

CES ifo

**12093
2026**

Working Papers

Original Version: August 2025
This Version: August 2026

Unpacking the Distributional Implications of the Energy Crisis: Lessons from the Iberian Electricity Market

Natalia Fabra, Clément Leblanc, Mateus Souza

CES ifo

Imprint:

CESifo Working Papers

ISSN 2364-1428 (digital)

Publisher and distributor: Munich Society for the Promotion
of Economic Research - CESifo GmbH

Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740

Email office@cesifo.de
<https://www.cesifo.org>

Editor: Clemens Fuest

An electronic version of the paper may be downloaded free of charge

- from the CESifo website: www.ifo.de/en/cesifo/publications/cesifo-working-papers
- from the SSRN website: www.ssrn.com/index.cfm/en/cesifo/
- from the RePEc website: <https://ideas.repec.org/s/ces/ceswps.html>

Unpacking the Distributional Implications of the Energy Crisis:

Lessons from the Iberian Electricity Market*

Natalia Fabra, Clément Leblanc, and Mateus Souza

August 2026

Abstract

The 2021–2023 European energy crisis caused by the war in Ukraine prompted broad policy interventions across energy markets. In contrast to the retail-side subsidies and public transfers implemented elsewhere, Spain and Portugal targeted the wholesale electricity market through the so-called Iberian solution. Combining wholesale market simulations with machine learning techniques, we quantify the distributional implications of the crisis and this intervention on Spanish electricity firms and consumers. The crisis shifted substantial wealth from consumers to generators: power plant profits rose by €29 Billion. Even accounting for the Iberian solution and the accompanying energy-saving measures, the crisis remained highly regressive – the lowest income quintile lost 3.4% of disposable income, three times the relative burden of the highest quintile. Without interventions, the crisis would have been considerably more costly for all households. Further, the policy package delivered proportionally larger relief to lower-income groups, at no net fiscal cost.

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: Nova School of Business and Economics, Universidade NOVA de Lisboa. We thank Michael Duarte, Catarina Pintassilgo, Álvaro Oballe, Lorenzo Cattabriga, and Wisse Rutgers 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), the Mannheim Conference on Energy and the Environment, the World Congress of Environmental and Resource Economists (Lisbon), Universidade NOVA de Lisboa, and the 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 gratefully acknowledges support from the German Research Foundation (DFG) through CRC TR 224 (Project B07), and from Portuguese National Funds through FCT – Fundação para a Ciência e a Tecnologia (projects UIDB/00124/2020 and UIDP/00124/2020). 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.

²While 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 Evidence from the UK (Feger, Pavanini, and Radulescu, 2022), Germany (Dertwinkel-Kalt and Wey, 2025), and the Netherlands (Haan and Schinkel, 2023) shows that some of these retail interventions weakened competition, raised average prices relative to what they would have been, and increased fiscal costs. Beyond these effects, 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.

The scheme was not without controversy. Germany and the Netherlands criticized it for distorting the single market and for potentially increasing gas consumption by lowering the effective price of gas for power generation. This paper also quantifies these claims.

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. In-

³In 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).

terestingly, 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 €47B between July 2021 and June 2023, i.e., a 2.4-fold increase, while their profits rose by €29B over the same period.⁴ Taken together, the interventions reduced consumer payments by €15.3B (€10.6B from the price intervention and €4B from the savings measures).

The interventions also affected electricity exports and emissions. In particular, the price intervention alone increased carbon emissions in Iberia by around 2.7%, as lower prices led to higher exports to France, which in turn required more (less) gas-fired generation in Iberia (France). The savings intervention more than offset this effect, so that the combined intervention was approximately emissions-neutral. This is an important finding, as it highlights the complementarity between the two policies.

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 8,564 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

⁴These 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.

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 \(2026\)](#) 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 to climate-related shocks and decarbonization efforts, understanding the distributional effects of alternative policy instruments becomes essential to ensure broad social support.

Several 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, Gazze, and Bishop \(2024\)](#) 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, Gazze, and Bishop \(2024\)](#) – 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 estimated elasticity is much lower, so 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 Küçük, 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 us-

⁵Relatedly, 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.

age, 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. 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 electricity 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 qualitatively 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 reducing 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

⁶As 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.

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.

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

⁷Kroger, 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.


⁸According 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.

⁹Since all thermal plants, regardless of their actual efficiency rates or technology, received the same subsidy, the merit order between them was not distorted. This is formally explained in Section 3.2.

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.

electricity bills.¹⁰

The primary objective of the mechanism was to incentivize gas- and coal-fired plants

¹⁰Specifically, 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.

– 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 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 Theoretical framework

We begin with a highly stylized representation of the wholesale electricity market. Consider first the demand side. Demand is denoted by $D(p; \theta)$. It can be thought of as arising from a mix of consumers: those exposed to real-time prices, whose consumption responds to prices, and those on fixed-price contracts, whose consumption is effectively price-inelastic. A larger inelastic share simply makes the demand curve steeper, without altering the qualitative results below. Demand is also affected by shocks, captured by θ , which reflect seasonal or stochastic fluctuations as well as non-price policy-related factors. These could include energy-saving measures or the adoption of behind-the-meter solar generation, which reduces the quantity of electricity purchased from the grid. An

¹¹The 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.

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

increase, $\bar{\theta} > \theta_F$, corresponds to an outward shift in the demand curve (e.g., if energy-saving measures are not put in place) relative to baseline demand with θ_F .

Consider next the supply side. Let k denote the capacity of low-cost, carbon-free generation technologies. Assuming that k is never sufficient to meet total demand, this capacity is always inframarginal, so its offers do not affect the market-clearing price. The remaining supply, denoted by $S(p; \eta)$, is weakly increasing in price and may therefore be either flat or upward sloping (e.g., depending on whether the price-setting gas-fired plants have the same or different marginal costs). The supply shifter η captures seasonal or random factors affecting fuel prices and the policy environment. We distinguish three values for it. The value η_F corresponds to the baseline crisis-level gas prices, but with the price intervention in place, such that producers internalize a lower effective gas price. The value $\bar{\eta}$ corresponds to crisis-level gas prices in the absence of the price intervention. The value $\underline{\eta}$ corresponds to pre-crisis supply conditions, with low gas prices. Accordingly, the supply curve $S(p; \bar{\eta})$ lies above baseline supply $S(p; \eta_F)$, which in turn lies above $S(p; \underline{\eta})$.

The market equilibrium is determined by the intersection of the relevant demand and supply curves, and thus varies across the demand and supply scenarios described below, and illustrated in Figure 2.

Scenarios. We consider five scenarios. The first is the baseline scenario, referred to as *Factual*. It is based on observed demand and supply during the crisis and with the real policy environment, including both the energy-saving measures and the gas price cap. This scenario is associated with demand and supply parameters (θ_F, η_F) , and its equilibrium outcome is denoted by F in the figure.

The remaining four scenarios are counterfactual. The *No-Energy-Crisis* scenario assumes that the energy crisis never occurred. In that case, demand and supply would have remained at pre-crisis levels, with neither energy-saving measures nor high gas prices. This scenario is therefore characterized by $(\bar{\theta}, \underline{\eta})$, and its equilibrium outcome is denoted by NC .

The *Business-as-Usual* scenario assumes that the crisis occurred, but that no policy interventions were introduced. Demand is therefore not affected by energy-saving measures, and supply reflects crisis-level gas prices in the absence of the price cap. This scenario is characterized by $(\bar{\theta}, \bar{\eta})$, and its equilibrium outcome is denoted by BaU .

The *Savings-Intervention-Only* scenario retains the energy-saving measures but removes the gas price cap. It is therefore associated with $(\theta_F, \bar{\eta})$, and its equilibrium outcome is denoted by SI . Conversely, the *Price-Intervention-Only* scenario retains the gas price cap but assumes demand in the absence of energy-saving measures. It is associated with $(\bar{\theta}, \eta_F)$, and its equilibrium outcome is denoted by PI .

Price, quantity and distributional effects. For each component of the policy environment, we define its “price effect” as the difference between the price under the *Factual* scenario, $p_F^* \equiv p^*(\theta_F, \eta_F)$, and the price that would prevail under the relevant counterfactual scenarios, p_X^* for $X \in \{NC, BaU, SI, PI\}$. Four effects are of interest. The *energy-crisis effect*, $p_F^* - p_{NC}^* > 0$, captures the impact of the crisis (under the factual interventions). The *full-intervention effect*, $p_F^* - p_{BaU}^* < 0$, captures the joint impact of the two policy measures. The *price-intervention effect*, $p_F^* - p_{SI}^* < 0$, isolates the marginal contribution of the gas price cap, holding the savings measures fixed. The *savings-intervention effect*, $p_F^* - p_{PI}^* \leq 0$, isolates the marginal contribution of the savings measures, holding the price cap fixed.

Symmetrically, we define the “quantity effect” as the difference between equilibrium demand under the *Factual* scenario, $q_F^* \equiv q^*(\theta_F, \eta_F)$, and demand under the relevant counterfactual, q_X^* . The *energy-crisis effect* on quantity, $q_F^* - q_{NC}^* < 0$, reflects the reduced electricity consumption during the crisis, primarily through the non-price energy-saving measures rather than through direct price responses, given the low short-run elasticity of demand. The *full-intervention effect*, $q_F^* - q_{BaU}^*$ is also negative as long as the demand is sufficiently inelastic, so that the reduction induced by the savings measures, $q_F^* - q_{PI}^*$, dominates over the demand increase induced by the price reduction, $q_F^* - q_{SI}^*$.

In the empirically relevant case in which demand is highly inelastic and supply is relatively flat in the relevant range, shifts in supply translate into large price changes but muted quantity effects, while shifts in demand have little effect on prices but large quantity effects. Accordingly, the price effects are dominated by the price intervention – whose contribution coincides with the full-intervention effect in the limit of perfectly flat supply – while the quantity effects are shaped almost entirely by the savings measures. The two instruments are therefore complementary: the price cap compressed wholesale prices and inframarginal rents, while the savings measures reduced the volume of electricity purchased at those elevated prices.

The price and quantity effects also have direct distributional implications. Changes in wholesale prices translate one-for-one into changes in the inframarginal rents earned by the low-cost capacity k . In addition, they alter the quasi-rents of thermal generators, captured by the area between the equilibrium price and their supply curve. The total change in aggregate producer surplus between the *Factual* scenario and counterfactual scenarios $X \in \{NC, BaU, SI, PI\}$ is therefore

$$\Delta\Pi(X) = \underbrace{k [p_F^* - p_X^*]}_{\text{inframarginal rents of } k} + \underbrace{\int_{p_F^*}^{p_X^*} S(p; \eta_F) dp - \int_{p_F^*}^{p_X^*} S(p; \eta_X) dp}_{\text{change in thermal quasi-rents}}. \quad (1)$$

The first term captures the windfall gain (or loss) of inframarginal producers whose costs are unaffected by the shock. The second term captures the change in quasi-rents of




Figure 2: A stylized framework for analyzing wholesale market effects

Notes: This figure presents a stylized demand–supply framework for analyzing the effects of the energy crisis and policy interventions in the wholesale electricity market. The outward demand curve, $D(p; \bar{\theta})$, corresponds to demand in the absence of energy-saving measures, while the inward demand curve, $D(p; \theta_F)$, reflects the demand reduction induced by those measures. The highest residual supply curve, $S(p; \bar{\eta})$, represents crisis-period supply in the absence of a price intervention. The intermediate supply curve, $S(p; \eta_F)$, reflects supply under the gas price cap. The lowest supply curve, $S(p; \underline{\eta})$, represents the counterfactual supply that would prevail in the absence of the energy crisis. The inframarginal low-cost, non-gas generation capacity is denoted by k . The intersections between the relevant demand and supply curves determine the equilibrium under each scenario: F (*Factual*), SI (*Savings-Intervention-Only*), PI (*Price-Intervention-Only*), BaU (*Business-as-Usual*), and NC (*No-Energy-Crisis*).

thermal generators, which depends on both the price change and the shift in the supply curve itself. When supply is perfectly flat – the empirically relevant case – thermal quasi-rents are zero and (1) collapses to the expression $k[p_F^* - p_X^*]$. More generally, for an upward-sloping thermal supply curve, the second term is positive for the energy-crisis effect and negative for the intervention effects.

Our welfare accounting is based on producer surplus and consumer electricity bills; it does not take a stand on whether the demand contraction induced by the savings measures reflected a binding, privately costly constraint on households or a largely costless adjustment (e.g., via solar self-consumption or discretionary curtailment households would have undertaken regardless). To the extent that the mandate was binding for some households, our estimates of the benefits of the savings intervention are best read as an upper bound on the associated welfare gain.

3.2 Simulating the wholesale electricity market

Equipped with this stylized framework, our next step is to characterize the actual demand and supply schedules, and the corresponding equilibrium outcomes, using data

from the Spanish electricity market under the factual and counterfactual scenarios described above. Our simulations, which assume competitive behavior,¹³ are performed at the hourly level over the 17,520 hours spanning the sample period from 1st July 2021 to 30th June 2023.

Supply. To model the supply side, we draw on granular data describing the technical characteristics of actual power plants operating in the Spanish electricity market (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 CO₂ allowances, using either the realized values (Bloomberg, 2025; SENDECO2, 2024) or counterfactual values in the absence of the crisis. The methods used to construct gas and CO₂ prices in the *No-Energy-Crisis* scenario are described in detail below.

Using 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 CO₂ emission factor; τ_t is the CO₂ price 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.

Because plant-specific efficiency rates for gas-fired plants are not publicly available, we assume in the simulations that all such plants operate at the baseline efficiency rate of 55%. Under this assumption, the competitive supply curve is essentially flat whenever gas-fired plants are marginal, with only minor differences arising from heterogeneity in operation and maintenance costs.

To construct the industry’s hourly competitive supply curve in the absence of the intervention, we order the estimated marginal costs from lowest to highest. Low-carbon inframarginal capacity is measured using the available nuclear and renewable capacity in each hour, and hydro production is allocated efficiently across time, shaving the peaks of demand net of renewable and nuclear generation.

The price intervention shifted the competitive supply curve by subsidizing thermal generation. Specifically, a path of reference gas prices, p_t^R , was set, and whenever the actual gas price, p_t^G , exceeded that level, generators were compensated for the difference

¹³Under the assumption of strategic bidding (modeled following the framework in de Frutos and Fabra, 2012), the simulation results remain quantitatively similar. See Appendix F.2 for details.

based on the amount of gas that a representative gas-fired plant with the baseline efficiency rate of 55% would need to purchase. Formally, when $p_t^G > p_t^R$, all thermal plants receive a per-unit subsidy for electricity generated in period t , given by

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

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, because all thermal generators received the same subsidy, their relative position in the marginal cost ranking was unchanged.

Demand. Demand under the *Factual* scenario is measured using the realized hourly demand series from [REE \(2024\)](#). However, we allow for medium-run demand adjustments through energy-saving responses and the adoption of solar panels that reduce electricity needs from the grid. Incorporating this requires predicting how demand would have evolved in the absence of the crisis, as well as under crisis conditions, but without the interventions. We obtain these counterfactual demand trajectories using machine-learning methods, as detailed below.

Interconnections. 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. Table 1 summarizes the assumptions on the supply and demand variables used to construct our five scenarios considered. The methods for constructing the counterfactual simulation inputs are described below.¹⁵

¹⁴The 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.

¹⁵Appendixes A and B 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 Price Int.				✓	
	Crisis w/o Savings Int.			✓		✓
Gas Prices	Realized	✓				✓
	No Energy Crisis		✓			
	Crisis w/o Price Int.			✓	✓	
CO₂ & French Prices	Realized	✓		✓	✓	✓
	No Energy Crisis		✓			


Notes: This table summarizes the assumptions on demand, gas prices, CO₂ prices, and French electricity prices used under the five simulated scenarios.

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

Counterfactual CO₂ prices. An increase in gas prices improves the relative competitiveness of coal-fired power plants, thereby increasing the demand for CO₂ allowances and exerting upward pressure on CO₂ prices. Accordingly, in a scenario without elevated gas prices, CO₂ prices would have remained lower. To construct the counterfactual CO₂ price in the absence of the crisis, we remove the component of the CO₂ price that is 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 CO₂ prices (EU-ETS allowances) over the 2015–2023 period, allowing for distinct time trends across different EU-ETS phases.¹⁶ Using this model, we predict CO₂ prices under the assumption that natural gas prices would have followed the pre-crisis trends described above. As seen in Figure 3, CO₂ prices would have increased even without the energy crisis, but not as much.

Figure 3: Gas and CO₂ prices – Historic and *No-Energy-Crisis* counterfactual




Notes: This figure shows the evolution of the realized TTF natural gas prices and CO₂ prices (solid lines) versus the estimated counterfactual prices in the absence of the crisis (dashed line, from July 2021).

¹⁶The model (equation A.1) and the estimation results (Table A.1) are reported in Appendix.

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 CO₂ price, the French electricity residual demand (net of renewable and nuclear production), and time-fixed effects.¹⁷ Using this model, we predict French electricity prices under the counterfactual scenario where natural gas and CO₂ prices reflect the hypothetical *No-Energy-Crisis* conditions. Figure 4 plots the resulting prices, which follow a similar trend as the pre-crisis prices.

Figure 4: 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).

Counterfactual wholesale electricity demand. Aggregate electricity demand is another key input for our simulations of the Spanish electricity market. Naturally, in the 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.¹⁸

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

¹⁷The regression results are summarized in Table A.2 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.

¹⁸In the case of Spain, in May 2022, a regulation stipulated that public buildings could not be cooled below 27°C in the warm season or heated above 19°C in the cold season (BOE, 2022b). Similar measures were extended to commercial establishments in August of the same year (BOE, 2022a).

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 B.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 demand data in Spain from 2015 to 2020.¹⁹ 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 B.3). The algorithm with the lowest root mean squared error (RMSE) was a variant of gradient-boosted regression trees (more details in Appendix B.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 47 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 CO₂ prices (as well as lag1, lag7, lag14, lag30, and lag60 versions of those).²⁰

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

¹⁹Our 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 B.3.

²⁰Aggregate 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 CO₂ prices were provided by [SENDECO2 \(2024\)](#).

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., [Albalade and Fageda, 2021](#)), in Appendix Figure B.2 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.²¹

Natural gas and CO₂ prices were included as proxies to capture demand elasticity to electricity prices.²² 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 CO₂ 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.

The demand series used in the simulations follow the assumptions in Table 1. For the *Business-as-Usual* and *Price-Intervention-Only* scenarios, in which the crisis occurs but energy-saving measures are not implemented, we use aggregate demand predicted with the actual TTF gas and CO₂ prices observed during the crisis (the “high CO₂ and gas prices” prediction, in red in Figure 5). For the *No-Energy-Crisis* scenario, we use demand predicted with the counterfactual, no-crisis prices (the “low CO₂ and gas prices” prediction, in blue), which captures the limited but non-zero demand response to prices. The *Factual* scenario uses realized demand (the black curve). Finally, for the *Savings-Intervention-Only* scenario, which retains the energy-saving measures but removes the gas price cap, we start from realized demand and remove the demand effect of the cap. That is, we take the difference between the blue and the red curves, and apply that to the factual demand from June 2022 onward, when the Iberian Solution took effect. The green curve therefore captures crisis demand with savings but without the price intervention.

Figure 5 plots the 30-day rolling averages of these series. The two ML-based counterfactuals (red and blue) are close to one another, reflecting a low short-run price elasticity. For those two curves, we also assess model-based uncertainty and present confidence bands computed via block-residual bootstrapping (see Appendix B.6).²³ Predicted counterfac-

²¹Our 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).

²²Own-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 CO₂ prices.

²³We further run wholesale market simulations with the upper- and lower-bound variants of those demand series, with results presented in Appendix F.2.

tuals are indistinguishable from factual demand during most of 2021. Yet, both curves (and their confidence bands) lie well above realized demand once the crisis intensifies, because consumers curtailed grid demand by more than prices alone would predict.

Figure 5 also shows that the green series sits just below realized demand, indicating that the price cap’s own effect on the quantity demanded was small.²⁴ Realized and predicted demand overlap for most of 2021; discrepancies emerge towards the end of 2021, as electricity and gas prices soared, and become more pronounced in 2022 and 2023. We find that electricity demand would have been significantly higher in the absence of the crisis.²⁵ For instance, actual demand in April and May 2023 was nearly 11% lower than the counterfactual predictions (see Appendix Table B.2), despite wholesale gas and electricity prices in Spain having begun to stabilize by that time. The reduction was also unevenly distributed within the day: it was twice as large during peak hours as during night hours, consistent with the design of the energy-saving measures (see Appendix F.1 for a detailed characterization of its timing). These results suggest that the energy crisis triggered strong and persistent reductions in electricity consumption, effects that cannot be ignored when simulating the wholesale market impacts.

3.3 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.²⁶ In the following sub-sections, we summarize the key results.


Simulated electricity market prices across scenarios. Figure 6 summarizes the wholesale market simulations. Panels (a)-(c) of Figure 6 plot the daily simulated wholesale price series under the *Factual* scenario alongside the four counterfactuals. Panel (d) plots the four price effects defined above by taking the differences between the factual and counterfactual prices. Table 2 reports the corresponding demand-weighted average prices for each half-year and scenario. Appendix Table F.3 presents the associated confidence intervals, obtained by simulating prices under the lower- and upper-bound demand estimates described above. The results indicate that the estimated prices are highly robust

²⁴We do not model uncertainty for this series, as it is not based on ML predictions, but rather on an implicit assumption of low price elasticity. In Appendix D we further discuss this assumption and estimate elasticities with two independent datasets and methods. The resulting energy price elasticities are close to -0.01.

²⁵In 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.

²⁶Throughout the analysis, we use the simulated *Factual* prices – based on realized demand and input values – rather than the actual realized prices, ensuring that any potential biases in the simulations cancel out when comparing across scenarios. As shown in Table F.2 in the Appendix, the simulated *Factual* prices and generation mix are close to the observed ones, providing confidence in the use of the model for simulating counterfactual scenarios. Taking the realized electricity prices as the baseline would not materially affect our conclusions.

Figure 5: Actual versus predicted aggregate electricity demand



Notes: This figure illustrates the historical and predicted counterfactual Spanish aggregate electricity demand from January 2019 to June 2023. Predictions in red use real CO₂ prices and real TTF natural gas prices. Predictions in blue use CO₂ and TTF price projections assuming that the crisis had not happened. The green curve, plotted from June 2022, shows counterfactual demand in the absence of the Iberian solution’s price cap, obtained by applying the model-implied price response to realized demand. Displaying 30-day rolling averages, computed after dropping outliers (top and bottom 0.5% values of full series). Shaded areas represent confidence bands computed via block-residual bootstrapping. Residuals for this procedure are based on out-of-sample predictions (see Appendix B.6).

to uncertainty in the demand estimates.

Three main messages emerge, which are consistent with the theoretical framework presented above. First, the energy crisis had a large and persistent impact on wholesale electricity prices. Second, once activated, the price intervention substantially mitigated this impact. Third, the savings measures – although highly effective at curbing consumption – had only a marginal direct effect on prices.

By the second half of 2021, as gas prices began to rise, the simulated *Factual* price climbed to a half-year average of €169/MWh, nearly three times the €63/MWh predicted by the *No-Energy-Crisis* scenario (Table 2). The energy-crisis effect (blue series in Panel (b)) widens rapidly throughout the period, exceeding €100/MWh on several days.²⁷

The full-scale invasion of Ukraine in late February 2022 triggered an extreme spike

²⁷Because the price intervention had not yet been introduced, the *Factual*, *Business-as-Usual*, and *Savings-Intervention-Only* curves lie essentially on top of each other. Accordingly, the price-intervention effect is not defined for this period (shown as “–” in Table 2), and both the price- and full-intervention effects in Panel (d) are zero.

in gas prices, which peaked in March 2022 at over €240/MWh and pushed simulated wholesale electricity prices above €500/MWh under *Business-as-Usual* (Panel (a)). Over the half-year, the *Business-as-Usual* average reached €220/MWh, against €66/MWh in the absence of the crisis. The introduction of the Iberian Solution on June 14, 2022 marks a clear inflection point: from this date onward, the *Factual* and *Price-Intervention-Only* curves separate from the *Business-as-Usual* and *Savings-Intervention-Only* curves. Although the price intervention was active for only the last two weeks of the half-year, it was already enough to bring the *Factual* half-year average price below the *Business-as-Usual* level (€201 vs. €220/MWh).


The contrast between scenarios is the starkest during the second half of 2022. In the absence of any intervention, the simulated wholesale price would have averaged €256/MWh under *Business-as-Usual*, and €252/MWh under *Savings-Intervention-Only*. With the price intervention in place, the *Factual* average price drops to €120/MWh – a reduction of more than half. Panel (d) makes this very visible: the price-intervention effect (and, almost coincident with it, the full-intervention effect) becomes deeply negative, frequently exceeding -€200/MWh on a daily basis. Over this period, since the savings-intervention effect (green series) remains close to zero, the full-intervention effect is fully explained by the price-intervention effect.

Gas prices remained elevated in the first two months of 2023, but then declined rapidly. Combined with the periodic upward revisions of the trigger gas price, this meant that the Iberian Solution was rarely binding over this half-year. Consistently, in Panel (d) the price-intervention effect is virtually nil, and Table 2 shows that the *Factual* (€102/MWh) and *Savings-Intervention-Only* (€105/MWh) prices nearly coincide. As expected, all scenarios converge in Panel (a) towards the end of the sample.


A consistent feature across all sub-periods is that the *Savings-Intervention-Only* curve (green) tracks the *Business-as-Usual* curve (gray) very closely, and that the savings-intervention effect hovers around zero. Quantitatively, the savings measures reduced the demand-weighted average wholesale price by no more than a few €/MWh in any given half-year (Table 2). The reason is structural: gas-fired plants remained the price-setting technology throughout the crisis, so the demand reduction induced by the savings measures translated into small movements along an essentially flat segment of the supply curve, with price changes essentially limited to switches between gas units of similar marginal cost. As we show in the next subsection, this does *not* mean that the savings measures were inconsequential: by lowering the quantity of electricity consumed at still-elevated prices, they generated sizable savings in consumer payments.

Distributional effects between consumers and electricity generators. The price effects summarized in Figure 6 led to substantial wealth transfers from consumers to electricity generators. Aggregate effects on power plants’ earnings, costs, and profits are


Figure 6: Simulated Wholesale Electricity Prices and Price Effects across Scenarios




(a) Factual *vs* Business-As-Usual



(b) Factual *vs* Savings-Only and Price-Only



(c) Factual *vs* No Crisis



(d) Simulated price effects

Notes: Panel (a) reports the daily demand-weighted average simulated wholesale electricity prices from July 2021 through June 2023 under the *Factual* and *Business-as-Usual* (BaU) scenarios. Panel (b) compares the *Factual* scenario with the *Price-Intervention-Only* and *Savings-Intervention-Only* scenarios. Panel (c) compares the *Factual* scenario with the *No Energy Crisis* scenario. Finally, panel (d) reports the simulated price effects, computed as the differences between prices under the *Factual* scenario and each of the counterfactual scenarios.

Table 2: Demand-weighted Average Prices and Implied Effects – Competitive Scenario (€/MWh)

	2021 (Q3-4)	2022 (Q1-2)	2022 (Q3-4)	2023 (Q1-2)
<i>Panel A. Price levels</i>				
Factual	169	201	120	102
No Energy Crisis	63	66	71	65
Business-as-usual	171	220	256	114
Savings intervention only	169	215	252	105
Price intervention only	171	207	122	110
<i>Panel B. Effects</i>				
Energy-crisis effect	106	135	50	37
Full-intervention effect	-1	-19	-136	-11
Price-intervention effect	–	-13	-132	-3
Savings-intervention effect	-1	-6	-2	-7

Notes: Panel A reports the simulated demand-weighted average prices in the Spanish wholesale market, assuming competitive bidding, during each period under each scenario. The price intervention took place in 2022. Panel B reports the implied effects, computed as the difference between the factual scenario and the relevant counterfactual.

summarized in Table 3; the confidence bands for these effects are reported in Appendix Table F.4.

Relative to the *No-Energy-Crisis* scenario, earnings by Spanish power plants increased by €47B over the two-year period (July 2021 to June 2023), representing a 2.4-fold increase. The surge in the costs of marginal gas-fired plants, which were passed through to wholesale electricity prices, increased firm profits by €29.2B. These were primarily driven by the growth in inframarginal rents earned by nuclear, hydro, and renewable energy resources.²⁸ 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, €17.8B, 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 €15.3B, which consumers

²⁸ 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.

saved after factoring in the subsidies to gas- and coal-fired plants (paid by consumers in their electricity bills). In particular, consumers collectively saved €10.6B due to the price intervention and €4B from the savings measures, and €15.3B by the combination of both. Firms' costs did not change substantially relative to BaU, given that the increase in generation costs due to the price intervention was approximately of the same magnitude as the cost reduction due to the demand reduction associated with the savings measures. 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 – Competitive Scenario (M€, July 2021-July 2023)

	Earnings	Costs	Profits
Energy-Crisis Effect	46,953	17,773	29,186
Full-Intervention Effect	-15,322	323	-15,644
Price-Intervention Effect	-10,574	2,368	-12,948
Savings-Intervention Effect	-4,080	-1,969	-2,105

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. Earnings 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 changes in the trade balance, thermal generation, and carbon emissions in Spain. As shown in Table 4, the price intervention reduced Spain's net import position (imports minus exports) by 37.4% relative to the *Factual* scenario, reflecting a substantial increase in net exports driven by the lower Iberian prices. Had only the savings measures been implemented, net imports would have fallen by a lower amount, 18.7%. The combined effect of both policies was to reduce net imports by 59.4%. The increase in exports associated with the price intervention required additional thermal generation, raising domestic output from gas-fired plants by 7.2% and increasing Spain's power-sector CO₂ emissions by 2.7%. The savings measures, by contrast, curtailed domestic consumption, reducing thermal generation by 7.5% and cutting emissions by 7.1%. On balance, the combined effect of both policies was to reduce power-sector CO₂ emissions by 4.8%.²⁹

This pattern also points to an efficiency cost of the price intervention: by compressing Iberian wholesale prices, it made Spanish generation relatively cheap compared to the French market, pulling in additional exports that required additional domestic thermal

²⁹Appendix Table F.5 presents the confidence bands for these results. Most remain statistically significant, except for the energy crisis and full-intervention effects on CO₂ emissions.

generation in place of imports. We do not convert this into a monetary deadweight-loss estimate, but note that this margin is more than offset by the accompanying demand-reduction measures.

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

	Imp-Exp	Thermal Generation	CO2 Emissions	Demand
Energy-Crisis Effect	-28.5	-7.6	-2.9	-4.1
Full-Intervention Effect	-59.4	-0.7	-4.8	-3.5
Price-Intervention Effect	-37.4	7.2	2.7	0.3
Savings-Intervention Effect	-18.7	-7.5	-7.1	-3.5

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 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.

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.³⁰

Hourly residential electricity consumption data at the zip code level has been provided by [Datadis \(2023\)](#). The raw data covers 11,153 zip codes. We restrict the sample to continental Spain and drop zip codes with unreliable consumption records: those with more than 1% of hourly values missing, or with anomalous ratios of contracts to resident households or of consumption per household. Importantly, this filtering operates at the zip code level only. No hours are excluded, so all price episodes, including the extreme spikes, are fully reflected in the analysis. We are left with a balanced panel of $8,564 \times 17,520$ zip code-by-hour observations, covering the period from July 2021 until June 2023 and roughly 16.2 million households (more than 90% of the population of peninsular

³⁰We restrict our analysis to continental Spain, excluding the Balearic Islands, the Canary Islands, Ceuta, and Melilla, where specific electricity market regulations apply.

Spain). Besides, using the total number of resident households, we compute the average household’s hourly electricity consumption in each zip code.³¹

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 2017.³²

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 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 counterfactual residential electricity demand at the province level (more details in subsection 4.2 below), we average the weather observations across the stations of each province. 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*.³³

Counterfactual retail prices for households. Household electricity bills are computed by assuming that retail prices are equal to the simulated wholesale electricity prices

³¹We 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.

³²Income 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.

³³Although 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.

obtained in Section 3, augmented by network access charges, taxes, and other components typically included in household bills. Note that only the energy component of the bill —the simulated wholesale price applied to each household’s hourly consumption— varies across our five scenarios. Network access charges, the capacity-based charge tied to contracted power, and taxes are held at their actual, observed levels throughout, since these are regulated independently of the wholesale market outcomes. Because our analysis focuses on bill changes relative to the Factual scenario, the non-volumetric charges are identical across scenarios and cancel out of every comparison. Components that apply proportionally to the energy price, such as grid losses and taxes, are held constant and scale the price differences across scenarios.

We further apply the same counterfactual energy price to all households regardless of contract type; this is exact for those on the regulated real-time tariff, and an approximation for those on fixed-price contracts, which Appendix Figure F.2 shows adjust to wholesale price changes within roughly one renewal cycle.

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.

Counterfactual demand for households. Our simulations for household bill impacts take into account the fact that, as described in Section 3.2, 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 3.2. The zip code-level data described in Section 4.1 is not available for years before 2021, such that it cannot be used for this exercise. Instead, we use the hourly “low voltage” electricity demand of each province, made available by REE (2024) from 2018 onward, as the outcome variable. This series covers consumption under the 2.0TD access tariff, the regulated tariff for low voltage supply points with contracted power of up to 15 kW, which applies to virtually all households; it is therefore a close measure of residential demand.³⁴ We train one gradient-boosting model per province, for the 47 peninsular provinces, using the same classes of predictors as in Section 3.2 (calendar variables, province-level weather, wind and solar generation forecasts, and mobility and pandemic-stringency measures), with hyperparameters tuned separately for each province. Models are trained on data from January 2018 to June 2021. Out-of-fold prediction errors and the estimated crisis effects for each province are

³⁴The 2.0TD tariff was introduced in June 2021, unifying the previous domestic tariff categories (2.0A, 2.0DHA, and 2.1A). For earlier years, the series reports the consumption of the same group of consumers under those predecessor tariffs. Some very small commercial consumers with contracted power below 15 kW are also included.

reported in Appendix Table E.2.

Because wholesale electricity prices and their gas and CO₂ drivers are deliberately excluded from the predictors, the models' predictions represent demand under pre-crisis price conditions: this is the *No-Energy-Crisis* counterfactual demand. The demand series for the remaining scenarios in Table 1 are obtained by adjusting these predictions for the response of demand to electricity prices.

We estimate the short-run price elasticity of residential electricity demand in two independent ways: from hourly household-level smart-meter data, instrumenting the real-time retail price with cost shifters (day-ahead renewable generation forecasts), and as the elasticity implied by our aggregate demand model, dividing its predicted demand response to the crisis by the corresponding change in simulated wholesale prices. The two approaches converge to nearly identical values, -0.0104 and -0.0106 respectively, and we use the household-level estimate in what follows.³⁵ For the *Business-as-Usual* and *Price-Intervention-Only* scenarios, in which the crisis occurs but the energy-saving response does not, we scale the *No-Energy-Crisis* prediction down by the elasticity applied to the hourly gap between the simulated *Factual* and *No-Energy-Crisis* wholesale prices.³⁶ For the *Savings-Intervention-Only* scenario, we instead start from observed consumption and remove the demand-boosting effect of the gas price cap, applying the elasticity to the hourly gap between the *Price-Intervention-Only* and *Factual* retail prices from June 14, 2022 onward.

Note that we focus on short-run demand responses within the two-year crisis window. Gaps between observed consumption and the counterfactuals without savings therefore bundle changes in consumption behavior with structural adaptation. Most notably, the adoption of solar self-consumption can reduce grid consumption without reducing energy use. Appendix F.6 relates our estimated bill effects to the expansion of self-consumption capacity. We find that this expansion had only a limited impact on household bills during the early phase of the crisis period and was more pronounced later on.³⁷

Figure 7 depicts our counterfactual predictions, aggregated into four broad regions in the country. In all regions, observed consumption (in black) is generally lower than the counterfactual predictions without energy savings, especially from the second half

³⁵Appendix D details the household-level instrumental-variables estimation and the implied-elasticity calculation.

³⁶The adjustment is applied only in hours when the simulated *Factual* price exceeds the *No-Energy-Crisis* price. In the remaining hours, concentrated in the first half of 2023 when the price intervention and falling gas prices pushed factual prices below the no-crisis level, demand is left at the predicted baseline rather than scaled up. This choice is conservative, as it slightly understates the estimated savings effect at the end of the sample period.


³⁷We find that regions where capacity per household grew faster experienced significantly smaller crisis effects on electricity bills. However, as shown in Appendix Figure F.7, most of this capacity was installed from late 2022 onward, after the peak of the price shock: in the months when the crisis effect on bills was strongest, the capacity accumulated up to that point could offset only about 2% of it. Structural adaptation thus played a minor role during the crisis itself, although the capacity in place by 2023 can be expected to reduce households' exposure to future price shocks.

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. The green curve, which removes the price cap from observed consumption, lies only slightly below the black curve, reflecting the small short-run price response. Comparing these results with the aggregate demand reductions reported in Figure 5 shows that the reductions were smaller at low voltage. This finding is consistent with the notion that high-voltage consumers (primarily industrial users) have greater flexibility to adjust their electricity consumption than low-voltage consumers—predominantly households.

We then construct counterfactual consumption specific to each zip code based on the demand counterfactuals of its own province. For each hour, we compute the percent difference between the province’s factual and counterfactual residential demand, and apply it to the consumption data of the zip codes in that province. Consistent with the assumptions presented in Table 1, we use the counterfactual without savings for the *BaU* and *Price-Intervention-Only* scenarios, the *No-Energy-Crisis* counterfactual for the *No-Energy-Crisis* scenario, and, for the *Savings-Intervention-Only* scenario, observed consumption net of the price-cap adjustment described above.

With this procedure, we capture heterogeneity in demand responses across provinces. Some bias may remain because we do not directly predict zip code-level counterfactuals. However, we argue that this is unlikely to bias our distributional analyses. 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 Appendix Table F.11). If anything, low-income households seemed to be less aware or less responsive. Therefore, any potential biases from our procedure would most likely be in the direction of underestimating bill increases for the least wealthy households. Further, because we apply *proportional* demand adjustments, zip codes with similar percentage responses can still face very different bill impacts. This is because our approach accommodates high-frequency price changes, as shown in Figure 6, which affect households differently depending on their consumption patterns over time. We expand on this point in the sub-section below.

Figure 7: Actual versus predicted low voltage electricity demand, by region



Notes: This figure illustrates the factual and predicted counterfactual “low voltage” electricity demand from January 2019 to June 2023, summing the provinces of each region (Northwest: Galicia, Asturias, Cantabria, and Castilla y León; Northeast: the Basque Country, Navarra, La Rioja, Aragón, and Cataluña; Center: Madrid, Castilla-La Mancha, and Extremadura; South and Mediterranean: Andalucía, Murcia, and the Comunidad Valenciana). Counterfactuals without savings are the province-level machine-learning predictions (low CO₂ and gas prices), adjusted by the estimated price elasticity for the high-price variant; the counterfactual without the price intervention applies the same elasticity to the retail price gap induced by the gas price cap. Displaying 30-day rolling averages.

4.3 Distributional implications of the energy crisis and policy interventions


We begin by illustrating the heterogeneous price impacts over time. Figure 8, 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 8, 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 8, Panel (a)).

Our detailed consumption data at the zip code level allow 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.

³⁸This pattern was also visible in Figure 6, when comparing wholesale prices under the *Factual* and *BaU* scenarios.

Figure 8: 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 $h \in [1, \dots, 24]$, and month-of-year $m \in [1, \dots, 12]$; 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. Note that the clusters are formed on consumption profiles and location characteristics, not on income. Substantial income variation remains within each cluster, which we exploit in Section 4.3.2 below when estimating how bill impacts vary with household income.

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 9, where Panel (a) plots the average consumption of households by hour of the day, and Panel (b) plots average consumption

³⁹The 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 C presents more details about the clustering algorithm and the performance metrics.

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 zip codes belonging to each cluster.

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

Cluster nº	1	2	3	4	5	6	Full Sample
Electricity Cons. [kWh/year.HH]	2940.0	2540.0	3131.0	2735.0	3311.0	3496.0	3216.0
Avg. Elec Price [c€/kWh]	24.8	25.0	25.0	24.3	25.1	25.1	25.0
Income [kEUR/year.HH]	33.4	40.1	42.0	27.4	34.1	32.4	36.3
Household Size	2.5	2.4	2.5	2.3	2.5	2.6	2.5
Electric Heating [%]	23.6	15.1	14.7	17.1	19.3	43.1	27.3
Air Conditioning [%]	7.1	7.5	50.5	27.5	51.8	68.6	50.2
Cities [%]	42.7	58.4	68.7	11.4	37.9	53.2	55.5
Towns & semi-dense areas [%]	33.5	25.6	19.7	14.6	37.1	36.3	29.7
Rural [%]	23.8	16.0	11.6	74.0	25.0	10.5	14.8
Min Temp. [°C]	5.4	4.1	1.6	1.6	1.9	5.4	3.7
Max Temp. [°C]	26.8	26.3	32.5	31.0	33.6	33.7	32.0
Number of ZIP codes	967	1147	2053	1276	1152	1969	8564
Total Households (thousands)	992	1776	4878	251	1780	6547	16224

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. Appendix Tables F.9 and F.10 explore the determinants of household consumption and average prices at the zip code level. 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.

From the intraday patterns, we observe that all clusters reach their maximum demand in the evening, around 8–9 pm, although the peak is most pronounced in Cluster 6 and flattest in Cluster 1. The clusters differ more sharply in their seasonal profiles and in the location and equipment of their households, as shown in Table 5 and the map in Panel (c).


Cluster 6, by far the largest with 6.5 million households, groups zip codes along the Mediterranean coast and in Andalucía (Valencia, Alicante, Málaga, Sevilla, and the Barcelona area). It records the highest consumption per household, sustained by hot summers and the highest air conditioning share (69%), and electric heating is also most widespread there (43%), such that demand is elevated both in summer and in winter. Cluster 3 comprises the large metropolitan areas, with Madrid and Barcelona accounting for over 60% of its households. It includes the wealthiest and most urban zip codes, and its demand also peaks in both seasons, with about half of households equipped with air conditioning.

Cluster 5 covers central and southern inland areas, including parts of the Madrid region, Castilla-La Mancha, and inland Andalucía. Its households exhibit the most pronounced winter peaks of all clusters, with January consumption about 70% above the cluster’s annual average, reflecting cold winters, complemented by a secondary summer peak driven by air conditioning.


Clusters 1 and 2 are located in the north. Cluster 1 comprises the Atlantic northwest (Galicia and Asturias), where the mild climate makes air conditioning virtually absent (7%) and demand tilts moderately towards winter. Cluster 2, covering the Cantabrian coast, the Basque Country, and Navarra, groups relatively wealthy and urban zip codes with the lowest consumption per household and the flattest seasonal profile, consistent with mild summers and winters.

Finally, Cluster 4 consists of small, predominantly rural zip codes (74%) scattered across the northern interior, with the lowest average income of all clusters. These households face the lowest average electricity prices, indicating that their consumption is concentrated in low-price periods. The cluster displays a striking August peak, about 70% above its annual average, which may be attributed to air conditioning together with the higher occupancy of rural areas during the summer.


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 9: 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, 2026](#)), 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 as a 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)

	1	2	3	4	5	6	Full Sample
Energy-crisis effect [EUR/year.HH]	333	290	368	313	390	411	376
[% income]	1.02	0.76	0.91	1.19	1.2	1.33	1.1
Full-intervention effect [EUR/year.HH]	-157	-126	-148	-145	-161	-174	-158
[% income]	-0.48	-0.33	-0.37	-0.55	-0.5	-0.56	-0.47
Price intervention [EUR/year.HH]	-77	-68	-91	-89	-91	-106	-94
[% income]	-0.24	-0.18	-0.22	-0.34	-0.28	-0.34	-0.28
Savings intervention [EUR/year.HH]	-66	-48	-46	-44	-57	-54	-52
[% income]	-0.2	-0.13	-0.12	-0.17	-0.18	-0.17	-0.15
Interaction [EUR/year.HH]	-13	-10	-11	-12	-13	-13	-12
[% income]	-0.04	-0.03	-0.03	-0.05	-0.04	-0.04	-0.04

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, households in Clusters 6 and 5 suffered from the largest bill increase due to the crisis (+ €411 and + €390/year) while also benefiting the most from the policy interventions (− €174 and − €161/year). Relative to disposable income, the crisis burden was also heaviest in these two clusters and in Cluster 4 (1.3%, 1.2%, and 1.2%, respectively), the latter because of its low income levels despite a smaller absolute increase. At the other end, Cluster 2, which comprises relatively wealthy zip codes along the northern coast, was spared by the crisis effect (+ €290/year, representing 0.8% of disposable income) and benefited less from the policy interventions (− €126/year).

Appendix Figures F.3 and F.4 further illustrate the distribution of these effects across Spanish provinces. Tarragona experienced the largest increase in annual electricity

bills due to the crisis (+ €466/year), while Jaén obtained the strongest relief from the interventions (− €247/year). However, when measured relative to income, the southern provinces were hit hardest. In Granada, Córdoba, Toledo, Málaga, and Alicante, the bill increase exceeded 1.5% of households’ disposable income. These same regions also benefited most from the interventions, with reductions amounting to as much as 0.9% 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 available only at the zip code level, which masks income heterogeneity among households within the same area. As illustrated in Table 7, grouping zip codes – rather than households – into national income quintiles substantially compresses the income distribution. Average income in the lowest zip code quintile is €24,773, against €9,481 for the lowest household quintile. Conversely, the richest zip codes average €50,316, against €57,485 for the top household quintile. The gap between the bottom and top quintiles is thus more than twice as large at the household level (€48,004) as at the zip-code level (€25,543). A naïve zip-code-level approach would therefore underestimate the true gap in bill impacts between low- and high-income households.⁴⁰

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 more accurately capture the differentiated impacts across household-level income

⁴⁰For completeness, we present results of the naïve approach in Appendix F.5. Reguant, Fabra, and Wang (2026) show that such an approach obscures much of the true distributional impact.

quintiles, we proceed in two steps.⁴¹

Step 1: Estimating the income gradient of bill changes at the zip code level.

We first estimate how bill changes vary with income across zip codes, controlling for other determinants of electricity consumption. Specifically, for each zip code i , we regress the log of the bill change on the log of average household income, as well as controls for average household size, the shares of electric heating and air conditioning, and the degree of urbanization:

$$\log(\Delta Bill_i) = \beta_0^c + \beta_{Inc}^c \log(Income_i) + \beta^c \mathbf{X}_i + \epsilon_i. \quad (2)$$

This regression is estimated separately for each consumption cluster c , allowing the income gradient β_{Inc}^c to vary across regions with different consumption profiles (Tables F.12–F.15 in the Appendix). Table 8 reports the pooled results across all zip codes.

Step 2: Extrapolating to household-level income quintiles.

We then use the estimated coefficients to simulate what bill changes would look like for households at each point of the national income distribution. For each zip code i in cluster c , we replace the observed zip code income with the national average income of quintile N , while holding all other zip code characteristics $\mathbf{X}_i$ fixed at their observed values:

$$\widehat{\Delta Bill}_i^N = \exp\left(\hat{\beta}_0^c + \hat{\beta}_{Inc}^c \log(\overline{Income}^N) + \hat{\beta}^c \mathbf{X}_i\right). \quad (3)$$

This yields $8,564 \times 5$ predicted bill changes – one per zip code per quintile – that capture both the effect of belonging to income quintile N and the geographic and equipment characteristics of each zip code. We then aggregate to the national level for each quintile, weighting by the number of households from quintile N residing in each zip code:


$$\widehat{\Delta Bill}^N = \frac{1}{pop^N} \sum_i \widehat{\Delta Bill}_i^N \times pop_i^N, \quad (4)$$

where pop_i^N is the number of quintile- N households in zip code i and pop^N is their national total. By weighting in this way, the aggregation accounts for the fact that low-income households are not uniformly distributed across the country but tend to be concentrated in particular types of areas — a feature that matters for the final distributional estimates.

Table 8 reports the regression results from Step 1. Household income is positively and significantly correlated with bill changes across all specifications. The pooled estimate implies that a 1% increase in average zip code income is associated with a 0.4% larger

⁴¹Reguant, Fabra, and Wang (2026) 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.

Figure 10: 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.

bill change – whether due to the crisis or to the policy interventions – and this gradient ranges from 0.3% to 0.6% across clusters. Beyond income, larger household size, higher penetration of electric heating or air conditioning, and location in towns or rural areas are all associated with larger bill impacts, consistent with these factors driving higher electricity consumption (see Appendix Table F.9).

Figure 10 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 €323 for households in the first (lowest) income quintile – equivalent to 3.4% of their disposable income. In contrast, households in the fifth (highest) income quintile experienced an average increase of €648, corresponding to 1.1% 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 larger share attributable to the price intervention. Without the interventions, bills for the lowest-income households would

have increased by an additional €138, reaching a total impact of 4.9% of their income. For the highest-income households, the interventions reduced bills by €253 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 1.6% 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 impacts 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 F.10 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 0.006%.

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

€29 billion over the two-year duration of the energy crisis.

Policy interventions, such as the “Iberian Solution” and energy-saving measures, mitigated part of this transfer by lowering wholesale prices, thereby reducing the infra-marginal gains by almost €16 billion and alleviating the financial burden on consumers. Without the interventions, the increase in firms’ profits would have reached approximately €45 billion relative to the no-crisis scenario.

At the *retail level*, the energy crisis had regressive effects, even accounting for the intervention. 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 3.4% 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, thus partly mitigating the regressivity of the crisis. 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 1.5% 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. While the Iberian solution was not targeted by design, a more explicit targeting would have improved its distributional effects.

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.

- Albalade, 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 Küçük. 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. 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, University 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 Commission.
- 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.
- Feger, Fabian, Nicola Pavanini, and Doina Radulescu. 2022. “Welfare and Redistribution in Residential Electricity Markets with Solar Power.” *The Review of Economic Studies* 89 (6):3267–3302. URL <https://doi.org/10.1093/restud/rdac005>.
- 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, Ludovica Gazze, and Menna Bishop. 2024. “Distributional and climate implications of policy responses to energy price shocks.” *Economic Policy* 39 (120):711–756. URL <https://doi.org/10.1093/epolic/eiae038>.
- 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>.

- Gelaro, Ronald, Will McCarty, Max J. Suárez et al. 2017. “The Modern-Era Retrospective Analysis for Research and Applications, Version 2 (MERRA-2).” *Journal of Climate* 30 (14):5419–5454.
- 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.
- Künsch, Hans R. 1989. “The jackknife and the bootstrap for general stationary observations.” *The annals of Statistics* :1217–1241.
- 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>.
- Lundberg, Scott M, Gabriel Erion, Hugh Chen, Alex DeGrave, Jordan M Prutkin, Bala Nair, Ronit Katz, Jonathan Himmelfarb, Nisha Bansal, and Su-In Lee. 2020. “From local explanations to global understanding with explainable AI for trees.” *Nature Machine Intelligence* 2 (1):56–67.
- Lundberg, Scott M and Su-In Lee. 2017. “A unified approach to interpreting model predictions.” *Advances in Neural Information Processing Systems* 30.
- 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/>.
- Politis, Dimitris N and Halbert White. 2004. “Automatic block-length selection for the dependent bootstrap.” *Econometric reviews* 23 (1):53–70.
- 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 Eléctrica de España. 2024. “Generation and Consumption Data.” *ESIOS-Sistema de Información del Operador del Sistema* URL <https://www.esios.ree.es/en/generation-and-consumption>.
- Reguant, Mar, Natalia Fabra, and Jingyuan Wang. 2026. “The Distributional Impacts of Time-Varying Electricity Pricing: A Novel Approach to Estimating Household Income.” *Journal of the European Economic Association* 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.”

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

	Energy Crisis	Full Intervention	Energy Savings Intervention	Price Intervention
$\log(\text{Income})$	0.416*** (0.016)	0.362*** (0.017)	0.317*** (0.021)	0.407*** (0.016)
<i>Household Size</i>	0.279*** (0.017)	0.279*** (0.018)	0.276*** (0.022)	0.284*** (0.017)
<i>Electric Heating</i>	0.413*** (0.026)	0.523*** (0.028)	0.464*** (0.035)	0.487*** (0.027)
<i>Air Conditioning</i>	0.338*** (0.015)	0.206*** (0.016)	-0.138*** (0.021)	0.447*** (0.016)
<i>Degree of urbanization</i> (ref.: Cities)				
Towns and Semi-dense areas	0.252*** (0.008)	0.258*** (0.009)	0.247*** (0.011)	0.260*** (0.008)
Rural areas	0.299*** (0.011)	0.333*** (0.012)	0.324*** (0.015)	0.337*** (0.011)
(Intercept)	0.808*** (0.169)	0.472*** (0.182)	-0.030 (0.230)	-0.601*** (0.175)
R ²	0.263	0.228	0.119	0.313
Adj. R ²	0.262	0.228	0.119	0.312
Num. obs.	8564	8564	8564	8564

Notes: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Each observation corresponds to one zip code. The dependent variable $\log(|\Delta \text{Bill}|)$ is the log of the absolute increase (decrease) in households' annual bill due to the crisis (interventions). *Income* stands for average disposable income per household per year; *Household Size* for the mean number of people per household; *Electric Heating* and *Air Conditioning* for the share of households with each appliance.

Appendix – For Online Publication

A Counterfactual simulation inputs

A.1 Natural Gas and CO₂ 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 CO₂ 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 CO₂ prices, we estimate the linear regression model below, capturing the effect of natural gas price (TTF) on CO₂ prices, while allowing for different linear time trends in different periods of the EU-ETS:

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

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

A.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

Table A.1: Linear model of monthly CO₂ prices (2015-2023)

$Y =$	$\log(CO_2 \text{ price})$ [EUR/t]
$\log(TTF \text{ price})$ [EUR/MWh]	0.18*** (0.03)
<i>EU-ETS Period</i> (ref.: Period A)	
Period B	-0.39** (0.15)
Period C	-0.35** (0.16)
Period D	0.55** (0.23)
<i>trend</i> $\times$ Period A	0.02 (0.01)
<i>trend</i> $\times$ Period B	-0.01 (0.01)
<i>trend</i> $\times$ Period C	0.01 (0.01)
<i>trend</i> $\times$ Period D	-0.00 (0.01)
(Intercept)	1.37*** (0.14)
R^2	0.98
Adj. R^2	0.98
Num. obs.	108

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

function of the natural gas price (TTF), the CO₂ price, electricity demand in France net of renewable and nuclear production (*Net Load*), a quadratic term of *Net Load*, interaction terms between gas and CO₂ 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 A.2. The model explains a substantial share of the variance in electricity prices ($R^2 = 0.9$), and both gas and CO₂ 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 4). Specifically, we assume TTF gas prices are equal to the monthly average observed during 2017–2019 and use the counterfactual CO₂ price series described in Appendix A.1.

Table A.2: Linear model estimates for prediction of French electricity prices

$Y =$	French Electricity Price [EUR/MWh]
<i>Net Load</i> [GW]	0.4355*** (0.0347)
$(\textit{Net Load} \text{ [GW]})^2$	0.0027 (0.0023)
<i>TTF Price</i> [EUR/MWh]	1.0580*** (0.0065)
<i>CO₂ Price</i> [EUR/t]	0.5293*** (0.0061)
<i>Net Load</i> [GW] $\times$ <i>TTF Price</i> [EUR/MWh]	0.0439*** (0.0011)
<i>Net Load</i> [GW] $\times$ <i>CO₂ Price</i> [EUR/t]	0.0151*** (0.0008)
$(\textit{Net Load} \text{ [GW]})^2 \times \textit{TTF Price}$ [EUR/MWh]	0.0032*** (0.0001)
$(\textit{Net Load} \text{ [GW]})^2 \times \textit{CO}_2 \text{ Price}$ [EUR/t]	-0.0022*** (0.0001)
(Intercept)	-10.6258*** (0.7028)
R ²	0.9050
Adj. R ²	0.9050
Num. obs.	77978

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

B Predicting counterfactual electricity demand

B.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 $g(\cdot)$ that drives no-crisis counterfactual aggregate electricity demand.

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

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

We refer to $g(\cdot)$ 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.

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

Taken together, assumptions 1 and 2 imply:

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

We train flexible machine learning algorithms using pre-crisis data ($Y_{pre}(0)$ and $\mathbf{X}_{pre}(0)$), aiming to estimate the counterfactual function $\hat{g}(\cdot)$. 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 $\hat{g}(\cdot)$, where the outcome variable is aggregate hourly electricity demand in Spain.

B.2 Machine learning algorithms

For estimating the counterfactual function $\hat{g}(\cdot)$, 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 B.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.

B.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 B.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. Rather than treating the number of trees as a hyperparameter fixed

⁴²The 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 B.2).

⁴³We 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 B.2).

by the grid, we select it through *early stopping*. For each candidate configuration, we add boosting iterations (trees) sequentially and monitor the cross-validated RMSE; once it fails to improve for 50 consecutive iterations, training halts and the iteration achieving the lowest CV RMSE is retained. Some configurations are capped at 2,000 trees, while others are capped at 3,000. The “Number of Trees” column in Table B.1 reports the early-stopped iteration count selected for each configuration.

For training, we use data from 2015 to 2020. The last two columns of Table B.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 counterfactual predictions shown in Section 3.2. Table B.2 presents comparisons between realized versus counterfactual predicted demand for each month of the sample.

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 B.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 3.2, the discrepancies between real and predicted demand observed from 2020 onward may be attributed to unexpected crises (COVID and the energy crisis itself).

Table B.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	291.63	520.57
2	2,000	10	0.1	10	17.07	494.10
3	2,000	5	0.2	10	182.25	511.17
4	1,985	10	0.2	10	1.47	540.85
5	2,000	10	0.1	5	6.03	518.06
6	2,996	10	0.1	5	0.99	517.92
7	1,800	20	0.1	5	0.05	640.72
8	1,800	20	0.1	5	0.05	640.72
9	2,997	10	0.1	10	4.96	493.67
10	1,998	20	0.1	10	1.21	566.77
11	2,061	20	0.1	10	1.17	566.77
12	2,000	10	0.1	20	45.38	465.31
13	2,998	10	0.1	20	20.50	463.99
14	2,000	20	0.1	20	4.68	485.70
15	2,522	20	0.1	20	3.51	485.67
16	1,500	30	0.1	20	7.68	514.63
17	2,997	20	0.1	30	6.13	456.44
18	2,349	30	0.1	30	9.31	477.38


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.

Table B.2: Average monthly electricity demand vs. counterfactual

Month	Pred. (real)	Reduction (%)	Pred. (stable)	Reduction (%)	Actual
Jan 2021	29,958	-3.7	29,935	-3.8	31,060
Feb 2021	28,766	0.0	28,815	0.2	28,758
Mar 2021	28,003	-0.2	28,040	-0.1	28,057
Apr 2021	26,276	-0.2	26,241	-0.3	26,320
May 2021	26,275	1.2	26,322	1.4	25,951
Jun 2021	27,423	0.1	27,391	-0.0	27,403
Jul 2021	28,538	-2.4	28,693	-1.9	29,231
Aug 2021	28,259	0.8	28,273	0.8	28,034
Sep 2021	27,438	-0.0	27,557	0.4	27,446
Oct 2021	26,326	2.5	26,394	2.8	25,662
Nov 2021	28,515	0.6	28,755	1.4	28,341
Dec 2021	28,362	1.1	28,612	1.9	28,060
Jan 2022	28,904	-0.1	29,353	1.4	28,935
Feb 2022	28,785	1.1	29,162	2.4	28,460
Mar 2022	28,236	3.2	28,318	3.5	27,329
Apr 2022	26,644	4.1	26,794	4.6	25,563
May 2022	26,456	3.1	26,554	3.5	25,633
Jun 2022	27,966	0.7	28,049	1.0	27,781
Jul 2022	29,433	-0.7	29,581	-0.2	29,630
Aug 2022	29,070	5.1	29,103	5.2	27,593
Sep 2022	27,555	4.0	27,662	4.4	26,454
Oct 2022	26,524	7.4	26,533	7.5	24,548
Nov 2022	27,502	7.1	27,618	7.4	25,562
Dec 2022	27,833	6.8	27,948	7.2	25,937
Jan 2023	29,393	4.4	29,794	5.7	28,091
Feb 2023	28,983	0.7	29,434	2.2	28,780
Mar 2023	27,583	5.4	27,735	6.0	26,081
Apr 2023	25,823	11.0	25,919	11.4	22,977
May 2023	26,134	10.3	26,295	10.8	23,448
Jun 2023	27,755	6.5	27,787	6.6	25,948
Jul 2023	29,095	1.9	29,179	2.1	28,555
Aug 2023	29,085	6.6	29,055	6.5	27,170
Sep 2023	27,389	7.1	27,509	7.5	25,450
Oct 2023	26,837	7.5	26,847	7.5	24,828
Nov 2023	27,492	5.1	27,613	5.6	26,077
Dec 2023	27,922	4.0	28,132	4.7	26,806

Notes: Monthly averages over the forecast period. “Reduction” is $100 \times (\text{predicted} - \text{actual})/\text{predicted}$ for each counterfactual without savings (high vs. low CO₂ and gas prices); predicted and actual demand in MWh.

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




Notes: This figure illustrates the factual 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.

B.4 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 B.2, 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 B.2: 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).

B.5 Variable importance

To understand which predictors drive our counterfactual demand model, we examine variable importance for the best-performing XGBoost configuration (Model ID 17 in Table B.1). We rely on two complementary metrics. The first is the standard *gain*-based importance, which measures the total reduction in squared-error loss attributable to splits on each variable across all trees of the XGBoost ensemble. The second builds on *SHapley Additive exPlanations* (SHAP) (Lundberg and Lee, 2017), computed exactly for tree ensembles via TreeSHAP (Lundberg et al., 2020). SHAP decomposes each individual prediction into additive per-variable contributions, expressed in the units of the outcome (MWh): the average of their absolute values, $|\overline{\text{SHAP}}|$, yields a global importance ranking, while the *distribution* of the signed contributions reveals the direction and heterogeneity of each variable’s effect. We compute SHAP values on a random subsample of 15,000 hours from the pre-crisis training period.


Figure B.3 ranks the predictors by both metrics, which agree closely. As expected for hourly electricity demand, the calendar variables dominate: the hour-of-day is by far the most important predictor, followed by the day-of-week and by finer seasonal controls (week-of-year and quarter). Weather variables form the next tier, followed by our mobility proxy (road traffic fatalities; see Appendix B.4), the public-holiday indicator, the daily time trend, and lagged natural-gas (TTF) prices. The prominence of these demand fundamentals, together with the negligible role of idiosyncratic predictors, supports the validity of the model.

Figure B.4 reports, for the most important *numeric* predictors, the SHAP value of every observation in the subsample. Each point is one hour from a random subsample of 15,000 pre-crisis hours. The color encodes the variable’s value, rescaled to $[0, 1]$ by min–max normalization within each variable (blue = low, red = high). The graph thus conveys both the magnitude and the direction of each variable’s influence across hours. We note that some temperature controls exhibit a non-monotonic effect, as both unusually cold and unusually warm hours raise demand. Holiday hours are associated with lower demand predictions, as expected. Also, low numbers of traffic fatalities lead to demand reduction, helping to capture COVID lockdown measures.

Finally, the calendar variables are categorical, so a low-to-high color encoding is not meaningful for them. We therefore display their effects in Figure B.5 as the distribution of SHAP values at each level. These panels recover well-known demand regularities: a pronounced diurnal cycle, systematically higher demand on weekdays than on weekends, and a seasonal pattern with elevated winter and summer demand relative to the shoulder months. The first and last weeks of the year carry strongly negative contributions, reflecting the Christmas–New-Year holiday lull. Some low SHAP values for March, as well as Q2, are most likely reflecting COVID lockdowns. Week 33 of the year, also as-


sociated with low SHAP, includes August 15th, which is an important holiday in Spain (Assumption of the Virgin Mary). Again, the SHAP values align with expected patterns and demand fundamentals.

Figure B.3: Variable importance – aggregate electricity demand model




Notes: Variable importance for the best-performing XGBoost configuration (Model ID 17 in Table B.1), trained on aggregate Peninsular-Spain electricity demand, 2015–2020. Left: gain-based importance (share of total split-gain). Right: mean absolute SHAP value (share of total). One-hot indicators of the calendar variables are aggregated back to their parent variable. The top 20 variables are shown.

Figure B.4: SHAP values for the leading numeric predictors



Notes: Each point is one hour from a random subsample of 15,000 pre-crisis hours; the horizontal position is that hour's SHAP value (contribution to predicted demand, in MWh) and the color is the min-max scaled value of the variable (blue = low, red = high). The figure is restricted to the leading *numeric* predictors; the categorical calendar variables are shown in Figure B.5. Province-level weather variables are labeled with the province name in parentheses; "National avg." denotes the population-weighted national series.

Figure B.5: SHAP values by level for the calendar variables



Notes: Distribution of SHAP values (contribution to predicted demand, in MWh) at each level of the categorical/cyclical calendar variables, over a random subsample of 15,000 pre-crisis hours. Boxes show the interquartile range and median; whiskers the $1.5 \times \text{IQR}$ range.

B.6 Uncertainty bands for the counterfactual predictions

The shaded bands around the counterfactual demand paths in Figure 5 quantify the uncertainty in our prediction of $\hat{Y}_{post}(0) = \hat{g}(\mathbf{X}_{post})$ that is attributable to the model’s forecast error. We construct them with a *block-residual bootstrap* that holds the estimated function $\hat{g}()$ fixed and resamples its prediction errors. The procedure has four steps, as follows.

Step 1. Forecast-error residuals. We cannot observe the model’s error during the crisis, because the realized residual there, $Y_{post}(1) - \hat{g}(\mathbf{X}_{post})$, confounds the model’s prediction error with the very crisis effect we seek to estimate. We therefore characterize the error distribution from the *pre-crisis* period, where the realized residual is a clean draw of forecast error. To avoid in-sample optimism, we use *out-of-fold* residuals, $e_t = Y_t(0) - \hat{g}_{-k(t)}(\mathbf{X}_t)$, where $\hat{g}_{-k(t)}$ is fitted on the cross-validation folds that exclude the fold containing observation t . More specifically, each fold for this procedure is one calendar year. We restrict the residual pool to 2015–2019, dropping 2020 so that the anomalous pandemic errors do not contaminate the bands.

Step 2. Seasonal heteroskedasticity. The forecast error is strongly seasonal: its standard deviation is more than twice as large in winter (e.g., January) as in the mild shoulder months (e.g., May and October; see Figure B.6), reflecting the greater volatility of temperature-driven heating and cooling demand. A bootstrap that pooled residuals across the whole year would therefore mis-size the band—too narrow in winter and too wide in spring and autumn. We accordingly stratify the residuals by calendar month and, within each month, center them (subtracting the monthly mean), so that the resampled errors are mean-zero and the band stays centered on the point prediction while its *width* varies by season.


Step 3. Block resampling. Hourly demand errors are serially correlated and carry strong diurnal and weekly structure, so resampling individual residuals would severely understate uncertainty. We instead use a circular block bootstrap (Künsch, 1989; Politis and White, 2004): each forecast month is filled with contiguous blocks of 168 consecutive residuals (one week), drawn only from the matching pre-crisis calendar month. Thus, each block preserves the within-week error structure.

Step 4. From bootstrap replicates to seasonal scenario bands. Steps 1 to 3 yield $B = 1,000$ bootstrap replicates of the hour-by-hour residual series; adding each to the point counterfactual $\hat{g}(\mathbf{X}_{post})$ produces B internally coherent counterfactual demand paths. A naïve *pointwise* band takes the 2.5th and 97.5th percentiles across replicates separately at each hour. The resulting lower (or upper) curve is problematic: because a different replicate can sit at the relevant percentile in each hour, the curve is stitched together across replicates and is not a demand path that any single replicate follows. It is an outer boundary of the set of simulated paths rather than a path itself. Because

forecast errors partially cancel when summed across hours, using such a boundary as if it were a demand series greatly overstates the uncertainty of any temporal aggregate. Summed over the crisis window, its lower curve even falls below realized demand, which would spuriously reverse the sign of the estimated crisis effect. This boundary is therefore unsuitable both as an input to the wholesale-market simulation, which requires coherent hourly series, and for reporting aggregate quantities.

We instead summarize the bootstrap distribution at the level of the calendar month, which is the natural unit of the resampling scheme, since each forecast month is filled with its own blocks (Steps 2 and 3). For each year-month m and each replicate b , we sum the replicate’s hourly counterfactual over the hours of m to obtain a bootstrap monthly total $S_m^{(b)}$. We take the 2.5th and 97.5th percentiles of $\{S_m^{(b)}\}_{b=1}^B$ and divide them by the point prediction’s monthly total, $\sum_{t \in m} \hat{g}(\mathbf{X}_t)$, giving a pair of scaling factors $\underline{k}_m$ and $\bar{k}_m$. The lower and upper bands are then the point counterfactual scaled hour by hour within each month, $\underline{k}_m \hat{g}(\mathbf{X}_t)$ and $\bar{k}_m \hat{g}(\mathbf{X}_t)$ for $t \in m$. By construction, each band is itself a coherent hourly demand path: it preserves the diurnal and weekly *shape* of the point prediction and shifts only its *level*, and its monthly total reproduces the corresponding percentile of the bootstrap distribution. The width of the band therefore inherits the seasonal heteroskedasticity of Step 2, being wider in winter than in the shoulder months. The same residual pools feed both the real-price and the no-crisis (stable-price) counterfactuals, since the forecast error is a property of the model rather than of the price scenario. Figure 5 displays the 30-day moving average of each of these series (point, lower, and upper).

Figure B.6: Seasonality of the demand model’s forecast error



Notes: Standard deviation of the out-of-fold prediction residuals of the aggregate electricity demand model (Model ID 17 in Table B.1), by calendar month, computed over the pre-crisis period 2015–2019. The dashed line marks the overall standard deviation. The pronounced seasonality motivates the month-stratified block bootstrap.

C Clustering algorithm and performance metrics

C.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 8,564 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) $\mathbf{C} = \{C_1, C_2, \dots, C_6\}$, such that:

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

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.

C.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 (C.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 8,564 zip codes in our sample, we take the average.

Results for total within-cluster sum of squares are presented in Panel (a) of Figure C.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 C.1: Performance metrics to determine optimal number of clusters

D Elasticity to real-time electricity prices

D.1 Estimates with cost shifters

Our aggregate counterfactual for electricity demand (Appendix B) captures the price response only indirectly, through variation in gas and CO₂ prices. To quantify how households respond to *real-time* electricity prices directly, we estimate a short-run own-price elasticity of demand and use it to adjust the retail counterfactual for the price-response channel. The estimate is also of independent interest, as it speaks to how much of the demand reduction during the crisis can be attributed to prices as opposed to non-price energy-saving behavior.

Data. We use household-level, hourly smart-meter consumption records from a large Spanish electricity retailer, matched to each household’s contract information, over the period January 2021–July 2023. We restrict attention to households on an *indexed* (“*Indexado*”) low-voltage residential tariff (2.0TD/2.0A), whose bills track the hourly regulated price (PVPC). These households therefore face the real-time price directly, which makes them the relevant population for identifying the response to hourly prices. The estimation sample contains 7,193 households and roughly 66 million household-hour observations. We merge the consumption data with the hourly PVPC price, day-ahead wind and solar generation forecasts, and local weather (temperature and wind speed) from the MERRA-2 reanalysis (Gelaro et al., 2017), aggregated to the municipality level.

Specification. Because prices and quantities are determined jointly, we estimate the elasticity by two-stage least squares (2SLS), instrumenting the price with the *day-ahead wind generation forecast*, as in Fabra et al. (2021). Let i index households and t index hours. The second stage is

$$\text{asinh}(q_{i,t}) = \beta \text{asinh}(p_t) + f(\text{temp}_{i,t}, \text{wind}_{i,t}) + \alpha_{i,w,h} + \gamma_{z,d} + \varepsilon_{i,t},$$

where $q_{i,t}$ is hourly consumption, p_t is the hourly PVPC price, and both enter through the inverse hyperbolic sine (IHS) transform, so that β is interpretable as an elasticity (the IHS behaves like a logarithm away from zero while accommodating zeros). The term $f(\cdot)$ collects second-order polynomials in temperature and in local wind speed. We saturate the specification with week-of-year $\times$ hour-of-day $\times$ household fixed effects $\alpha_{i,w,h}$ — which absorb each household’s typical consumption at a given hour and season — and day-of-sample $\times$ zip code fixed effects $\gamma_{z,d}$, which absorb any local, day-specific shock common to a postcode. Standard errors are clustered at the household level. We estimate the model on the full sample and separately for each of the six postcode clusters defined from aggregate consumption profiles (Section 4.3.1).

Identification. The day-ahead wind forecast is a cost/supply shifter: wind generation displaces more expensive marginal technologies and lowers the wholesale price, and hence the PVPC, but is plausibly unrelated to any individual household’s idiosyncratic demand shock once weather and the rich fixed effects are controlled for. The first stage is strong and precisely estimated: a higher forecast lowers the price (coefficient -0.095 , Table D.2), with a first-stage F -statistic of about 9.6×10^5 in the full sample. The reduced form (Table D.3) shows the mirror-image pattern: a higher wind forecast is associated with *higher* consumption ($+0.0010$, significant at the 1% level), as expected when demand slopes down in price.

Results. Table D.1 reports the 2SLS estimates. In the full sample, the own-price elasticity of hourly electricity demand is -0.0104 (standard error 0.0022, significant at the 1% level): a 1% increase in the real-time price reduces consumption by about one hundredth of a percent within the hour. This confirms that the short-run demand response to prices is real but very small, consistent with the low short-run elasticities documented in the literature (Labandeira, Labeaga, and López-Otero, 2017; Deryugina, MacKay, and Reif, 2020). Across clusters the estimate is negative throughout and largest (in absolute value) for Cluster 1 (-0.0160 , significant at 1%); it is marginally significant for Cluster 5 (-0.0081) and negative but imprecisely estimated for the remaining clusters, which are smaller. We adopt the full-sample point estimate, -0.0104 , as the price-elasticity parameter when adjusting our retail counterfactual demand for the price-response channel (Section 4.1). Its small magnitude further justifies attributing the bulk of the observed demand reduction during the crisis to non-price energy-saving measures rather than to the price increase itself.

Table D.1: Elasticity Estimates – Day-ahead wind generation forecast as instrument

Dependent Variable:	IHS Hourly Consumption						
	Full Sample	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Variables</i>							
IHS Electricity Price	-0.0104*** (0.0022)	-0.0160*** (0.0033)	-0.0101 (0.0094)	-0.0072 (0.0071)	-0.0057 (0.0046)	-0.0081* (0.0045)	-0.0144 (0.0133)
Weather Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day of Sample by Zipcode FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Week of Year by Hour by Household FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	66,083,003	28,965,828	4,152,890	5,156,723	17,324,921	9,943,990	538,651
Households	7,193	3,152	410	522	2,028	1,018	63
Adjusted R ²	0.688	0.657	0.761	0.742	0.693	0.686	0.739

Notes: The dependent variable is the inverse hyperbolic sine (IHS) of hourly electricity consumption. The IHS of the day-ahead wind generation forecast was used as an instrument for IHS electricity prices. Standard errors clustered at the household level. Significance codes: ***: 0.01, **: 0.05, *: 0.1

Table D.2: First Stage results

Dependent Variable: Model:	IHS Electricity Price						
	Full Sample (1)	Cluster 1 (2)	Cluster 2 (3)	Cluster 3 (4)	Cluster 4 (5)	Cluster 5 (6)	Cluster 6 (7)
<i>Variables</i>							
IHS Day-Ahead Wind Generation Forecast	-0.0949*** (0.0002)	-0.0955*** (0.0003)	-0.0886*** (0.0010)	-0.0961*** (0.0009)	-0.0925*** (0.0003)	-0.0925*** (0.0008)	-0.0840*** (0.0040)
Weather Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day of Sample by Zipcode FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Week of Year by Hour by Household FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	66,083,003	28,965,828	4,152,890	5,156,723	17,324,921	9,943,990	538,651
Households	7,193	3,152	410	522	2,028	1,018	63
Adjusted R ²	0.845	0.844	0.842	0.843	0.846	0.847	0.838
F-test (1st stage)	964,725.2	428,485.7	49,545.6	72,815.3	248,580.6	132,266.8	5,490.9

Notes: The dependent variable is the inverse hyperbolic sine (IHS) of hourly electricity prices. Standard errors clustered at the household level.
Significance codes: ***: 0.01, **: 0.05, *: 0.1

Table D.3: Reduced Form results

Dependent Variable:	IHS Hourly Consumption						
	Full Sample	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Variables</i>							
IHS Day-Ahead Wind Generation Forecast	0.0010*** (0.0002)	0.0015*** (0.0003)	0.0009 (0.0008)	0.0007 (0.0007)	0.0005 (0.0004)	0.0008* (0.0004)	0.0012 (0.0011)
Weather Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day of Sample by Zipcode FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Week of Year by Hour by Household FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	66,083,003	28,965,828	4,152,890	5,156,723	17,324,921	9,943,990	538,651
Households	7,193	3,152	410	522	2,028	1,018	63
Adjusted R ²	0.688	0.657	0.761	0.742	0.693	0.686	0.739

Notes: The dependent variable is the inverse hyperbolic sine (IHS) of hourly electricity consumption. Standard errors clustered at the household level. Significance codes: ***: 0.01, **: 0.05, *: 0.1

D.2 Implied price elasticity from the aggregate demand model

The household estimates above identify the response to real-time prices within a specific retail sample. As an independent, market-level check, we can recover the price elasticity implied by the aggregate demand model of Appendix B, even though that model deliberately excludes the electricity price from its predictors. The model instead uses the TTF gas price and the CO₂ price as proxies for the crisis’s effect on power prices. Our two counterfactual demand paths (Figure 5) differ only in those fuel prices, so the gap between them is the model’s demand response to the joint move from crisis to no-crisis price levels.

Because the fuels enter the model only as proxies for the electricity price, this response can be expressed as an elasticity with respect to the electricity price itself. We take the average demand response and divide it by the average change in the wholesale electricity price between the *Factual* and *No-Energy-Crisis* simulations, using the demand-weighted price over the period July 2021 to June 2023. Table D.4 reports the result. The model’s mean demand response is small, about +154 MWh/h, or 0.55% of average demand, against a wholesale price that is more than twice as high under the crisis (€144 versus €68 per MWh). The implied elasticity of electricity demand with respect to the wholesale price is -0.011 .

This estimate is remarkably close to the household IV estimate of -0.010 (Table D.1). Two very different approaches, a household-level instrumental-variables regression on real-time prices and an aggregate machine-learning model that never observes the electricity price, agree on a very inelastic short-run response of about -0.01 . This reinforces our conclusion that only a small part of the demand reduction during the crisis is attributable to prices, with the remainder reflecting non-price energy-saving behavior.

Table D.4 also reports the response expressed against each fuel price separately. Because the gas and CO₂ prices move together across the two scenarios, each of those figures attributes the entire joint demand response to a single fuel; they are not *ceteris-paribus* partial elasticities and should not be added. Similar to the case of wholesale electricity prices, we find small demand elasticities to natural gas and CO₂ prices, even during the crisis period.

Table D.4: Implied electricity-demand elasticity from the aggregate demand model

Price series	Real (crisis)	No-crisis	Implied elasticity
Natural gas (TTF, €/MWh)	90.5	17.9	−0.0069
CO ₂ (€/t)	77.7	57.7	−0.0215
Wholesale electricity (€/MWh)	144.2	68.4	−0.0106

Notes: The demand response is the gap between the two counterfactual predictions of aggregate demand (real minus no-crisis fuel prices), averaged over July 2021 to June 2023; it equals about 0.55% of average demand. Prices are demand-weighted averages of the simulated wholesale price under the *Factual* and *No-Energy-Crisis* scenarios. Each elasticity is the percentage demand response divided by the percentage price change. The gas and CO₂ rows attribute the full joint response to a single fuel and are not partial elasticities.

E Predicting counterfactual residential electricity demand

For predicting counterfactual residential demand, we train and tune separate ML models for each of the 47 peninsular provinces (Section 4.2). For each province, we evaluate a grid of 32 XGBoost configurations: maximum tree depth in $\{4, 6, 8, 10\}$, minimum child weight in $\{10, 50\}$, minimum split loss γ in $\{0, 1\}$, and learning rate η in $\{0.05, 0.1\}$, with row and column subsampling fixed at 0.8. The number of boosting rounds is not part of the grid but is chosen by early stopping (up to 3,000 rounds, stopping after 50 rounds without improvement). Cross-validation uses calendar-year folds within the training period (2018, 2019, 2020, and the first half of 2021), so that each fold holds out a full year of data and performance reflects genuine out-of-sample prediction across years.

Table E.1 reports, for each province, the selected configuration (the one minimizing cross-validated RMSE) and its out-of-sample performance, expressed as the CV RMSE relative to the province’s mean demand. The selected configurations are quite stable across provinces: all favor the lower learning rate ($\eta = 0.05$) and no split-loss penalty ($\gamma = 0$), most select the larger minimum child weight and moderate depths of 6 or 8, consistent with the smaller per-province sample sizes relative to the aggregate model. Normalized CV errors range from 4.6 to 8.6 percent of mean demand (median 6.2 percent). As expected, these are larger than the corresponding error of the aggregate national model (about 1.6 percent), since province-level demand is a noisier series.

Table E.2 complements these results with the models’ out-of-fold prediction errors over the training period and the estimated crisis effects on demand. Mean prediction errors are small and centered close to zero across provinces, indicating no systematic bias. The last two columns report the demand effects implied by the counterfactuals of Table 1: the difference between realized demand and the *No-Energy-Crisis* prediction (the total effect of the crisis on low voltage demand), and the difference between realized demand and the counterfactual without energy savings (the effect net of the price response, obtained with the elasticity described in Appendix D).

Table E.1: Hyperparameter tuning – province-level low voltage demand

Province	Max. depth	Min. child weight	η	Boosting rounds	CV RMSE (% of mean)
Araba/Álava	6	50	0.05	1,050	5.41
Albacete	8	50	0.05	798	5.71
Alicante	6	50	0.05	1,131	6.38
Almería	8	50	0.05	968	6.11
Ávila	8	50	0.05	592	7.72
Badajoz	6	50	0.05	1,035	7.72
Barcelona	6	10	0.05	651	6.34
Burgos	6	50	0.05	1,290	5.02
Cáceres	4	50	0.05	1,799	6.72
Cádiz	6	50	0.05	1,592	6.79
Castellón	6	50	0.05	1,267	7.46
Ciudad Real	6	10	0.05	1,417	5.68
Córdoba	8	50	0.05	1,207	8.05
A Coruña	6	50	0.05	1,559	5.92
Cuenca	6	10	0.05	741	6.17
Girona	10	50	0.05	388	6.53
Granada	8	50	0.05	612	5.99
Guadalajara	6	50	0.05	1,010	6.39
Gipuzkoa	8	50	0.05	736	6.33
Huelva	6	50	0.05	1,303	7.37
Huesca	10	50	0.05	570	6.47
Jaén	10	50	0.05	728	6.92
León	6	50	0.05	1,174	4.56
Lleida	6	10	0.05	540	6.16
La Rioja	6	50	0.05	1,489	4.86
Lugo	6	50	0.05	1,102	4.57
Madrid	6	50	0.05	849	6.57
Málaga	8	50	0.05	878	5.90
Murcia	6	50	0.05	1,051	6.51
Navarra	6	10	0.05	623	5.15
Ourense	4	10	0.05	1,588	5.41
Asturias	8	50	0.05	449	5.73
Palencia	6	50	0.05	1,125	4.61
Pontevedra	8	50	0.05	1,282	5.50
Salamanca	6	10	0.05	850	5.24
Cantabria	8	50	0.05	728	5.23
Segovia	8	50	0.05	1,018	6.07
Sevilla	6	50	0.05	1,263	8.62
Soria	8	50	0.05	1,132	6.18
Tarragona	10	50	0.05	929	6.67
Teruel	8	50	0.05	751	6.88
Toledo	6	50	0.05	1,244	6.57
Valencia	6	50	0.05	948	7.72
Valladolid	6	50	0.05	1,026	4.80
Bizkaia	6	10	0.05	865	6.37
Zamora	6	50	0.05	985	5.38
Zaragoza	8	50	0.05	794	5.97

Notes: Selected XGBoost configuration per province (minimum cross-validated RMSE over a grid of 32 configurations; row and column subsampling fixed at 0.8). The minimum split loss γ was tuned over $\{0, 1\}$ but never selected, with all provinces choosing $\gamma = 0$, and is omitted from the table. Boosting rounds chosen by early stopping. Cross-validation uses calendar-year folds within the 2018 to June 2021 training period. The last column reports the cross-validated RMSE as a percentage of the province's mean hourly demand.

Table E.2: Prediction accuracy and estimated crisis effects, by province

Province	Prediction error (%)	Demand effect (%), realized vs. counterfactual	
		<i>No-Energy-Crisis</i>	Crisis w/o Savings Int.
Araba/Álava	-0.91 (0.39)	-3.54 (0.46)	-3.04 (0.47)
Albacete	-1.07 (0.45)	-5.49 (0.50)	-5.01 (0.52)
Alicante	-0.42 (0.51)	-1.86 (0.58)	-1.36 (0.60)
Almería	-1.02 (0.44)	-3.31 (0.54)	-2.82 (0.55)
Ávila	0.23 (0.75)	-5.22 (0.65)	-4.73 (0.67)
Badajoz	-1.26 (0.65)	-3.81 (0.61)	-3.31 (0.62)
Barcelona	-0.78 (0.45)	-4.84 (0.48)	-4.36 (0.49)
Burgos	-0.26 (0.43)	-3.49 (0.52)	-3.00 (0.54)
Cáceres	-1.00 (0.64)	-4.68 (0.91)	-4.19 (0.92)
Cádiz	-0.48 (0.50)	-5.22 (0.50)	-4.73 (0.52)
Castellón	-0.70 (0.58)	-4.51 (0.72)	-4.02 (0.74)
Ciudad Real	-1.08 (0.52)	-8.69 (0.69)	-8.22 (0.71)
Córdoba	-1.36 (0.66)	-5.55 (0.45)	-5.07 (0.46)
A Coruña	-0.47 (0.51)	-5.78 (0.79)	-5.30 (0.80)
Cuenca	-0.83 (0.54)	-4.92 (0.68)	-4.44 (0.68)
Girona	-0.33 (0.51)	-3.95 (0.72)	-3.46 (0.74)
Granada	-0.66 (0.40)	-3.76 (0.46)	-3.27 (0.47)
Guadalajara	-0.91 (0.65)	-3.52 (0.65)	-3.02 (0.67)
Gipuzkoa	-0.66 (0.52)	-4.71 (0.74)	-4.22 (0.76)
Huelva	-0.75 (0.60)	-5.27 (0.52)	-4.79 (0.53)
Huesca	-0.75 (0.52)	-2.95 (1.22)	-2.45 (1.24)
Jaén	-0.98 (0.63)	-8.54 (0.66)	-8.07 (0.67)
León	-0.13 (0.38)	-4.44 (0.54)	-3.95 (0.56)
Lleida	-0.67 (0.49)	-3.92 (0.86)	-3.43 (0.88)
La Rioja	-0.23 (0.42)	-3.12 (0.52)	-2.63 (0.53)
Lugo	-0.29 (0.35)	-5.47 (0.53)	-4.98 (0.54)
Madrid	-0.84 (0.56)	-2.63 (0.54)	-2.13 (0.55)
Málaga	-0.41 (0.54)	-0.93 (0.33)	-0.43 (0.34)
Murcia	-1.09 (0.45)	-2.68 (0.33)	-2.19 (0.35)
Navarra	-0.43 (0.28)	-3.47 (0.39)	-2.98 (0.42)
Ourense	-0.57 (0.42)	-6.15 (0.74)	-5.67 (0.75)
Asturias	-0.14 (0.49)	-6.19 (0.74)	-5.71 (0.76)
Palencia	-0.12 (0.36)	-3.84 (0.58)	-3.35 (0.60)
Pontevedra	-1.04 (0.50)	-6.06 (0.77)	-5.58 (0.78)
Salamanca	-0.62 (0.43)	-5.25 (0.57)	-4.77 (0.58)
Cantabria	-0.60 (0.38)	-5.34 (0.56)	-4.85 (0.57)
Segovia	-0.28 (0.55)	-6.36 (0.67)	-5.88 (0.69)
Sevilla	-1.24 (0.68)	-4.06 (0.53)	-3.56 (0.54)
Soria	-0.33 (0.57)	-5.70 (0.71)	-5.21 (0.72)
Tarragona	-0.63 (0.49)	-4.83 (0.69)	-4.34 (0.70)
Teruel	-0.78 (0.67)	-5.82 (0.82)	-5.34 (0.83)
Toledo	-1.62 (0.58)	-5.47 (0.61)	-4.99 (0.62)
Valencia	-1.36 (0.49)	-3.96 (0.57)	-3.46 (0.58)
Valladolid	-0.27 (0.34)	-3.81 (0.44)	-3.32 (0.46)
Bizkaia	-1.16 (0.52)	-4.98 (0.57)	-4.49 (0.58)
Zamora	-0.21 (0.48)	-6.11 (0.74)	-5.63 (0.75)
Zaragoza	-0.53 (0.35)	-5.02 (0.43)	-4.54 (0.44)

Notes: The first column reports the mean out-of-fold prediction error over the training period (2018 to June 2021), excluding the COVID-19 lockdown periods and storm Filomena. The last two columns report the mean percent difference between realized demand and each counterfactual over July 2021 to June 2023: the *No-Energy-Crisis* prediction, and the counterfactual without the energy-savings response (the prediction adjusted by the price elasticity of Appendix D). Standard errors, clustered by month, in parentheses.

F Additional Results

F.1 Timing of the crisis-induced demand reduction

The counterfactual predictions described in Section 3.2 deliver an hourly estimate of the demand reduction induced by the crisis: the difference between predicted demand under the *No-Energy-Crisis* scenario (the “low CO₂ and gas prices” series in Figure 5) and realized demand, expressed as a percentage of the counterfactual. In this appendix, we characterize when this reduction took place. Because the short-run price response is very small (Appendix D), the patterns documented here essentially reflect the non-price energy-saving response.


Over the crisis window (July 2021 to June 2023), the reduction averages 3.9% of counterfactual demand, or 19.2 TWh in total. To describe its timing, we classify each hour of the sample into the three blocks: peak hours (10:00–14:00 and 18:00–22:00), night hours (00:00–08:00), and flat hours in between. We complement this with a workday indicator (weekdays excluding national holidays) and with bins of the national population-weighted daily mean temperature.

Figure F.1 plots the average reduction by hour of day, separately for each half-year. Three patterns stand out. First, the reduction is far from uniform within the day: over the full window it averages 5.0% during peak hours, twice the 2.5% observed during night hours, peaking at around 7% in the midday hours and falling to essentially zero in the early morning (6–7am). Second, the entire hourly profile shifts upward as the crisis unfolds: midday reductions grow from about 3% in the second half of 2021 to about 12% in the first half of 2023. We argue that part of the demand reduction reflects the adoption of solar self-consumption, particularly in 2023 (more details in Appendix F.6). Solar self-consumption reduces grid purchases, without necessarily reducing energy use. Since this capacity generates in daylight hours, it contributes to the midday concentration of the estimated reduction. However, since the reduction is not confined to daylight hours, we argue that a substantial share of the gap reflects genuine energy-saving behavior. The estimated savings are thus best understood as a combination of behavioral savings, present throughout the day from the onset of the crisis, and solar self-consumption, which grows over time and is concentrated in midday hours.

Table F.1 summarizes these patterns in a regression of the hourly percentage reduction on the covariates above, with standard errors clustered by day. Column (2) adds year-month fixed effects, so that coefficients capture within-month variation. The peak- and flat-hour coefficients confirm the intraday tilt of the reduction. Two further results emerge. The reduction is systematically larger on weekends and holidays than on workdays (by about 4.8 percentage points), and larger on mild days than on very cold or very hot days. Both patterns suggest that the savings response operated primarily

through changes in household consumption behavior: they were willing to forego heating and cooling when dispensable. Reductions did not accrue exclusively through the mandates imposed on public buildings and commercial establishments, which would otherwise predict reductions concentrated on workdays.

Figure F.1: Demand reduction by hour of day



Notes: This figure shows the average estimated demand reduction by hour of day, separately for each half-year of the crisis window (July 2021 to June 2023). The reduction is computed as the difference between the counterfactual demand prediction under *No-Energy-Crisis* prices and realized demand, as a percentage of the counterfactual. Shaded areas mark peak periods (10:00–14:00 and 18:00–22:00).

Table F.1: Demand reduction and covariates

Dependent Variable: Model:	Demand gap (% of counterfactual)	
	(1)	(2)
<i>Variables</i>		
Constant	8.32*** (0.437)	
Flat hours (8-10h, 14-18h, 22-24h)	1.94*** (0.115)	1.94*** (0.115)
Peak hours (10-14h, 18-22h)	2.51*** (0.119)	2.51*** (0.119)
Workday	-4.77*** (0.384)	-4.76*** (0.303)
Temperature <10°C	-3.90*** (0.483)	-2.35*** (0.710)
Temperature 10-15°C	-1.01** (0.499)	-0.493 (0.517)
Temperature 20-25°C	-4.39*** (0.464)	-1.28*** (0.361)
Temperature >25°C	-4.51*** (0.638)	-1.48** (0.726)
<i>Fixed-effects</i>		
Year-month		Yes
<i>Fit statistics</i>		
Observations	17,494	17,494
R ²	0.2594	0.4747
Within R ²		0.2439

Notes: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. The dependent variable is the hourly demand reduction – counterfactual demand under *No-Energy-Crisis* prices minus realized demand – as a percentage of the counterfactual, over July 2021 to June 2023. Reference categories: night hours (00:00–08:00), weekend or national holiday, and daily mean temperature between 15 and 20°C. Temperature is the national population-weighted daily mean. Standard errors are clustered by day.

F.2 Wholesale Market Simulations

Table F.2 assesses the fit of the simulation model by reporting ratios of simulated to realized outcomes under the Factual scenario, over the full sample period and by half-year, and under both competitive and strategic bidding. The first panel compares demand-weighted average prices. Under competitive bidding, the simulated average price is 99% of the realized one over the full period. The fit is tightest before the price intervention; afterwards, the simulation underpredicts prices in the second half of 2022 and overpredicts them in the first half of 2023, deviations in opposite directions that largely offset over the full window. The second panel compares dispatch by technology. Nuclear, hydro, wind, and solar output are matched almost exactly. Thermal generation, the technology most sensitive to our cost and bidding assumptions, is matched within 3% over the full period, although with larger deviations in the post-intervention subperiods. The main discrepancy concerns imports, which the simulation overstates, while exports are matched within 2%. Overall, the competitive simulations track realized prices and dispatch slightly more closely than the strategic ones, which is why we retain competitive bidding as our baseline. Moreover, since our estimated effects are computed as differences between simulated scenarios, rather than between simulations and realized outcomes, any remaining level biases largely net out of the comparisons.

The counterfactual demand series that enter the simulations carry uncertainty (Appendix B.6). Tables F.3, F.4, and F.5 reproduce the main simulation tables, adding the effects obtained under the low- and high-demand series, showing that the estimated effects are robust to this uncertainty.

Table F.2: Validating the Simulation Model: Ratios to Realized Outcomes

Measure	Bidding	Full period	2021 (Q3-4)	2022 (Q1-2)	2022 (Q3-4)	2023 (Q1-2)
<i>Prices</i>						
Demand-weighted average prices	Competitive	0.99	1.02	0.97	0.91	1.12
	Strategic	1.02	1.04	1.00	0.92	1.19
<i>Generation by technology</i>						
Thermal generation	Competitive	0.97	0.95	0.99	1.09	0.82
	Strategic	0.97	0.94	0.97	1.09	0.81
Nuclear	Competitive	0.99	0.99	1.00	0.98	0.99
	Strategic	0.98	0.99	0.98	0.98	0.96
Hydro	Competitive	0.99	1.00	1.00	1.00	0.98
	Strategic	0.99	1.00	1.00	1.00	0.98
Wind	Competitive	1.00	1.00	1.00	1.00	0.99
	Strategic	0.99	0.99	0.99	1.00	0.98
Solar PV	Competitive	1.00	1.00	1.00	1.00	1.00
	Strategic	1.00	1.00	1.00	1.00	0.99
Imports	Competitive	1.15	1.36	1.16	0.46	1.26
	Strategic	1.31	1.45	1.37	0.52	1.54
Exports	Competitive	0.98	0.91	1.02	1.11	0.83
	Strategic	0.96	0.89	1.00	1.11	0.80

Notes: This table reports ratios of simulated outcomes to realized outcomes under the factual scenario. Competitive and Strategic refer to the factual simulations under competitive and strategic bidding, respectively.

Table F.3: Demand-weighted Average Prices and Implied Effects – Competitive Scenario, with Demand Bounds (€/MWh)

	2021 (Q3-4)	2022 (Q1-2)	2022 (Q3-4)	2023 (Q1-2)
<i>Panel A. Price levels</i>				
Factual	169	201	120	102
No Energy Crisis	63 [63; 64]	66 [66; 67]	71 [70; 71]	65 [64; 66]
Business-as-usual	171 [169; 172]	220 [217; 222]	256 [255; 257]	114 [111; 116]
Savings intervention only	169	215	252	105
Price intervention only	171 [169; 172]	207 [204; 209]	122 [122; 122]	110 [108; 112]
<i>Panel B. Effects</i>				
Energy-crisis effect	106 [107; 105]	135 [136; 135]	50 [50; 50]	37 [39; 36]
Full-intervention effect	-1 [1; -3]	-19 [-15; -20]	-136 [-135; -137]	-11 [-9; -14]
Price-intervention effect	—	-13	-132	-3
Savings-intervention effect	-1 [1; -3]	-6 [-2; -7]	-2 [-1; -2]	-7 [-5; -9]

Notes: Panel A reports the simulated demand-weighted average prices in the Spanish wholesale market, assuming competitive bidding, during each period under each scenario. The price intervention took place in 2022. Panel B reports the implied effects, computed as the difference between the factual scenario and the relevant counterfactual. Values in square brackets report low- and high-demand bounds. Cells without square brackets (including Factual, Savings intervention only, and Price-intervention effect) use demand series without additional low-/high-demand simulations.

Table F.4: Policy Intervention and Crisis Effects on Earnings, Costs, and Profits – Competitive Scenario, with Demand Bounds (M€, July 2021-July 2023)

	Earnings	Costs	Profits
Energy-Crisis Effect	46,953 [47,843; 46,069]	17,773 [18,344; 17,216]	29,186 [29,501; 28,863]
Full-Intervention Effect	-15,322 [-12,547; -17,822]	323 [2,131; -1,429]	-15,644 [-14,682; -16,388]
Price-Intervention Effect	-10,574	2,368	-12,948
Savings-Intervention Effect	-4,080 [-1,373; -6,485]	-1,969 [-147; -3,728]	-2,105 [-1,224; -2,746]

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. Earnings and Costs include the gas compensation in the *Factual* and *Price-Intervention-Only* scenarios. Values in square brackets report low- and high-demand bounds. The Price-Intervention Effect row has no bounds because the *Savings-Intervention-Only* scenario has no additional low-/high-demand simulations.

Table F.5: Policy Intervention and Crisis Effects on Imports, Exports, Thermal Generation, and Demand – Competitive Scenario, with Demand Bounds (% , July 2021-July 2023)

	Imp-Exp	Thermal Generation	CO2 Emissions	Demand
Energy-Crisis Effect	-28.5 [-19.7; -36.0]	-7.6 [-2.1; -13.0]	-2.9 [2.0; -7.7]	-4.1 [-1.8; -6.2]
Full-Intervention Effect	-59.4 [-52.2; -65.3]	-0.7 [5.1; -6.3]	-4.8 [0.5; -9.9]	-3.5 [-1.2; -5.6]
Price-Intervention Effect	-37.4	7.2	2.7	0.3
Savings-Intervention Effect	-18.7 [-12.3; -23.9]	-7.5 [-1.6; -13.3]	-7.1 [-1.7; -12.3]	-3.5 [-1.2; -5.6]

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. Values in square brackets report low- and high-demand bounds. The Price-Intervention Effect row has no bounds because the *Savings-Intervention-Only* scenario has no additional low-/high-demand simulations.

Tables F.6, F.7, and F.8 replicate the simulation results under strategic bidding, modeled following de Frutos and Fabra (2012), including the low- and high-demand bounds. The estimated price, distributional, and quantity effects are very similar to those obtained under competitive bidding, confirming that the main conclusions do not hinge on the assumption of competitive behavior.

Table F.6: Demand-weighted Average Prices and Implied Effects – Strategic Scenario, with Demand Bounds (€/MWh)

	2021 (Q3-4)	2022 (Q1-2)	2022 (Q3-4)	2023 (Q1-2)
<i>Panel A. Price levels</i>				
Factual	173	208	121	108
No Energy Crisis	65 [64; 65]	67 [67; 68]	71 [71; 71]	67 [66; 68]
Business-as-usual	175 [173; 176]	225 [223; 226]	258 [257; 258]	121 [119; 122]
Savings intervention only	173	221	256	113
Price intervention only	175 [173; 176]	211 [210; 213]	122 [122; 122]	115 [114; 117]
<i>Panel B. Effects</i>				
Energy-crisis effect	109 [110; 108]	141 [141; 140]	51 [51; 50]	41 [42; 40]
Full-intervention effect	-2 [0; -3]	-17 [-15; -18]	-137 [-136; -137]	-13 [-11; -14]
Price-intervention effect	—	-13	-135	-4
Savings-intervention effect	-2 [0; -3]	-4 [-2; -5]	-1 [-1; -1]	-7 [-5; -8]

Notes: Panel A reports the simulated demand-weighted average prices in the Spanish wholesale market, assuming strategic bidding, during each period under each scenario. The price intervention took place in 2022. Panel B reports the implied effects, computed as the difference between the factual scenario and the relevant counterfactual. Values in square brackets report low- and high-demand bounds. Cells without square brackets (including Factual, Savings intervention only, and Price-intervention effect) use demand series without additional low-/high-demand simulations.

Table F.7: Policy Intervention and Crisis Effects on Earnings, Costs, and Profits – Strategic Scenario, with Demand Bounds (M€, July 2021-July 2023)

	Earnings	Costs	Profits
Energy-Crisis Effect	48,200 [49,128; 47,389]	17,739 [18,320; 17,168]	30,468 [30,810; 30,232]
Full-Intervention Effect	-15,360 [-12,802; -17,735]	329 [2,147; -1,439]	-15,688 [-14,952; -16,290]
Price-Intervention Effect	-11,042	2,375	-13,422
Savings-Intervention Effect	-3,805 [-1,323; -6,128]	-1,988 [-152; -3,755]	-1,810 [-1,169; -2,361]

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, under strategic bidding. Earnings and Costs include the gas compensation in the *Factual* and *Price-Intervention-Only* scenarios. Values in square brackets report low- and high-demand bounds. The Price-Intervention Effect row has no bounds because the *Savings-Intervention-Only* scenario has no additional low-/high-demand simulations.


Table F.8: Policy Intervention and Crisis Effects on Imports, Exports, Thermal Generation, and Demand – Strategic Scenario, with Demand Bounds (% , July 2021-July 2023)

	Imp-Exp	Thermal Generation	CO2 Emissions	Demand
Energy-Crisis Effect	-30.7 [-24.1; -36.4]	-7.7 [-2.0; -13.3]	-3.2 [1.9; -8.2]	-4.1 [-1.8; -6.2]
Full-Intervention Effect	-66.7 [-60.7; -71.2]	-0.9 [5.1; -6.6]	-4.8 [0.6; -10.1]	-3.5 [-1.2; -5.6]
Price-Intervention Effect	-46.7	7.3	2.9	0.3
Savings-Intervention Effect	-16.4 [-11.5; -20.6]	-7.8 [-1.7; -13.7]	-7.3 [-1.8; -12.6]	-3.5 [-1.2; -5.6]

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, under strategic bidding. Note that Demand is equal to imports-exports plus in-home generation. Values in square brackets report low- and high-demand bounds. The Price-Intervention Effect row has no bounds because the *Savings-Intervention-Only* scenario has no additional low-/high-demand simulations.

F.3 Retail Market

Figure F.2: 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.

F.4 Descriptive statistics across zip codes

Table F.9: Determinants of household electricity consumption

	Electricity Cons. – $\log(KWh/year.HH)$	
	(1)	(2)
$\log(Income)$	0.073*** (0.013)	0.406*** (0.016)
<i>Household Size</i>	0.445*** (0.017)	0.284*** (0.016)
<i>Electric Heating</i>		0.422*** (0.025)
<i>Air Conditioning</i>		0.290*** (0.015)
<i>Degree of urbanization</i> (ref.: Cities)		
Towns and Semi-dense areas		0.246*** (0.008)
Rural areas		0.297*** (0.011)
(Intercept)	6.211*** (0.149)	3.063*** (0.167)
R ²	0.077	0.252
Adj. R ²	0.077	0.252
Num. obs.	8565	8564

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 F.10: Determinants of the average electricity price paid by households

	Average Price per unit [EUR/MWh]				
	(1) $\log(p)$	(2) $\log(\Delta p)$	(3) $\log(-\Delta p)$	(4) $\log(-\Delta p)$	(5) $\log(-\Delta p)$
$\log(\text{Income})$	0.006*** (0.001)	0.004*** (0.001)	-0.023*** (0.003)	0.001 (0.003)	-0.102*** (0.004)
<i>Household Size</i>	-0.004*** (0.001)	-0.005*** (0.001)	-0.002 (0.003)	-0.000 (0.003)	-0.018*** (0.004)
<i>Electric Heating</i>	0.007*** (0.001)	-0.003 (0.002)	0.078*** (0.004)	0.065*** (0.005)	0.011* (0.006)
<i>Air Conditioning</i>	0.012*** (0.001)	0.021*** (0.001)	0.074*** (0.002)	0.156*** (0.003)	-0.173*** (0.004)
<i>Degree of urbanization</i> (ref.: Cities)					
Towns and Semi-dense areas	0.002*** (0.000)	0.006*** (0.001)	0.010*** (0.001)	0.013*** (0.002)	-0.025*** (0.002)
Rural areas	-0.002*** (0.000)	0.003*** (0.001)	0.033*** (0.002)	0.040*** (0.002)	-0.021*** (0.003)
(Intercept)	5.471*** (0.006)	4.764*** (0.012)	3.769*** (0.027)	3.244*** (0.034)	2.878*** (0.042)
R ²	0.131	0.061	0.330	0.406	0.305
Adj. R ²	0.130	0.060	0.330	0.405	0.304
Num. obs.	8564	8564	8564	8564	8564


Notes: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. 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 F.11: Household 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 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 F.3: Energy-crisis effect on household bills by province




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



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 F.5: Average electricity consumption and income per household



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

Table F.12: 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
$\log(Income)$	0.416*** (0.016)	0.599*** (0.074)	0.571*** (0.054)	0.485*** (0.025)	0.430*** (0.078)	0.295*** (0.049)	0.388*** (0.032)
<i>Household Size</i>	0.279*** (0.017)	0.302*** (0.058)	0.437*** (0.055)	0.267*** (0.027)	0.057 (0.063)	0.304*** (0.045)	0.308*** (0.035)
<i>Electric Heating</i>	0.413*** (0.026)	0.602*** (0.163)	0.418*** (0.157)	0.713*** (0.096)	−0.824*** (0.233)	0.654*** (0.133)	0.371*** (0.058)
<i>Air Conditioning</i>	0.338*** (0.015)	−0.260 (0.225)	−0.349** (0.139)	0.144*** (0.036)	0.521*** (0.138)	0.026 (0.052)	0.258*** (0.067)
<i>Degree of urbanization</i> (ref.: Cities)							
Towns and Semi-dense areas	0.252*** (0.008)	0.146*** (0.032)	0.041* (0.025)	0.260*** (0.015)	1.661*** (0.062)	0.182*** (0.024)	0.270*** (0.015)
Rural areas	0.299*** (0.011)	0.135*** (0.040)	0.277*** (0.035)	0.339*** (0.019)	1.518*** (0.058)	0.262*** (0.027)	0.258*** (0.024)
(Intercept)	0.808*** (0.169)	−0.897 (0.771)	−1.062* (0.554)	0.228 (0.269)	0.016 (0.765)	2.090*** (0.517)	1.073*** (0.354)
R ²	0.263	0.101	0.149	0.328	0.451	0.188	0.218
Adj. R ²	0.262	0.096	0.145	0.326	0.448	0.184	0.215
Num. obs.	8564	967	1147	2053	1276	1152	1969

Notes: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. 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 F.13: Full-intervention effect and household characteristics within each cluster

$Y = \log(-\Delta \text{Bill})$	Full Sample	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
$\log(\text{Income})$	0.362*** (0.017)	0.568*** (0.074)	0.431*** (0.054)	0.411*** (0.028)	0.403*** (0.083)	0.089* (0.052)	0.421*** (0.035)
<i>Household Size</i>	0.279*** (0.018)	0.343*** (0.057)	0.380*** (0.055)	0.226*** (0.029)	0.019 (0.067)	0.279*** (0.048)	0.395*** (0.038)
<i>Electric Heating</i>	0.523*** (0.028)	0.472*** (0.162)	0.689*** (0.158)	1.313*** (0.106)	-0.845*** (0.247)	0.524*** (0.141)	-0.083 (0.062)
<i>Air Conditioning</i>	0.206*** (0.016)	-0.835*** (0.224)	-0.493*** (0.139)	-0.060 (0.039)	0.416*** (0.147)	0.016 (0.056)	0.757*** (0.073)
<i>Degree of urbanization</i> (ref.: Cities)							
Towns and Semi-dense areas	0.258*** (0.009)	0.106*** (0.032)	0.051** (0.025)	0.277*** (0.017)	1.707*** (0.065)	0.258*** (0.026)	0.284*** (0.016)
Rural areas	0.333*** (0.012)	0.136*** (0.040)	0.301*** (0.036)	0.358*** (0.021)	1.618*** (0.062)	0.325*** (0.029)	0.305*** (0.026)
(Intercept)	0.472*** (0.182)	-1.369* (0.768)	-0.432 (0.557)	0.141 (0.297)	-0.490 (0.810)	3.202*** (0.550)	-0.488 (0.382)
R ²	0.228	0.103	0.134	0.312	0.461	0.263	0.220
Adj. R ²	0.228	0.098	0.130	0.310	0.458	0.259	0.218
Num. obs.	8564	967	1147	2053	1276	1152	1969

Notes: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Each observation corresponds to one zip code. The dependent variable $\log(-\Delta \text{Bill})$ 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 F.14: 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
$\log(Income)$	0.317*** (0.021)	0.583*** (0.076)	0.328*** (0.055)	0.299*** (0.034)	0.169** (0.086)	-0.088 (0.067)	0.521*** (0.045)
<i>Household Size</i>	0.276*** (0.022)	0.338*** (0.059)	0.302*** (0.056)	0.153*** (0.036)	0.136* (0.069)	0.264*** (0.062)	0.553*** (0.049)
<i>Electric Heating</i>	0.464*** (0.035)	0.583*** (0.167)	1.330*** (0.161)	2.080*** (0.129)	-0.758*** (0.257)	0.676*** (0.181)	-1.092*** (0.081)
<i>Air Conditioning</i>	-0.138*** (0.021)	-1.440*** (0.231)	-0.824*** (0.141)	-0.519*** (0.048)	0.038 (0.152)	-0.306*** (0.071)	1.625*** (0.094)
<i>Degree of urbanization</i> (ref.: Cities)							
Towns and Semi-dense areas	0.247*** (0.011)	0.054 (0.033)	0.053** (0.025)	0.257*** (0.020)	1.509*** (0.068)	0.331*** (0.033)	0.304*** (0.020)
Rural areas	0.324*** (0.015)	0.069* (0.041)	0.259*** (0.036)	0.295*** (0.026)	1.391*** (0.064)	0.379*** (0.037)	0.383*** (0.033)
(Intercept)	-0.030 (0.230)	-2.318*** (0.791)	-0.286 (0.566)	0.361 (0.361)	0.529 (0.842)	3.944*** (0.706)	-3.208*** (0.495)
R ²	0.119	0.127	0.138	0.283	0.406	0.294	0.249
Adj. R ²	0.119	0.121	0.134	0.280	0.403	0.291	0.247
Num. obs.	8564	967	1147	2053	1276	1152	1969

Notes: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. 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 F.15: Price-intervention effect and household characteristics within each cluster

$Y = \log(-\Delta \text{Bill})$	Full Sample	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
$\log(\text{Income})$	0.407*** (0.016)	0.559*** (0.073)	0.519*** (0.055)	0.478*** (0.026)	0.509*** (0.083)	0.222*** (0.048)	0.387*** (0.034)
<i>Household Size</i>	0.284*** (0.017)	0.353*** (0.057)	0.444*** (0.056)	0.266*** (0.028)	-0.025 (0.067)	0.293*** (0.045)	0.309*** (0.037)
<i>Electric Heating</i>	0.487*** (0.027)	0.385** (0.160)	0.283* (0.161)	0.900*** (0.100)	-0.895*** (0.249)	0.487*** (0.132)	0.381*** (0.061)
<i>Air Conditioning</i>	0.447*** (0.016)	-0.457** (0.222)	-0.305** (0.141)	0.214*** (0.037)	0.590*** (0.148)	0.205*** (0.052)	0.358*** (0.071)
<i>Degree of urbanization</i> (ref.: Cities)							
Towns and Semi-dense areas	0.260*** (0.008)	0.143*** (0.032)	0.049* (0.025)	0.282*** (0.016)	1.775*** (0.066)	0.201*** (0.024)	0.271*** (0.015)
Rural areas	0.337*** (0.011)	0.187*** (0.039)	0.327*** (0.036)	0.382*** (0.020)	1.698*** (0.062)	0.286*** (0.027)	0.263*** (0.025)
(Intercept)	-0.601*** (0.175)	-2.047*** (0.759)	-2.026*** (0.565)	-1.182*** (0.279)	-1.987** (0.816)	1.286** (0.514)	-0.355 (0.373)
R ²	0.313	0.102	0.156	0.351	0.470	0.219	0.210
Adj. R ²	0.312	0.096	0.151	0.349	0.467	0.215	0.208
Num. obs.	8564	967	1147	2053	1276	1152	1969

Notes: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Each observation corresponds to one zip code. The dependent variable $\log(-\Delta \text{Bill})$ 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.

F.5 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 F.6.

Table F.16: Summary statistics per zip code-level income quintile

Quintile	Q1	Q2	Q3	Q4	Q5	Full Sample
Electricity Cons. [kWh/year.HH]	3315.0	3233.0	3098.0	3049.0	3384.0	3216.0
Avg. Elec Price [c€/kWh]	25.0	25.0	25.0	25.0	25.1	25.0
Income [kEUR/year.HH]	24.7	30.9	35.5	40.0	50.3	36.3
Household Size	2.5	2.6	2.5	2.5	2.5	2.5
Electric Heating [%]	36.6	33.0	23.6	22.8	20.4	27.3
Air Conditioning [%]	54.6	51.1	46.3	48.5	50.7	50.2
Cities [%]	18.9	44.2	63.2	71.4	79.5	55.5
Towns & semi-dense areas [%]	37.5	41.3	28.8	23.0	18.1	29.7
Rural [%]	43.6	14.5	8.0	5.5	2.4	14.8
Min Temp. [°C]	4.7	4.5	3.3	3.0	2.7	3.7
Max Temp. [°C]	33.4	32.1	31.3	31.6	31.8	32.0
Number of ZIP codes	4748	1607	909	742	558	8564
Total Households (thousands)	3244	3244	3218	3272	3246	16224

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




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


effects of the crisis and the interventions appear more contrasted across quintiles. The 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.

F.6 Bill effects and solar self-consumption

Solar self-consumption capacity grew rapidly in Spain during the crisis period, and households with self-consumption are partially insulated from retail price increases. We assess whether our estimated bill effects are related to this expansion. Cumulative installed self-consumption capacity is published by REE (2024) at the monthly level for each autonomous community (CCAA), which is the finest aggregation available. We aggregate our bill estimates accordingly: hourly zip code consumption is summed to the province level, monthly bills under each scenario are computed as in Section 4.2 (applying the province-level counterfactual demand adjustments), and bills are then summed to the CCAA level.⁴⁴ Monthly bill effects per household are defined as the same scenario differences as in the main analysis.

Figure F.7 shows the evolution of capacity per household by CCAA: penetration was low and flat before the crisis, accelerated after mid-2022, and grew fastest in the southern and inland regions, with Castilla-La Mancha and Murcia reaching more than four times the capacity of the northern coastal CCAAs by the end of 2023.

Figure F.7: Solar self-consumption capacity by autonomous community



Notes: Cumulative installed solar self-consumption capacity per household, by autonomous community (CCAA), from January 2019 to December 2023. The regional breakdown is published from early 2020 onward; earlier months are shown as zero. Dashed vertical lines mark the onset of the energy crisis and the start of the Iberian Mechanism. The legend orders CCAAs by capacity at the end of 2023.

⁴⁴The capacity data covers all consumer segments, not only households, and is therefore a proxy for residential self-consumption penetration. Since cumulative installed capacity cannot decline, we correct transient reporting spikes in the raw series by capping each month at the minimum of all subsequent values. This only affects two observations: Castilla-La Mancha in May and June 2023.

We regress each monthly bill effect on the cumulative solar self-consumption capacity per household,

$$\Delta Bill_{c,m} = \beta Solar_{c,m} + \alpha_c + \gamma_m + \varepsilon_{c,m}, \quad (\text{F.1})$$

where $Solar_{c,m}$ is measured in watts (W) per household, α_c and γ_m are CCAA and month fixed effects, respectively. Observations are weighted by the number of households, and standard errors are clustered by CCAA. The CCAA fixed effects absorb the fixed exposure of each region (climate, equipment, consumption profiles), and the month fixed effects absorb the common evolution of prices, so β is identified from differences in the timing of capacity growth across CCAAs. Still, these estimates may be considered as mostly descriptive: capacity growth is not randomly assigned, and the identifying variation cannot rule out region-specific trends correlated with installation timing.

Table F.17 reports the results. For completeness, the odd-numbered columns include month fixed effects only: in this cross-region comparison the association between capacity and the crisis effect is positive, reflecting that the regions facing the largest bill increases also expanded self-consumption capacity the most. Indeed, across CCAAs, cumulative capacity additions over the sample period correlate strongly with the average crisis effect on bills (correlation coefficient of 0.81): sunnier regions both face higher returns to solar investment and experienced larger bill increases. Adding CCAA fixed effects (even-numbered columns) absorbs these cross-region differences, so that the negative coefficient is identified from within-region capacity growth.

Results from column (2) of Table F.17 reveal that the Regions where self-consumption capacity expanded faster experienced smaller energy-crisis effects on household bills. However, the timing of the expansion limited its mitigating role. In the months where the crisis effect on bills was strongest (December 2021 to March 2022 and July to August 2022, averaging €58 per household-month), only about 21 W per household had been added since the start of the crisis, as most of the capacity arrived in late 2022 and 2023. A back-of-the-envelope calculation using our estimate implies a mitigation of roughly €1.2 per household-month in those peak months, about 2% of the effect that would have occurred without the added capacity. Self-consumption thus expanded largely as a delayed response to the crisis, shielding households mainly once the worst of the price shock had passed.

Table F.17: Monthly bill effects and solar self-consumption capacity

Dependent Variables:	Energy crisis		Full intervention		Price intervention		Savings intervention	
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Variables</i>								
Solar self-consumption capacity (W/HH)	0.0420*** (0.0127)	-0.0542*** (0.0128)	-0.0275*** (0.0052)	0.0003 (0.0053)	-0.0167*** (0.0040)	-0.0004 (0.0025)	-0.0077 (0.0046)	0.0018 (0.0021)
<i>Fixed-effects</i>								
Month	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
CCAA		Yes		Yes		Yes		Yes
<i>Fit statistics</i>								
Observations	360	360	360	360	360	360	360	360
R ²	0.9583	0.9820	0.9479	0.9545	0.9603	0.9638	0.6263	0.7486
Within R ²	0.1468	0.1510	0.1122	0.0000	0.0726	0.0000	0.0565	0.0013

Notes: Each observation is a CCAA $\times$ month (15 peninsular CCAAs, July 2021 to June 2023). The dependent variables are the monthly bill effects per household (€/household/month); the regressor is the cumulative installed solar self-consumption capacity in W per household. Weighted OLS (household weights). Odd-numbered columns include month fixed effects only, so that identification comes from cross-region differences in capacity levels; even-numbered columns add CCAA fixed effects, so that identification comes from within-region capacity growth. Standard errors, clustered by CCAA, in parentheses. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.