Unpacking the Distributional Implications of the Energy Crisis

Lessons from the Iberian Electricity Market*

Natalia Fabra, Clément Leblanc, and Mateus Souza

November 2025

Abstract

The 2021-2023 European energy crisis, triggered by the war in Ukraine, led to broad policy interventions in energy markets. In contrast to the retail-side measures and public transfers implemented elsewhere, Spain and Portugal targeted the wholesale electricity market through the so-called *Iberian solution*. We quantify the distributional implications of the crisis and this market intervention on Spanish electricity firms and across consumer groups. We find that the crisis shifted substantial wealth from consumers to generators, with regressive impacts among consumers. Conversely, the policy's relief was progressive, delivering larger gains to lower-income groups.

Keywords: energy crisis, electricity markets, distributional implications, machine learning, policy evaluation.

*Natalia Fabra: CEMFI and CEPR, natalia.fabra@cemfi.es. Clément Leblanc: Universidad Carlos III de Madrid; Mateus Souza: Department of Economics, University of Mannheim. We thank Michael Duarte, Catarina Pintassilgo, and Álvaro Oballe for their outstanding research support. We are also grateful for the comments and feedback received from Iivo Vehviläinen and seminar participants at the Congress of the Spanish Economic Association (Salamanca), the CESifo Area Conference on Energy and Climate Economics (Munich), the UB International Academic Symposium (Barcelona), the Conference of the Yrjö Jahnsson Foundation on Climate change and inequality (Helsinki), Universidade NOVA de Lisboa, and Geneva Graduate Institute. Fabra has benefited from generous funding by the European Union (ERC grant *ENERGY-IN-TRANSITION*, 101142583). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. Souza also gratefully acknowledges support from the German Research Foundation (DFG) through CRC TR 224 (Project B07). The authors declare that they have no financial or personal conflicts of interest that could have influenced the work reported in this paper.

1 Introduction

Few instances in recent history have seen the price of an essential commodity – such as energy – increase as much as tenfold within a short period. Triggered by the lead-up to the war in Ukraine, the energy crisis that unfolded in Europe between 2021 and 2023 stands out as one such episode. It thus offers a unique opportunity to study the economic consequences of a large and sudden shock to energy markets. In this paper, we investigate the distributional consequences of this energy crisis and the subsequent policy interventions designed to cushion the surge in electricity bills. We focus on the direct impacts on electricity producers and consumers in the context of Spain.

While the macroeconomic consequences of energy crises – such as GDP losses, productivity slowdowns, rising inflation and, as a consequence, tighter monetary policy – are well documented (see Bachmann et al. (2024) and Krebs and Weber (2024), among others), their distributional effects have received comparatively less attention.[1] Yet, one cannot ignore the heterogeneous impacts across sectors, income strata, or regions. For example, firms in energy-intensive sectors might suffer from reduced competitiveness when faced with high energy costs, as emphasized in Draghi (2024)'s report. For households, energy is a basic necessity, and price increases tend to disproportionately burden low-income households that face tighter budget constraints and are less able to absorb or adapt to higher energy costs (Ahlvik, Liski, and Makimattila, 2024; Ahlvik et al., 2025). This exacerbates concerns related to energy poverty, which can have adverse physical and mental health effects (Lee and Yuan, 2024; Bentleya et al., 2024), among others.

As a response to the energy crisis, several policies were implemented in Europe – ranging from voluntary demand reduction incentives to market interventions – in order to cushion the effects of rising energy costs (European Commission, 2025). The design of these policies is crucial for their effectiveness in reducing energy bills without distorting market efficiency or inducing adverse environmental consequences (Levell, O'Connell, and Smith, 2025). Moreover, such policies frequently carry significant distributional implications, as they influence how the economic burden and benefits are shared between energy companies, consumers, and taxpayers. While heterogeneous impacts are expected, estimates of their relative magnitudes are still scarce. In this paper, we help fill this gap by leveraging the unique response from Spain and Portugal – the only countries in Europe to address the energy crisis by intervening directly in their wholesale electricity market.[2]

[1] Bachmann et al. (2024) and Krebs and Weber (2024) study the context of Germany. They show that the one-year GDP loss resulting from the energy crisis was comparable in magnitude to the short-run output declines observed during the COVID-19 crisis in 2020 and the global financial crisis in 2008. In addition, inflation surged and real wages fell more sharply in 2022 than in any other year in post-war Germany. They also provide evidence of the lasting economic harm of the energy crisis, as reflected in the sluggish recovery of the German economy following the energy shock.

[2] While Spain and Portugal had a unified response, sharing a well-connected electricity market, our analyses focus on Spanish consumers and producers, due to data limitations.

Our insights could be helpful for the effective design of measures in response to future energy crises.

Most European nations, other than Spain and Portugal, mainly relied on subsidies, tax cuts, retail price caps, or direct public transfers to households and firms. While these helped to weather the crisis, they entailed substantial fiscal costs ultimately borne by taxpayers.[3] Importantly, such measures left the profits of energy producers largely untouched. Yet, one of the most significant distributional imbalances during the energy crisis emerged at the level of electricity generation. As the President of the European Commission noted: *"Low-carbon energy sources are making... revenues they never dreamt of... that do not reflect their production costs"* (von der Leyen, 2022). This occurred because surging gas prices drove up wholesale electricity prices, substantially increasing the inframarginal rents of low-carbon producers (nuclear, hydro, and renewable energy plants) – whose costs remained largely unchanged. As the preceding quote suggests, this profit increase was largely perceived as a windfall.

The market intervention implemented by Spain and Portugal – the so-called *Iberian solution* – was designed precisely to limit the pass-through of inflated gas prices to wholesale electricity prices, mitigating the increase in inframarginal rents, and thus reducing energy costs for consumers without fiscal costs. As a result, the *Iberian solution* had distributional consequences both at the wholesale level – by altering the allocation of surplus between electricity producers and consumers – and at the retail level, where the price impacts varied across consumers depending on their consumption levels and patterns over time. This paper quantifies both of these effects.

We begin by simulating counterfactual wholesale market outcomes under five distinct scenarios – including a no-crisis scenario and variations that help identify the effects of the *Iberian solution* (which we refer to as the price intervention) and the accompanying energy-saving measures. To build these scenarios, we estimate input prices in the absence of the crisis (fossil fuel prices, carbon prices, and electricity prices in France), and estimate counterfactual electricity demand using machine learning techniques. Interestingly, we find that the crisis led to a decline in electricity demand – driven more by the implementation of non-price energy-saving measures, rather than by direct demand response to rising energy prices.

Our simulations indicate that, due to the crisis, Spanish power plant earnings increased by almost €50B between July 2021 and June 2023, i.e., a 250% increase, while

[3] In Europe, Ferdinandusse and Delgado-Téllez (2024) estimate that fiscal measures to support households and firms in response to the energy price shock amounted to 1.8% of GDP in 2022 and 1.3% in 2023. See also Sgaravatti et al. (2023).

their profits rose by €27B over the same period, i.e., a 224% increase.[4] Taken together, the interventions reduced consumer payments by €4B (€1.6B from the price intervention and €2.3B from the savings measures).

We then combine the simulated wholesale prices with hourly household electricity consumption data at the zip code level to compute counterfactual electricity bills for households across 8,390 zip codes in Spain. We find that, on average, the crisis increased household electricity bills by €338 per year, equivalent to about 1% of their disposable income. Without the interventions, this increase would have been €487 per year, given that the measures reduced average annual bills by €149 (€81 due to the price intervention and €56 due to the energy-saving measures), avoiding 30% of the bill increase that would otherwise have occurred.

These averages, however, mask substantial heterogeneity across households. Using data on weather, socio-demographics, and electricity consumption, we group zip codes into six clusters with distinct consumption profiles. This approach helps us to identify the factors driving the intensity of the crisis and policy intervention effects. Results suggest that household income, as well as higher shares of electric heating or air conditioning, are positively correlated with bill changes – consistent with these factors driving higher electricity consumption during periods when crisis-level prices peaked and the policy interventions were most effective.

We also uncover heterogeneity across (household-level) income quintiles by extrapolating the drivers of bill variation across zip codes. We find that the electricity price shock had regressive effects. While higher-income households experienced larger absolute increases in bills, lower-income households bore a greater burden relative to their disposable income. Importantly, policy interventions helped mitigate these regressive impacts. Although the monetary savings were larger for wealthier households in absolute terms, the policy measures – particularly the price intervention – delivered proportionally greater relief to lower-income groups, thus exhibiting progressive impacts.

Related Literature. This paper contributes to a growing literature emphasizing the importance of equity in the design of energy and climate policies (see, for instance, Deryugina, Fullerton, and Pizer (2019) for academic contributions; OECD (2024) for the policy perspective; and Reguant, Fabra, and Wang (2025) and Borenstein, Sallee, and Fowlie (2022) for analyses focused on the distributional consequences of electricity pricing). As the energy transition accelerates and prices become more volatile due

[4]These figures are computed under the assumption that all power plants were remunerated at market prices and should be interpreted as an upper bound on the actual effects due to two caveats. First, plants owned by vertically integrated utilities may not have directly benefited from elevated wholesale prices due to internal transfer pricing, but likely captured additional profits through their retail margins. Second, certain regulated renewable power plants received fixed payments rather than full market remuneration, although the share of such plants was already relatively small.

to climate-related shocks and decarbonization efforts, understanding the distributional effects of alternative policy instruments becomes essential to ensure broad social support.

A number of recent papers have evaluated the support measures adopted by European governments to mitigate the impact of rising energy costs.[5] For example, the UK government subsidized energy producers to maintain retail prices at a cap below market levels. In their analysis, Fetzer (2014) find that, due to the positive correlation between income and energy consumption, high-income households received higher bill reductions in absolute terms. However, their data do not allow for an assessment of energy expenditures as a share of income, a key metric for evaluating the progressivity or regressivity of a policy (Borenstein, Sallee, and Fowlie, 2022). Under the *Iberian solution*, we show that although high-income households experienced greater savings in absolute terms – consistent with Fetzer (2014) – the relative benefit was greater for low-income households, rendering the policy progressive.

Levell, O'Connell, and Smith (2025) also study the UK's relief package and estimate that, absent the intervention, average household welfare losses would have amounted to 6% of their income. They document significant demand reductions in response to price increases, with an average elasticity of 0.31, and stronger responses among heavy energy users. In contrast, our results attribute most of the demand reduction to non-price energy-saving measures, which is consistent with findings in the German context (Behr, Köveker, and Kücük, 2025).

Germany and the Netherlands implemented schemes that shielded consumers from retail prices above a guaranteed threshold. Dertwinkel-Kalt and Wey (2025) and Haan and Schinkel (2023) find that these measures weakened retail market competition, raised retail prices, and increased fiscal costs for the government. Notably, because subsidies were tied to baseline consumption levels, consumers still had incentives to reduce usage, underscoring the importance of policy design in balancing protection, conservation incentives, and fiscal sustainability.

Unlike the UK, Germany, and the Netherlands – whose measures primarily involved retail market interventions or direct transfers – the distinguishing feature of the *Iberian solution* is that it targeted the generation segment without imposing fiscal costs. We are aware of three other recent papers that have examined the effects of this intervention on the Iberian electricity market. None of these have assessed distributional implications at the wholesale or retail levels. Ruiz, Schult, and Wunder (2024) employ a synthetic control approach and find that the policy reduced average spot prices by approximately 40% between July 2022 and June 2023. Similarly, by building a counterfactual supply, Lou et al. (2025) estimate a 35.3% spot price reduction between June 2022 and February 2023.

[5]Relatedly, other work examines structural electricity market-design reforms intended to prevent a recurrence of such episodes. See Fabra (2023), Gerlagh, Liski, and Vehvilainen (2022) and Polo et al. (2023), among others.

However, they also highlight an unintended consequence: the intervention led to a 19.2% increase in gas consumption in the Iberian Peninsula, driven by rising exports, which in turn contributed to higher carbon emissions. Finally, using Bayesian structural time series models, Hidalgo-Pérez et al. (2024) find that the intervention reduced household prices between 20% and 28%, accompanied by increased gas consumption and more exports to France. Our estimates regarding the aggregate effects of the *Iberian solution* are broadly consistent with these.

The remainder of the paper is organized as follows. Section 2 provides a detailed overview of the energy crisis and the policy response implemented in Iberia. Section 3 examines the effects at the wholesale level by simulating electricity market outcomes under counterfactual scenarios – with and without the crisis, and with and without the policy interventions. Section 4 quantifies the distributional impacts on Spanish households and investigates the channels driving changes in their electricity bills. Section 5 concludes. Additional supporting material is provided in the Appendix.

2 The Energy Crisis and the Policy Response

Starting in January 2021, gas prices on European exchanges began to rise sharply (Panel (a) in Figure 1). At the time, the underlying cause remained unclear: Gazprom, the Russian state-owned gas company, had begun to reduce gas exports to Europe in a strategic move to inflate prices ahead of the invasion of Ukraine (Keliauskaitė et al., 2024). By January 2022 – just one month before the invasion – European gas prices had reached unprecedented levels, surpassing €200/MWh, a tenfold increase relative to the pre-crisis average.[6] This escalation reflected mounting concerns that winter gas supplies would be insufficient to avoid curtailments. The outbreak of the war further exacerbated the situation. In the summer of 2022, gas prices soared to €300/MWh, prompting a series of policy interventions that eventually helped ease market pressure. By the winter of 2023, gas prices had declined but remained roughly twice their pre-crisis levels.[7]

The surge in gas prices was directly transmitted to wholesale electricity markets (Panel (b) in Figure 1), where gas-fired power plants frequently set the marginal price. This led to a sustained period of exceptionally high electricity prices. Additional factors – such as severe heatwaves across Europe, low hydro and wind output, and widespread outages in the French nuclear fleet – amplified the upward pressure on electricity prices, contributing to multiple extreme price spikes.

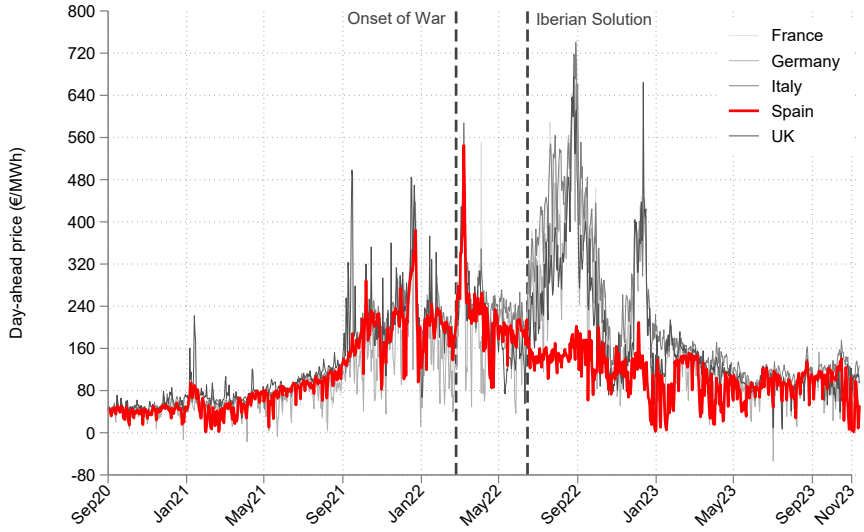
[6] As shown in Figure 1 Panel (a), gas prices in the Iberian market remained slightly below those in the TTF. This divergence reflects both the limited interconnection capacity between the Iberian Peninsula and the rest of Europe, and the greater LNG regasification capacity along the Iberian coast.

[7] Kroger, Neuhoff, and Schwenen (2025) document that the price escalation caused a dramatic decline in traded contract volumes, particularly in bilateral trading, and prompted a shift towards centralized markets for liquidity provision.

Figure 1: Wholesale Gas and Electricity Prices in Europe



(a) Wholesale gas prices



(b) Wholesale electricity prices

Notes: Panel (a) shows the evolution of natural gas prices in the wholesale exchanges of the Netherlands (TTF) and Iberia (MIBGAS). Panel (b) shows the evolution of electricity prices in several European exchanges. In both figures, which span from September 2020 until November 2023, prices in Iberia are highlighted in red.

The energy crisis prompted European governments to adopt a range of price and non-price measures aimed at cushioning the impact of rising energy bills on consumers. On the one hand, the EU issued emergency rules requiring member states to implement mandatory or voluntary energy-saving targets – such as reducing overall electricity consumption by 10% and enforcing a minimum 5% reduction during peak hours (European Commission, 2023). Governments also ran public campaigns to encourage energy-saving behaviors, such as turning down heating/air conditioning, shifting usage outside of peak

hours, and adopting efficient practices in homes and businesses.

Regarding price measures, most governments relied on a combination of tax exemptions, retail price caps, and targeted support for vulnerable groups.[8] Spain and Portugal, as already mentioned, deviated from this standard approach by intervening directly in the Iberian wholesale electricity market (MIBEL).

In June 2022, the European Commission authorized both countries to implement a one-year mechanism designed to cap the input costs of fossil fuel power plants – commonly referred to as the "Iberian Solution" or "Iberian Exception." Under this mechanism, gas and coal power plants were compensated for the difference between the cost of producing one MWh of electricity under the actual market price of natural gas and an administratively set cap, assuming an efficiency rate of 55% for converting gas into electricity.[9] The cap started at €40/MWh for the first six months and subsequently increased by €5/MWh each month, reaching €70/MWh by the end of the one-year period. This compensation did not have a fiscal cost, as it was passed on to consumers through a surcharge on their electricity bills.[10]

The primary objective of the mechanism was to incentivize gas- and coal-fired plants – often the price-setting market units – to internalize the subsidy in their bids.[11] This would lower wholesale electricity prices, thereby reducing the inframarginal rents of producers such as nuclear, hydro, and renewables, which were not exposed to the increase in fuel costs. As illustrated in Panel (b) of Figure 1, the *Iberian Solution* proved effective in decoupling electricity prices from gas prices, creating a visible wedge between Iberian wholesale electricity prices and those in the rest of continental Europe. This, in turn, allowed Spain to weather the energy crisis more effectively than many other EU countries, with the *Iberian Solution* contributing to significantly reducing its inflation rate (Roldán, Comajuncosa, and Hidalgo, 2023; Ruiz, Schult, and Wunder, 2024).

The implementation of the *Iberian Solution* sparked controversy within the European Union. Countries such as Germany and the Netherlands criticized the temporary price cap, arguing that it undermined the integrity of the EU's internal electricity market and risked distorting cross-border competition. A key concern was that lower Iberian electricity prices would spur exports to France, saturating interconnection capacity and effectively subsidizing French consumers. Analysts also warned that the cap could lead

[8] According to Sgaravatti et al. (2023), the price measures amounted to cumulative fiscal expenditures of almost €651 billion across the EU between September 2021 and June 2023.

[9] Since all thermal plants, regardless of their actual efficiency rates or technology, received the same subsidy, the merit order between them was not distorted.

[10] Specifically, subsidies to thermal generators were financed in part by higher congestion rents on the France-Spain interconnector arising from the widened price spread. The remainder was allocated pro rata across demand exposed to wholesale prices – including customers on the default real-time tariff and fixed-price customers whose contracts had been renewed after the intervention – via a per-kWh surcharge added to electricity bills.

[11] The price of coal also increased during the energy crisis. The main purpose of subsidizing coal generation was to avoid distortions in the cost ranking across the price-setting plants.

to increased gas-fired generation and higher gas consumption in the Iberian Peninsula, counteracting EU-wide efforts to reduce fossil fuel dependence and emissions. Finally, critics noted that the price relief was not targeted by income level, raising concerns that low-income households would receive less support.[12]

This paper offers a rigorous quantification of the distributional effects of the energy crisis and the Iberian market intervention on Spanish electricity consumers, analyzing impacts at both the wholesale (Section 3) and retail levels (Section 4).

3 Wholesale Market Impacts

3.1 Simulating the wholesale electricity market

Our first step is to conduct simulations of the Spanish wholesale electricity market under counterfactual scenarios for electricity supply and demand – with and without the energy crisis, and with and without the price and non-price policy interventions. Simulations are based on a model originally developed by de Frutos and Fabra (2012), assuming competitive supply.[13] All simulations are performed at the hourly level over the 17,520 hours spanning the sample period from 1st July 2021 to 30th June 2023.

To model the supply side, we draw on granular data describing the technical characteristics of actual power plants operating in the Spanish electricity market – including their capacity, thermal efficiency, and emissions rates (Global Energy Monitor, 2025). We also incorporate the observed hourly availability of renewable resources (REE, 2024), which remains unaffected by the crisis or any policy interventions. These inputs are combined with daily prices for fossil fuels and EU-ETS CO_2 allowances, using either the realized values (Bloomberg, 2025) or counterfactual values in the absence of the crisis. The methods used to construct gas and CO_2 prices in the *No-Energy-Crisis* scenario are described in detail below.

Equipped with this information, we estimate marginal generation costs following standard methodologies (see, for example, Fabra and Imelda, 2023). For renewable generation, marginal costs are assumed to be equal to operation and maintenance (O&M) costs. For a thermal plant i, marginal costs in period t also depend on fossil fuel prices as follows:

$$c_{it} = \frac{p_t^f + \tau_t \epsilon^f}{e_i} + om_i,$$

where f denotes the fossil fuel used by the power plant, either gas or coal, $f = \{G, C\}$, p_t^f denotes its price in period t and ϵ^f denotes its CO_2 emission factor; τ_t is the CO_2 price

[12] See Corbeau, Farfan, and Orozco (2023) and Patel (2022) for a summary of these critiques.

[13] Under the assumption of strategic bidding (modeled following the framework in de Frutos and Fabra (2012)), the simulation results remain quantitatively similar. Simulation results are available upon request.

in period t; e_i is the plant-specific efficiency rate in converting fuel into electricity; and om_i stands for operation and maintenance costs.

The price intervention mechanism is binding when the gas price p_t^G exceeds the reference price p_t^R. In such a case, all thermal plants receive a per-unit subsidy for electricity generated in period t, defined as:

$$s_t = \frac{p_t^G - p_t^R}{0.55},$$

where 0.55 is the benchmark thermal efficiency.

Accordingly, the plant's marginal cost net of the subsidy becomes:

$$c_{it} - s_t = \frac{p_t^R}{0.55} + \left(\frac{p_t^f}{e_i} - \frac{p_t^G}{0.55} \right) + \frac{\tau_t \epsilon^f}{e_i} + om_i.$$

Importantly, the intervention preserves the marginal cost ranking, or merit order, among generating units. Only gas-fired plants operating at the reference efficiency of 0.55 exhibit a net marginal cost equal to the cap plus the emissions costs. By contrast, coal plants and gas plants with lower efficiency levels display net marginal costs that exceed that level, consistent with their underlying cost characteristics.

To construct the industry's hourly competitive supply curve, we use either the estimated marginal costs in the absence of the market intervention, or the marginal costs net of the subsidy when the intervention is in place.

To model the demand side, we use either the realized hourly demand values (REE, 2024) or one of two counterfactual predictions generated using machine learning techniques: (i) demand in the absence of the crisis, and (ii) demand assuming that the crisis occurred so that higher prices triggered some demand response, but excluding the effects from non-price energy-saving measures. The methods used to construct these counterfactual scenarios are detailed below.

Finally, since the Iberian electricity market is interconnected with France, French electricity prices influence supply and demand conditions in Iberia, depending on whether France is importing from or exporting to the region.[14] Below, we describe the methodology used to compute French electricity prices under the counterfactual scenario without the energy crisis.

Matching actual or counterfactual market demand with actual or counterfactual supply enables us to determine the competitive hourly price and production allocation across plants in the Spanish electricity market under our five scenarios, as described next.

[14]The Spanish market is closely linked to the Portuguese market, together forming the Iberian electricity market. This interconnection is rarely congested, resulting in nearly identical electricity prices in Spain and Portugal. Furthermore, both markets have been similarly affected by the crisis, as they are influenced by natural gas prices in the Iberian Peninsula and both have been subject to the *Iberian Solution*. Hence, for the purposes of this analysis, we assume that the Spanish and Portuguese markets are always coupled.

Scenarios. Table 1 summarizes the assumptions on the supply and demand variables used to construct our scenarios. The first is the *Factual* scenario, which serves as our benchmark. It uses the realized values for demand, gas, CO_2, and French electricity prices, under the existing market interventions, including the energy-saving measures and price cap on gas.[15]

The other four scenarios are counterfactual. The *No-Energy-Crisis* scenario assumes that the energy crisis did not occur (i.e., both demand and supply would have followed paths similar to those observed prior to the crisis). Comparing it to the *Factual* scenario allows us to evaluate the *energy-crisis effect*.

The *Business-as-Usual* (BaU) scenario assumes that the crisis happened, but neither energy-saving measures nor price interventions were implemented. Comparing it with the *Factual* scenario allows us to quantify the combined effects of both interventions, which we call the *full-intervention effect*. The *Savings-Intervention-Only* scenario removes the price control measure from gas prices, isolating the *price-intervention effect* by comparison with the *Factual* scenario. Conversely, the *Price-Intervention-Only* scenario assumes demand without energy-saving measures, enabling us to separately quantify the *savings-intervention effect*.[16]

Below, we outline the methodology used to construct the counterfactual inputs for the relevant supply and demand variables used in the simulations.[17]

Counterfactual gas prices. We assume that, in the absence of the crisis, gas prices would have likely remained at their average pre-crisis levels. Specifically, we calculate the average over the 2017–2019 period, excluding 2020 and 2021 due to significant deviations from long-term trends caused by the COVID-19 pandemic. To account for seasonal variation, we compute twelve month-specific average gas prices, one for each calendar month.

Counterfactual CO_2 prices. An increase in gas prices improves the relative competitiveness of coal-fired power plants, thereby increasing the demand for CO_2 allowances and exerting upward pressure on CO_2 prices. Accordingly, in a scenario without elevated gas prices, CO_2 prices would have remained lower. To construct the counterfactual CO_2 price in the absence of the crisis, we remove the component of the CO_2 price that is

[15] Note that the *Factual* scenario contains simulated electricity prices – based on realized demand and input values – rather than the actual realized prices. This approach ensures that any potential biases in our simulations cancel out when comparing across scenarios. Nevertheless, taking the realized electricity prices as the baseline scenario would not materially affect our conclusions, since the simulated prices closely track the observed ones.

[16] Due to the interaction between the price intervention and the energy savings intervention, the sum of the *price-intervention effect* and the *savings-intervention effect* need not be equal to the *full-intervention effect*. In practice, we find this difference to be small.

[17] Appendixes B and C contain further details on the methods.

Table 1: Simulation Scenarios

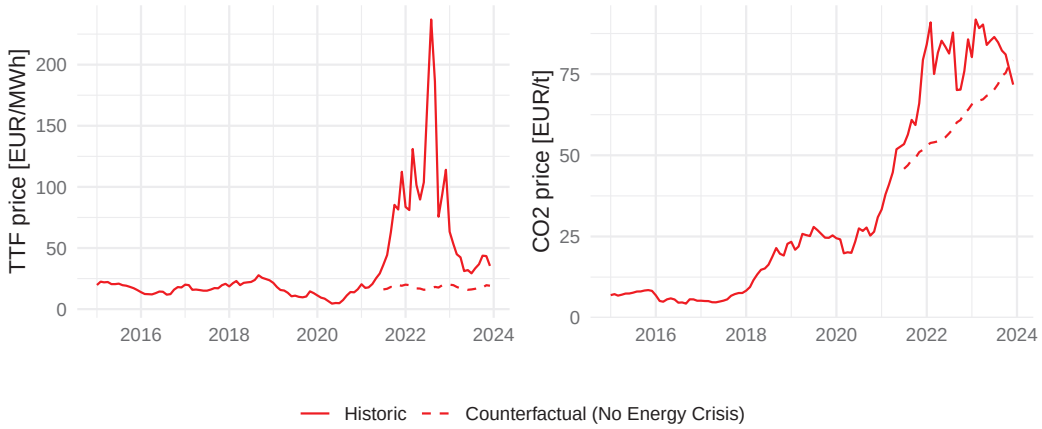
Variables	Assumed Value	Factual	No Energy Crisis	BaU	Savings Int. Only	Price Int. Only
Demand	Realized	✓			✓	
	No Energy Crisis		✓			
	Crisis w/o Savings Int.			✓		✓
Gas Prices	Realized	✓				✓
	No Energy Crisis		✓			
	Crisis w/o Price Int.			✓	✓	
CO_2 & French Prices	Realized	✓		✓	✓	✓
	No Energy Crisis		✓			

Notes: This table summarizes the assumptions on demand, gas prices, CO_2 prices, and French electricity prices used under the five simulated scenarios.

attributable to the gas price increase, while preserving the underlying upward trend associated with structural reforms to the EU-ETS – particularly, the launch of Phase IV in 2021 – that would have taken place without the crisis.

To this end, we estimate a linear model capturing the relationship between natural gas prices (TTF) and CO_2 prices (EU-ETS allowances) over the 2015–2023 period, allowing for distinct time trends across different EU-ETS phases.[18] Using this model, we predict CO_2 prices under the assumption that natural gas prices would have followed the pre-crisis trends described above. As seen in Figure 2, CO_2 prices would have increased even without the energy crisis, but not as much.

Figure 2: Gas and CO_2 prices – Historic and *No-Energy-Crisis* counterfactual



Notes: This figure shows the evolution of the realized TTF natural gas prices and CO_2 prices (solid lines) versus the estimated counterfactual prices in the absence of the crisis (dashed line, from July 2021).

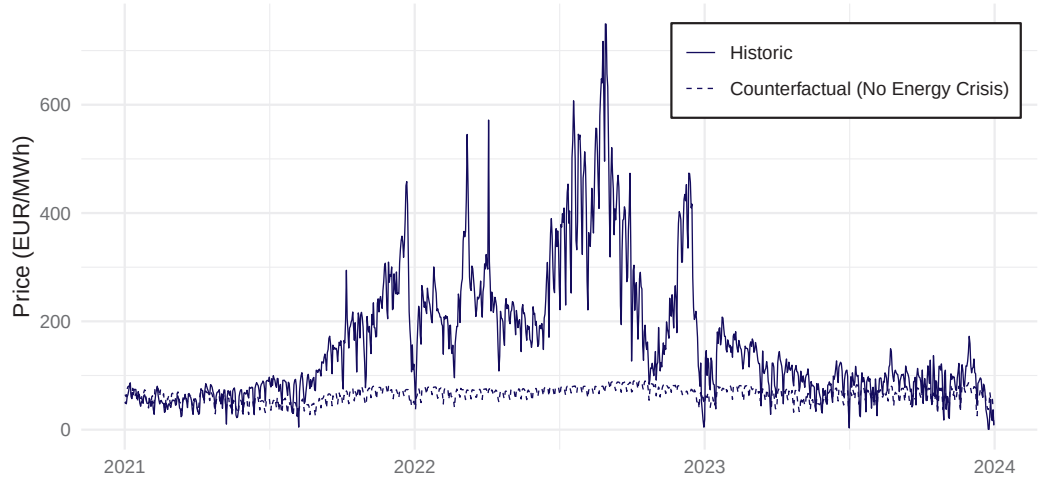
Counterfactual French electricity prices. To construct counterfactual French electricity prices, we first estimate a regression model using historical data from 2015 to 2023. The model explains the French electricity price as a function of the natural gas price (TTF), the CO_2 price, the French electricity residual demand (net of renewable and nuclear production), and time-fixed effects.[19] Using this model, we predict French electricity prices under the counterfactual scenario where natural gas and CO_2 prices reflect the hypothetical *No-Energy-Crisis* conditions. Figure 3 plots the resulting prices, which follow a similar trend as the pre-crisis prices.

Counterfactual wholesale electricity demand. Aggregate electricity demand is another key input for our simulations of the Spanish electricity market. Naturally, in the

[18] The model (equation B.1) and the estimation results (Table B.1) are reported in Appendix.

[19] The regression results are summarized in Table B.1 in the Appendix. The high value of the $R^2 = 0.9$ shows that a high proportion of the price variation is explained by the model. Yet, note that we do not account for changes in electricity demand in neighboring countries that may have occurred in reaction to the crisis, neither do we account for policy measures implemented outside the Iberian market.

Figure 3: French electricity prices – Historic and *No-Energy-Crisis* counterfactual



Notes: This figure shows the evolution of the realized French electricity prices (solid line) versus the estimated counterfactual prices in the absence of the crisis (dashed line).

absence of the crisis, one would expect higher electricity demand than observed, both due to lower electricity prices and the absence of the energy-saving measures put in place by the government during the crisis.[20]

To predict this counterfactual demand, we train machine learning algorithms. Prior literature has used similar machine learning strategies to analyze, for example, the impact of the COVID-19 pandemic on electricity markets (e.g., Graf, Quaglia, and Wolak, 2021; Fabra, Lacuesta, and Souza, 2022; Benatia and Gingras, 2023). These studies have shown that accurate predictions can be obtained through the implementation of flexible algorithms with high-dimensional (preferably hourly) demand data.

We caveat that, by nature, the true counterfactual can never be observed. The validity of our predictions therefore relies on two key assumptions: (i) that the *No-Energy-Crisis* counterfactual aggregate electricity demand is driven by a "stable" regression function that we can estimate; and (ii) that the variables used to predict counterfactual demand are independent of the energy crisis itself (more details in Fabra, Lacuesta, and Souza, 2022). In Appendix C.1 we formalize these assumptions, based on the Neyman-Rubin potential outcomes framework (Neyman, 1923; Rubin, 1974).

The key idea from our approach is to use pre-crisis data to learn patterns that would be expected for the unobserved counterfactual demand in case the crisis had not happened. We therefore train machine learning algorithms using hourly electricity de-

[20] In the case of Spain, in May 2022, a regulation stipulated that public buildings could not be cooled below 27^{o}C in the warm season or heated above 19^{o}C in the cold season (BOE, 2022b). Similar measures were extended to commercial establishments in August of the same year (BOE, 2022a).

mand data in Spain from 2015 to 2020.[21] The best-performing ML algorithm is selected based on cross-validation strategies that allow us to assess out-of-sample prediction accuracy (details in Appendix C.3). The algorithm with the lowest root mean squared error (RMSE) was a variant of gradient-boosted regression trees (more details in Appendix C.2 and in Chen and Guestrin, 2016).

Our algorithms incorporate the following predictor variables: daily national average (population weighted) and squared average of minimum, maximum, and median surface temperatures; daily temperatures for the 50 provinces of Peninsular Spain; daily national average (population weighted) and squared average of sunshine hours and average wind speeds; daily sunshine hours and average wind speeds for the 5 most populous provinces (Madrid, Barcelona, Valencia, Seville, and Alicante); day-ahead hourly forecasts of wind electricity generation (and the logged version of that variable); hour of day, day of week, week of year, month of year, quarter of year and year fixed effects; holiday indicators; a daily trend; monthly nationwide road traffic fatalities; daily Dutch TTF natural gas prices and daily EU-ETS CO_2 prices (as well as lag1, lag7, lag14, lag30, and lag60 versions of those).[22]

The latter three variables warrant further explanation. The response to the COVID-19 pandemic – through lockdowns and other restrictions – triggered a significant shock to electricity consumption. While this period only partially overlaps with the energy crisis, it substantially affected much of the data used to train the algorithms. Traffic fatalities are included as a proxy to flexibly capture reduced urban mobility due to COVID-19 lockdowns and restrictions. Even though the relationship between urban mobility and traffic fatalities is complex (e.g., Albalate and Fageda, 2021), in Appendix Figure C.3 we show that aggregate road fatalities in Spain were somewhat stable from 2015 to early 2020 (excluding seasonality patterns), but then reduced substantially from March to June 2020, and again during the end of 2020/beginning of 2021 (when a second state of emergency was in effect). We prefer this measure over economic indicators (such as GDP) that might also have been impacted by the energy crisis itself.[23]

Natural gas and CO_2 prices were included as proxies to capture demand elasticity

[21] Our preferred specification is trained with data also from 2020, when strict lockdown measures were in place due to COVID-19. We argue that data from 2020 can be useful to capture general trends in the relationship between economic activity and electricity demand, even if demand in that year was much lower than usual. Regardless, prediction results are similar if we train the algorithms with data only up to 2019, as shown in Appendix C.3.

[22] Aggregate electricity demand and day-ahead wind generation forecasts for Spain are publicly available through the Spanish System Operator website (REE, 2024). Weather data was obtained through the State Meteorological Agency's API (AEMET, 2024). Traffic fatality data are available through DGT and EpData (2024). Daily Dutch TTF Natural Gas prices were obtained from Bloomberg (2024). Historical EU-ETS CO_2 prices were provided by SENDECO2 (2024).

[23] Our choice is also justified by evidence that driving behavior is quite inelastic to energy price changes in the short run (e.g., Hughes, Knittel, and Sperling, 2008). We also considered including a more direct measure of aggregate traffic conditions in Spain over time (such as travelers per month). Unfortunately, such a measure is not reliably available during the whole sample period (2015 to 2023).

to electricity prices.[24] We do not use wholesale electricity prices from the Spanish market directly as predictors because they are themselves influenced by electricity demand in Spain, which raises a simultaneity issue. Instead, we use the TTF natural gas price and the EU-ETS CO_2 price as proxies. These Europe-wide prices capture the impact of the crisis on Spanish electricity prices while being less likely to be significantly affected by the Spanish electricity demand.

To construct the demand that is used to simulate the *Business-as-Usual* and *Price-Intervention-Only* scenarios – where the crisis occurs but energy savings are not implemented – we simulate counterfactual electricity demand using the actual TTF and CO_2 prices observed during the crisis (labeled "real TTF and CO_2 prices" in Figure 4). To construct the demand that is used to simulate the *No-Energy-Crisis* scenario, we use the counterfactual TTF and CO_2 prices (labeled "stable TTF and CO_2 prices" in Figure 4). This approach allows us to capture the limited but non-zero demand response to energy prices during the crisis.[25]

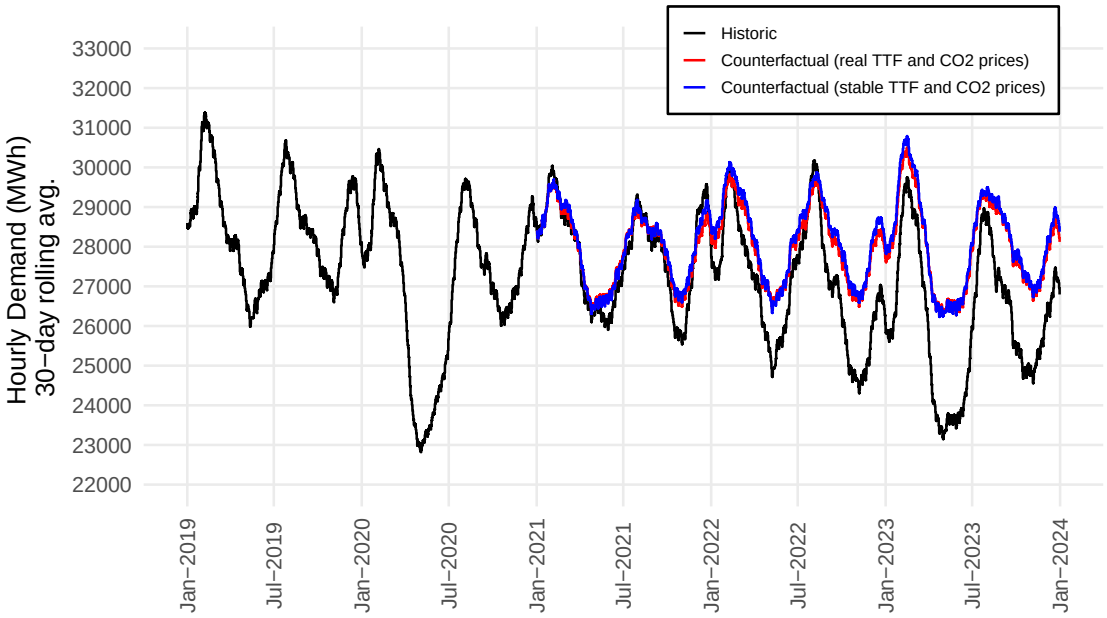
Figure 4 depicts the time series (30-day rolling averages) of our counterfactual predictions for aggregate electricity demand in Spain. The predictions in red use actual gas and CO_2 prices, while the predictions in blue are for the scenario with counterfactual prices. The two predicted curves look similar, which may be attributed to a relatively low short-term price elasticity. Realized demand is depicted by the black curve. There is substantial overlap between realized and predicted demand for most of 2021. Discrepancies start to appear towards the end of 2021, when electricity and gas prices started to soar. Deviations are even more prominent in 2022 and 2023. We find that electricity demand would have been significantly higher in the absence of the crisis.[26] For instance, actual demand in May 2023 was nearly 11% lower than the counterfactual predictions, despite the fact that wholesale gas and electricity prices in Spain had begun to stabilize by that time. These results suggest that the energy crisis triggered strong and persistent reductions in electricity consumption. Such effects cannot be ignored when simulating the wholesale market impacts.

[24] Own-price elasticity of electricity demand is typically low, especially in the short-run (e.g., Labandeira, Labeaga, and López-Otero, 2017; Deryugina, MacKay, and Reif, 2020). Rather than ignoring this channel for our counterfactual predictions, we attempt to incorporate it through variations of natural gas and CO_2 prices.

[25] In the simulations, we use the same counterfactual demand regardless of whether the price intervention was in place, although it could slightly affect demand via price elasticity. Since the overall demand response to prices is very small, we consider the omission negligible.

[26] In the context of the Italian gas market, Polo and Roccuzzo (2025) also find a significant drop in demand, but they attribute it mostly to an increase in price elasticity.

Figure 4: Actual versus predicted aggregate electricity demand



Notes: This figure illustrates the historical and predicted counterfactual Spanish aggregate electricity demand from Jan. 2019 to Dec. 2023. Predictions in red use real CO_2 prices and real TTF natural gas prices. Predictions in blue use CO_2 and TTF price projections assuming that the crisis had not happened.

3.2 Results from the wholesale market simulations

Combining the elements described above, we have all the inputs required to simulate the Spanish wholesale electricity market under our five scenarios of interest.[27] In the following sub-sections, we summarize the key results.

Simulated electricity market prices across scenarios. Figure 5 plots the time series of the simulated prices, comparing factual prices against the simulated electricity prices in all counterfactual scenarios, and Table 2 reports the average electricity prices under each scenario. As can be seen, the energy crisis had a significant impact on prices, which was only partially offset by the price intervention. In contrast, while the savings measures were effective in curbing consumption, they had little to no downward impact on prices.

During the second half of 2021, gas prices began to rise, driving electricity prices to an average of €163/MWh. This represents a 246% increase compared to the coun-

[27] To assess the validity of the electricity market model, Figure A.1 in the Appendix compares the daily averages of the simulated prices under the *Factual* scenario with actual market prices. As shown, the model accurately captures price variations, providing confidence in its use for simulating counterfactual scenarios. In any case, when assessing the impact of the energy crisis and the policy responses, we compare the counterfactual scenarios against the simulated *Factual* scenario, ensuring that any potential biases cancels out across scenarios.

terfactual electricity prices in the absence of the energy crisis.[28] The surge in gas prices continued throughout 2022, accompanied by extreme volatility. The peak occurred in March 2022, when gas prices soared to a record high of over €240/MWh, pushing electricity prices above €500/MWh. The price intervention, which became effective on June 14, mitigated the pass-through of gas price spikes to electricity prices from that point onward. In the second half of 2022, this intervention reduced the average wholesale market price from €237/MWh (under the *Savings-Intervention-Only* scenario) to €116/MWh. During 2023, gas prices remained relatively high in the first two months but then began to decline rapidly. This, combined with the increase in the trigger gas price for the intervention, meant that the *Iberian Solution* was rarely implemented over this time. Indeed, during the first half of 2023, the average price without the intervention (€97/MWh) was very close to the factual price (€93/MWh). The green curve in Figure 5, representing the *Savings-Intervention-Only* scenario, reveals that the energy-saving measures had a limited impact on market prices. The small differences observed are primarily due to variations in the marginal costs of different gas plants, which continued to be the price-setting technology throughout the period.

Table 2: Demand-weighted Average Prices across Scenarios (€/MWh)

Scenario	2021 (Q3-4)	2022 (Q1-2)	2022 (Q3-4)	2023 (Q1-2)
Factual	163	200	116	93
Business as Usual	163	213	241	97
Savings Interv. Only	163	212	237	95
Price Interv. Only	–	201	118	95
No Energy Crisis	66	68	71	72

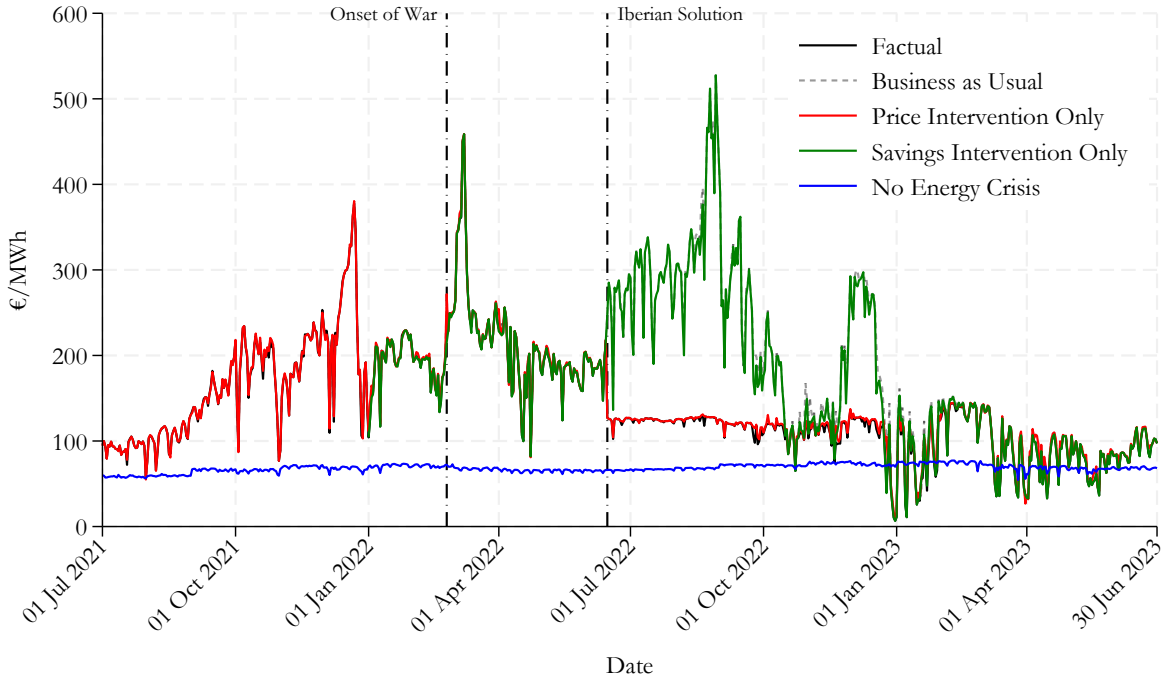
Notes: This table reports the simulated demand-weighted average prices in the Spanish wholesale market during each period under each scenario. The price intervention took place in 2022.

Distributional effects between consumers and electricity generators. The price effects summarized in Figure 5 led to substantial wealth transfers from consumers to electricity generators. Aggregate effects on power plants' earnings, costs, and profits are summarized in Table 3.

Relative to the *No-Energy-Crisis* scenario, earnings by Spanish power plants increased by €49,287M over the two-year period (July 2021 to June 2023), representing a 250% increase. The surge in the costs of marginal gas-fired plants, which were passed through to wholesale electricity prices, led to a 224% increase in firm profits (€27,090M),

[28] Note that during this period, no price intervention was in place. As a result, the *Factual* scenario and the *Savings-Intervention-Only* scenario yield the same wholesale electricity prices.

Figure 5: Simulated Wholesale Electricity Prices across Scenarios



Notes: This figure reports the daily demand-weighted averages of the simulated electricity wholesale prices from July 2021 until the end of June 2023 under the various scenarios. The market intervention (Iberian Solution) was implemented in June 2022. The *Factual* and the *Price-Intervention-Only* scenarios overlap until mid 2022 (when the price intervention had not yet been implemented) and during the first half of 2023 (when the reference gas price was not binding). The *Business-as-Usual* and *Savings-Intervention-Only* scenarios overlap on most dates, given that the price impacts of the saving measures were limited.

compared to the *No-Energy-Crisis* scenario. Profit increases were primarily driven by the growth in inframarginal rents earned by nuclear, hydro, and renewable energy resources.[29] Since their costs remained unchanged during the crisis, yet their output prices more than doubled – even after the market intervention – these generators benefited significantly from the crisis, at the expense of consumers. The overall increase in costs, €22,189M, is mainly explained by the gas price increase caused by the crisis.

The price and savings interventions helped curb wealth transfers from consumers to firms by partially reducing inframarginal rents. Had the interventions not been implemented, the crisis would have increased firm earnings by 262% – instead of 250% – and the increase in firms' profits would have reached 282% – instead of 224%. After factoring in the subsidies to gas- and coal-fired plants (paid by consumers in their electricity bills), consumers collectively saved €1,605M due to the price intervention and €2,346M

[29] A fraction of renewable energy assets receives regulated payments – specifically, they earn the market price supplemented by a regulated top-up. This top-up is revised every two years, based on the market revenues earned during the preceding regulatory period. As a result, the excess earnings accrued during the energy crisis led to a reduction in regulated payments in the subsequent period. This adjustment is not computed in our estimated effects.

from the savings measures. Together, these interventions reduced consumer payments by €4,052M. In sum, while the interventions successfully softened the impact of the crisis, they only partially mitigated its effects.

Table 3: Policy Intervention and Crisis Effects on Earnings, Costs, and Profits (M€, July 2021-July 2023)

	Earnings	**Costs**	**Profits**
Energy-Crisis effect	49,287	22,189	27,090
Full-Intervention effect	-4,052	8,585	-12,651
Price-Intervention effect	-1,605	9,721	-11,344
Savings-Intervention effect	-2,346	-1,480	-861

Notes: This table reports the changes in earnings by power generators located in Spain, their production costs, and their profits due to the policy interventions and the energy crisis. Payments and Costs include the gas compensation in the *Factual* and *Price-Intervention-Only* scenarios.

Effects on other variables of interest. The electricity price effects of the interventions also triggered an increase in exports, which in turn led to higher thermal production and associated emissions in Spain. As shown in Table 4, Spain increased its exports by 120% due to the price intervention. Had only the savings measures been implemented, the exports would have increased by 36% relative to the *Factual* scenario due to the slight price reduction. The combined effect of both policies was to increase exports by 162%. The increase in exports had to be met with an increase in thermal generation, which contributed to higher emissions even though the savings measures helped mitigate this effect. More specifically, the price intervention increased emissions by 23%, while the energy savings measures allowed for a 6% reduction. The combined effect of both policies was to increase the power-sector carbon emissions by 18%.[30]

4 Retail Market Impacts

Thus far, our analysis has focused on the aggregate wealth transfers from consumers to electricity producers induced by the energy crisis, along with the mitigating role of policy interventions. In this section, by analyzing retail market impacts, we explore the heterogeneity across households and provide further distributional insights.

[30]However, it is important to note that the increase in emissions in Spain might have been partially offset by a reduction in emissions in France, to the extent that Spanish electricity exports helped reduce gas-fired generation in France. Since carbon is a global pollutant, only the net cross-border effect should matter.

Table 4: Policy Intervention and Crisis Effects on Imports, Exports, Thermal Generation, and Demand (%, July 2021-July 2023)

	Imp-Exp	Thermal Generation	CO2 Emissions	Demand
Energy-Crisis Effect	-140.7	18.4	20.7	-4.8
Full-Intervention Effect	-161.8	23.8	17.8	-4.2
Price-Intervention Effect	-120.5	28.6	22.6	0.0
Savings-Intervention Effect	-36.1	-6.0	-5.8	-4.2

Notes: This table reports the changes in percentage terms in the trade balance (imports-exports), thermal generation, carbon emissions, and total demand in Spain due to the crisis and the policy interventions. Note that Demand is equal to imports-exports plus in-home generation.

4.1 Background and data

We begin by constructing a database of residential electricity consumption at the zip code level. This dataset is then merged with detailed characteristics of each zip code, as well as socio-demographic attributes of the households residing within them.[31]

Hourly residential electricity consumption data at the zip code level has been provided by Datadis (2023). After restricting to continental Spain and filtering out missing data and outliers, we are left with 8,390 × 17,520 zip code-by-hour observations, covering the period from July 2021 until June 2023 and roughly 16.3 million households (more than 90% of the population of peninsular Spain). Besides, using the total number of resident households, we compute the average household's hourly electricity consumption in each zip code.[32]

We merge the demand data with demographic characteristics, also at the zip code level, compiled by MB-Research (2023) based on information from the Spanish National Statistics Institute (INE) and Eurostat. This includes total population, total number of households, average household size, average purchasing power (€/pc), and the number of households in each national-level income quintile, for the year 2017.[33]

In addition, we obtain data on the share of households equipped with electric heating and air conditioning, based on the 2021 survey from INE (2023). These shares are reported at the province level for all provinces, and at the municipality level for provincial capitals and municipalities with more than 50,000 inhabitants. Therefore, we are able to

[31] We restrict our analysis to continental Spain, excluding the Balearic Islands, the Canary Islands, Ceuta, and Melilla, where specific electricity market regulations apply.

[32] We also know the number of electricity contracts registered in each zip code, which typically exceeds the number of resident households. In some zip codes, the discrepancy can be large, explained mostly by a large number of second homes. For this reason, we prefer to consider the number of resident households in each zip code and we exclude those where the discrepancy with the number of contracts is too large.

[33] Income quintile thresholds are defined using the full Spanish population. The dataset then reports, for each zip code, the number of households belonging to each quintile.

compute equipment shares separately for these cities and for the remaining areas within each province.

We classify zip codes into "cities", "towns and semi-dense areas" and "rural areas" following Dijkstra et al. (2021). Finally, we obtained weather data from AEMET (2024), consisting of hourly observations for hundreds of stations located across Spain. For our models to predict aggregate counterfactual low voltage electricity demand (more details in subsection 4.2 below), we aggregate weather data to the national level, using population-weighted averages. We also separately include measurements from the five most populous provinces (Madrid, Barcelona, Valencia, Seville, and Alicante). For the clustering algorithm described in subsection 4.3.1, we match hourly zip code-level residential electricity demand with average weather observations at the province level.

4.2 Computing household electricity bills

To assess the impact of the energy crisis and the associated policy interventions on households, we estimate electricity bills at the zip code level under each of the five simulated scenarios. We then compute bill changes across scenarios to quantify four distinct effects: the *full-intervention effect*, the *savings-intervention effect*, the *price-intervention effect*, and the *energy-crisis effect*.[34]

Counterfactual retail prices for households. Household electricity bills are computed by assuming that retail prices are equal to the simulated wholesale electricity prices obtained in Section 3, augmented by network charges, taxes, and other components typically included in household bills. This assumption is motivated by the fact that a large share of Spanish households are exposed to real-time pricing, where retail prices closely track wholesale market prices (Fabra et al., 2021). The remaining households typically purchase electricity through fixed-price contracts negotiated in the retail market, which are usually set for one year. Empirical evidence suggests that these fixed-price contracts adjust to wholesale price changes upon renewal, thereby reflecting the crisis-induced price increases over time.[35]

Two clarifications are in order. First, in both the *Factual* and *Price-Intervention-Only* scenarios, household bills include the observed per-kWh surcharge used to finance subsidies to thermal generators. Second, we abstract from any other tax changes that occurred during our sample period, as our focus is on bill variations due to the crisis and market interventions.

[34] Although we do not offer a formal utilitarian assessment of these outcomes, it is worth emphasizing that the bill effects are likely to be more relevant for low-income households, who often face tighter budget constraints. Given that electricity expenses comprise a larger portion of their total spending, and their consumption is primarily focused on essential needs, the marginal utility of income in this context is expected to be higher for low-income households.

[35] See Figure A.2 in the Appendix.

Counterfactual demand for households. Our simulations for household bill impacts take into account the fact that, as described in Section 3.1, electricity demand in Spain was affected by price and non-price interventions. To do so, we predict households' counterfactual demand in the absence of the crisis with machine learning techniques analogous to those used in Section 3.1. The zip code-level data described in Section 4.1 is not available for years prior to 2021, such that it cannot be used for this exercise. Instead, we use the aggregate "low voltage" hourly electricity demand made available by REE (2024) as the outcome variable. This can be considered as a proxy for residential demand, even though small commercial establishments may be included in this category.[36] The predictor variables are the same as described at the end of Section 3.1. We also perform the same tuning and cross-validation steps (validation results are described in Appendix C.4).

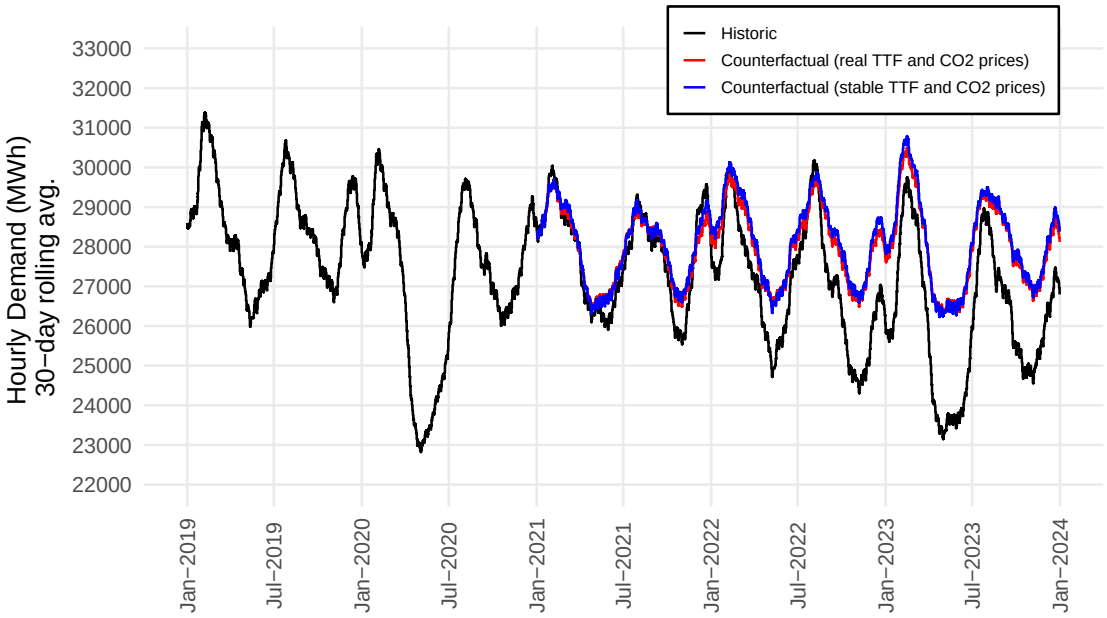
Our counterfactual predictions are depicted in Figure 6. The blue curve is for predictions using stable TTF and CO_2 prices. Comparison with the red curve (using real TTF/CO_2 prices) suggests that short-term residential demand is somewhat inelastic. Observed historic consumption (in black) is generally lower than our counterfactual predictions, especially from the second half of 2022 onward. Differences at the end of 2021 may also be attributed to the salience effects of the crisis, as the stark price increases were already frequently reported in the news. We also note some slight discrepancies at the beginning of 2021, which may be attributed to COVID-19 containment measures that were still in place during that period.[37]

We then construct counterfactual consumption specific to each zip code based on these aggregate low voltage demand counterfactuals. For each hour, we compute the percent reduction in consumption corresponding to the difference between factual and predicted counterfactual low voltage demand. We then apply those percent reductions to the consumption data at the zip code level. Consistently with the assumptions presented in Table 1, we consider the counterfactual demand with real TTF and CO_2 prices for the *BaU* and *Price-Intervention-Only* scenarios, and we consider the counterfactual demand with stable TTF and CO_2 prices in the case of the *No-Energy-Crisis* scenario.

With this procedure, we assume that residential electricity consumption decreases homogeneously across zip codes (in proportion to observed consumption). In particular, we assume away heterogeneous responses to the crisis across income groups. This may introduce some bias. However, a survey of 5,000 Spanish households conducted during our study period suggests limited discrepancies between social classes in their stated sensitivity to changes in electricity prices (see Table A.3). If anything, low-income households seemed to be less aware or less responsive. Therefore, any potential biases from our homogeneity assumption would most likely be in the direction of underestimating bill

[36] We use the "E0" category, which corresponds to consumption at voltage lower than one kV.
[37] The official State of Emergency due to COVID-19 was still effective until 9th of May, 2021 (BOE, 2020).

Figure 6: Actual versus predicted low voltage electricity demand



Notes: This figure illustrates the historical and predicted counterfactual "low voltage" electricity demand in Spain from Jan. 2019 to Dec. 2023. Predictions in red use real CO_2 prices and real TTF natural gas prices. Predictions in blue use CO_2 and TTF price projections assuming that the crisis had not happened.

increases for the least wealthy households.

It is important to note, however, that bill impacts can vary substantially between zip codes, even with a homogeneous decrease in consumption. This is because our approach still accommodates high-frequency price changes, as shown in Figure 5, which affect households differently depending on their consumption patterns over time. We expand on this point in the sub-section below.

4.3 Distributional implications of the energy crisis and policy interventions

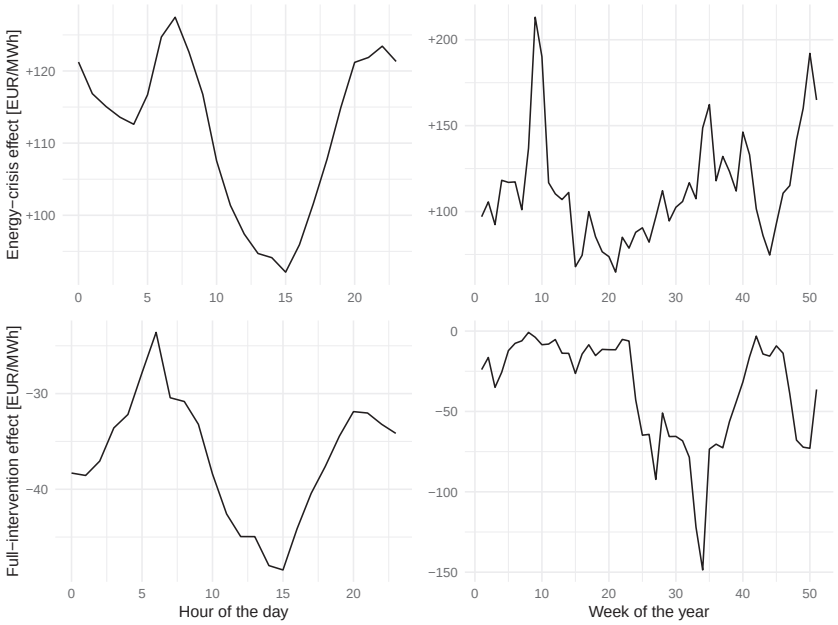
We begin by illustrating the heterogeneous price impacts over time. Figure 7, Panel (a), displays the evolution of retail electricity prices under the energy crisis and the price intervention, broken down by hour of the day (left) and by week of the year (right). Clear intraday patterns emerge. For example, price increases are mitigated during sunny hours due to the availability of solar generation. In contrast, prices are most strongly affected by the energy crisis during the winter months, as evident in the right-hand side of the figure. Conversely, the price intervention appears to have been more effective in suppressing prices during the summer.[38]

The total impact of these price changes on household electricity bills depends on consumption patterns, which vary substantially across zip codes. Figure 7, Panel (b), illustrates the consumption patterns of two zip codes: Elche (Alicante), located in a sunny region on the Mediterranean coast, and Ponferrada (León), located in a mountainous area in northern Spain. The substantially higher electricity consumption in Elche is largely driven by increased daytime usage during the summer months, primarily due to the more widespread and intensive use of air conditioning relative to Ponferrada. In this case, the within-day and across-month differences in consumption profiles between Elche and Ponferrada work in favor of Elche, which was – in per-unit terms – relatively less affected by the energy crisis and benefited relatively more from the price intervention (as can be seen from the lower figures in Figure 7, Panel (a)).

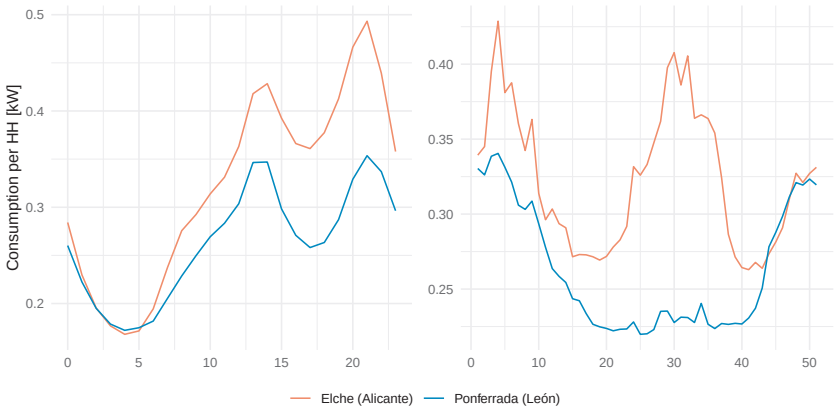
Our detailed consumption data at the zip code level allows for multiple comparisons analogous to those presented above. To deal with the high dimensionality of the data, we group zip codes based on the similarity of their average consumption profiles using *k-means* clustering, as described in the following section. This method allows us to strike a balance between capturing heterogeneity and maintaining interpretability.

[38] This pattern was also visible in Figure 5, when comparing wholesale prices under the *Factual* and *BaU* scenarios.

Figure 7: Illustration of the heterogeneous price effects across time and households



(a) Energy-crisis and full-intervention effects on retail prices



(b) Electricity consumption patterns for two illustrative zip codes

Notes: These figures use data from two zip codes – Elche (Alicante) in orange and Ponferrada (León) in blue – to illustrate how time-varying effects of the energy crisis and the policy intervention (Panel (a)) can lead to heterogeneous effects on households explained by differences in consumption patterns across hours of the day and weeks of the year (Panel (b)).

4.3.1 Clustering algorithm for consumption profiles

We use *k-means* clustering to group zip codes according to their electricity consumption profiles. Recall that we observe zip code-level consumption for each hour from July 2021 until June 2023. Clustering based on untransformed electricity consumption would be somewhat uninformative, as the resulting separations would be driven primarily by differences in the population sizes of zip codes. Instead, the main variables included in our clustering algorithm are hour-by-month consumption "shares" calculated as follows:

$$s_{z,h,m} = \frac{\bar{e}_{z,h,m}}{\bar{e}_z},$$

where $\bar{e}_{z,h,m}$ is the average electricity consumption for zip code z, in hour-of-day $h \in [1, ..., 24]$, and month-of-year $m \in [1, ..., 12]$; and $\bar{e}_z$ is the unconditional average hourly consumption of a given zip code. This results in a total of $24 \times 12 = 288$ shares for each zip code. These shares mute the level differences between zip codes, while retaining seasonal and intraday variations in electricity consumption.

We also include average consumption by number of households and average by number of electricity contracts, both to capture differences in intensity of electricity usage. In terms of weather variables, we include province-level min., max., and median temperature averages by month-of-year. To account for geographic differences, we include min. and max. altitudes at the province level. Finally, the algorithm also incorporates zip code-level indicators for the degree of urbanization. In total, the clustering is based on 375 variables.

We test the performance of clustering algorithms that construct 1 to 20 groups of zip codes. We aim to find the "optimal" number of groups that reduces the dimensionality of our data while retaining meaningful and interpretable heterogeneity in energy consumption profiles. To assess performance, we compute the total within-cluster sum of squares (Krzanowski and Lai, 1988) and the average "silhouette" score or width (Rousseeuw, 1987) resulting from each cluster allocation.[39] We determine the optimal number of clusters to be 6.

Table 5 provides some descriptive statistics for the zip codes allocated to each of the 6 clusters. Consumption profiles are presented in Figure 8, where Panel (a) plots the average consumption of households by hour of the day, and Panel (b) plots average consumption by month of year. Figures on the left correspond to demand in levels (kWh per household-year), while figures on the right depict demand shares. The demand patterns should be interpreted alongside the map on Panel (c), which shows the geographic location of the

[39] The within-cluster sum of squares measures the similarity (distance) between observations and the "centroid" of a given cluster. Lower values imply more similarity. The silhouette score is a metric that incorporates both within- and between-cluster variability. Silhouette scores close to 1 imply high separation and silhouette scores close to 0 indicate a high degree of overlap between clusters. Appendix D presents more details about the clustering algorithm and the performance metrics.

zip codes belonging to each cluster.

Table 5: Descriptive statistics of zip codes and households grouped in each cluster

Cluster	**1**	**2**	**3**	**4**	**5**	**6**	**Full Sample**
Electricity Cons. [kWh/year.HH]	3131	3717	3294	3691	2625	2946	**3229**
Avg. Electricity Price [c€/kWh]	24.5	24.2	24.3	24.5	24.3	23.8	**24.4**
Income [k€/year.HH]	38.3	37.5	33.9	31.8	39.1	26.4	**36.3**
Household Size	2.5	2.5	2.4	2.6	2.4	2.3	**2.5**
Electric Heating [%]	22.6	13.6	19.4	49.7	14.4	12.5	**27.6**
Air Conditioning [%]	61.4	33.5	22	71.6	10.2	19.8	**50.5**
Cities [%]	65.7	11.1	34.9	54.1	58.7	0	**55.9**
Towns & semi-dense areas [%]	25.5	36.1	40	35.8	25.3	14.4	**29.6**
Rural [%]	8.8	52.8	25.1	10.1	16	85.6	**14.5**
Min Temp. [ºC]	1.7	0.5	2.1	5.2	2.8	0.2	**2.8**
Max Temp. [ºC]	31.9	30.4	28.7	33.1	26.8	29.8	**31**
Number of zip codes	1762	1812	826	1304	1788	898	**8390**
Total households (thousands)	7008	889	1018	4331	2942	138	**16325**

Notes: All values are averages weighted by zip code population. The average electricity price is the price faced by households in the *Factual* scenario depending on the timing of their consumption, with different consumption patterns leading to different average prices.[a] The share of electric heating, the share of air conditioning are based on province-level data or city-level data (for large cities). Average min. and max. temperatures are based on province-level data. Average min. temperatures correspond to those of the month of January, while average max. temperatures correspond to those of the month of August.

[a]Tables A.1 and A.2 in the Appendix explore the determinants of household consumption and average prices at the zip code level.

From the intraday patterns, we observe that households in Cluster 6 generally exhibit lower electricity demand than those in any other cluster and tend to consume relatively more electricity during periods when prices are lower. As shown in Table 5, these households are predominantly located in small rural zip codes. The demand peak observed in Cluster 6 during August may be attributed to increased air conditioning usage.

Households in Cluster 5 also show relatively low electricity demand across all hours of the day and throughout the year. However, their average electricity prices are more closely aligned with the national average. The map indicates that Cluster 5 primarily includes zip codes in Northern Spain, where the climate is milder and air conditioning is seldom necessary, though some electric heating may be used during the winter months.

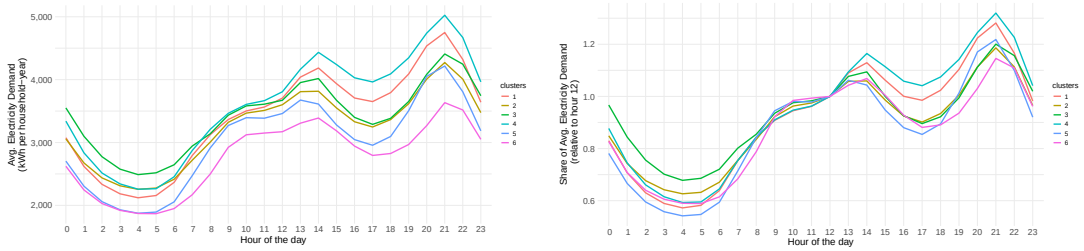
Cluster 4 is characterized by the highest electricity demand, particularly during daytime hours and the summer months. Households in this cluster also face above-average electricity prices. Notably, Cluster 4 includes some of the southernmost zip codes in Spain, where extreme summer heat drives high air conditioning usage.

Cluster 1 comprises zip codes primarily located in central regions which also ex-

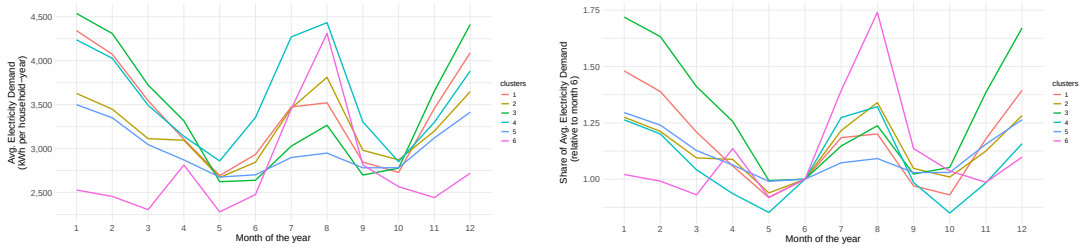
perience hot climates and significant AC usage, although with less pronounced summer peaks than Cluster 4. These households similarly encounter above-average electricity prices. Importantly, Cluster 1 includes several of Spain's most populous and affluent cities, such as Madrid and Barcelona.

Cluster 2 consists primarily of rural zip codes located in colder regions. Its electricity consumption patterns resemble those of Cluster 5 (in the Northern region), though Cluster 2 exhibits a more pronounced peak during the summer months. Cluster 3 also includes colder zip codes, likely situated in more mountainous areas where access to gas infrastructure is often limited. This may account for the cluster's high electricity demand during the winter months.

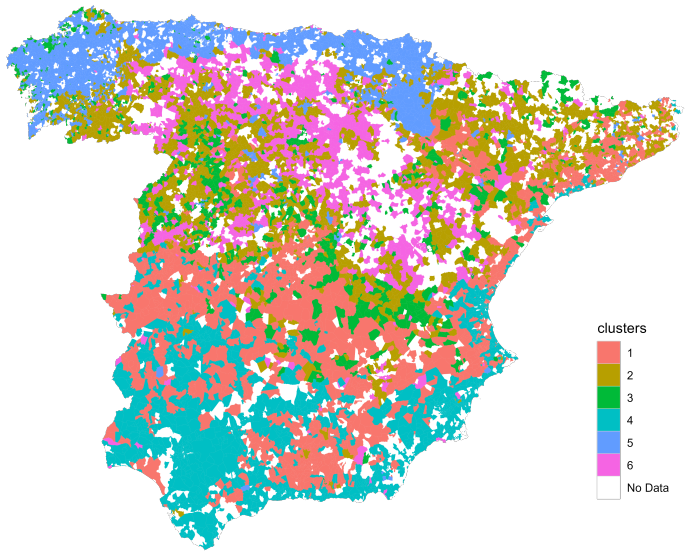
We conclude that the data-driven clustering algorithm yields a meaningful and interpretable segmentation of households. In the following sub-sections, we estimate how the impact of the energy crisis varied across clusters. Significant differences are to be expected, given the variability in electricity prices and the extent to which demand is influenced by climate and sociodemographic factors.



(a) Average household electricity demand by hour of the day



(b) Average household electricity demand by month of year



(c) Geographic location of zip codes by cluster

Figure 8: Electricity consumption profiles across clusters and their geographic locations

Notes: Panels (a) and (b) show the average electricity consumption by hour of the day and by month of year of Spanish households grouped in the six clusters. The left panels show the results in levels, while the right panels show the consumption shares. Geo-locations belonging to each cluster are illustrated in Panel (c).

4.3.2 Distributional implications of the bill changes

We are now equipped to quantify the bill impacts of the energy crisis and associated policy interventions across households depending on their location and income levels.

Across clusters and locations. As discussed above (and further studied by Reguant, Fabra, and Wang, 2025), a household's location is strongly correlated with its installed electrical equipment (mainly electric heating and AC), which affects both the level and the timing of electricity demand. Consistent with this, our clustering procedure groups households located in similar areas, reflecting shared characteristics. The intensity of the crisis and the mitigating effects of the interventions varied substantially across these clusters. Table 6 reports estimates of these effects on households' annual bill, in absolute terms and in percentage of their disposable income.

Table 6: Simulated Effects of the Energy Crisis and the Policy Interventions on Household Electricity Bills (€/year, July 2021-June 2023)

Cluster nº	**1**	**2**	**3**	**4**	**5**	**6**	**Full Sample**
Energy Crisis	329.2	386	344	387.7	269.9	297.2	**337.7**
% of disp. income	*0.91*	*1.09*	*1.04*	*1.27*	*0.71*	*1.16*	***0.99***
Full Intervention	-143.3	-174	-142.2	-176.2	-116.4	-146	**-148.8**
% of disp. income	*-0.40*	*-0.49*	*-0.43*	*-0.58*	*-0.31*	*-0.57*	***-0.44***
Savings Intervention	-54.6	-65.2	-58.6	-63.2	-47	-51.7	**-56.3**
% of disp. income	*-0.15*	*-0.18*	*-0.18*	*-0.21*	*-0.12*	*-0.2*	***-0.17***
Price Intervention	-77.9	-95.5	-72.6	-99.9	-60.4	-82.9	**-81.3**
% of disp. income	*-0.22*	*-0.27*	*-0.22*	*-0.33*	*-0.16*	*-0.32*	***-0.24***

Notes: This table reports changes in annual household electricity bills resulting from the energy crisis and the associated policy interventions, both in levels (€) and relative to households' disposable income (%). The *full-intervention effect* includes both the *savings-intervention effect*, the *price-intervention effect*, and their interaction.

In absolute terms, Clusters 4 and 2 were those who suffered from the largest bill increase due to the crisis (+ €388 and + €386/year) and who benefited the most from the policy intervention (− €176 and − €174/year). Yet in Cluster 4, which comprises less wealthy zip codes located in the south of Spain, these effects represented a larger share of households' disposable income. At the other end, Cluster 5, which comprises rather wealthy zip codes located in northern Spain, was relatively spared by the crisis effect (+ €270/year, representing 0.7% of their disposable income) and benefited less from the policy interventions (− €116/year).

Appendix Figures A.3 and A.4 further illustrate the distribution of these effects across Spanish provinces. In absolute terms, both the crisis and the policy interventions had stronger impacts in the south and east of Spain, while the mildest effects were

observed in the north. Tarragona experienced both the largest increase in annual electricity bills due to the crisis (+ €425/year) and the strongest relief from the intervention (− €192/year). However, when measured relative to income, the southern provinces were hit hardest. In Granada, Córdoba, Jaén, and Toledo, the bill increase reached 1.4% of households' disposable income. These same regions also benefited most from the interventions, with reductions amounting to as much as 0.6% of households' disposable income.

Across income bins. Households were affected by the crisis and policy interventions to varying degrees depending on their income levels. We now assess the average bill impact across income quintiles.

A key challenge is that our income data is only available at the zip code level, which masks income heterogeneity among households within the same area. A naïve approach would be to classify zip codes into national income quintiles based on their average income.[40] However, this method averages out within-zip code variation, leading to a substantial underestimation of the gap between low- and high-income households, as illustrated in Table 7. To more accurately capture the differentiated impacts of the crisis and interventions across household-level income quintiles, some extrapolation beyond zip code-level income data is needed.

Table 7: Average income per quintile: Household level or zip code level

Income Quintile	Average Income (household-level)	Average Income (zip code-level)
1	9,481	24,773
2	17,980	30,945
3	25,622	35,457
4	34,926	40,022
5	57,485	50,316

Notes: This table reports average disposable income per household within each quintile, depending on the level of aggregation used to construct the quintiles. At the *household level*, all individual households in Spain are ranked by disposable income and grouped into five quintiles, within which the average income is computed. At the *zip code level*, zip codes are ranked by the average income of their resident households and grouped into five quintiles with equal total household population; average income is then computed within each group using population weights. The average income by household-level income quintile is obtained from Eurostat (2025) for the year 2020. The average income by zip code-level income quintile is computed with the dataset described in Section 4.1.

To address this, we develop a procedure that models the relationship between income and the estimated effects of the crisis and policy interventions that we observe at the zip

[40] We develop this alternative in Appendix A.2. Reguant, Fabra, and Wang (2025) show that such an approach obscures much of the true distributional impact.

code level, before extrapolating this relationship to fully reflect the household-level income heterogeneity in Spain.[41]

First, we assess the correlation between bill changes and the average income of resident households in a given zip code i, while controlling for average household size, the share of electric heating and air conditioning, and the degree of urbanization. Specifically, we estimate the following regression:

$$log(\Delta Bill_i) = \beta_0^c + \beta_{Inc}^c log(Income_i) + \beta^c X_i + \epsilon_i. \quad (1)$$

Table 8 presents the results from estimating the model using the full sample of zip codes. To allow for heterogeneity in the effects across different types of areas, we also estimate the model separately for each cluster c (Tables A.4–A.7 in the Appendix).

We find that household income is positively correlated with bill changes resulting from the crisis and policy interventions, with relatively little variation across clusters. A 1% increase in average income is associated with a 0.4% higher bill increase on average, ranging from 0.31% in cluster 1 to 0.52% in cluster 5. Similar magnitudes are observed for the policy intervention effects, with coefficients ranging between 0.34% and 0.48%. In addition, larger household size, higher shares of electric heating or air conditioning, and zip code locations in towns or rural areas are all associated with greater bill impacts. These patterns are consistent with higher average electricity consumption in such areas (see Appendix Table A.1).

[41] Reguant, Fabra, and Wang (2025) propose a method to infer household-level income from zip code aggregates using individual electricity consumption data. However, we cannot apply it here, as our electricity data is also aggregated at the zip code level.

Table 8: Determinants of the Crisis and Interventions Effects for households

$Y = log(\lvert\Delta Bill\rvert)$	**Energy-Crisis effect**	**Full-Intervention effect**	**Savings-Intervention effect**	**Price-Intervention effect**
log(*Income*)	0.40^{***}	0.39^{***}	0.38^{***}	0.40^{***}
	(0.02)	(0.02)	(0.01)	(0.02)
Household Size	0.27^{***}	0.27^{***}	0.27^{***}	0.28^{***}
	(0.02)	(0.02)	(0.02)	(0.02)
Electric Heating	0.43^{***}	0.46^{***}	0.38^{***}	0.51^{***}
	(0.02)	(0.02)	(0.02)	(0.03)
Air Conditioning	0.32^{***}	0.36^{***}	0.24^{***}	0.46^{***}
	(0.02)	(0.02)	(0.01)	(0.02)
Degree of urbanization (ref.: Cities)				
Towns and Semi-dense areas	0.25^{***}	0.25^{***}	0.23^{***}	0.26^{***}
	(0.01)	(0.01)	(0.01)	(0.01)
Rural areas	0.29^{***}	0.31^{***}	0.28^{***}	0.33^{***}
	(0.01)	(0.01)	(0.01)	(0.01)
(Intercept)	0.84^{***}	0.08	-0.70^{***}	-0.71^{***}
	(0.17)	(0.17)	(0.16)	(0.17)
R^2	0.27	0.29	0.24	0.34
Adj. R^2	0.27	0.29	0.24	0.34
Num. obs.	8,390	8,390	8,390	8,390

Notes: $^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$. Each observation corresponds to one zip code. The dependent variable $log(|\Delta Bill|)$ is the log of the absolute increase in households' annual bill in the zip code due to the crisis, or the log of the absolute decrease in households' annual bill due to the interventions. *Income* stands for the average disposable income per household per year, *Household Size* for the mean number of people per household. *Electric Heating* and *Air Conditioning* denote the share of households with electric heating and air conditioning, respectively.

We next use the estimated models to extrapolate the effects across household-level income quintiles. To compute the average effect for households in income quintile N, we proceed as follows. First, for each zip code i, we simulate the counterfactual effect, denoted $\Delta\widehat{Bill}_i^N$, by assuming that the average income in zip code i equals the national average income for quintile N. The objective is to estimate the change in household bills for residents of zip code i who belong to income quintile N, accounting for both the income level of that quintile and other characteristics specific to zip code i. To do so, we use the estimated parameters β_0^c, β_{Inc}^c, and β^c, depending on the corresponding cluster c to which zip code i belongs. In this computation, we hold all control variables $\mathbf{X}_i$ at their observed values for zip code i, replacing only the variable $Income_i$ with the national average income for quintile N.

Equipped with these $8,390 \times 5$ (zip code $\times$ quintile) projected effects $\Delta\widehat{Bill}_i^N$, we compute five national average effects $\Delta\widehat{Bill}^N$ – one for each household income quintile N – as follows:

$$\Delta\widehat{Bill}^N = \frac{1}{pop^N} \sum_i \Delta\widehat{Bill}_i^N \times pop_i^N$$

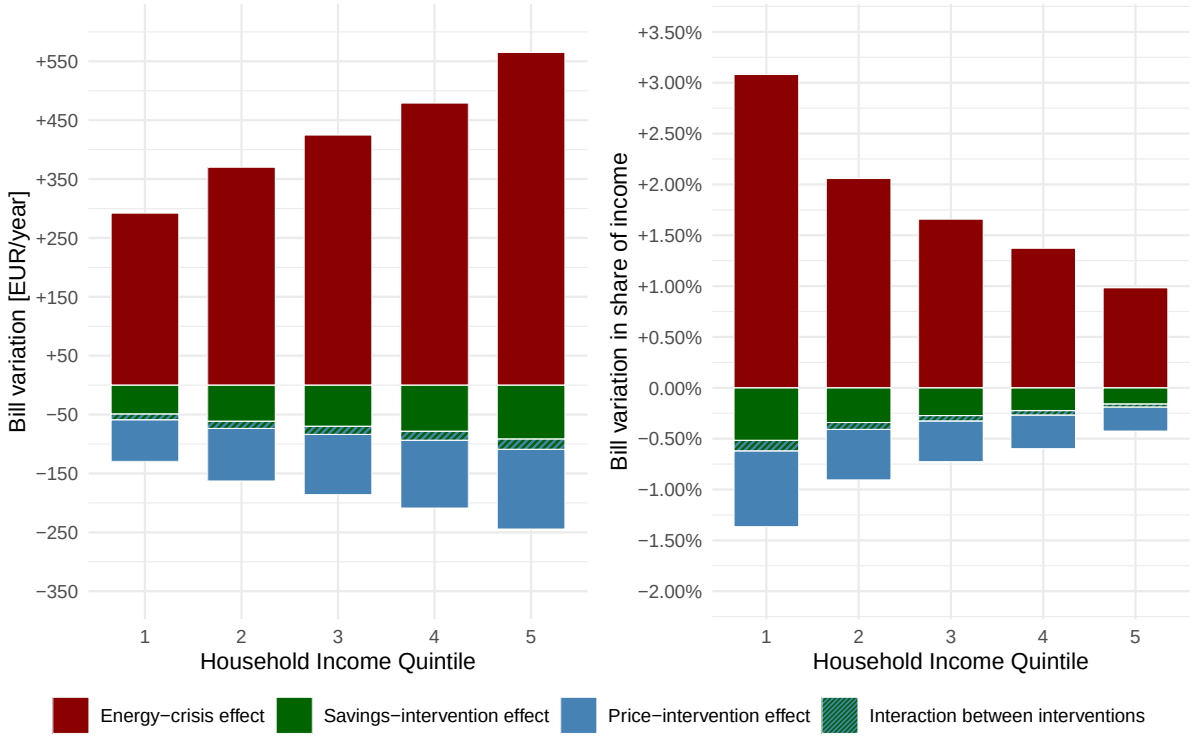
where pop_i^N denotes the number of households in income quintile N residing in zip code i, and pop^N is the total number of households in quintile N nationwide. This weighted average captures the national-level effect for households in income quintile N, appropriately accounting for their geographic distribution across zip codes.

Figure 9 shows the average bill impacts across household income quintiles. The left panel presents the effects in euros per year, while the right panel expresses them relative to disposable income. As indicated by the red bars, the energy crisis raised annual electricity bills by €292 for households in the first (lowest) income quintile – equivalent to 3.1% of their disposable income. In contrast, households in the fifth (highest) income quintile experienced an average increase of €565, corresponding to 1.0% of their income.

Absent policy interventions, these effects would have been significantly larger. The mitigation effects of the price and savings interventions are shown by the blue and green bars, respectively. Both were substantial, with a slightly larger share attributable to the price intervention. Without the interventions, bills for the lowest-income households would have increased by an additional €129, reaching a total impact of 4.45% of their income. For the highest-income households, the interventions reduced bills by €244 in absolute terms – roughly twice as much as in the first quintile, but this represented only 0.4% of their disposable income. For these households, the total impact without intervention would have been limited to 1.4% of their disposable income.

We note some differentiated effects of the crisis and policy interventions across income quintiles, stemming from two main channels: (1) differences in aggregate electricity consumption, and (2) differences in consumption patterns over time, which may align positively or negatively with the timing of price changes. The larger absolute bill im-

Figure 9: Crisis and policy intervention effects by household income quintile



Notes: The bars depict the average effect of the crisis (*red*) and policy interventions (*green* and *blue*) on households' bills for households belonging to each income quintile, expressed in (a) absolute level in €/year and (b) relative to households' disposable income. The *full-intervention effect* is divided into the *savings-intervention effect* (*green*), the *price-intervention effect* (*blue*) and the interaction of both (*hatched green and blue*). The cumulative height of the red, green, blue and hatched bars measures the effect the crisis would have had in the absence of any policy interventions.

pacts observed for higher-income households are primarily driven by the first channel, as electricity consumption is strongly correlated with household income.

To assess the importance of the second channel, we examine the average electricity price paid by households and how this price was affected by the crisis and policy interventions. The linear regression results in Appendix Table A.2 show that households in higher-income zip codes tend to pay slightly higher average prices, and that these prices increased marginally more during the crisis compared to lower-income areas. However, the magnitude of these effects is economically small: a 1% increase in income is associated with an absolute change in the average price paid of approximately 0.005%.

Therefore, we conclude that the larger bill increases observed among high-income households were primarily driven by their higher levels of electricity consumption. However, bill increases were more burdensome for low-income households, as they represented a larger share of their disposable income. In this sense, the energy crisis had regressive effects.

Policy interventions partially mitigated these adverse effects. While low-income households saw smaller absolute reductions in their bills, the relief was relatively larger in proportional terms, given their lower incomes. Thus, the policy interventions had progressive effects across the income distribution.

5 Conclusions

This paper quantifies the distributional impacts of the 2021–2023 energy crisis and the associated policy interventions in Spain, analyzing effects at both the wholesale and retail levels.

At the *wholesale level*, we employ electricity market simulations, combined with machine learning techniques for counterfactual demand prediction, to estimate the wealth transfers induced by the crisis. We find that the surge in gas prices led to a sharp increase in electricity prices, boosting the profits of inframarginal producers not directly affected by the cost shock (nuclear, hydro, and renewable power plants). As a result, a significant share of surplus was redistributed from consumers to producers. Specifically, our estimates indicate that the profits of Spanish electricity firms increased by approximately €27 billion over the two-year duration of the energy crisis.

Policy interventions, such as the "Iberian Solution" and targeted energy-saving measures, mitigated part of this transfer by lowering wholesale prices, thereby reducing the inframarginal gains by almost €13 billion and alleviating the financial burden on consumers.

At the *retail level*, the energy crisis had regressive effects. While high-income households experienced larger increases in electricity bills in absolute terms due to higher electricity consumption, the relative burden was greater for low-income households, for

whom electricity expenditures represent a larger share of income. On average, households in the lowest income quintile lost 3.1% of their disposable income due to higher electricity bills – three times the relative burden experienced by households in the highest quintile.

Policy interventions, by contrast, exhibited progressive effects. Although the absolute bill reductions were larger for higher-income households, the relief was proportionally greater for lower-income groups. On average, households in the lowest income quintile saw a reduction amounting to 1.4% of their disposable income – more than three times larger than the relative benefit received by those in the highest quintile.

Finally, regional disparities in equipment ownership – such as electric heating and air conditioning – amplified the effects due to income differences. Southern Spain, where such technologies are more prevalent and where household income is below the national average, was more severely impacted by the crisis but also benefited relatively more from the interventions.

Our findings highlight the importance of incorporating distributional considerations into the design of crisis response policies.

References

Agencia Estatal de Meteorología. 2024. "Climatologías Diarias." *AEMET OpenData* URL https://opendata.aemet.es/centrodedescargas/inicio.

Ahlvik, L, T Kaariaho, M Liski, and I Vehviläinen. 2025. "Household-Level Responses to the European Energy Crisis." *CEPR Discussion Paper* 19972. URL https://cepr.org/publications/dp19972.

Ahlvik, Lassi, Matti Liski, and Mikael Makimattila. 2024. "Pigouvian Income Taxation." CESifo Working Paper Series 11174, CESifo. URL https://ideas.repec.org/p/ces/ceswps/_11174.html.

Albalate, Daniel and Xavier Fageda. 2021. "On the relationship between congestion and road safety in cities." *Transport Policy* 105:145–152. URL https://www.sciencedirect.com/science/article/pii/S0967070X21000731.

Bachmann, Ruediger, David Baqaee, Christian Bayer, Moritz Kuhn, Andreas Löschel, Benjamin Moll, Andreas Peichl, Karen Pittel, and Moritz Schularick. 2024. "What if? The macroeconomic and distributional effects for Germany of a stop of energy imports from Russia." *Economica* 91 (364):1157–1200.

Behr, Sophie M, Till Köveker, and Merve Kücük. 2025. "Understanding energy savings in a crisis: The role of prices and non-monetary factors." Tech. rep., DIW Discussion Papers. Number 2112.

BELab. Banco de España Data Laboratory. 2025. "CNMC Household Panel (CNMC-PH)." URL doi.org/10.48719/belab.cnmc.ph1922_02. Remote access.

Benatia, David and Samuel Gingras. 2023. "Reaching new lows? The pandemic's consequences for electricity markets." *The Energy Journal* 44 (4):195–221.

Bentleya, Rebecca, Lyrian Danielb, Yuxi Lia, Emma Bakerb, and Ang Li. 2024. "The effect of energy poverty on mental health, cardiovascular disease and respiratory health: a longitudinal analysis." *The Lancet Regional Health – Western Pacific* 35:100734.

Bloomberg. 2024. "Dutch TTF Natural Gas Prices." *Markets–Energy* URL https://www.bloomberg.com/energy.

Bloomberg. 2025. "Terminal Bloomberg." URL https://www.bloomberg.com/professional/products/data/.

Boletín Oficial del Estado. 2020. "Real Decreto 956/2020, de 3 de noviembre, por el que se prorroga el estado de alarma declarado por el Real Decreto 926/2020, de 25 de octubre, por el que se declara el estado de alarma para contener la propagación de infecciones causadas por el SARS-CoV-2." *Ministerio de la Presidencia, Relaciones con las Cortes y Memoria Democrática* BOE-A-2020-13494. URL https://www.boe.es/eli/es/rd/2020/11/03/956.

———. 2022a. "Plan de Choque de Ahorro y Gestión Energética en Climatización." *Jefatura del Estado* BOE-A-2022-12925. URL https://www.boe.es/eli/es/rdl/2022/08/01/14.

———. 2022b. "Plan de Medidas de Ahorro y Eficiencia Energética de la Administración General del Estado y las Entidades del Sector Público Institucional Estatal." *Ministerio de la Presidencia, Relaciones con las Cortes y Memoria Democrática* BOE-A-2022-8563. URL https://www.boe.es/eli/es/o/2022/05/25/pcm466.

Borenstein, Severin, James M. Sallee, and Meredith Fowlie. 2022. "Paying for Electricity in California: How Residential Rate Design Impacts Equity and Electrification." URL https://energyathaas.wordpress.com/2022/09/12/paying-for-electricity-in-california/. Working Paper.

Chen, Tianqi and Carlos Guestrin. 2016. "Xgboost: A scalable tree boosting system." In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 785–794.

Corbeau, Anne-Sophie, Juan Camilo Farfan, and Sebastian Orozco. 2023. "The Iberian Exception and Its Impact." Tech. rep., Center on Global Energy Policy, SIPA, Univer-

sity of Columbia. URL https://www.energypolicy.columbia.edu/publications/the-iberian-exception-and-its-impact/.

Datadis. 2023. *Datadis API*. URL https://datadis.es/private-api.

de Frutos, María-Ángeles and Natalia Fabra. 2012. "How to Allocate Forward Contracts: The Case of Electricity Markets." *European Economic Review* 56 (3):451–469.

Dertwinkel-Kalt, Markus and Christian Wey. 2025. "Why Energy Price Brakes Encourage Moral Hazard, Raise Energy Prices, and Reinforce Energy Savings." *RAND Journal of Economics* (56(2)):129–144.

Deryugina, Tatyana, Don Fullerton, and William Pizer. 2019. "An Introduction to Energy Policy Trade-Offs between Economic Efficiency and Distributional Equity." *Journal of the Association of Environmental and Resource Economists* 6 (S1):S1 – S6. URL https://EconPapers.repec.org/RePEc:ucp:jaerec:doi:10.1086/701515.

Deryugina, Tatyana, Alexander MacKay, and Julian Reif. 2020. "The Long-Run Dynamics of Electricity Demand: Evidence from Municipal Aggregation." *American Economic Journal: Applied Economics* 12 (1):86–114. URL https://www.aeaweb.org/articles?id=10.1257/app.20180256.

Dijkstra, Lewis, Aneta J Florczyk, Sergio Freire, Thomas Kemper, Michele Melchiorri, Martino Pesaresi, and Marcello Schiavina. 2021. "Applying the degree of urbanisation to the globe: A new harmonised definition reveals a different picture of global urbanisation." *Journal of Urban Economics* 125:103312.

Dirección General de Tráfico and EpData. 2024. "Muertos en accidentes de tráfico en los últimos 10 años." *Datos de víctimas mortales en accidentes de tráfico* URL https://www.epdata.es/datos/accidentes-trafico-datos-estadisticas/65/espana/106.

Draghi, Mario. 2024. "The Future of European Competitiveness." Tech. rep., European Commision.

European Commission. 2023. "Report from the Commission to the European Parliament and the Council on the review of emergency interventions to address high energy prices in accordance with Council Regulation (EU) 2022/1854." Tech. rep. URL https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52023DC0302.

———. 2025. "Timeline - Energy prices and security of supply." Tech. rep.

Eurostat. 2025. "Mean and median economic resources of households by income, consumption and wealth quantiles [icw_res_02]." URL doi.org/10.2908/ICW_RES_02.

Fabra, Natalia. 2023. "Reforming European electricity markets: Lessons from the energy crisis." *Energy Economics* 126 (C). URL https://ideas.repec.org/a/eee/eneeco/v126y2023ics0140988323004619.html.

Fabra, Natalia and Imelda. 2023. "Market power and price exposure: learning from changes in renewable energy regulation." *American Economic Journal: Economic Policy* 15 (4):323–358.

Fabra, Natalia, Aitor Lacuesta, and Mateus Souza. 2022. "The implicit cost of carbon abatement during the COVID-19 pandemic." *European Economic Review* 147:104165. URL https://www.sciencedirect.com/science/article/pii/S0014292122000939.

Fabra, Natalia, Mar Reguant, David Rapson, and Jinyuan Wang. 2021. "Estimating the Elasticity to Real Time Pricing: Evidence from the Spanish Electricity Market." *AEA Papers and Proceedings, American Economic Association* 111:425–29.

Ferdinandusse, Marien and Mar Delgado-Téllez. 2024. "Fiscal policy measures in response to the energy and inflation shock and climate change." Tech. rep., European Central Bank.

Fetzer, Thiemo. 2014. "Fracking Growth." Cep discussion papers, Centre for Economic Performance, LSE. URL https://EconPapers.repec.org/RePEc:cep:cepdps:dp1278.

Frey, Brendan J and Delbert Dueck. 2007. "Clustering by passing messages between data points." *science* 315 (5814):972–976.

Friedman, Jerome H. 2002. "Stochastic gradient boosting." *Computational Statistics and Data Analysis* 38 (4):367–378. URL https://www.sciencedirect.com/science/article/pii/S0167947301000652.

Gerlagh, Reyer, Matti Liski, and Iivo Vehvilainen. 2022. "Stabilizing the EU Electricity Market: Mandatory Demand Reduction and a Lower Price Cap." *EconPol Forum* 23 (06):8–12. URL https://ideas.repec.org/a/ces/epofor/v23y2022i06p8-12.html.

Global Energy Monitor. 2025. "Building an open guide to the world's energy system." URL https://globalenergymonitor.org/.

Graf, Christoph, Federico Quaglia, and Frank A. Wolak. 2021. "(Machine) learning from the COVID-19 lockdown about electricity market performance with a large share of renewables." *Journal of Environmental Economics and Management* 105:102398. URL https://www.sciencedirect.com/science/article/pii/S0095069620301212.

Haan, Marco and Maarten Pieter Schinkel. 2023. "Energy price ceilings with partial cover: A Dutch master?" *Journal of the European Court of Auditors* April.

Hidalgo-Pérez, Manuel, Natalia Collado, Jorge Galindo, and Ramón Mateo. 2024. "The Iberian exception: Estimating the impact of a cap on gas prices for electricity generation on consumer prices and market dynamics." *Energy Policy* 188:114092.

Hughes, Jonathan E, Christopher R Knittel, and Daniel Sperling. 2008. "Evidence of a shift in the short-run price elasticity of gasoline demand." *The Energy Journal* 29 (1):113–134.

Hyndman, Rob J and George Athanasopoulos. 2018. *Forecasting: principles and practice*. OTexts. URL https://otexts.com/fpp3/tscv.html.

Keliauskaitė, Ugnė, Ben McWilliams, Giovanni Sgaravatti, and Georg Zachmann. 2024. "European natural gas imports." URL https://www.bruegel.org/dataset/european-natural-gas-imports.

Krebs, Tom and Isabella Weber. 2024. "Can Price Controls Be Optimal? The Economics of the Energy Shock in Germany." IZA Discussion Papers 17043, Institute of Labor Economics (IZA). URL https://ideas.repec.org/p/iza/izadps/dp17043.html.

Kroger, Mats, Karsten Neuhoff, and Sebastian Schwenen. 2025. "Contracts in Crisis: The War in Ukraine and Long-Term Contracts in Energy Markets." Discussion Papers of DIW Berlin 2118, DIW Berlin, German Institute for Economic Research. URL https://ideas.repec.org/p/diw/diwwpp/dp2118.html.

Krzanowski, Wojtek J and YT Lai. 1988. "A criterion for determining the number of groups in a data set using sum-of-squares clustering." *Biometrics* :23–34.

Labandeira, Xavier, José M. Labeaga, and Xiral López-Otero. 2017. "A meta-analysis on the price elasticity of energy demand." *Energy Policy* 102:549 – 568. URL http://www.sciencedirect.com/science/article/pii/S0301421517300022.

Lee, Chien-Chiang and Zihao Yuan. 2024. "Impact of energy poverty on public health: A non-linear study from an international perspective." *World Development* 174:106444. URL https://www.sciencedirect.com/science/article/pii/S0305750X23002620.

Levell, Peter, Martin O'Connell, and Kate Smith. 2025. "The welfare effects of price shocks and household relief packages: Evidence from the European Energy Crisis." *CEPR Discussion Paper* 19939. URL https://cepr.org/publications/dp19939.

Lou, H. K., M. G. Pollitt, R. Robinson, and A. V. Arcos. 2025. "The Iberian Exception: What was the Cost of Distorting Electricity Markets During the 2021-23 European Energy Crisis?" Cambridge Working Papers in Economics 2535, Faculty of Economics, University of Cambridge. URL https://ideas.repec.org/p/cam/camdae/2535.html.

MB-Research. 2023. *Market Data and Digital Boundaries for Spain*. URL https://www.english.mb-research.de/spain.html.

Mouselimis, Lampros. 2024. *ClusterR: Gaussian Mixture Models, K-Means, Mini-Batch-Kmeans, K-Medoids and Affinity Propagation Clustering*. URL https://CRAN.R-project.org/package=ClusterR. R package version 1.3.3.

Neyman, J. 1923. "On the Application of Probability Theory to Agricultural Experiments. Essay on Principles. Section 9." *Roczniki Nauk Rolniczych Tom X [Polish]; translated in Statistical Science (1990)* 5:465–480.

OECD. 2024. "Distributional impacts of energy transition pathways and climate change." Tech. rep., OECD, Economics Department.

Patel, Chetan. 2022. "What lessons can be learned from the Iberian price cap?" Tech. rep., ICIS. URL https://www.icis.com/explore/resources/news/2022/09/06/10802840/what-lessons-can-be-learned-from-the-iberian-price-cap/.

Polo, Michele, Mar Reguant, Karsten Neuhoff, David Newbery, Matti Liski, Gerard Llobet, Reyer Gerlagh, Albert Banal-Estanol, Francesco Decarolis, Natalia Fabra, Anna Creti, Claude Crampes, Estelle Cantillon, Sebastian Schwenen, Camile Landais, Iivo Vehviläinen, and Stefan Ambec. 2023. "Electricity market design: Views from European Economists." (CEPR Policy Insight, 120).

Polo, Michele and Tommaso Roccuzzo. 2025. "And Yet It Moves: A Study of Natural Gas Consumption in Italy During the 2022 Energy Crisis." Tech. Rep. mimeo, Bocconi University.

Red Elétrica de España. 2024. "Generation and Consumption Data." *ESIOS–Sistema de Informacíon del Operador del Sistema* URL https://www.esios.ree.es/en/generation-and-consumption.

Reguant, Mar, Natalia Fabra, and Jingyuan Wang. 2025. "The Distributional Impacts of Time-Varying Electricity Pricing: A Novel Approach to Estimating Household Income." CEPR Discussion Papers 17200, C.E.P.R. Discussion Papers. URL https://ideas.repec.org/p/cpr/ceprdp/17200.html.

Roldán, Toni, Josep M. Comajuncosa, and Manuel Hidalgo. 2023. "Spain's economy will slow down in 2023 but not enter a recession." Tech. rep., ESADE. URL https://www.esade.edu/en/news/spains-economy-will-slow-down-in-2023-not-enter-a-recession.

Rousseeuw, Peter J. 1987. "Silhouettes: a graphical aid to the interpretation and validation of cluster analysis." *Journal of Computational and Applied Mathematics* 20:53–65.

Rubin, Donald B. 1974. "Estimating causal effects of treatments in randomized and nonrandomized studies." *Journal of Educational Psychology* 66 (5):668–701.

Ruiz, Miguel Haro, Christoph Schult, and Christoph Wunder. 2024. "The effects of the Iberian exception mechanism on wholesale electricity prices and consumer inflation: a synthetic-controls approach." *Applied Economics Letters* 0 (0):1–7. URL https://doi.org/10.1080/13504851.2024.2425834.

Sgaravatti, Giovanni, Simone Tagliapietra, Cecilia Trasi, and Georg Zachmann. 2023. "National fiscal policy responses to the energy crisis." URL https://www.bruegel.org/dataset/national-policies-shield-consumers-rising-energy-prices.

Sistema europeo de negociación de CO2. 2024. "Precios CO2." *Histórico de Precios 2010-2024* URL https://www.sendeco2.com/es/precios-co2.

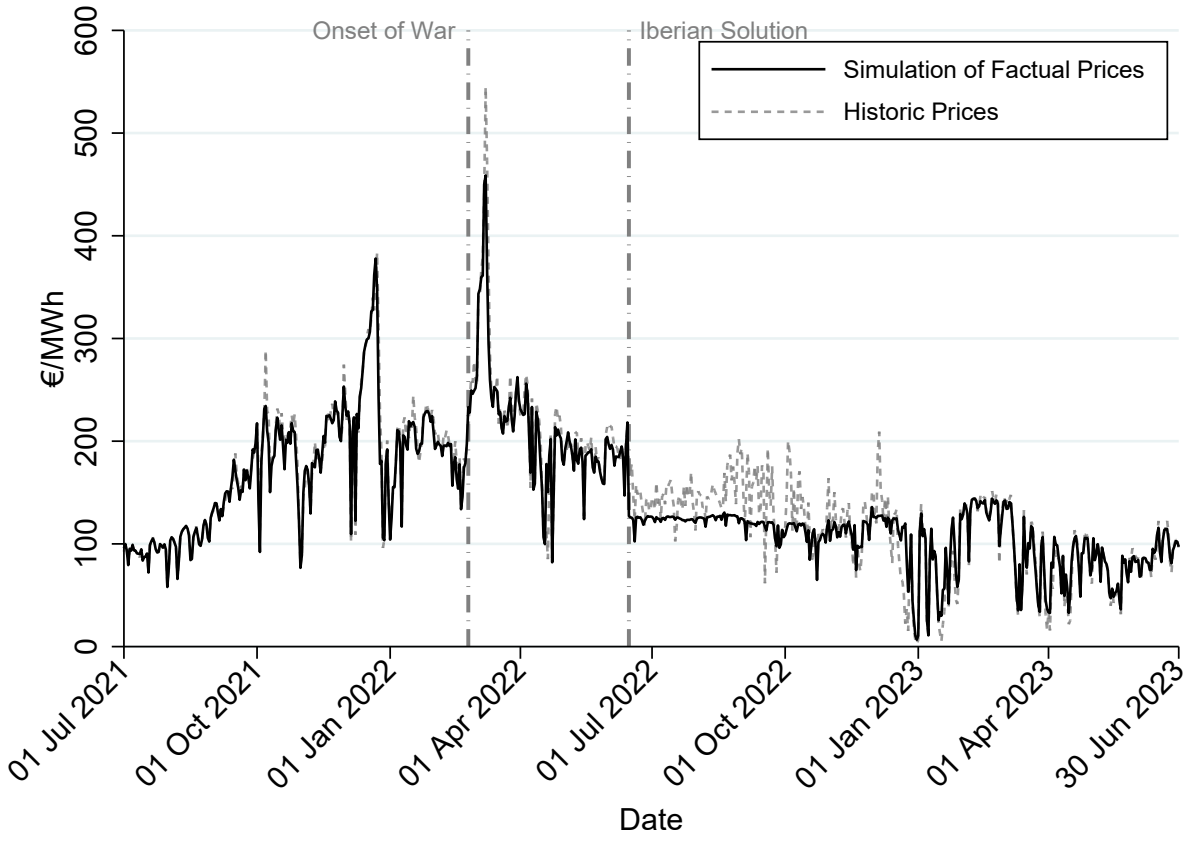
Spanish National Statistics Institute. 2023. *Survey on Essential Characteristics of Population and Housing. Year 2021*. URL https://ine.es/dynt3/inebase/en/index.htm?padre=8981.

von der Leyen, Ursula. 2022. "Statement by President von der Leyen on energy, 7 September 2022."

Appendix – For Online Publication

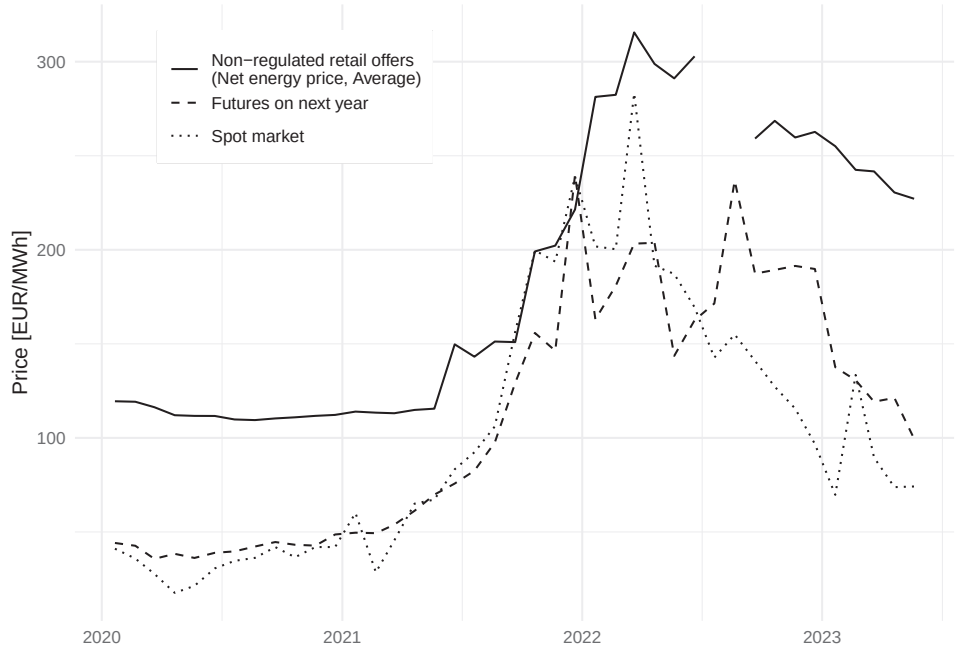
A Additional Results

Figure A.1: Validating the Simulation Model



Notes: The figure displays the historic prices in the Spanish electricity market versus the simulated ones under the *Factual* scenario.

Figure A.2: Retail prices in the retail market: futures & spot vs. non-regulated retail tariffs



Notes: The figure displays the trajectory of the energy component of non-regulated retail tariffs (solid lines) – excluding distribution and transmission charges, as well as taxes – compared to spot (dotted line) and futures prices (dashed line). The data reveal that the rise in spot prices during the crisis was largely reflected in futures prices and, subsequently, in retail tariffs. Nevertheless, the compression of retail margins suggests that suppliers absorbed part of the cost increase.

A.1 Descriptive statistics across zip codes

Table A.1: Determinants of household electricity consumption

$Y = log(Electricity\ Cons.)$	(1)	(2)
log(*Income*)	0.067***	0.396***
	(0.013)	(0.015)
Household Size	0.439***	0.276***
	(0.017)	(0.016)
Electric Heating		0.426***
		(0.024)
Air Conditioning		0.288***
		(0.015)
Degree of urbanization (ref.: Cities)		
Towns and Semi-dense areas		0.243***
		(0.008)
Rural areas		0.290***
		(0.011)
(Intercept)	6.292***	3.187***
	(0.146)	(0.162)
R^2	0.080	0.265
Adj. R^2	0.080	0.264
Num. obs.	8393	8392

Notes: $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$. Each observation corresponds to one zip code. The dependent variable is the log of the average electricity consumption in the zip code expressed in kWh per household per year. *Income* stands for the mean disposable income per household per year, *Household Size* for the mean number of people per household. *Electric Heating* and *Air Conditioning* denote the share of households with electric heating and air conditioning, respectively.

Table A.2: Determinants of the average electricity price paid by households

	Average Price [€/MWh]				
	(1) $log(p)$	(2) $log(\Delta p)$	(3) $log(-\Delta p)$	(4) $log(-\Delta p)$	(5) $log(-\Delta p)$
log(*Income*)	0.005***	0.007***	−0.003	0.009**	−0.103***
	(0.001)	(0.001)	(0.003)	(0.003)	(0.003)
Household Size	−0.004***	−0.005***	−0.003	0.000	−0.026***
	(0.001)	(0.001)	(0.003)	(0.004)	(0.003)
Electric Heating	0.007***	0.005***	0.067***	0.082***	−0.082***
	(0.001)	(0.002)	(0.005)	(0.006)	(0.005)
Air Conditioning	0.013***	0.026***	0.145***	0.175***	−0.102***
	(0.001)	(0.001)	(0.003)	(0.003)	(0.003)
Degree of urbanization (ref.: Cities)					
Towns and semi-dense	0.001***	0.006***	0.011***	0.014***	−0.019***
	(0.000)	(0.001)	(0.002)	(0.002)	(0.001)
Rural areas	−0.003***	0.003***	0.034***	0.041***	−0.020***
	(0.000)	(0.001)	(0.002)	(0.002)	(0.002)
(Intercept)	5.449***	4.658***	3.254***	3.008***	2.171***
	(0.006)	(0.014)	(0.033)	(0.037)	(0.031)
R^2	0.162	0.089	0.396	0.427	0.353
Adj. R^2	0.161	0.089	0.396	0.426	0.353
Num. obs.	8392	8392	8392	8392	8392

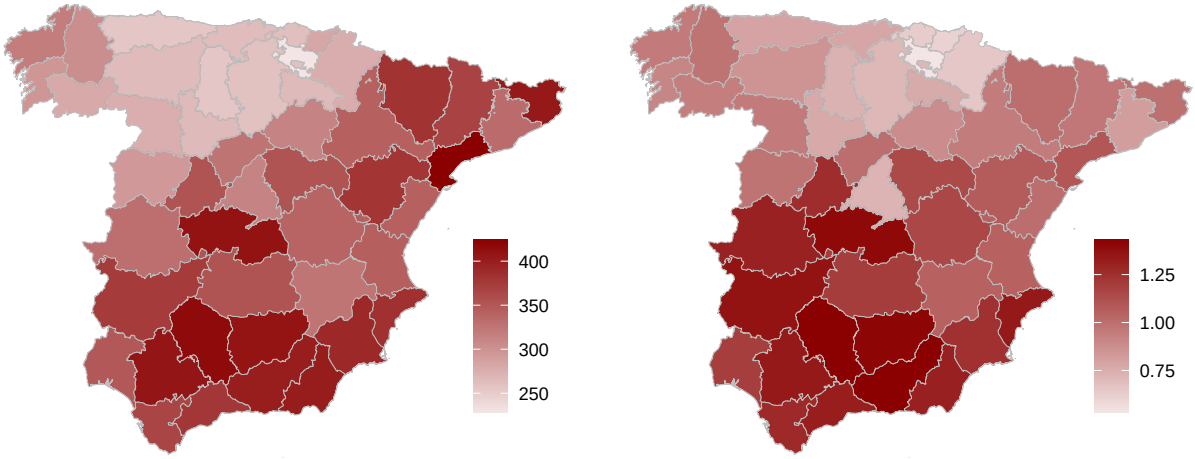
Notes: *p<0.1; **p<0.05; ***p<0.01. Each observation corresponds to one zip code. The dependent variables are logs of (1) the average electricity price paid by households in the *Factual* scenario, (2) the change in this average price due to the crisis relative to the *Factual* scenario, and the change in this average price relative to the *Factual* scenario due to (3) the *full-intervention effect*, (4) the *savings-intervention effect*, (5) the *price-intervention effect*. *Income* stands for the average disposable income per household per year, *Household Size* for the mean number of people per household. *Electric Heating* and *Air Conditioning* denote the share of households with electric heating and air conditioning, respectively.

Table A.3: Households responsiveness to electricity prices across social classes

	Upper Class / Middle Upper Class	Middle Class	Lower Middle Class	Lower Class	**Total**
Have you changed your electricity consumption habits in the last twelve months following the information campaigns on bill changes?					
Yes	56,5%	57,9%	55,3%	56,3%	**56,8%**
No, I have not changed them	40,8%	38,3%	40,1%	35,9%	**39,4%**
I was not aware of the bill changes	2,7%	3,9%	4,6%	7,8%	**3,7%**
Do you take into account the difference in prices at different times of the day in your consumption habits, for example, to put on the washing machine or turn on the electric heating or air conditioning?					
Yes, quite a lot	40,9%	41,6%	41,5%	33,5%	**41,0%**
Yes, a little	25,6%	24,1%	26,4%	38,3%	**25,7%**
No	33,4%	34,2%	32,0%	28,3%	**33,3%**

Notes: Survey conducted in the first half of 2022 among 4,907 households located in continental Spain (5000 households in the whole country). *At that time, the regulated market in Spain exposed consumers to real time pricing, i.e., the the hourly wholesale market price + fees and taxes. Source: BELab. Banco de España Data Laboratory, Panel de Hogares CNMC. DOI: 10.48719/Belab.CNMC.PH1922_02, 2025, own calculations.

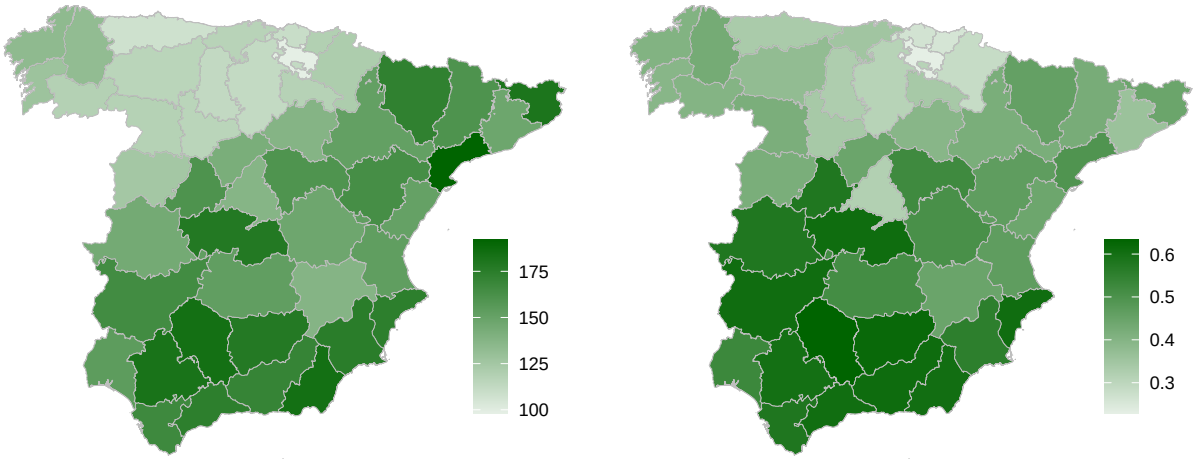
Figure A.3: Energy-crisis effect on household bills by province



(a) Bill increase in level [€/year] (b) Bill increase relative to income [%]

Notes: The maps depict the average effect of the crisis on households' bills in each province expressed in (a) absolute level in €/year and in (b) percentage of households' disposable income.

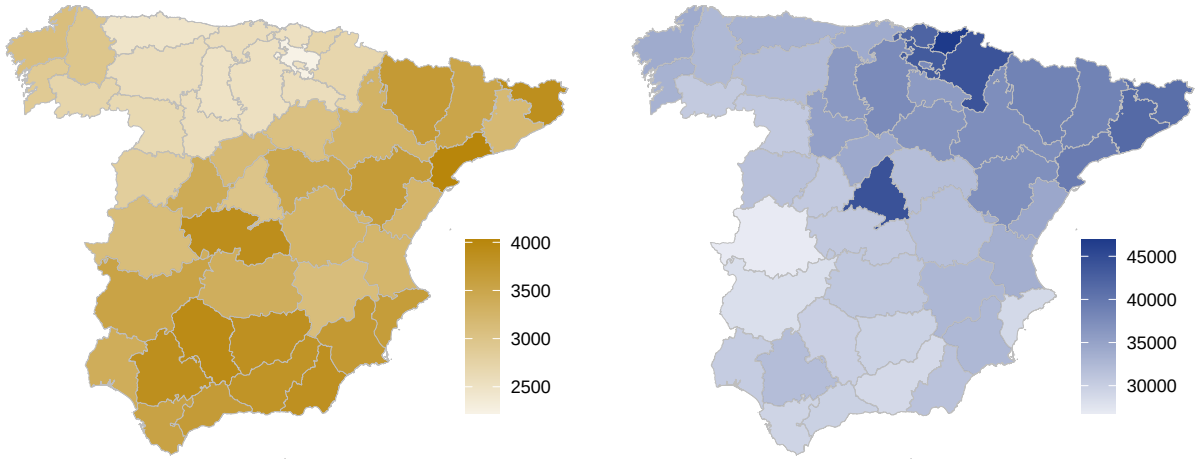
Figure A.4: Full-intervention effect on household bills by province



(a) Bill decrease in level [€/year] (b) Bill decrease relative to income [%]

Notes: The maps depict the average effect of the policy intervention (price intervention and energy savings) on households' bills in each province expressed in (a) absolute level in €/year and in (b) percentage of households' disposable income.

Figure A.5: Average electricity consumption and income per household



(a) Electricity Consumption [kWh/year] (b) Disposable Income [€/year]

Notes: The maps depict the average per-household electricity consumption in kWh/year (a) and disposable income in €/year (b), for each province.

Table A.4: Energy-crisis effect and household characteristics within each cluster

$Y = log(\Delta Bill)$	Full Sample	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
log(*Income*)	0.40***	0.36***	0.43***	0.39***	0.45***	0.52***	0.35***
	(0.02)	(0.03)	(0.05)	(0.06)	(0.04)	(0.04)	(0.10)
Household Size	0.27***	0.26***	0.18***	0.12	0.24***	0.40***	−0.06
	(0.02)	(0.03)	(0.04)	(0.06)	(0.04)	(0.04)	(0.08)
Electric Heating	0.43***	0.40***	−0.24	0.30**	0.28**	0.46***	−1.49***
	(0.02)	(0.06)	(0.14)	(0.11)	(0.10)	(0.10)	(0.32)
Air Conditioning	0.32***	0.09	0.73***	0.31***	0.16	0.29***	0.44*
	(0.02)	(0.06)	(0.07)	(0.05)	(0.09)	(0.07)	(0.17)
Degree of urbanization (ref.: Cities)							
Towns and Semi-dense areas	0.25***	0.25***	0.26***	0.21***	0.29***	0.04*	
	(0.01)	(0.01)	(0.03)	(0.03)	(0.02)	(0.02)	
Rural areas	0.29***	0.29***	0.33***	0.25***	0.25***	0.20***	−0.27***
	(0.01)	(0.02)	(0.04)	(0.03)	(0.03)	(0.03)	(0.06)
(Intercept)	0.84***	1.38***	0.85	1.57*	0.67	−0.57	2.76**
	(0.17)	(0.30)	(0.48)	(0.64)	(0.43)	(0.42)	(0.98)
R^2	0.27	0.25	0.26	0.18	0.20	0.16	0.06
Adj. R^2	0.27	0.25	0.26	0.18	0.19	0.15	0.06
Num. obs.	8390	1762	1812	826	1304	1788	898

Notes: $^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$. Each observation corresponds to one zip code. The dependent variable $log(\Delta Bill)$ is the log of the absolute increase in households' annual bill in the zip code due to the crisis. *Income* stands for the mean disposable income per household per year, *Household Size* for the mean number of people per household. *Electric Heating* and *Air Conditioning* denote the share of households with electric heating and air conditioning, respectively.

Table A.5: Full-intervention effect and household characteristics within each cluster

$Y = log(-\Delta Bill)$	Full Sample	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
log(*Income*)	0.39***	0.37***	0.41***	0.35***	0.44***	0.47***	0.38***
	(0.02)	(0.03)	(0.05)	(0.06)	(0.04)	(0.04)	(0.10)
Household Size	0.27***	0.28***	0.10*	0.16**	0.21***	0.41***	−0.14
	(0.02)	(0.03)	(0.04)	(0.06)	(0.04)	(0.04)	(0.08)
Electric Heating	0.46***	0.34***	−0.44**	0.28**	0.30**	0.39***	−1.62***
	(0.02)	(0.06)	(0.14)	(0.11)	(0.10)	(0.10)	(0.32)
Air Conditioning	0.36***	0.10	0.77***	0.29***	0.09	0.48***	0.48**
	(0.02)	(0.06)	(0.06)	(0.05)	(0.09)	(0.07)	(0.18)
Degree of urbanization (ref.: Cities)							
Towns and Semi-dense areas	0.25***	0.23***	0.27***	0.20***	0.29***	0.05**	
	(0.01)	(0.01)	(0.03)	(0.03)	(0.02)	(0.02)	
Rural areas	0.31***	0.28***	0.32***	0.26***	0.24***	0.23***	−0.23***
	(0.01)	(0.02)	(0.04)	(0.03)	(0.03)	(0.02)	(0.06)
(Intercept)	0.08	0.48	0.41	0.93	0.09	−0.97*	1.93
	(0.17)	(0.30)	(0.47)	(0.63)	(0.44)	(0.41)	(0.99)
R^2	0.29	0.23	0.25	0.19	0.18	0.18	0.06
Adj. R^2	0.29	0.23	0.25	0.19	0.18	0.18	0.06
Num. obs.	8390	1762	1812	826	1304	1788	898

Notes: $^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$. Each observation corresponds to one zip code. The dependent variable $log(-\Delta Bill)$ is the log of the absolute decrease in households' annual bill in the zip code thanks to the full intervention (price intervention and energy savings intervention). *Income* stands for the mean disposable income per household per year, *Household Size* for the mean number of people per household. *Electric Heating* and *Air Conditioning* denote the share of households with electric heating and air conditioning, respectively.

Table A.6: Savings-intervention effect and household characteristics within each cluster

$Y = log(-\Delta Bill)$	Full Sample	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
log(*Income*)	0.38***	0.34***	0.41***	0.37***	0.44***	0.47***	0.35***
	(0.01)	(0.03)	(0.05)	(0.06)	(0.04)	(0.04)	(0.10)
Household Size	0.27***	0.26***	0.15***	0.13*	0.24***	0.39***	−0.04
	(0.02)	(0.03)	(0.04)	(0.06)	(0.04)	(0.04)	(0.08)
Electric Heating	0.38***	0.28***	−0.35**	0.30**	0.34***	0.41***	−1.41***
	(0.02)	(0.06)	(0.14)	(0.11)	(0.10)	(0.09)	(0.30)
Air Conditioning	0.24***	0.05	0.70***	0.22***	0.00	0.35***	0.40*
	(0.01)	(0.05)	(0.06)	(0.05)	(0.09)	(0.06)	(0.16)
Degree of urbanization (ref.: Cities)							
Towns and Semi-dense areas	0.23***	0.23***	0.27***	0.20***	0.27***	0.04*	
	(0.01)	(0.01)	(0.03)	(0.03)	(0.02)	(0.02)	
Rural areas	0.28***	0.28***	0.34***	0.24***	0.23***	0.20***	−0.26***
	(0.01)	(0.02)	(0.04)	(0.03)	(0.03)	(0.02)	(0.06)
(Intercept)	−0.70***	−0.12	−0.69	−0.08	−0.91*	−1.78***	0.96
	(0.16)	(0.29)	(0.46)	(0.63)	(0.43)	(0.40)	(0.93)
R^2	0.24	0.23	0.25	0.16	0.18	0.16	0.06
Adj. R^2	0.24	0.23	0.24	0.15	0.18	0.16	0.06
Num. obs.	8390	1762	1812	826	1304	1788	898

Notes: $^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$. Each observation corresponds to one zip code. The dependent variable $log(-\Delta Bill)$ is the log of the absolute decrease in households' annual bill in the zip code thanks to the energy savings intervention. *Income* stands for the mean disposable income per household per year, *Household Size* for the mean number of people per household. *Electric Heating* and *Air Conditioning* denote the share of households with electric heating and air conditioning, respectively.

Table A.7: Price-intervention effect and household characteristics within each cluster

$Y = log(-\Delta Bill)$	Full Sample	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
log(*Income*)	0.40^{***}	0.39^{***}	0.42^{***}	0.34^{***}	0.45^{***}	0.48^{***}	0.40^{***}
	(0.02)	(0.03)	(0.05)	(0.06)	(0.04)	(0.04)	(0.11)
Household Size	0.28^{***}	0.29^{***}	0.07	0.18^{**}	0.20^{***}	0.42^{***}	-0.21^{*}
	(0.02)	(0.03)	(0.04)	(0.06)	(0.04)	(0.04)	(0.09)
Electric Heating	0.51^{***}	0.39^{***}	-0.49^{***}	0.27^{*}	0.27^{**}	0.36^{***}	-1.84^{***}
	(0.03)	(0.06)	(0.14)	(0.11)	(0.10)	(0.10)	(0.34)
Air Conditioning	0.46^{***}	0.13^{*}	0.83^{***}	0.35^{***}	0.16	0.57^{***}	0.54^{**}
	(0.02)	(0.06)	(0.07)	(0.05)	(0.10)	(0.07)	(0.19)
Degree of urbanization (ref.: Cities)							
Towns and Semi-dense areas	0.26^{***}	0.23^{***}	0.26^{***}	0.21^{***}	0.30^{***}	0.06^{**}	
	(0.01)	(0.02)	(0.03)	(0.03)	(0.02)	(0.02)	
Rural areas	0.33^{***}	0.28^{***}	0.30^{***}	0.27^{***}	0.24^{***}	0.25^{***}	-0.21^{**}
	(0.01)	(0.02)	(0.04)	(0.03)	(0.03)	(0.03)	(0.07)
(Intercept)	-0.71^{***}	-0.44	-0.13	0.32	-0.52	-1.77^{***}	1.28
	(0.17)	(0.31)	(0.47)	(0.63)	(0.46)	(0.42)	(1.06)
R^2	0.34	0.23	0.26	0.22	0.18	0.20	0.07
Adj. R^2	0.34	0.22	0.26	0.22	0.18	0.20	0.06
Num. obs.	8390	1762	1812	826	1304	1788	898

Notes: $^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$. Each observation corresponds to one zip code. The dependent variable $log(-\Delta Bill)$ is the log of the absolute decrease in households' annual bill in the zip code thanks to the price intervention. *Income* stands for the mean disposable income per household per year, *Household Size* for the mean number of people per household. *Electric Heating* and *Air Conditioning* denote the share of households with electric heating and air conditioning, respectively.

A.2 Distributional Consequences – zip code-level approach

As an alternative to the approach developed in Section 4.3.2, we consider another approach that do not rely on any extrapolation and provides a lower bound of the gaps in crisis and intervention effects across income bins. We cannot directly group households into quintiles and estimate the effects on each group because we do not have household-level data. Thus, we instead group zip codes in quintiles of equal population (number of households) according to the average income of resident households, as depicted on the left column in Table 7. Then we compute the average effect of the crisis, the average effect of the price intervention and the average effect of the savings intervention for each zip code-level income quintile. The results are depicted in Figure A.6.

Table A.8: Summary statistics per zip code-level income quintile

Income Quintile	**Q1**	**Q2**	**Q3**	**Q4**	**Q5**	**Full Sample**
Electricity Cons. [kWh/year.HH]	3,332	3,267	3,097	3,063	3,385	**3,229**
Income [kEUR/year.HH]	24.8	30.9	35.5	40.0	50.3	**36.3**
Household Size	2.5	2.5	2.5	2.5	2.5	**2.5**
Electric Heating [%]	36.7	33.8	23.6	23.1	20.5	**27.6**
Air Conditioning [%]	54.9	51.9	46.3	48.8	50.8	**50.5**
Cities [%]	19.0	46.4	62.9	71.5	79.7	**55.9**
Towns & semi-dense areas [%]	38.1	39.8	29.1	23.1	18.0	**29.6**
Rural [%]	42.9	13.8	8.0	5.5	2.3	**14.5**
Min Temp. [ºC]	3.7	3.7	2.4	2.1	1.8	**2.8**
Max Temp. [ºC]	32.5	31.1	30.3	30.5	30.5	**31.0**
Number of ZIP codes	4,659	1,526	912	740	556	**8,393**
Total Households (thousands)	3,261	3,268	3,229	3,295	3,271	**16,325**

Notes: All values are averages weighted by zip code population. The share of electric heating, the share of air conditioning are based on province-level data or city-level data (for large cities). Average min. and max. temperatures are based on province-level data. Average min. temperatures correspond to those of the month of January, while average max. temperatures correspond to those of the month of August.

The differences between income quintiles at the zip code level are relatively small in absolute terms. Households in zip codes in the second and third quintiles experienced the largest impact of the crisis, with their average bill increasing by 356 €. However, this increase was only marginally higher than the smallest increase of 314 € experienced by households in zip codes in the 1st quintile. Similarly, the total average effect of the interventions ranged only from 170 € to 191 € across zip code income quintiles. Yet, when these bill variations are related to average disposable income of households, the effects of the crisis and the interventions appear more contrasted across quintiles. The

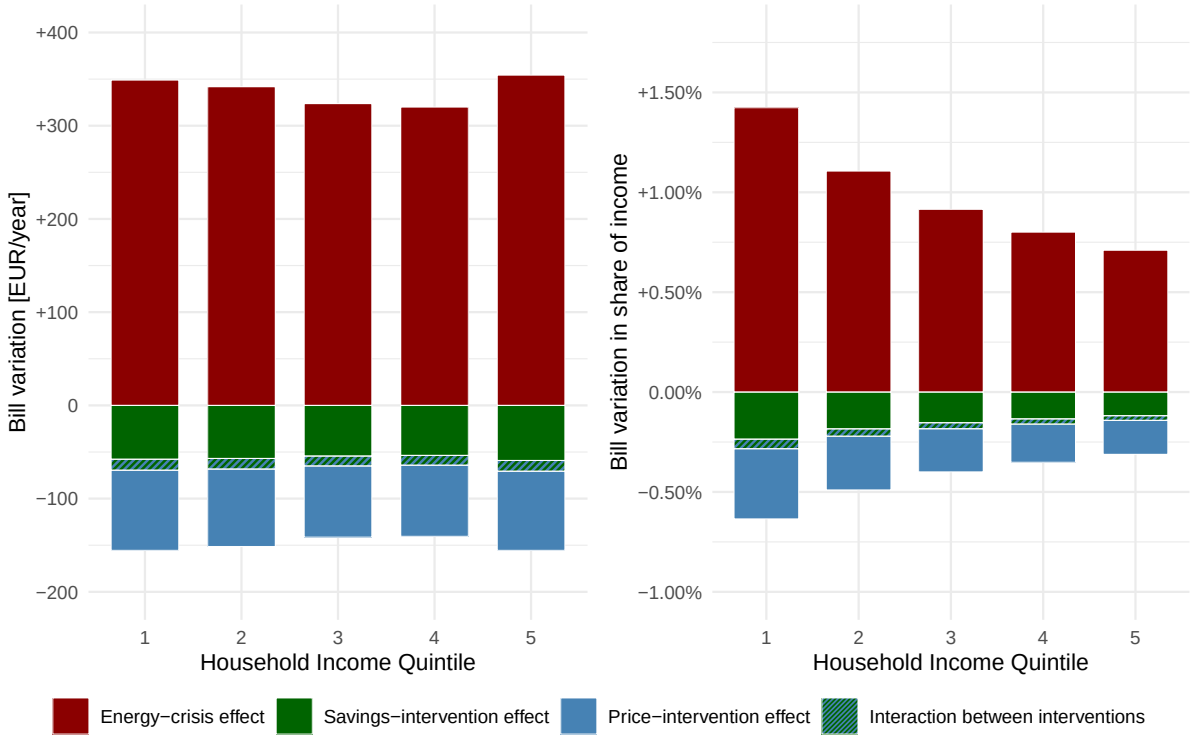


Figure A.6: Crisis and policy intervention effects by zip code-level income quintile

bill increase due to the crisis represented a share of their disposable income twice larger on average for households living in zip codes in the 1st quintile than for households living in zip codes in the 5th quintile. In addition, we must bear in mind that this gap is only a lower bound of the actual differences across (household level) income quintiles.

B Counterfactual simulation inputs

B.1 Natural Gas and CO_2 Prices

As a direct input to our simulation of the Spanish electricity market, we consider the reference price of gas on the Iberian market, that is the MIBGAS price. However, we use the TTF price as a predictor when assessing the indirect effect of the crisis through the CO_2 price, considering that the TTF price is the most relevant for most of the area covered by the EU-ETS.

We collect daily data from 2015 to 2023 that we aggregate at the monthly level. We build counterfactual natural gas prices for the period July 2021 - June 2023 by applying for each month the average price for the same month over the period 2017-2019 (e.g. January will take the average of prices observed in January 2017, January 2018, and January 2019). This allows us to preserve the seasonality of natural gas prices in the counterfactual scenario.

In order to build counterfactual CO_2 prices, we estimate the linear regression model below, capturing the effect of natural gas price (TTF) on CO_2 prices, while allowing for different linear time trends in different periods of the EU-ETS:

$$\log(\text{CO}_2 \text{ price}_t) = log(\text{TTF price}_t) + \text{trend} \cdot \text{factor(EU-ETS period)}_t \tag{B.1}$$

The index t denotes months of the period 2015-2023, and "factor(EU-ETS period)" denotes a categorical variable dividing the time period into four periods (A, B, C, and D) according to major policy shocks that affected the EU-ETS. Namely, the breakdown is based on the following events:

A→B: January 2016: Announcement of free allocation of allowances from the new entrants reserve (see climate.ec.europa.eu/news-your-voice/news/commission-publishes-status-update-new-entrants-reserve-2016-01-15_en)

B→C: January 2018: Market Stability Reserve amended by Directive (EU) 2018/410 so that a certain amount of permits inside the reserve would be canceled from 2023 onward (Directive passed on the 14 March 2018 but expected earlier)

C→D: January 2021: Beginning of phase IV of the ETS

B.2 Counterfactual electricity prices in France

To construct a set of counterfactual electricity prices in France under the *No-Energy-Crisis* scenario, we first estimate a linear model that explains electricity prices as a function of the natural gas price (TTF), the CO_2 price, electricity demand in France net

$Y =$	log(*CO2 price*) [EUR/t]
log(*TTF price*) [EUR/MWh]	0.18***
	(0.03)
EU-ETS Period (ref.: Period A)	
Period B	−0.39**
	(0.15)
Period C	−0.35**
	(0.16)
Period D	0.55**
	(0.23)
trend × Period A	0.02
	(0.01)
trend × Period B	−0.01
	(0.01)
trend × Period C	0.01
	(0.01)
trend × Period D	−0.00
	(0.01)
(Intercept)	1.37***
	(0.14)
R^2	0.98
Adj. R^2	0.98
Num. obs.	108

***$p < 0.01$; **$p < 0.05$; *$p < 0.1$

Table B.1: Linear model of monthly CO_2 prices (2015-2023)

of renewable and nuclear production (*Net Load*), a quadratic term of *Net Load*, interaction terms between gas and CO_2 prices and *Net Load*, as well as fixed effects for the month of the year and the hour of the day.

We include *Net Load* as an explanatory variable to account for the non-dispatchable nature of renewable generation (comprising wind, solar, and run-of-river hydro) and the limited flexibility of nuclear generation.

The model is estimated on hourly data spanning from 2015 to 2023. Estimation results are reported in Table B.1. The model explains a substantial share of the variance in electricity prices ($R^2 = 0.9$), and both gas and CO_2 prices are found to have a strong influence on electricity prices in France.

In the second step, we use the estimated model to predict electricity prices in France for the period from July 2021 to June 2023 under the *No-Energy-Crisis* scenario (see Figure 3). Specifically, we assume TTF gas prices are equal to the monthly average observed during 2017–2019 and use the counterfactual CO_2 price series described in Appendix B.1.

$Y =$	French Electricity Price [EUR/MWh]
Net Load [GW]	0.4355^{***}
	(0.0347)
(*Net Load* [GW])^2	0.0027
	(0.0023)
TTF Price [EUR/MWh]	1.0580^{***}
	(0.0065)
CO_2 Price [EUR/t]	0.5293^{***}
	(0.0061)
Net Load [GW] $\times$ *TTF Price* [EUR/MWh]	0.0439^{***}
	(0.0011)
Net Load [GW] $\times$ *CO_2 Price* [EUR/t]	0.0151^{***}
	(0.0008)
(*Net Load* [GW])^2 $\times$ *TTF Price* [EUR/MWh]	0.0032^{***}
	(0.0001)
(*Net Load* [GW])^2 $\times$ *CO_2 Price* [EUR/t]	-0.0022^{***}
	(0.0001)
(Intercept)	-10.6258^{***}
	(0.7028)
R^2	0.9050
Adj. R^2	0.9050
Num. obs.	77978

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table B.1: Linear model estimates for prediction of French electricity prices

C Predicting counterfactual electricity demand

C.1 Key assumptions for counterfactual predictions

Let Y_t denote aggregate electricity demand in hour t. Building on the Neyman-Rubin potential outcomes framework (Neyman, 1923; Rubin, 1974), let $Y_t(0)$ represent the potential outcome if the crisis had not happened. Then we assume:

Assumption 1. There exists a "stable" regression function $g()$ that drives no-crisis counterfactual aggregate electricity demand.

$$\begin{aligned} Y_t(0) &= g(\mathbf{X}_t(0)) + \varepsilon_t \\ \text{such that } \mathbb{E}[Y_t(0)|X_t(0)] &= g(\mathbf{X}_t(0)) \ , \end{aligned} \tag{Asm. 1}$$

where $\boldsymbol{X}_t$ is a vector of variables observed by the researcher, and ε_t is an idiosyncratic error term.

We refer to $g()$ as being a "stable" function in the sense that it explains potential outcomes $Y_t(0)$ both before and during the crisis. Let *pre* denote time periods before the energy crisis, and *post* denote time periods during the crisis. Ultimately, we are interested in predicting $Y_{post}(0)$ (i.e., demand during the crisis period under a counterfactual as if the crisis had not happened). In reality, we only observe $Y_{pre}(0)$, $\mathbf{X}_{pre}(0)$, $Y_{post}(1)$, and $\mathbf{X}_{post}(1)$. One complication is that, naturally, the predictor variables themselves might change during the crisis ($\mathbf{X}_{pre} \neq \mathbf{X}_{post}$). This warrants another assumption, namely:

Assumption 2: The predictor variables are independent of the energy crisis.

$$\mathbf{X}_t(0) = \mathbf{X}_t(1) = \mathbf{X}_t \ . \tag{Asm. 2}$$

Taken together, assumptions 1 and 2 imply:

$$\mathbb{E}[Y_{post}(0)|X_{post}(0)] = g(\mathbf{X}_{post}(0)) = g(\mathbf{X}_{post}) \ . \tag{C.1}$$

We train flexible machine learning algorithms using pre-crisis data ($Y_{pre}(0)$ and $\mathbf{X}_{pre}(0)$), aiming to estimate the counterfactual function $\hat{g}()$. Given this function, the counterfactual of interest can then be predicted as: $\hat{Y}_{post}(0) = \hat{g}(\mathbf{X}_{post})$. More specifically, we use hourly electricity demand data from 2015 to 2020 to estimate $\hat{g}()$, where the outcome variable is aggregate hourly electricity demand in Spain.

C.2 Machine learning algorithms

For estimating the counterfactual function $\hat{g}()$, we build on recent advances from the machine learning literature, both in regards to (i) ML algorithms and to (ii) cross-validation strategies.

We focus on tree-based machine learning algorithms. A nice feature of regression trees is that they inherently allow for non-linear relationships between the covariates and the outcome of interest. This is because, by design, the continuous covariates are transformed into categorical covariates through "branch splits." Also, interactions between variables can be easily incorporated by increasing the tree "depth" (i.e., the number of nodes at which branching is considered/allowed). Our algorithm of choice further models complexity in a data-driven manner by combining many "simple" trees into a single predictor via gradient boosting (Friedman, 2002).[42] Specifically, we implement XGBoost (Chen and Guestrin, 2016). We consider several configurations of the algorithm by varying: the minimum observations per terminal node; the maximum tree depth; the learning rate (which adjusts the weights assigned to each tree); and the total number of trees. The best performing configuration is selected via 4-fold cross-validation (CV), which allows us to assess *out-of-sample* prediction performance.[43]

We complement the 4-fold CV strategy with time series cross-validation (Hyndman and Athanasopoulos, 2018), which can be viewed as a form of 'placebo' test. We train our algorithm with data from 2015 to 2018, and assess its performance for 2019 to 2023. Results from that exercise are presented in Appendix Figure C.1. We show that prediction accuracy is high during all of 2019, which lends validity to our algorithm of choice. Performance in 2020 onward drops significantly, as expected due to the COVID-19 pandemic and later due to the energy crisis. During crisis periods, the model trained with data from 2015 to 2018 results in counterfactual predictions that are substantially higher than those from our baseline model (trained with data from 2015 to 2020). This suggests that the 2020 data might be particularly useful for the algorithm to capture potential structural demand changes triggered by the pandemic. Below, we provide further details.

C.3 Tuning and cross-validation for aggregate electricity demand

We implement 4-fold cross-validation for hyperparameter tuning of the XGBoost algorithm. As shown in Table C.1, we consider configurations that vary in the number of trees, maximum tree depth, shrinkage/learning rate, and minimum observations per terminal node. A total of 18 configurations were tested for a model to predict aggregate electricity demand. For training, we use data from 2015 to 2020. The last

[42] The algorithm starts with one (randomly defined) regression tree, then iteratively adds more trees to an ensemble model, assigning more weights to the trees that improve predictive performance (more details in Appendix C.2).

[43] We randomly split our sample into 4 equally sized sub-samples, use three of those to train the ML algorithm, and then test prediction accuracy in the fourth (hold-out) sub-sample. The process is repeated four times, such that all folds serve once as the hold-out sub-sample. We then select the configuration that results in the lowest root-mean squared error (RMSE) in the hold-out sub-samples (more details in Appendix C.2).

two columns of Table C.1 present the in-sample root-mean-squared errors (RMSE) and the cross-validated RMSE for each configuration. As expected, the in-sample errors are substantially lower than out-of-sample errors. The best-performing algorithm, based on having the lowest CV RMSE, is identified as Model ID 17 in the table. We use that configuration for our couterfactual predictions shown in Section 3.1.

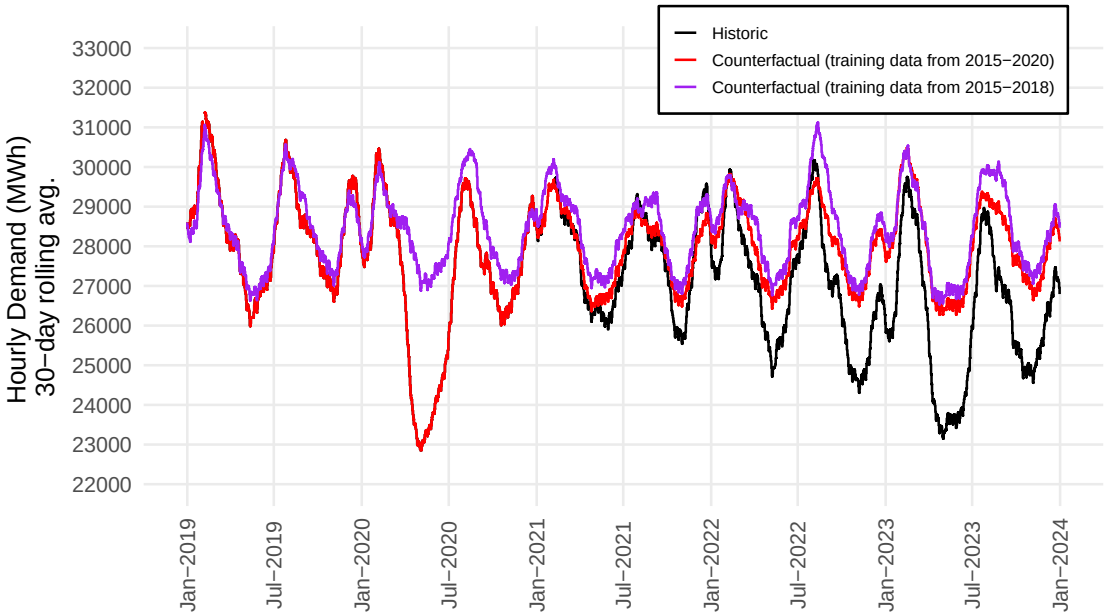
We also implement time series cross-validation to assess performance of our preferred algorithm (Hyndman and Athanasopoulos, 2018). For this, we train the model with data from 2015 to 2018, then assess prediction performance in 2019 and beyond. As shown in Figure C.1, comparing predictions with historic consumption, the model is accurate for 2019 and early 2020. We conclude that the model performs well in predicting demand for counterfactual scenarios without significant structural changes. Given the richness of the controls included, and under the assumptions stated in Section 3.1, the discrepancies between real and predicted demand observed from 2020 onward may be attributed to unexpected crises (COVID and the energy crisis itself).

Table C.1: Hyperparameter tuning – aggregate demand

Model ID	Number of Trees	Max. Tree Depth	Shrinkage	Min. Obs. per Node	In-sample RMSE	CV RMSE
1	2,000	5	0.1	10	307.50	534.62
2	2,000	10	0.1	10	30.46	490.39
3	2,000	5	0.2	10	208.53	528.43
4	2,000	10	0.2	10	4.06	532.76
5	2,000	10	0.1	5	14.97	509.99
6	3,000	10	0.1	5	4.00	509.73
7	2,000	20	0.1	5	0.01	514.93
8	3,000	20	0.1	5	0.01	514.93
9	3,000	10	0.1	10	11.57	485.83
10	2,000	20	0.1	10	0.90	471.20
11	3,000	20	0.1	10	0.74	471.20
12	2,000	10	0.1	20	64.28	468.64
13	3,000	10	0.1	20	34.30	466.83
14	2,000	20	0.1	20	5.99	435.19
15	3,000	20	0.1	20	2.50	435.14
16	3,000	30	0.1	20	1.94	438.85
17	3,000	20	0.1	30	5.71	434.33
18	3,000	30	0.1	30	3.83	435.86

Notes: Tuning hyperparameters for XGBoost (Chen and Guestrin, 2016). The outcome is aggregate electricity demand in Peninsular Spain. Cross-validated (CV) errors are based on 4-fold cross validation.

Figure C.1: Time series cross-validation – aggregate electricity demand



Notes: This figure illustrates the historical and predicted counterfactual Spanish aggregate electricity demand from Jan. 2019 to Dec. 2023. Predictions in red are from an algorithm trained with data from 2015 to 2020. Predictions in purple are from a model trained only with data from 2015 to 2018.

C.4 Tuning and cross-validation for low voltage electricity demand

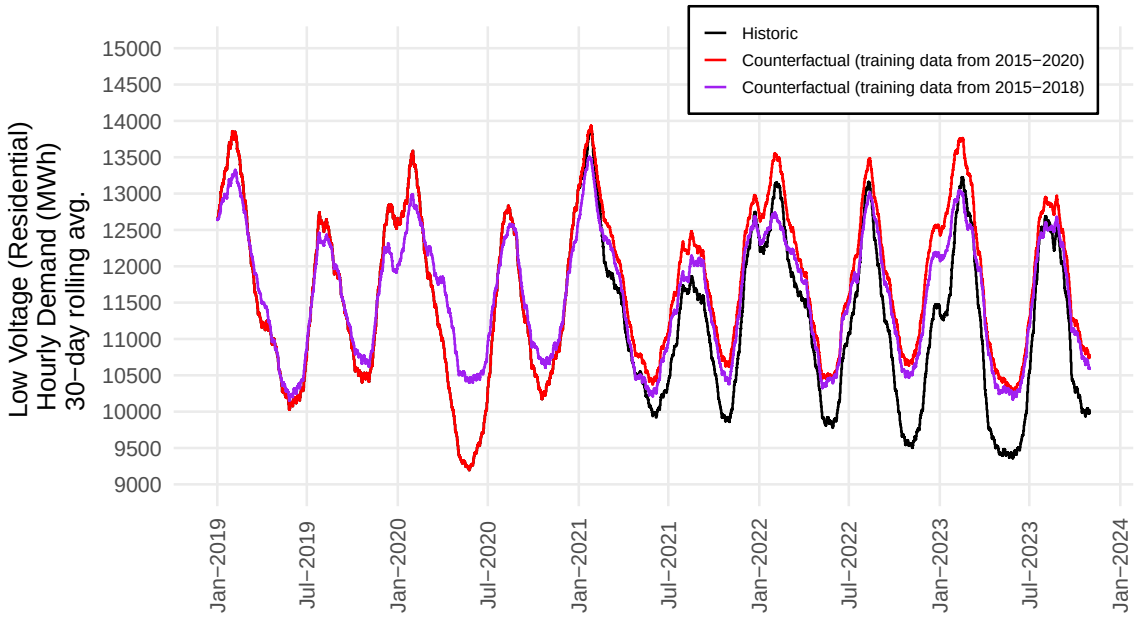
We also implement 4-fold CV for tuning the algorithm for prediction of low voltage electricity demand. Results are presented in Table C.2. For this case, Model ID 10 is considered the best performing, based on having lowest CV RMSE. Results are generally consistent with those reported in the subsection above (for aggregate electricity demand).

Table C.2: Hyperparameter tuning – low voltage demand

Model ID	Number of Trees	Max. Tree Depth	Shrinkage	Min. Obs. per Node	In-sample RMSE	CV RMSE
1	2,000	5	0.1	10	218.09	362.75
2	2,000	10	0.1	10	17.26	297.51
3	2,000	5	0.2	10	127.21	336.82
4	2,000	10	0.2	10	2.39	323.24
5	3,000	10	0.1	10	6.29	295.31
6	3,000	20	0.1	10	0.12	297.62
7	3,000	10	0.1	20	21.28	285.88
8	3,000	20	0.1	20	1.40	274.12
9	3,000	10	0.1	30	33.22	283.74
10	3,000	20	0.1	30	3.31	266.39

Notes: Tuning hyperparameters for XGBoost (Chen and Guestrin, 2016). The outcome is low voltage electricity demand in Peninsular Spain. Cross-validated (CV) errors are based on 4-fold cross validation.

Figure C.2: Time series cross-validation – low voltage electricity demand



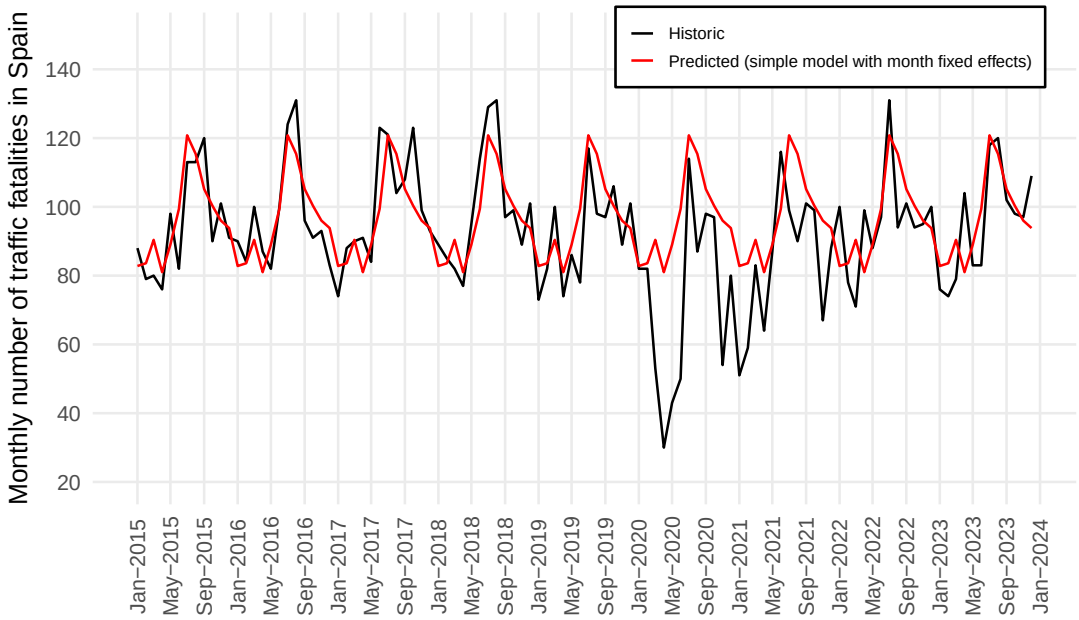
Notes: This figure illustrates the historical and predicted counterfactual "low voltage" electricity demand in Spain from Jan. 2019 to Dec. 2023. Predictions in red are from an algorithm trained with data from 2015 to 2020. Predictions in purple are from a model trained only with data from 2015 to 2018.

C.5 Traffic fatalities as a proxy for mobility intensity

We include monthly road traffic fatalities in our set of control variables to predict counterfactual electricity demand in Spain. Mobility by cars is not expected to directly influence electricity consumption in Spain, given the relatively low adoption of electric vehicles. However, an indirect relationship may arise, to the extent that a lack of mobility implies more people staying at home. We find that the patterns observed in traffic fatalities closely match the COVID-19 lockdown periods. As shown in Figure C.3, the seasonality patterns in fatalities were remarkably stable in Spain from 2015 to 2023, except during lockdown periods. The red curve in the Figure reveals that a simple regression specification with month-of-year fixed effects performs well in predicting fatalities.

We therefore argue that this measure can serve as a proxy for urban mobility. For the case of low-voltage demand, the correlation between mobility and electricity consumption is most likely negative (as lockdown restrictions may increase electricity at home). This relationship may be confounded by the fact that some small commercial establishments also use low-voltage electricity. Yet, we noted an improvement in cross-validated RMSE after the inclusion of traffic fatalities in our models, especially in those to predict residential demand (results without this control can be made available upon request). The improvements for predicting aggregate (residential plus industrial) electricity demand were negligible.

Figure C.3: Time series of road traffic fatalities in Spain



Notes: The predictions in red are based on a simple regression with month-of-year fixed effects plus a constant, estimated with data from 2015 to 2019 (pre-crisis).

D Clustering algorithm and performance metrics

D.1 *K-means* clustering

For the clustering algorithm, we consider each zip code of Peninsular Spain as the unit of observation, aggregating our key variables of interest to that level. As described in the main text, the key variables used for clustering capture seasonal and intraday electricity demand patterns. We implement *k-means* clustering with a final sample of 8,930 zip codes and 375 variables. The objective is to partition the zip codes into clusters, minimizing the within-cluster sum of squares (WCSS).

Let $\mathbf{z}$ denote the set of zip code observations, where each observation is a vector of dimension 375 (the number of variables). We partition the zip codes into 6 clusters (optimal according to the performance metrics) $\mathbf{C} = \{C_1, C_2, ..., C_6\}$, such that:

$$\begin{aligned} &\arg\min_{\mathbf{C}} \sum_{k=1}^{6} \sum_{\mathbf{z} \in C_k} ||\mathbf{z} - \boldsymbol{\mu_k}||^2, \\ &\text{where } \boldsymbol{\mu_k} = \frac{1}{|C_k|} \sum_{\mathbf{z} \in C_k} \mathbf{z}. \end{aligned} \tag{D.1}$$

Note that $|C_k|$ is the number of zip codes in cluster C_k, and $\boldsymbol{\mu_k}$ is the centroid of the cluster. We implement *k-means* clustering via the `ClusterR` package (Mouselimis, 2024; Frey and Dueck, 2007), initializing the algorithm with 5 different centroid seeds and keeping all other options at their default values.

D.2 Clustering performance metrics

One standard performance metric for clustering is the total within-cluster sum of squares. That simply consists of a summation, over all clusters, of the values resulting from the minimization problem in (D.1). However, we are also interested in analyzing the degree of separation *between* clusters. We therefore calculate silhouette scores (Rousseeuw, 1987). Let

$$a_z = \frac{1}{|C_k| - 1} \sum_{\mathbf{i} \in C_k, \mathbf{z} \neq \mathbf{i}} d(\mathbf{z}, \mathbf{i}),$$

where $d(\mathbf{z}, \mathbf{i})$ is the Euclidean distance between zip code $\mathbf{z}$ and all other zip codes $\mathbf{i}$ belonging to the same cluster C_k. Smaller values of a_z would indicate more similarity between the points contained in C_k. Additionally, let

$$b_z = \min_{l \neq k} \frac{1}{|C_l| - 1} \sum_{\mathbf{j} \in C_l} d(\mathbf{z}, \mathbf{j}),$$

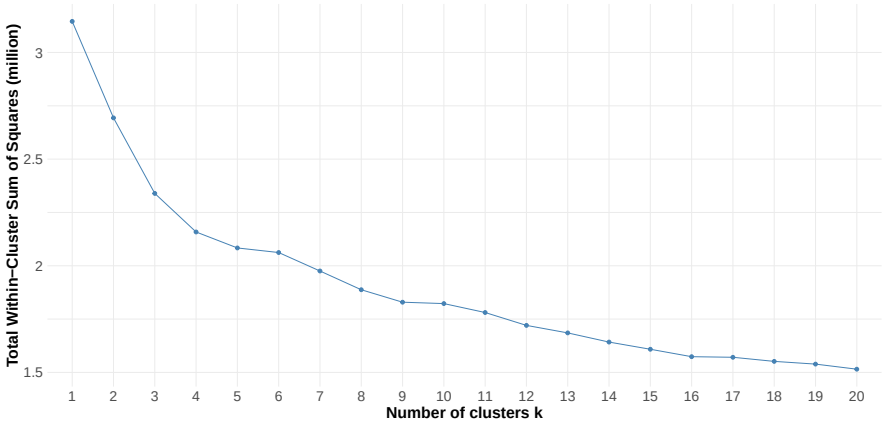
where $d(\mathbf{z}, \mathbf{j})$ is the distance between zip code $\mathbf{z}$ and all zip codes $\mathbf{j}$ from the nearest

neighboring cluster C_l. Higher values of b_z imply lower similarity between zip code $\mathbf{z}$ and points from its neighboring cluster. Finally, the silhouette score or width is defined as

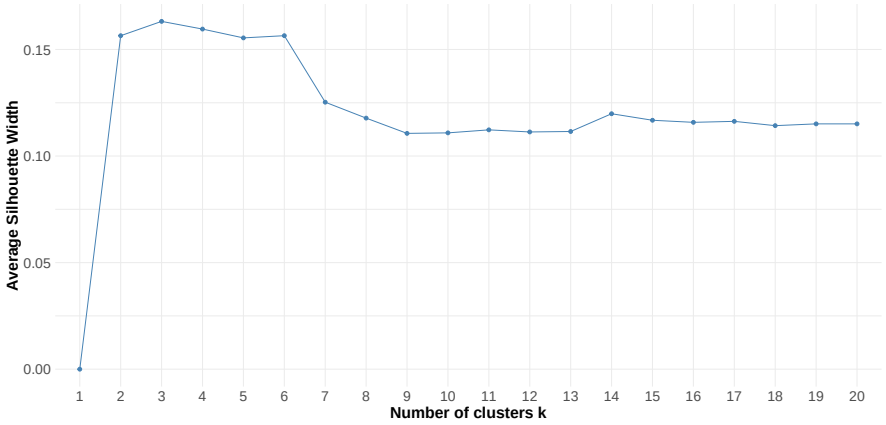
$$s_z = \begin{cases} 1 - \frac{a_z}{b_z} & \text{if } a_z < b_z, \\ 0 & \text{if } a_z = b_z, \\ \frac{b_z}{a_z} - 1 & \text{if } a_z > b_z. \end{cases}$$

Values of s_z close to 1 indicate high degree of separation between $\mathbf{z}$ and observations from other clusters. When s_z is 0, then zip code $\mathbf{z}$ is on the border of two clusters. If s_z is negative, then the zip code may have been allocated to the incorrect cluster (as it matches better with observations in the neighboring cluster). After computing silhouette scores for all 8,930 zip codes in our sample, we take the average.

Results for total within-cluster sum of squares are presented in Panel (a) of Figure D.1, while average silhouette scores are in Panel (b). We note that the total within-cluster sum of squares (Panel a) decreases as the number of clusters increases. There is a noticeable sharp decrease when moving from 6 to 7 clusters. However, the average silhouette scores reveals that any improvements beyond 6 clusters come at the expense of higher cluster overlap (or lower separation). Therefore, for the main analyses in our paper, we use 6 clusters to summarize the distributional consequences of the energy crisis.



(a) Total within-cluster sum of squares by number of clusters



(b) Average silhouette scores by number of clusters

Figure D.1: Performance metrics to determine optimal number of clusters