through a section is greater than what is physically possible. The demarcation of the four Swedish electricity areas follows three of the most common sections found in the electricity grid. There are four sections where bottlenecks occur frequently. Three of these cuts through the country in an east-west direction, and risk being congested when the electricity transmission goes in a north-south direction. The maximum transmission capacity across each section is not constant, but can vary from hour to hour and day to day, depending on the configuration of the grid (for example, if the grid is intact or if lines are disconnected for maintenance), production and consumption, and imports and exports.

The Swedish electricity market is deregulated, and consumers are free to choose between retailers and different type of retail contracts. In 2023, approximately 65% of all households had a retail contract with prices varying by month; 10% had a retail contract with prices varying by the hour, and the remaining households had a contract with prices fixed for a year or longer. In addition to the retail price per kWh, the consumer price of electricity includes the energy tax (roughly 0.05 €/kWh), the electricity certificate fee (0.002 €/kWh on average for 2022), sales tax (25%) and a transmission fee consisting of a variable part, which is roughly 0.01 €/kWh, and a fixed part, which varies between roughly €150 and €1500 per year, depending on the size of the household's fuse amp.

	2019	2020	2021	2022
Price area 1	0.026	0.013	0.039	0.057
Price area 2	0.036	0.013	0.039	0.06
Price area 3	0.036	0.020	0.06	0.124
Price area 4	0.038	0.024	0.074	0.146

Table A.1: Electricity price per price area and year, in €/kWh

None of these price components changed during the time of study, and except for the distribution price, they are identical across price areas.

Turning to the supply side of Sweden's electricity market, Sweden's electricity generation mix has undergone significant changes over the past 20 years, reflecting the country's commitment to sustainability and reducing carbon emissions. Historically, Sweden's electricity production was dominated by hydroelectric and nuclear power. Hydropower has been a cornerstone, consistently providing a substantial portion of the country's electricity due to Sweden's abundant water resources. Nuclear power has also played a crucial role, contributing significantly to the electricity supply.

In recent years, there has been a notable shift towards renewable energy sources. Wind power has seen remarkable growth, increasing its share in the electricity mix substantially. From a minor contributor in the early 2000s, wind power now accounts for a large portion of Sweden's electricity generation (approximately 20% as of 2024). Solar power, although still a smaller part of the mix, has also experienced rapid growth, especially during the 2010s. The use of biofuels has increased since the 1980s, and this trend has continued into the 21st century, further diversifying the energy mix3. Fossil-fuel based thermal power, once a more prominent part of the mix, has been largely phased out in favor of cleaner energy sources.

Overall, Sweden's electricity generation mix has evolved to become one of the cleanest in the world, with a strong emphasis on renewable energy and low-carbon sources. This transition reflects Sweden's ongoing efforts to decarbonize its energy system and lead by example in the global push towards sustainable energy.

For a more detailed description of the Swedish electricity market, see Holmberg and Tangerås (2023).

Appendix B Swedish policies related to EV demand and the electricity market

During the sample period, several policies took place that influenced both the market for EVs and the electricity market. However, as I will explain in detail below, none of these policies impact the results of my analysis.

First, the so-called Bonus-Malus policy in Sweden, which was introduced on July 1,

Figure A.1: Sweden and its four price areas

Note: i) SE1 is price area 1, SE2 is price area 2, SE3 is price area 3 and SE4 is price area 4.

2018, was designed to encourage the purchase of low-emission vehicles and discourage the acquisition of high-emission ones. Initially, this policy provided financial incentives, known as bonuses, for purchasing low-emission vehicles. For instance, vehicles emitting 0 grams of CO2 per kilometer could receive a bonus of up to SEK 70,000 (€6,900). The bonus amount decreased for vehicles emitting up to 60 grams of CO2 per kilometer. However, this bonus component was discontinued on November 8, 2022. While this policy might have influenced the demand for EVs, it should not affect the estimates presented in this paper. This is because the policy was applied uniformly across all price areas, and my analysis relies on discontinuities at the borders of these areas. To ensure robustness, I have also re-estimated all models in the paper excluding observations after November 2022, and this exclusion does not alter my results in any way.

Secondly, over the past few years, the Swedish government has implemented several policies aimed at mitigating the impact of high electricity prices. Up until 2024, there have been two distinct support schemes relevant to my sample period. The first support scheme was based on electricity consumption during the period from December 2021 to February 2022. Payments from this scheme were made to all households during the second half of 2022. Since this policy did not differentiate between price areas in terms of eligibility or the amount received, it does not affect my results.

The second support scheme was based on households' electricity consumption from October 2021 to September 2022 and specifically targeted households in price areas 3 and 4. This scheme also differentiated between these two price areas in terms of the amount of support provided, with households in SE4 receiving slightly more than those in SE3. However, this policy was decided and payments were made in the spring of 2023, which is after the end of my sample period. Therefore, this policy should not affect the results presented in this paper.

Appendix C Data

Appendix D RD and power

For the RD approach, I concentrate on plots situated close to the border between price areas 3 and 4. In 2022, the electricity price in price area 3 was €0.013 cents/kWh, while in price area 4, it was €0.016/kWh. This difference of 0.024 €/kWh is quite substantial and provides a clear basis for analysis. Although the price differential between price areas 2 and 3 was even larger, the number of households living near this border is relatively small. This region of Sweden is sparsely populated, predominantly covered by forests, and has very few inhabitants per square kilometer. In contrast, the border between price areas 3 and 4 is much more densely populated. This differences in population density are illustrated in Figure D.1, which shows the distribution of households near

EV make and model	Freq.	Percent (of all EVs)	Avg. price (in Euro)
All price areas			
Kia Niro (KG021)	189	5.92	49,227
Kia Ceed (KG017)	187	5.59	37,233
Volvo XC40 (VO029)	173	5.17	55,354
Nissan Leaf (NA032)	121	3.62	50,093
Volkswagen ID.4 (VW040)	119	3.56	57,942
Price area 1 & 2			
Kia Niro (KG021)	18	7.56	50,538
Kia Ceed (KG017)	11	4.62	38,748
Nissan Leaf (NA032)	8	3.36	39,115
Peugeot other (PG888)	8	3.36	52,279
Volvo XC40 (VO029)	8	3.36	59,622
Price area 3			
Kia Niro (KG021)	120	6.4	49,436
Kia Ceed (KG017)	114	5.65	37,076
Volvo XC40 (VO029)	99	4.91	55,368
Volkswagen ID.4 (VW040)	84	4.17	58,240
Nissan Leaf (NA032)	61	3.03	51,429
Price area 4			
Volvo XC40 (VO029)	66	6.06	54,571
Kia Ceed (KG017)	62	5.69	37,089
Nissan Leaf (NA032)	52	4.77	49,035
Kia Niro (KG012)	51	4.68	48,474
Ford Kuga (FO032)	38	3.49	48,720

Table C.1: Most common EVs per price area in 2022

Decile	Income (1000€)	Age	Commuting distance (km)	Vintage
1	0.001-0.099	18-30	0-0.498	1926-2005
2	0.099-0.220	31-35	0.499-1.655	2006-2008
3	0.221-0.284	36-41	1.656-3.157	2009-2010
4	0.285-0.322	42-45	3.158-6.260	2011-2012
5	0.323-0.353	46-49	6.261-10.587	2013-2013
6	0.354-0.384	50-53	10.588-15.836	2014-2014
7	0.385-0.419	54-56	15.837-22.166	2015-2015
8	0.419-0.465	57-61	22.167-30.364	2016-2016
9	0.466-0.544	62-66	30.365-53.101	2017-2018
10	0.544-10.171	67-103	53.105-1315.651	2019-2022

Table C.2: Income, age, commuting distance and vintage across deciles

the borders of the different price areas. Plots with zero inhabitants are not included in the data, and this is the reason for the few observations close to the border between price areas 2 and 3.

It is important to note that population density is not uniform across the running variable; instead, it exhibits a pattern of bunching, reflecting the locations of villages and towns. Additionally, there are relatively few inhabited plots exactly at the border, as these borders are situated in sparsely populated areas. In Section F, I utilize the panel structure of the data to demonstrate that these characteristics do not violate the identifying assumptions. For example, I show that the observed bunching in space is not driven by price differentials.

The limited number of observations near the border between price areas 2 and 3 leads to lower statistical power, making it challenging to detect significant effects. On the other hand, the higher number of observations near the border between price areas 3 and 4 results in much higher statistical power. This is demonstrated in Figure D.2, which presents power functions for these two subsets of data. The power functions are calculated using the approach outlined in Cattaneo et al. (2019) for computing statistical power in RD settings. In this figure, Tau (measured on the horizontal axis) represents the treatment effect and the vertical axis measures statistical power. The figure clearly shows that the statistical power is substantially higher for the border between price areas 3 and 4, and it increases even further when covariates are included in the analysis (see Figure E.1 in Section E in the Appendix for a power plot for my preferred RD specification, with covariates and winsorizing).

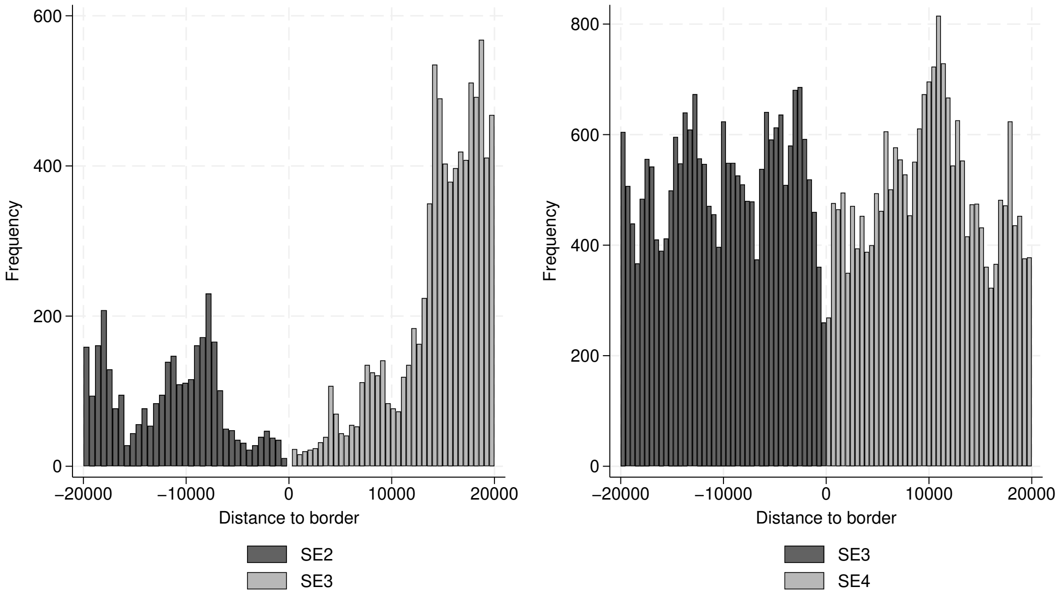

Figure D.1: Density of data close to price area borders for price areas 2 and 3, and for price areas 3 and 4.
Note: i) Observations are 1000×1000 meter plots in densely populated areas and 250×250 meter plots in sparsely populated areas (i.e., the same unit of observation as in the RD analysis. ii) Plots with zero inhabitants are not included in the data, and this is the reason for the difference in population density.

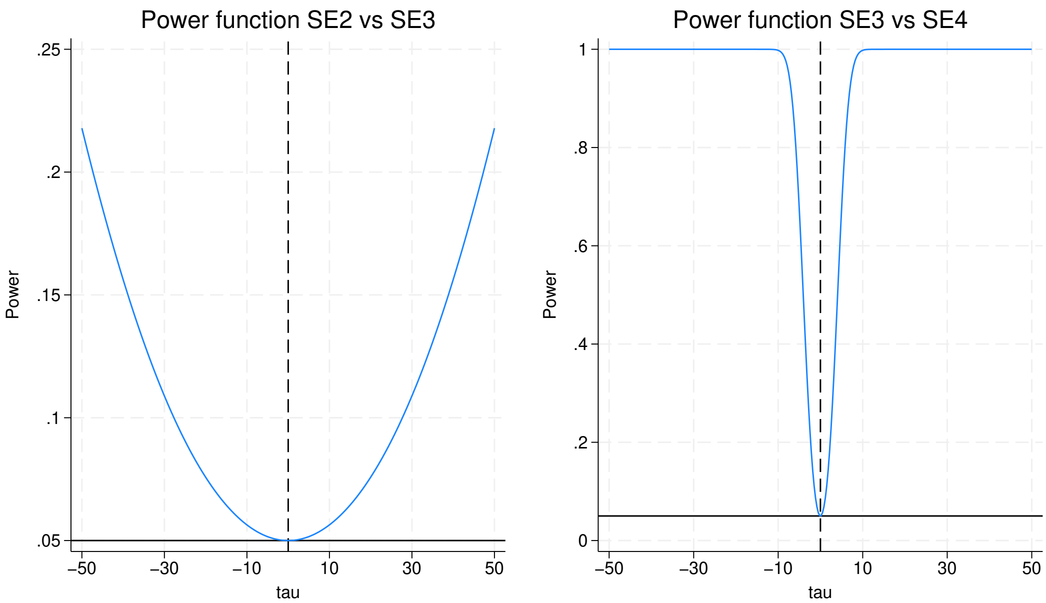

Figure D.2: Power functions for data close to price area borders for price areas 2 and 3, and for price areas 3 and 4.
Note: i) The statistical power is calculated using the approach outlined in Cattaneo et al. (2019), and without any covariates

	(1) No. of cars	(2) Commuting dist.	(3) Household inc.	(4) No. of kids	(5) Age	(6) Vintage
RD_Estimate	0.043	1517.097	100.813	-0.052	0.864	0.157
Robust 95% CI	[-.016 ; .119]	[-10612.766 ; 15117.013]	[-159.525 ; 473.753]	[-.153 ; .067]	[-1.23 ; 2.695]	[-.685 ; 1.001]
Kernel type	Triangular	Triangular	Triangular	Triangular	Triangular	Triangular
BW type	mserd	mserd	mserd	mserd	mserd	mserd
Obs.	10640	9616	10640	10640	11043	10621
Robust p-value	0.133	0.732	0.331	0.444	0.464	0.714
Order Loc. Poly. (p)	1.000	1.000	1.000	1.000	1.000	1.000
Order bias (q)	1.000	2.000	2.000	2.000	2.000	2.000

Table E.1: Placebo test on pre-treatment covariates

Note: The treatment effect is estimated using the approach detailed in Calonico et al. (2017), but using the pre-treatment covariates as the outcome variables.

	(1) EVs per 1000 hh.	(2) EVs per 1000 hh.	(3) EVs per 1000 hh.
RD_Estimate	-0.597	-1.370	-4.323*
Robust 95% CI	[-3.85 ; 3.169]	[-7.976 ; 3.817]	[-8.539 ; .858]
Kernel type	Triangular	Triangular	Triangular
BW type	mserd	mserd	mserd
Obs.	9792	9710	9645
Robust p-value	0.849	0.490	0.109
Order Loc. Poly. (p)	1.000	1.000	1.000
Order bias (q)	2.000	2.000	2.000
Price areas	3 & 4	3 & 4	3 & 4
Year	2019	2020	2021
Control variables	Yes	Yes	Yes
Winsorizing	No	No	No
Mean outcome, treatment	5.711	5.093	12.204
Mean outcome, control	4.657	6.086	13.501

Table F.1: RD estimation results for years 2019 to 2021

Note: In the first column, the outcome variable is number of EVs per 1000 households in 2019. In the second column, the outcome variable is the number of EVs per 1000 households in 2020, and in the third column, the outcome variable is the number of EVs per 1000 households in 2021. ii) For 2020 and 2021, I control for pre-treatment covariates using figures from 2019. For 2019, I use the covariates for that year. iii) Treatment effects are estimated using the robust RD approach suggested by Calonico et al. (2017) and with a triangular kernel. iv) Bandwidth is chosen using the data-driven MSE-optimal approach in Calonico et al. (2014). v) The order bias specifies the order of the local polynomial used to construct the bias correction. This bias correction helps to adjust the RD point estimator to account for the bias introduced by the polynomial approximation.

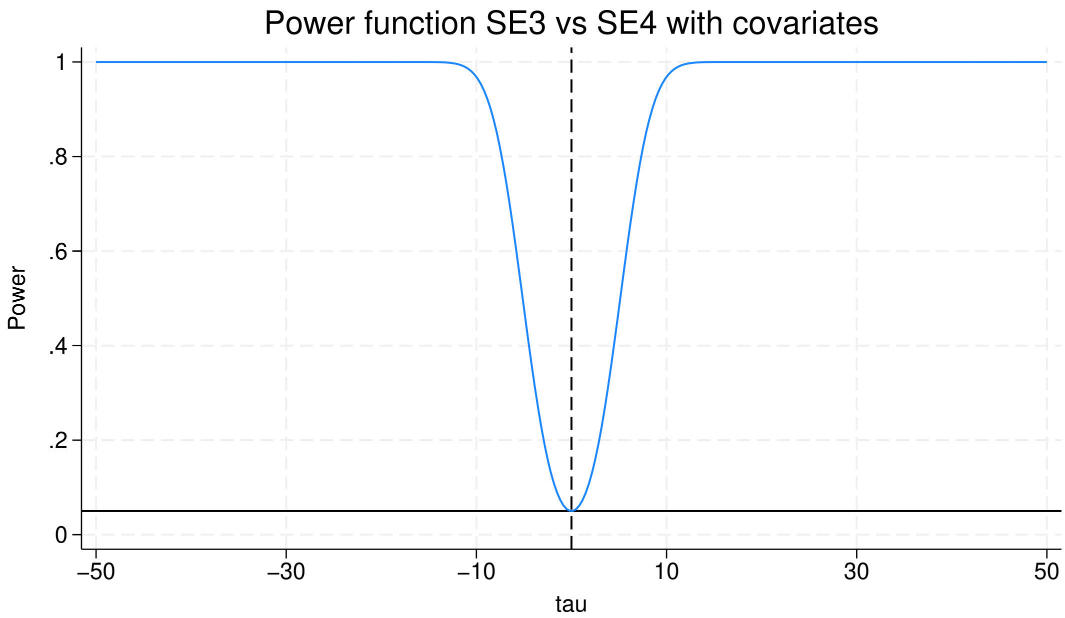

Figure E.1: Power plot, SE3 vs SE4, with covariates
Note: i) The statistical power is calculated using the approach outlined in Cattaneo et al. (2019), and with the covariates used in the RD model in Section 3.1.

Appendix E Additional RD results

Appendix F Sensitivity analysis, RD

Appendix G Additional PPML results

Appendix H Additional logit results

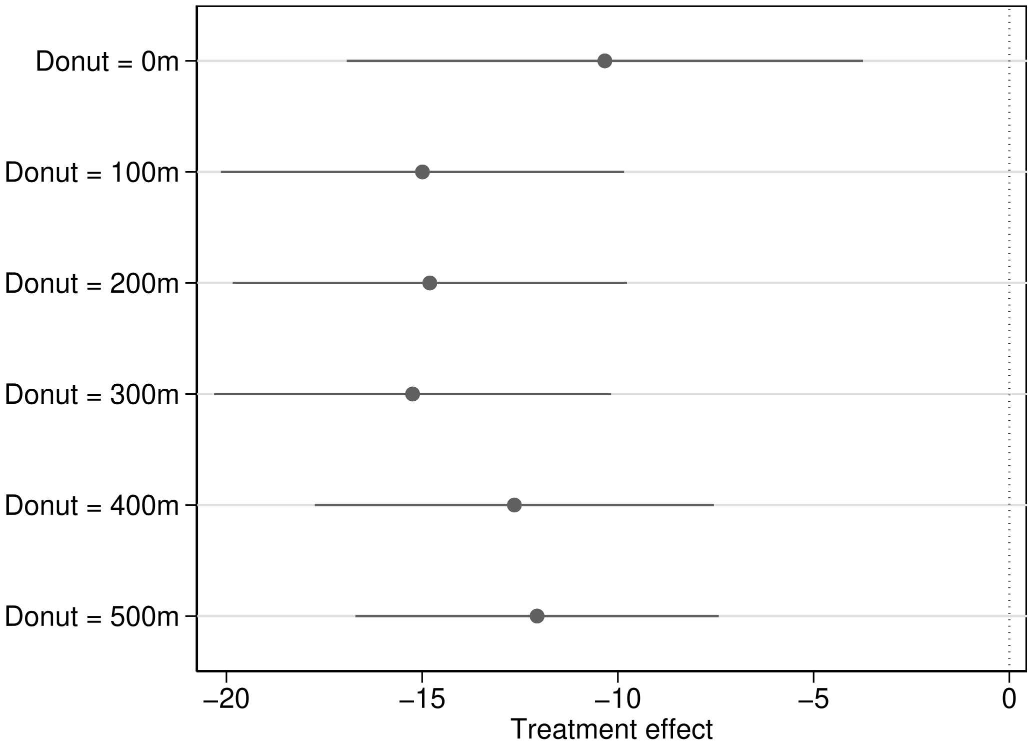

Figure E.2: Donut RD

Note: i) The treatment effect is estimated using the donut approach, where observations close to the cutoff is removed.

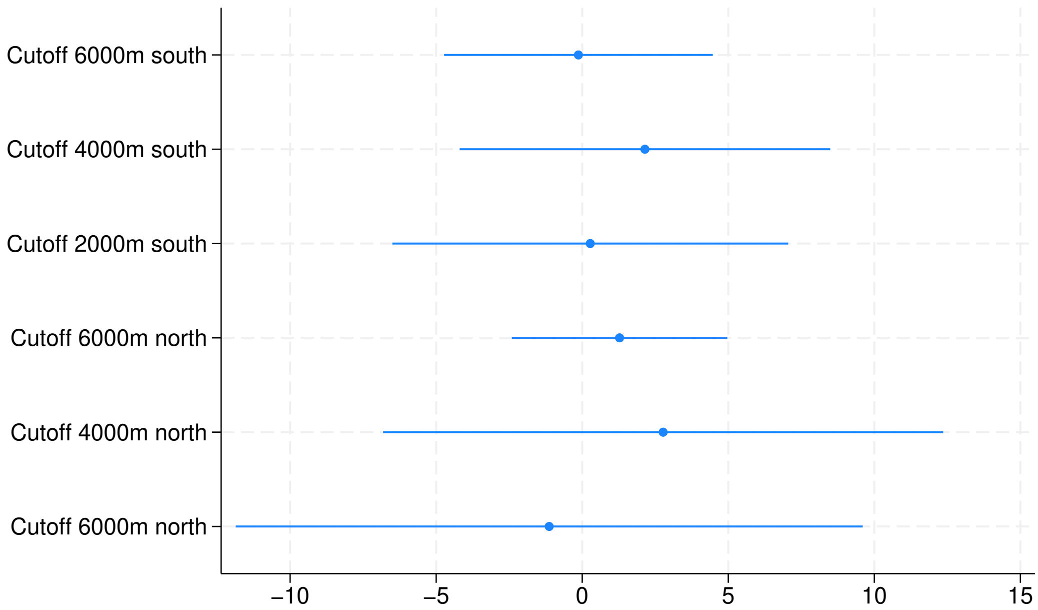

Figure F.1: RD results for placebo treatment away from the border

Note: The treatment effect is estimated using a placebo treatment away from the border ii) Price area borders are defined in the east-west direction, and the placebo treatments are either south or north of the true border. iii) Treatment effects are estimated using the robust RD approach suggested by Calonico et al. (2017) and with a triangular kernel. iv) Bandwidth is chosen using the data-driven MSE-optimal approach in Calonico et al. (2014). v) The order bias specifies the order of the local polynomial used to construct the bias correction. This bias correction helps to adjust the RD point estimator to account for the bias introduced by the polynomial approximation.

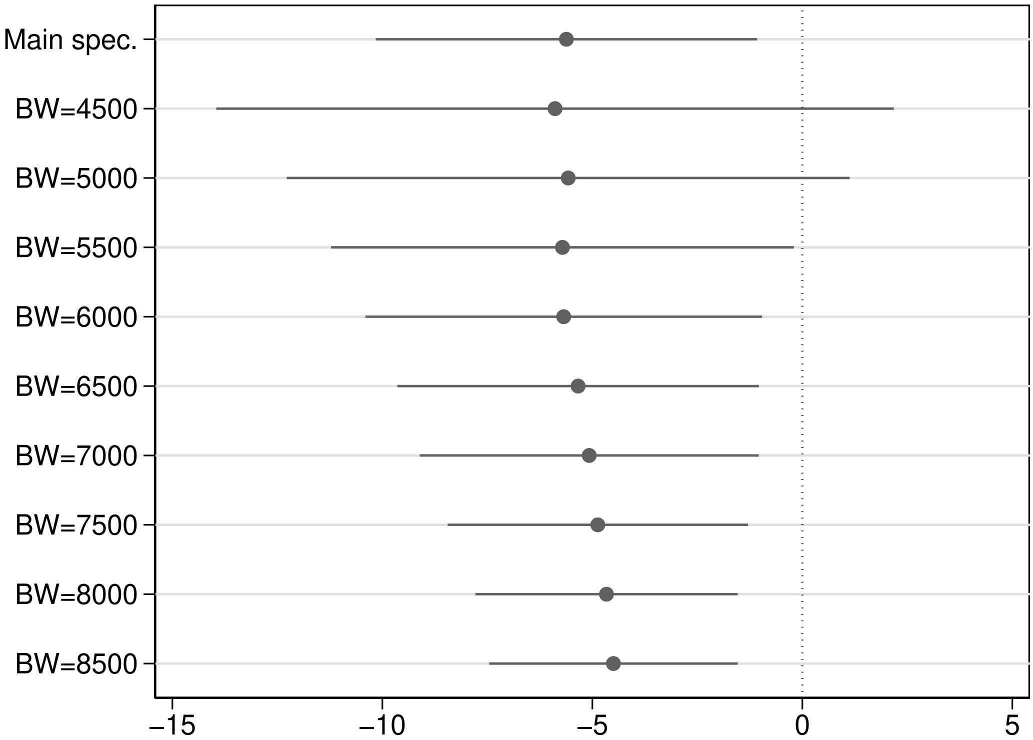

Figure F.2: RD results across bandwidth

Note: The treatment effect is estimated for different choices of bandwidth. ii) Treatment effects are estimated using the robust RD approach suggested by Calonico et al. (2017) and with a triangular kernel. iii) Bandwidth is chosen using the data-driven MSE-optimal approach in Calonico et al. (2014). iv) The order bias specifies the order of the local polynomial used to construct the bias correction. This bias correction helps to adjust the RD point estimator to account for the bias introduced by the polynomial approximation.

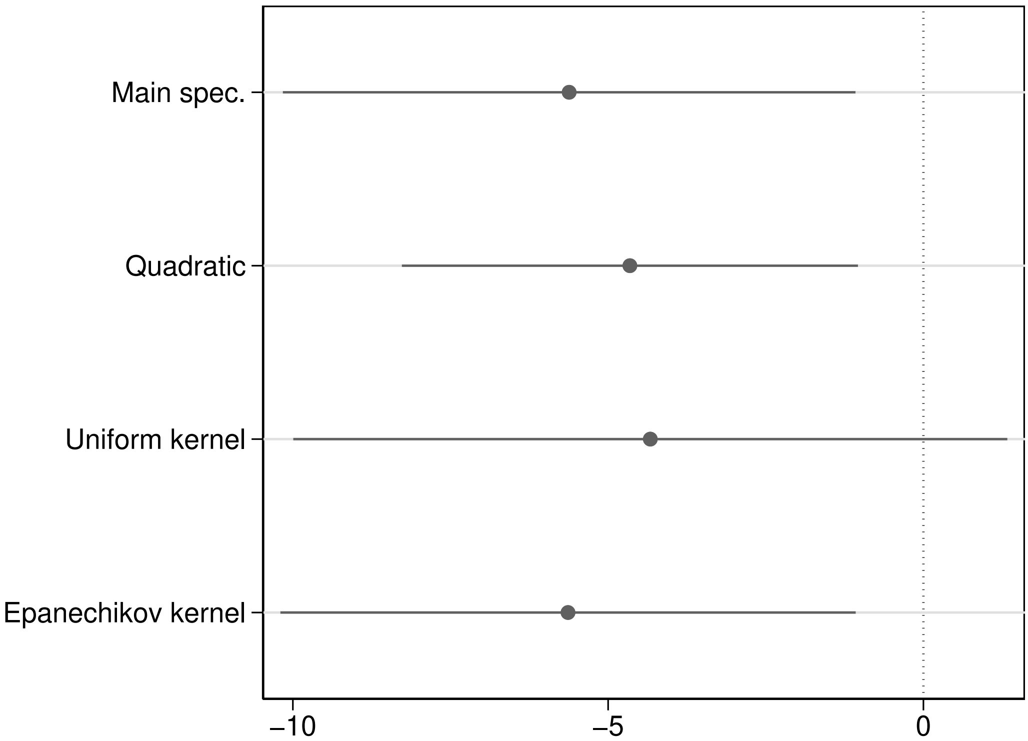

Figure F.3: RD results across polynomials and kernels

Note: The treatment effect is estimated for different choices of polynomial and kernels. ii) Treatment effects are estimated using the robust RD approach suggested by Calonico et al. (2017) and with a triangular kernel. iii) Bandwidth is chosen using the data-driven MSE-optimal approach in Calonico et al. (2014). iv) The order bias specifies the order of the local polynomial used to construct the bias correction. This bias correction helps to adjust the RD point estimator to account for the bias introduced by the polynomial approximation.

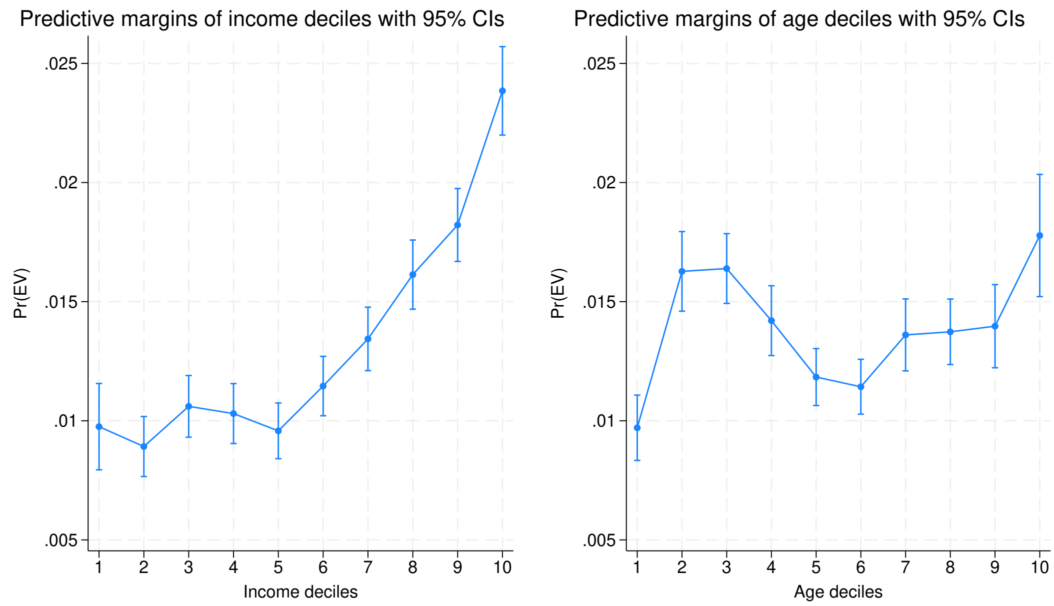

Figure H.1: Probability of purchasing an EV across income and age deciles

Note: i) The levels of the household characteristics across deciles are presented in Table C.2 in the Appendix

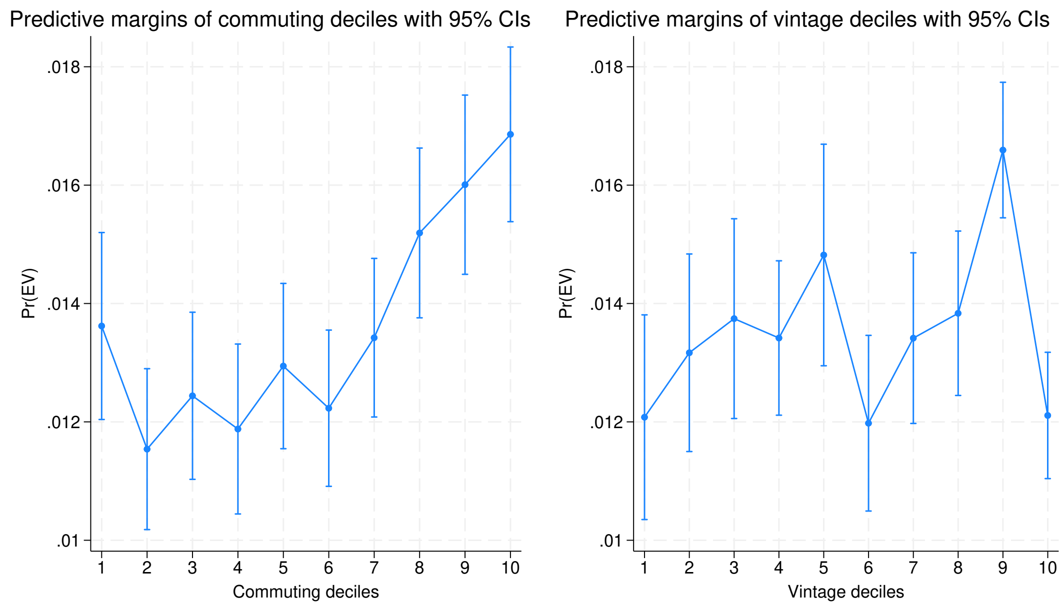

Figure H.2: Probability of purchasing an EV across commuting distance and vintage deciles

Note: i) The levels of the household characteristics across deciles are presented in Table C.2 in the Appendix

	2020	2021	2022	All years
Moving south				
SE1-SE2	21	21	15	57
SE2-SE3	46	57	59	162
SE3-SE4	240	238	216	694
SE1-SE3	0	3	0	3
SE2-SE4	1	0	0	1
SE1-SE4	0	1	1	2
Moving north				
SE4-SE3	185	208	235	628
SE3-SE2	76	44	61	181
SE2-SE1	13	21	12	46
SE4-SE2	1	1	1	3
SE3-SE1	3	1	0	4
SE4-SE1	1	1	1	3
Within price area				
SE1	109	131	98	338
SE2	288	301	280	869
SE3	3,196	3,687	3,295	10,178
SE4	1,894	2,133	1,822	5,849
No of households	134,222	135,445	135,414	

Table F.2: Number of households moving across and within price areas per year

Note: For each year, I calculate how many households who have moved from their previous price area price area further south, from their previous price area to a price area further north, and within (but not across) price areas.

	(1) EVs per 1000 hh.	(2) EVs per 1000 hh.	(3) EVs per 1000 hh.
ln(€/kWh)	-0.758***	-0.923***	-1.355***
	(0.214)	(0.290)	(0.466)
Household inc.	0.018	0.023	-0.016
	(0.021)	(0.025)	(0.037)
No. of cars	0.178***	0.159*	0.062
	(0.068)	(0.086)	(0.119)
Commuting dist.	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)
No. of kids	0.230***	0.206***	0.084
	(0.044)	(0.053)	(0.089)
Age	0.004	0.006	-0.001
	(0.004)	(0.005)	(0.007)
Vintage	0.289***	0.291***	0.318***
	(0.012)	(0.015)	(0.022)
Constant	-581.573***	-586.437***	-640.907***
	(23.619)	(29.733)	(44.842)
Spatial FE	Yes	Yes	Yes
Price areas	2-4	2-4	2-4
Observations	35911	21814	9528
Bandwidth (meters)	15000	10000	5000

Standard errors in parentheses
* p<0.10, ** p<0.05, *** p<0.01

Table G.1: PPML, EVs per 1000 households, for different bandwidths

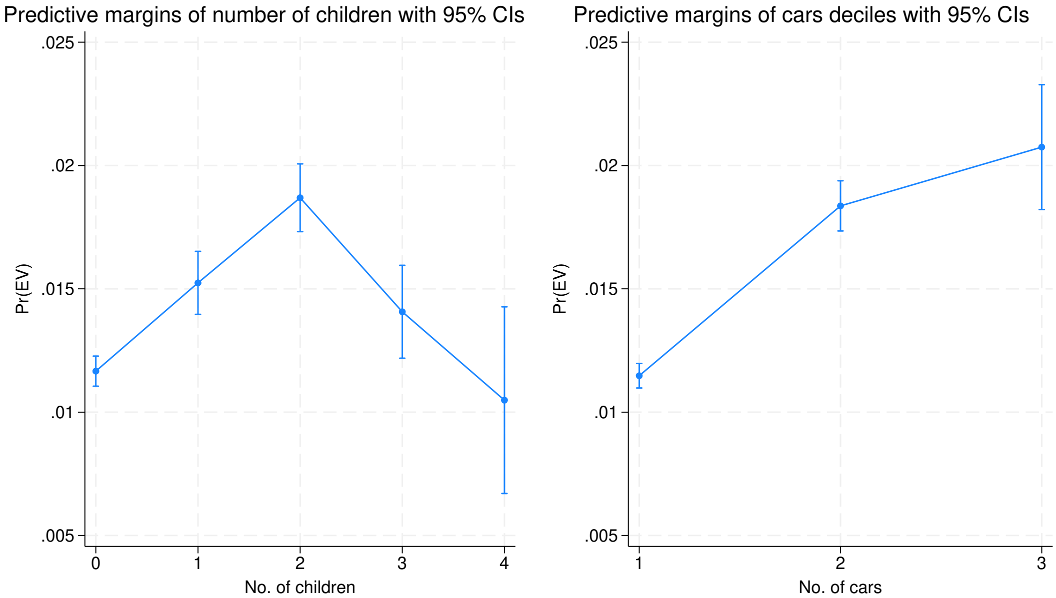

Figure H.3: Probability of purchasing an EV across number of children and cars
Note: i) The number of households purchasing an EV with zero existing cars are virtually non-existent in my sample.