Household-Level Responses to the European Energy Crisis

Lassi Ahlvik, Tuomas Kaariaho, Matti Liski, Iivo Vehviläinen*

Abstract

Cost-of-living shocks can have large and unequal effects, but causal evidence on households' adaptive responses is scarce. Using Finnish administrative microdata from the 2022 energy crisis, we exploit variation in electricity contract expirations to identify behavioral responses in electricity use, labor earnings, financial distress, and residual consumption and savings. We find notable heterogeneity in responses that shape the final incidence: high-income households primarily reduced electricity consumption, while low-income households reduced consumption less, responding instead through increased labor supply, more payment defaults, and cutbacks in spending and savings. Longer anticipation softens but does not eliminate these impacts.

Keywords: cost-of-living, energy crisis, financial shock, heterogeneity, inequality
JEL codes: D12, H23, J22, Q41

---

*Ahlvik (lassi.ahlvik@helsinki.fi): University of Helsinki, Helsinki GSE and University of Stavanger Business School. Kaariaho (tuomas.kaariaho@helsinki.fi): University of Helsinki, Helsinki GSE. Liski (matti.liski@aalto.fi) and Vehviläinen (iivo.vehvilainen@aalto.fi): Department of Economics of the Aalto University, and Helsinki GSE. This work has been carried out as a part of the VATT Datahuone Environment and Energy group (https://vattdatahuone.fi/). We thank Moritz Drupp, Martin Holm, Kristiina Huttunen, Sini Numminen, Olli Kauppi, Edson Severnini and participants at the CEPR Climate Change and the Environment Programme workshop in Paris, seminar participants at Umeå University, Turku University, Toulouse School of Economics, University of Copenhagen, Statistics Norway, Utrecht University, Wageningen University, the Finnish Competition and Consumer Authority, Transatlantic Seminar on Environmental and Energy Economics, and the Ministry of the Environment, as well as participants in EEA, EAERE and NAERE conferences. This project has been funded by Research Council of Finland COP27 (project 357065), the Finnish Cultural Foundation and the European Union (ERC GRETA, 101116659). Views and opinions expressed are those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency.

1 Introduction

How do households respond to cost-of-living shocks? Inflation in essential goods – energy, gasoline, and food – has revived concerns that rising costs place a disproportionate burden on low-income households. The distributional consequences depend not only on exposure to higher prices, but also on how households adjust consumption, labor earnings, and savings. There is limited causal evidence on such multi-margin responses across the income distribution. Identifying which margins households use, and how responses differ by income, is central to designing effective policy responses. These questions are especially important for energy prices because energy is a necessity, accounts for a larger share of expenditures among low-income households, and energy prices are directly shaped by climate policies such as carbon pricing.

This paper provides causal evidence on household adjustments to a large and unexpected cost-of-living shock created by the European energy crisis of 2022. We trace household responses across multiple margins using rich Finnish administrative microdata – metered electricity use, earnings, and court-reported payment defaults – together with an imputed residual measure capturing shifts in net spending and savings. Our identification strategy exploits predetermined expiration dates of fixed-price electricity contracts. Households whose contracts expired during the crisis faced an immediate price shock of up to eightfold, whereas others remained temporarily insulated by ongoing fixed-price contracts. This setting lends itself to a stacked difference-in-differences design, comparing treated and matched control households before and after contract expiration.

We find that households react to the price shock through multiple adjustment margins. On average, the energy crisis reduces electricity use by 8.8% – a moderate but statistically precise response consistent with electricity being a necessity. Average labor earnings rise by 0.6% at the intensive margin, with no detectable change in labor force participation. Financial distress increases: treated households experience a 10% increase in payment defaults relative to baseline, while residual consumption and savings decline by 2.8%.

These average effects mask large heterogeneity in household responses. High-income households respond mainly by reducing electricity consumption, with no detectable changes in earnings or defaults. At the other end of the distribution, low-income households reduce electricity use the least but increase labor earnings at the intensive margin by several times the average effect, as those already in the labor force work more to offset higher bills. Households with no labor earnings show no entry into work. These groups experience the sharpest increase in payment defaults and the steepest decline in residual consumption and savings. Middle-income households lie between these extremes: we find an increase in labor supply for this group, and also a rise in payment defaults, especially among those with a high debt-to-income ratio, consistent with the idea of a liquidity-constrained middle class (Kaplan et al., 2014).

These heterogeneous adjustments – reductions in electricity consumption among high-income households, income smoothing among lower-income working households, and financial strain among those at the bottom of the income distribution or with limited liquidity – show that aggregate averages mask large and distributionally uneven responses and impacts at the household level.

The variation in contract expiration times in our setting also allows us to study anticipation effects: households whose contracts expire later in the crisis have time to adjust their behavior prior to the price shock. By observing behavior in the months leading up to expiration, we find that households are forward-looking and begin reducing electricity use several months in advance. Like the treatment effect, the anticipation effect also grows with income, indicating heterogeneous capabilities to prepare. A longer anticipation period softens the financial impact and mitigates negative effects on residual consumption and savings but does not eliminate the increase in payment defaults.

Our paper connects to the literature on household responses to income and financial shocks, including negative earnings shocks in the labor market (Baker and Yannelis, 2017; Gelman et al., 2020), positive income shocks from lottery winnings (Cesarini et al., 2017; Picchio et al., 2018; Fagereng et al., 2021; Golosov et al., 2024), and temporary income transfers, such as tax rebates (Johnson et al., 2006; Parker et al., 2013, see also Boehm et al., 2025). Cost-of-living shocks, the focus of this paper, differ in that they alter relative prices as well as net incomes, and thereby substitution effects become important alongside income effects.

Research in energy and environmental economics has documented heterogeneity in exposure to policy-driven cost shocks both across incomes – low-income households are hit harder due to spending patterns – and within income groups (Grainger and Kolstad, 2010; Fischer and Pizer, 2019; Cronin et al., 2019; Pizer and Sexton, 2019; Douenne, 2020; Levinson and Silva, 2022). These studies typically measure static incidence, while we focus on active household responses: adjustments in energy use, labor supply, and financial decisions, that differ systematically across households. Our contribution is to provide empirical evidence on these broad behavioral responses, implying that inequality from price shocks, policy-induced or not, reflects not only initial exposure but also differences in households' capacity to adjust.

This value added also holds in comparison to studies that use quasi-experimental settings to examine demand responses to higher energy prices (e.g., Ito, 2014; Deryugina et al., 2020). A growing number of papers use microdata to estimate heterogeneous demand elasticities (e.g., Alberini et al., 2020; Rubin and Auffhammer, 2024; Çürük et al., 2025), consumption responses (Gelman et al., 2023), and the distributional impacts of price and policy changes (e.g., Burger et al., 2020; Hahn and Metcalfe, 2021; Fetzer et al., 2024; Reguant et al., 2025; Fabra et al., 2025; Levell et al., 2025). Our contribution is to document how households adjust jointly across multiple margins, and how these adjustments vary across the income distribution.

These income-dependent adjustment patterns matter for the design of relief as part of energy and climate policies. When governments subsidize energy bills, reduce energy taxes, or provide targeted transfers, the distributional and efficiency consequences depend on which households consume the subsidized goods and on how households choose to adjust when prices change. Classic public-finance results imply that if households across the income distribution respond similarly to relative price changes, redistribution can be achieved through the tax-and-transfer system, making commodity-specific relief redundant; if responses differ, commodity-specific interventions may be needed.[1] Our setting provides evidence relevant for this distinction by estimating income-dependent adjustment responses.

2 Background, data and empirical strategy

2.1 Background

On February 24, 2022, Russia invaded Ukraine. At the time, Russia supplied about 40 percent of the European Union's natural gas. Russia subsequently restricted exports to Europe, causing the 2022 European energy crisis during which energy prices reached record highs in the fall of that year. Figure 1a plots electricity prices in Finland for variable and two-year fixed-term contracts.[2] Fixed-term prices peaked in September 2022, while variable prices peaked in December. Both declined in spring 2023.

To illustrate our identification strategy, consider two households: one that signed a two-year fixed-price contract in October 2020, and another that did so in October 2021. Two years later, in October 2022, the first household's contract expired, forcing it to enter a new fixed- or variable-price contract amid the crisis and thereby generating within-household variation in prices. In contrast, the second household's contract maintained a low electricity price (below 10 c/kWh) throughout the crisis. We classify all households whose contracts expired in the middle of the crisis (between August 2022 and January 2023) as treated, while households whose contracts extended beyond our study period (June 2023) serve as controls. The crisis generated an average cost shock of about €90 per month, or roughly €1,000 over the eleven-month post-crisis study

[1] According to the Atkinson–Stiglitz theorem (1976), if marginal rates of substitution across goods are homogeneous among consumers, redistribution is in principle best achieved through income taxation or targeted transfers. If the conditions of the theorem do not hold, additional policy instruments beyond the tax-and-transfer system may be needed to address equity–efficiency trade-offs, as emphasized recently by Ahlvik et al. (2024); Ferey et al. (2024); Hellwig and Werquin (2025); Doligalski et al. (2025).

[2] The Finnish retail electricity market is fully deregulated, and households freely choose contracts. Before the crisis, 54% of households held fixed-price, fixed-term contracts, 9% held real-time contracts tied to hourly day-ahead spot prices, and 37% held variable tariffs tied to average spot levels (e.g., monthly or quarterly) but not hourly (Finnish Energy Authority, 2021). Contracts cover only electricity, as retailers do not sell other fuels. Residential gas heating is virtually absent; oil heating is more common but usually purchased from distributors without long-term contracts.

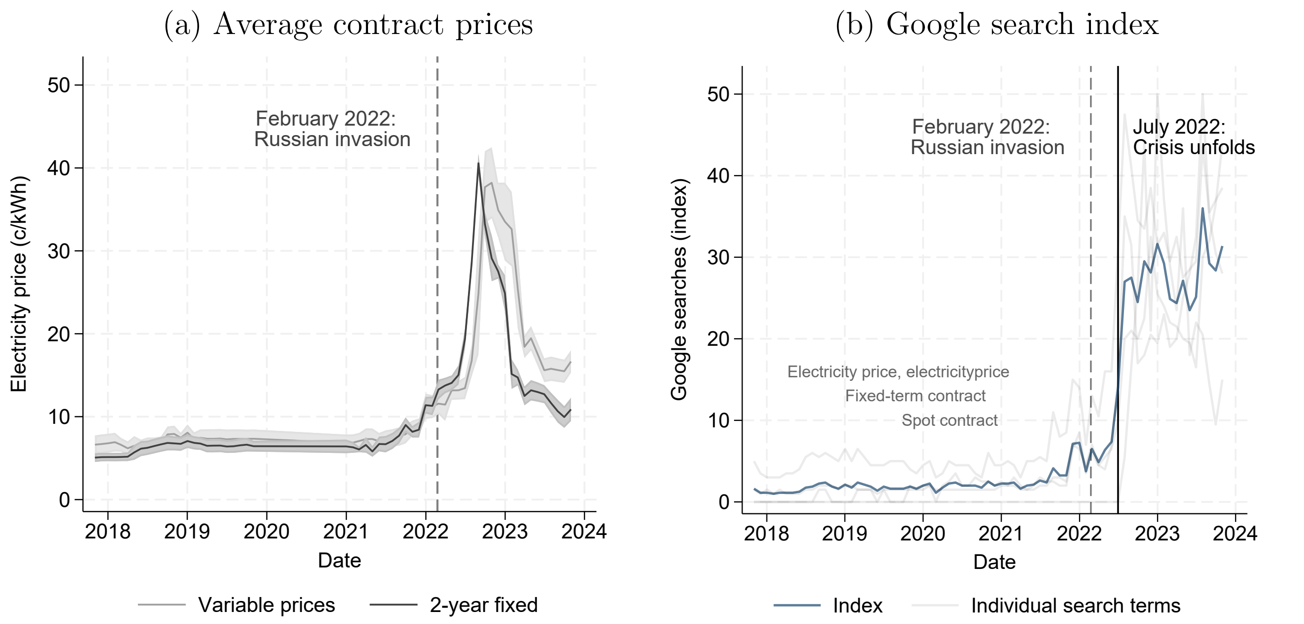

Figure 1: Development of the energy crisis in Finland

Notes: Figure (a) shows the development of average contract prices for variable-price contracts (gray line) and two-year fixed-term contracts (black line), with shaded areas indicating the cross-household range. Data are from the Finnish Energy Authority. Figure (b) gray lines show Google trends search results for selected keywords: (i) electricity price (*sähkönhinta*, *sähkön hinta*), (ii) fixed-term contract (*määräaikainen sähkösopimus*) and (iii) spot contract (*pörssisähkö*). The blue line is the simple average over these individual keywords.

period, with substantial variation across households. The shock peaked at an average of €170 in December.[3]

The crisis was unexpected, and electricity prices, including futures contracts, provided no advance warning. Figure 1b shows Google searches for individual electricity-related keywords (gray lines) and a combined index (blue line). The effect of the Russian invasion on electricity prices began to attract media and public attention in July 2022, as reflected in the sharp increase in the index. We treat July as the point at which the crisis unfolded and became widely salient.[4]

2.2 Data

Our initial dataset comprises monthly data on all households in Finland from March 2022 to June 2023. We consider the margins of household response using the following stylized budget

[3]These figures are derived by multiplying the average price difference between treated and control households, 14.2 c/kWh, by the average quantity consumed by the control group: 0.142 €/kWh × 651 kWh ≈ €90. For December, the price difference is 18.2 c/kWh, multiplied by the control group's electricity use of 939 kWh, yielding 0.182 €/kWh × 939 kWh ≈ €170.

[4]During the crisis, heightened awareness and information campaigns contributed to energy reductions, while broader economic conditions were also affected. Our estimated effects should be interpreted in this context. The identification strategy isolates price effects from non-monetary influences and broader economic conditions common to both treatment and control groups. See Prest (2020), Dertwinkel-Kalt et al. (2024) and Behr et al. (2025) on the role of non-price factors in energy savings.

constraint:

$$\underbrace{\text{Electricity use}}_{(1)} \times \text{Price} \; - \; \underbrace{\text{Defaulted payments}}_{(2)} \; - \; \underbrace{\text{Income}}_{(3)} \; + \; \underbrace{\text{Consumption} + \text{Savings}}_{(4)} = 0,$$

where each term represents a dependent variable in our analysis: (1) household-level electricity use from the transmission system operator Fingrid; (2) the total amount of unpaid debts and bills, sourced from the Legal Register Centre; (3) monthly gross income, including labor earnings, pensions, and government-paid benefits, as well as measures of the intensive margin (income per worker) and the share of working household members from the Finnish Tax Administration; and (4) consumption plus savings, which we compute mechanically as income net of electricity costs, defaulted payments, and taxes. For the heterogeneity analysis, we use pre-crisis household-level background information. These datasets are described below and a more detailed description of the final dataset is provided in the Supplemental Appendix, including summary statistics in Appendix Table B.1.

Household-level background data Household-level background data are sourced from Statistics Finland's Preliminary Population Statistics and the Basic Data and Income Data modules. The dataset provides annual individual-level information for 2021 – one year before the crisis – on employment status, disposable income, total debt, and family composition. We construct monthly household-level observations by grouping individuals who co-reside in the same dwelling and share the same electricity meter. Households that move during the sample period are excluded from the analysis. We divide households into income quintiles based on household disposable income per consumption unit in 2021, the year preceding the crisis. The upper quintile cutoffs for annual disposable incomes are €18,300 (1st quintile), €24,100 (2nd quintile), €29,900 (3rd quintile) and €38,100 (4th quintile).

Electricity consumption and contracts We obtain electricity consumption and contract data from Fingrid Datahub, the regulated, centralized information exchange system for Finland's electricity retail market. The dataset includes monthly household electricity consumption, contract type (fixed- or variable-price), contract start and end dates, and pseudonymized meter-, individual- and retailer-identifiers. We link electricity contracts to household-level records based on the contract holder identifier. Our data do not contain contract-level price information.[5] Instead, we use monthly average electricity prices by contract type and user type from the Finnish Energy Authority; these prices include both energy and network (distribution and transmission) charges.

[5]Because contract-level prices are not observed, our analysis does not directly capture potential heterogeneity in the pass-through of cost shocks across consumer characteristics (see Byrne et al., 2022; Gravert, 2024).

Earnings, pensions and benefits We obtain income and benefit data from the Finnish Tax Administration and calculate household-level monthly labor earnings, pensions, and other benefits by summing across all household members. Labor earnings are defined as wage income from paid employment only; they exclude all transfers and benefits and take the value zero when an individual is not working. During the crisis, the Finnish government implemented support schemes. Our analysis incorporates the most important ones when building the residual term: a temporary reduction in the value-added tax (VAT) on electricity, which is reflected in prices, and an electricity bill compensation (HE 324/2022), which was based on past electricity use and therefore did not affect the effective prices faced by households (November 2022–February 2023).[6] The compensation covered 50% of monthly electricity expenses exceeding €90, with a cap of €700 per month. Eligibility was restricted to individuals paying an electricity price above 10c/kWh. Compensation was paid automatically as reductions on electricity bills starting in March 2023. We calculate compensation amounts using electricity consumption data and the estimated prices.

Defaults Data on payment defaults are obtained from the Finnish Legal Register Centre. The dataset includes individual-level defaults enforced by district courts, which we link to households. The default data also include the subject matter of the default, the principal, interest, fees, penalties, collection expenses, the summons date, and the court decision date. We define the month of the summons date as the default time, which typically occurs one or two months after a missed payment. Our analysis includes defaults for all purposes (e.g., services, goods, debt), because we anticipate that a higher utility bill may lead to other types of defaults, for instance, via payday loans. We calculate the total amount of accumulated defaults for each household and construct an indicator variable that equals one if the household has at least one default recorded before a given month during the study period, and zero otherwise.

Residual consumption and savings Households may respond to the higher prices by switching to a more affordable consumption basket, using their savings, borrowing from relatives, and so forth.[7] These channels are unobservable to us but, as explained above, we can impute a "residual" term capturing consumption and savings decisions, which can be defined formally as: after-tax earnings (labor and pension) + after-tax benefits (including electricity support) - electricity bill

[6]The law was formally submitted to Parliament on February 1, 2023, and it entered into force on March 6, 2023. The other support programs, an income tax deduction for electricity expenses (HE 204/2022) and a subsidy through the social insurance institution (HE 234/2022) were small in size and underutilized; they totalled to only €3 million compared to €440 million through the VAT reduction and electricity support (YLE, 2023). We consider these other programs to be minor and exclude them from the analysis.

[7]For instance, Doremus et al. (2022) find that extreme weather events lead low-income households to cut back on food spending, while Steen et al. (2021) use household-level transaction data from a large Norwegian grocery chain to show that households shift to cheaper stores, bulk products, and sales in response to a regional income shock.

+ defaults. We calculate monthly residual consumption and savings for each household and treat it as a dependent variable in the analysis.

2.3 Empirical strategy

To identify causal effects, we use a difference-in-differences design comparing treated households whose contracts expire during the crisis (from August 2022 to January 2023) to a control group of households whose contracts end after our study period (after June 2023). The main sample includes only households with a two-year fixed-price contract at the beginning of our sample (March 2022). Since the energy crisis was unforeseen when these contracts were signed in 2020–2021, the timing of contract expirations during the crisis is exogenous to the crisis itself. However, the group of households that sign their contracts in a given month may not be random. For example, students tend to move during certain months, advertising campaigns may target households in specific areas, and certain types of contracts may gain popularity over time in specific groups. These factors can generate systematic differences between households whose contracts expire during the crisis and those whose contracts expire later. To address this, we use a matched difference-in-differences design, matching each household with its nearest neighbor based on heating technology, electricity use, default indicators, benefits, earnings, and earnings indicators in the five months preceding when the crisis unfolded (for the period March–July 2022).[8] The final matched dataset includes 231,491 households.

To estimate the effect of the energy crisis, we run the following stacked difference-in-differences model:

$$Y_{iht} = \underbrace{\beta_{tre}(D_i \times ContractEnds_{ht})}_{\text{Treatment effect}} + \underbrace{\beta_{ant}(D_i \times Anticipation_{ht})}_{\text{Anticipation effect}} + \alpha_{ih} + \gamma_{ht} + \epsilon_{iht}, \qquad (1)$$

where Y_{iht} is the dependent variable for household i in month $t = 1, ..., 16$ (1 = March 2022, 16 = June 2023), and $h = 6, ..., 11$ indexes cohorts whose contracts expire during the crisis months (from 6 = August 2022 to 11 = January 2023). Each matched control is assigned the same cohort label h as its treated match within that cohort.

The main treatment effect is captured by the interaction $D_i \times ContractEnds_{ht}$, where $D_i = 1$ for treated households and $D_i = 0$ for matched controls, and $ContractEnds_{ht} \equiv \mathbb{1}[t \geq h]$ equals one in and after the month when the treated households' contracts end. Coefficient β_{tre} therefore measures the average treatment effect of contract expiration (on the treated). Households whose contracts end later in the crisis know their expiration dates and may react in advance to the

[8] This approach follows Goldschmidt and Schmieder (2017) and Adams-Prassl et al. (2024). In the Supplemental Appendix we discuss the matching in more detail, demonstrate its performance in Table B.1, Figure B.6 and show results under an alternative matching approach in B.7.

anticipated price increase, for example by installing energy-saving appliances or building a financial buffer. Such anticipation could bias estimation if not accounted for (Malani and Reif, 2015). Based on the timing in Figure 1b, we conclude that the crisis became salient after July 2022. To capture potential pre-expiration responses, $Anticipation_{ht} \equiv \mathbb{1}[6 \leq t < h]$ equals one from August 2022 onward and for all periods before contract expiration. The interaction $D_i \times Anticipation_{ht}$ is one only for treated households in this window, and β_{ant} measures the corresponding anticipation effect.[9] There is no anticipation for the first cohort ($h = 6$), and the anticipation period is longest for the last cohort ($h = 11$).

The cohort-by-month fixed effects, γ_{ht}, absorb unobserved shocks shared by cohort h in calendar month t, such as information campaigns or broader economic conditions. Household–cohort fixed effects α_{ih} assign a separate effect to each household–cohort pair, capturing any unobservable time-invariant factors. Households are matched to controls cohort-by-cohort, and the same household may serve as a control for multiple treated units. The inclusion of household-cohort fixed effects allows the same household to serve as a control in multiple cohorts without contaminating comparisons. Importantly, this specification compares treated households only to matched control households as in Cengiz et al. (2019), thereby avoiding the potential bias identified in the recent literature (Goodman-Bacon, 2021; Roth et al., 2023). All standard errors are clustered at the household level.

For the dependent variables, we use the logarithms of electricity use, earnings per worker, and residual consumption and savings, while the default indicator and the share of workers per household enter the specification in levels. To capture both intensive and extensive margins of earnings, we also analyze treatment effects on transformed outcomes of the form $\mathbb{1}[Y > y]$ for different earnings thresholds y (see Chen and Roth, 2024).

To study dynamic treatment effects, we estimate models of the form:

$$Y_{iht} = \underbrace{\sum_{\tau=1}^{4} \beta_\tau (D_i \times I_\tau)}_{\text{Before crisis}} + \underbrace{\sum_{\tau=6}^{h-1} \beta_\tau^{ant} (D_i \times I_\tau)}_{\text{Anticipation effect}} + \underbrace{\sum_{\tau=h}^{16} \beta_\tau^{tre} (D_i \times I_\tau)}_{\text{Treatment effect}} + \alpha_{ih} + \gamma_{ht} + \epsilon_{iht}, \tag{2}$$

where t denotes the month in which the outcome is measured, and $I_\tau \equiv \mathbb{1}[t = \tau]$ is an indicator for month τ. We interact I_τ with treatment status D_i to obtain month-specific treatment effects. The first term selects pre-crisis months ($\tau < 5$, i.e. March–June 2022), with β_τ used to test parallel pretrends. We omit $\tau = 5$ (July 2022), so all results are shown relative to that month. The anticipation

[9] The control households have contracts expiring well into the future. In Appendix Figure B.8, we estimate a 'placebo' specification in which the households whose contracts expire after the crisis (June 2023) are defined as treated. We find no effects for this group, supporting the identifying assumption that the control group remains unaffected by the crisis.

effect β_τ^{ant} applies after the crisis but strictly before contract expiration ($6 \leq \tau \leq h-1$), while the treatment effect β_τ^{tre} applies beginning in the expiration month and thereafter ($\tau \geq h$).

3 Results

Figure 2 reports the estimated dynamic treatment effects for (a) electricity use, (b) payment defaults, (c) labor earnings, and (d) residual consumption and savings. For presentation, we pool cohorts by anticipation and treatment status, separately for early-treated ($h = 6, \ldots, 8$) and late-treated households ($h = 9, \ldots, 11$), when estimating equation (2). For the early group, we plot the treatment effects for all months, but do not show anticipation effects (left panel). For the late group, we plot the treatment effects and also anticipation effects for the first three months (right panel).[10] The figure provides graphical evidence consistent with common trends prior to the energy crisis (to the left of the vertical dashed line). Panel A of Table 1 summarizes the average treatment effects in numerical form.

To understand heterogeneity in adjustment channels, Figure 3 presents responses to contract expiration across income quintiles, based on household equivalent income. The figure aggregates all six treatment cohorts ($h = 6, ..., 11$) as in equation (1), but it estimates separate regressions for each subgroup. We show heterogeneity in both the treatment effect of contract expiration (brown bars), and for anticipation effect (light teal bars). Panel B of Table 1 presents heterogeneity in the results through four interaction terms: (i) household disposable income (in income quintiles), (ii) electricity use (below or above the income quintile median), (iii) liquidity constraints, proxied by their debt-to-income ratio (below or above the income quintile median), and (iv) the length of the anticipation period in months.

3.1 Electricity use

Figure 2a shows the effect on electricity use, the most direct adjustment margin when electricity prices increase. We find a clear effect for households whose contracts expire early in the crisis (left panel), where electricity use decreases by around 6-10% after households face a sharp rise in electricity prices. For households whose contracts end late (right panel), we find an anticipation effect after the crisis unfolds but before contract expiration. These anticipation effects grow over time, reducing electricity use by an additional 1–3%, or about one-fourth of the total effect.[11] In

[10] Separate event-study graphs for all cohorts are reported in Supplemental Appendix Figures B.1–B.5, including analysis for the worker share (extensive margin) in Figure B.4.

[11] In the Supplemental Appendix we provide more evidence that the anticipation effect is robust. Appendix Figure B.8 examines households with contracts expiring in June 2023 – eleven months after the crisis began – and shows no anticipatory behavior, consistent with adjustment occurring only when expiration is near and expected prices are high. Appendix Figure B.9 runs the analysis for households whose contracts end unexpectedly due to

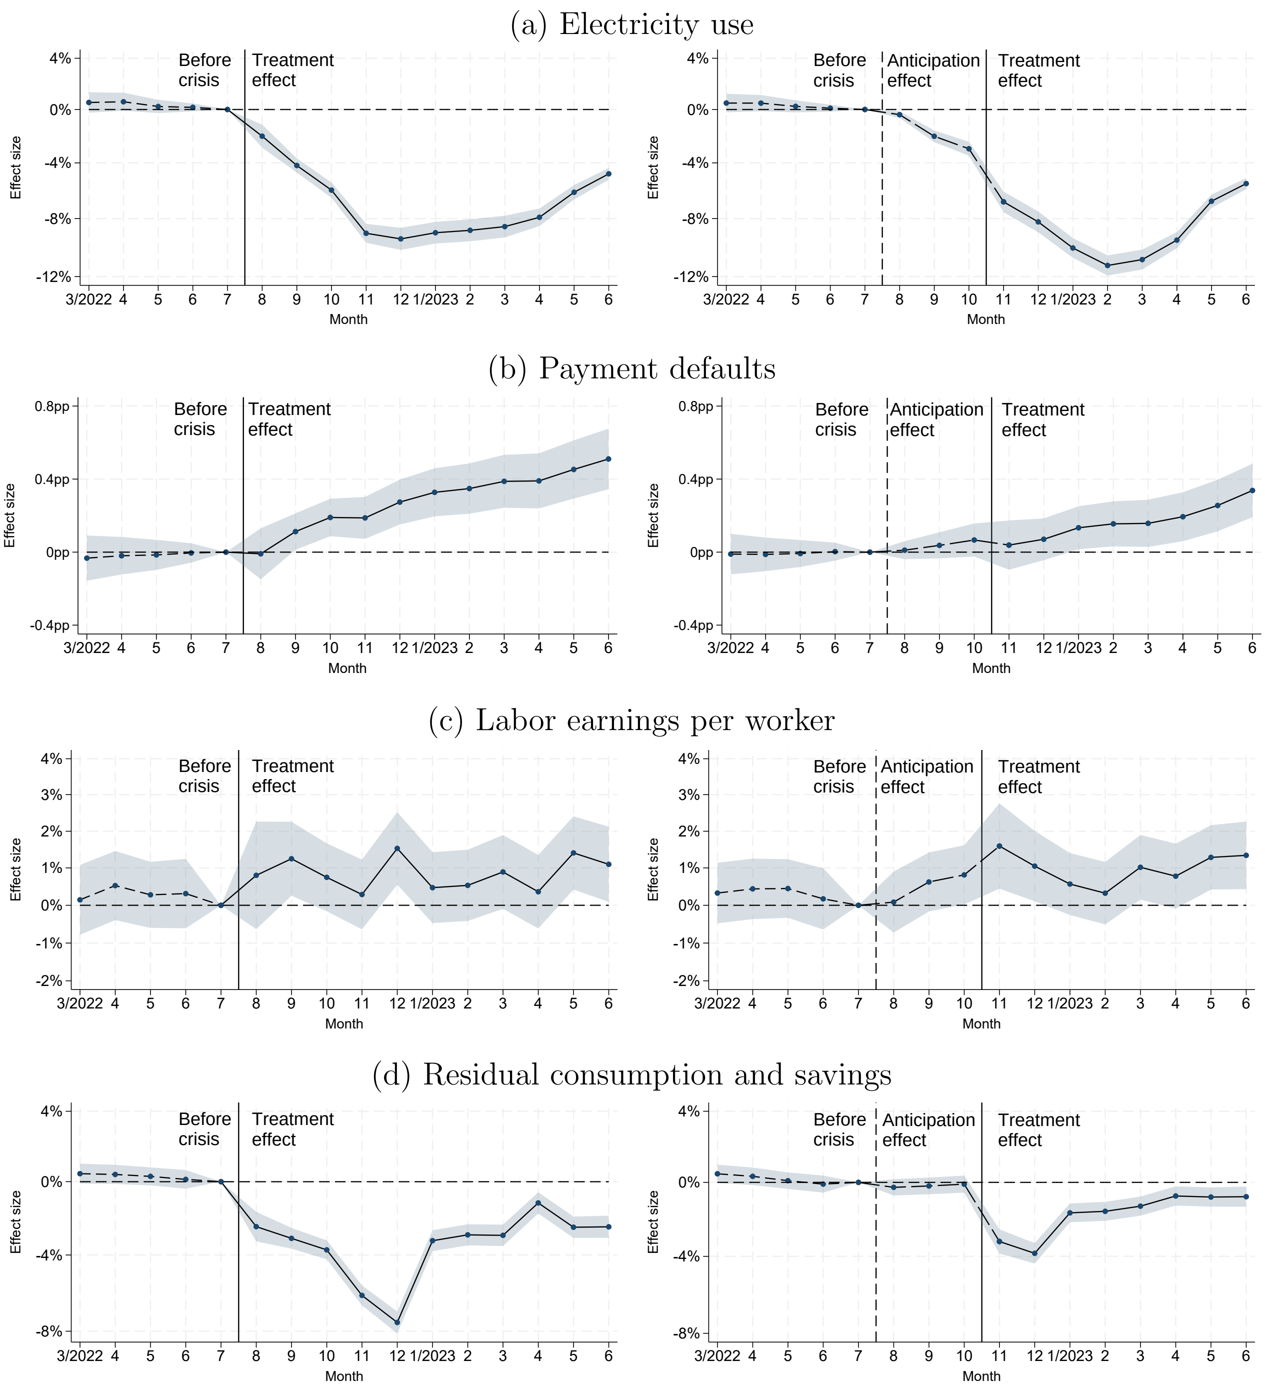

Figure 2: Timing of the treatment effects

Notes: Panels report dynamic coefficients from specification (2) by calendar month, relative to July 2022 (vertical dashed line). Left panel uses early-treated households ($h = 6, 7, 8$): August treatment coefficient corresponds to cohort $h = 6$, September coefficient corresponds to pooled cohorts $h = 6$ and $h = 7$, and from October onward coefficients pool cohorts $h = 6, 7, 8$; anticipation effects are not shown. Right panel uses late-treated households ($h = 9, 10, 11$): Anticipation effects for August-October are pooled over all cohorts $h = 9, 10, 11$; November treatment coefficient corresponds to cohort $h = 9$, December coefficient corresponds to pooled cohorts $h = 9$ and $h = 10$, from January onward coefficient corresponds to pooled cohorts $h = 9, 10, 11$. The vertical solid line marks the first month of contract expiration. Outcomes are: (a) log electricity use, (b) indicator for any payment default (effects in percentage points), (c) log labor earnings per worker, and (d) log residual consumption and savings. All estimates include household × cohort and cohort × month fixed effects. Shaded areas show 95% confidence intervals with standard errors clustered at the household level. Event study figures for individual cohorts are provided in the Supplemental Appendix.

the bankruptcy of their service provider. As expected, we find no ‘placebo anticipation effect’ for these households who are not aware of their contract expiration.

both panels, the effect grows gradually and fades toward the summer for two reasons: prices in variable-price contracts decrease over time, and the responsiveness is lower during the summer months.[12]

Column 1 of Table 1 reports the average treatment effects for the full sample and shows that a price increase lowers electricity use by 8.8%, with roughly a quarter of the reduction occurring prior to contract expiration due to anticipation effects. These effects may arise because forward-looking households make energy-saving investments, build cash buffers by cutting back on electricity expenses, or experiment with energy-saving measures before their contracts expire.

Figure 3a splits the treatment effects across household income quintiles. We find that high-income households are more responsive than low-income households. One potential mechanism is that low-income households are near their subsistence level of electricity use, making it more difficult to reduce consumption further. A similar pattern emerges for the anticipation effect: high-income households reduce electricity use more during the anticipation period, both in absolute terms and relative to their treatment effect. High-income households, being less liquidity-constrained, may be more able to make energy-saving investments. Column 1 of Table 1 shows that the effect becomes more negative by 0.7 percentage points per additional income quintile. The effect is also 4.8 percentage points larger among households with above-median electricity use, likely because such households are more likely to use electric heating and can reduce energy consumption by adjusting indoor temperatures.[13] Longer anticipation is associated with a larger reduction in electricity use, although this pattern may be partly confounded by seasonality.

3.2 Payment defaults

Figure 2b presents results for the cumulative probability of having at least one payment default in the study period. For households whose contracts end early (left panel), we find a treatment effect that kicks in after one month, likely due to the delay between billing and the recording of the default, and gradually increases after that. Column 2 of Table 1 reports an average treatment effect of 0.26 percentage points. In total, 2.5 percent of control households experienced at least one payment default on their bills in the crisis period, suggesting that the price shock increased the likelihood of default by roughly 10% on average. We find a similar pattern for households whose contracts expire later in the crisis: no anticipation effect but a positive treatment effect a few months after contract expiration (right panel). The anticipation period gives households time

[12]Appendix Figure B.10 disentangles these two channels by focusing on households that switch to fixed-price contracts, who face constant and persistently high prices during the post-expiration period. These households experience a small decrease in their response (in absolute value) toward the summer, around one-third from the winter level, indicating seasonal variation in price responsiveness as in Rubin and Auffhammer (2024).

[13]For reference, lowering the indoor temperature by 1°C reduces electricity use for households with electric heating by around 5% in winter (Motiva, 2015).

Table 1: Household-level responses to the energy crisis

	Electricity use	Payment defaults	Labor earnings per worker	Worker share	Residual consumption and savings
	(1)	(2)	(3)	(4)	(5)
Panel A: Average treatment effects					
Treatment	-0.0880***	0.0026***	0.0062***	-0.0002	-0.0277***
	(0.0013)	(0.0005)	(0.0017)	(0.0007)	(0.0012)
Anticipation	-0.0200***	0.0008*	0.0030	-0.0007	-0.0060***
	(0.0011)	(0.0004)	(0.0019)	(0.0006)	(0.0012)
N	4,114,627	4,116,737	2,343,273	4,116,737	3,977,310
Panel B: Heterogeneous effects					
Treatment	-0.0895***	0.0025***	0.0065***	-0.0003	-0.0269***
	(0.0012)	(0.0005)	(0.0018)	(0.0007)	(0.0012)
...x income quintile	-0.0066***	-0.0012***	-0.0055***	0.0014**	0.0093***
	(0.0009)	(0.0003)	(0.0015)	(0.0004)	(0.0009)
...x high-use	-0.0475***	0.0016	-0.0001	-0.0008	-0.0252***
	(0.0022)	(0.0008)	(0.0030)	(0.0011)	(0.0021)
...x high-debt	0.0043	0.0026**	0.0132***	0.0112***	0.0181***
	(0.0024)	(0.0008)	(0.0031)	(0.0012)	(0.0022)
...x months of anticipation	-0.0045***	-0.0005	0.0004	-0.0001	0.0070***
	(0.0006)	(0.0003)	(0.0009)	(0.0003)	(0.0006)
Anticipation	-0.0233***	0.0005	0.0031	-0.0008	-0.0023
	(0.0011)	(0.0004)	(0.0020)	(0.0006)	(0.0012)
N	4,091,561	4,093,659	2,329,908	4,093,659	3,954,981

Notes: The table reports regression coefficients from 2×5 separate regressions, 5 per panel. The dependent variables are the natural logarithm of electricity consumption (column 1), defaults indicator (column 2), logarithm of earnings per worker (column 3), share of household members with positive labor earnings (column 4), and logarithm of residual consumption and savings (column 5). In Panel A the first row estimates treatment effects, where *Treatment* is an indicator taking value one in or after the contract expiration month and zero otherwise and the second row presents the anticipation effects where indicator *Anticipation* takes value one between August 2022 and one month before contract expiration, and zero otherwise. Panel B presents heterogeneous treatment effects, interacting the treatment effect with demeaned background characteristics: income quintile, high electricity use (an indicator taking value one if electricity consumption is above median for a given income quintile), high-debt (an indicator taking value one if debt-to-income ratio is above quintile median) and months of anticipation defined as the difference between contract ending month and August 2022. All columns include household-cohort (ih) and cohort-month (ht) fixed effects. In Panel B all results also control for post-interactions (for example, *ContractEnds* × *high-use*). Standard errors, clustered by households, are shown in parentheses, with * p<.05, ** p<.01, and *** p<.001.

to make the needed investments and plan accordingly but, nevertheless, it does not eliminate their risk of payment default. According to column 2 of Table 1, each additional month of anticipation period lowers the post-expiration default risk by 0.05 percentage points, but this effect is not statistically significant.

Figure 3b presents the effect on payment defaults by income group. As expected, we find larger point estimates for the lowest-income households, which are more exposed to price increases and more likely to be credit constrained. These findings are consistent with Barreca et al. (2022), who document that higher energy use increases the likelihood of disconnection among low-income households in California. Surprisingly, the increase in defaults is observed among both low- and middle-income households. One possible explanation is that middle-income households may hold

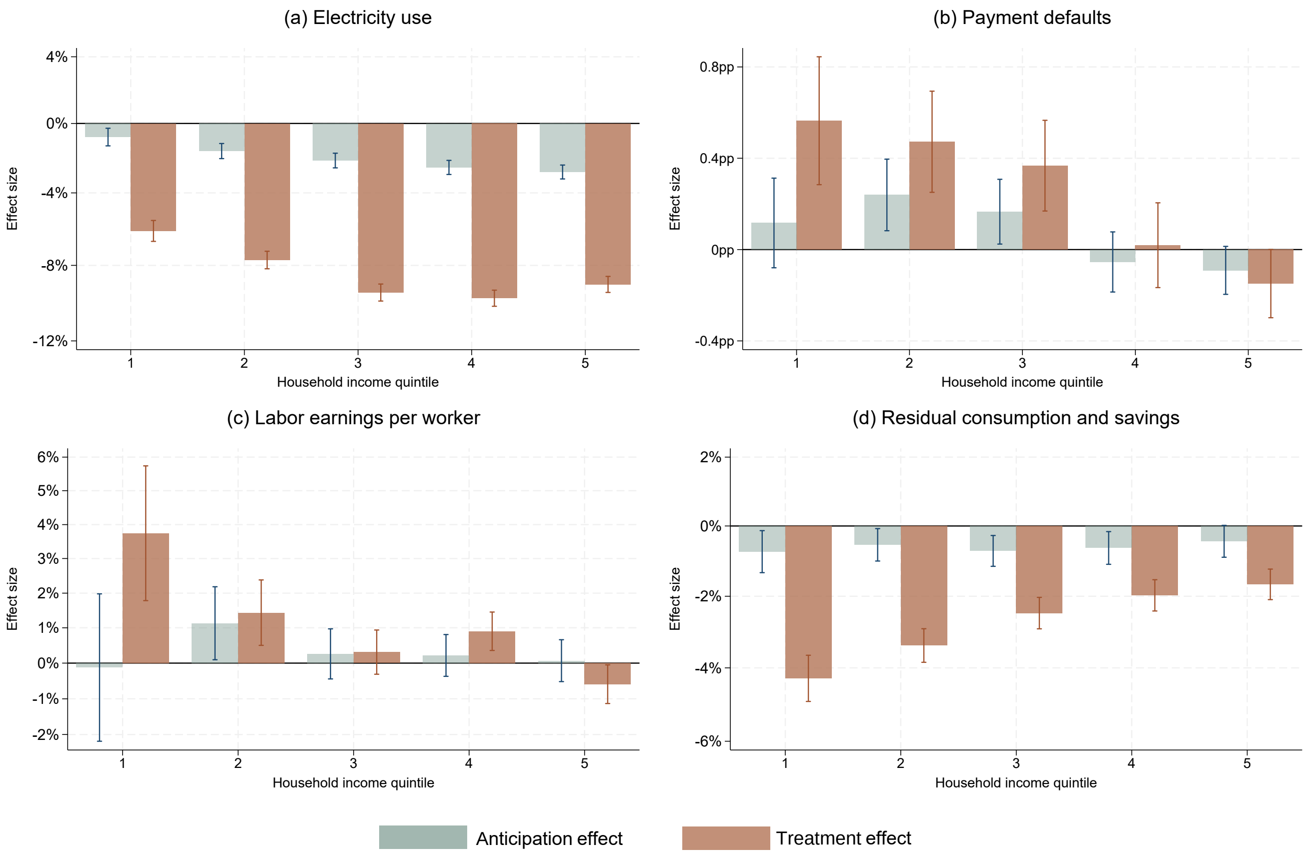

Figure 3: Heterogeneous responses to the energy crisis

Notes: The figure reports coefficients from the stacked regression equation (1) for different subsets of the data. The five categories on the x-axis estimate separate regressions for each household income quintile. Light teal bars represent the anticipation effect (indicator taking value one between August 2022 and one month before contract expiration), and brown bars are the treatment effects (indicator taking value one in or after contract expiration). All regressions include household-cohort (ih) and cohort-month (ht) fixed effects. Error bars denote 95% confidence intervals and standard errors are clustered at the household level.

mortgages or other debt, leaving them with limited liquidity despite relatively high incomes, consistent with the notion of 'wealthy hand-to-mouth' households (Kaplan et al., 2014; Campbell and Hercowitz, 2019). Table 1 shows further evidence on the mechanisms, indicating that the default increase is roughly twice as large for indebted households. We find no corresponding increase in defaults for the two highest income quintiles but, on average, default risk is income-dependent as for each income quintile default risk is 0.12 percentage points smaller.

3.3 Labor supply

We study the labor market effects on both the intensive and the extensive margins. Figure 2c shows labor supply effects on the intensive margin, measured as average labor earnings per worker. The average treatment effects are modest, but the aggregate effect in Table 1 is positive and statistically significant after the treatment expiration, indicating an increase of roughly 0.6% in labor earnings in response to contract expiration. Given an average income of €3,690 per worker per month, the treatment effect corresponds to approximately €22 per month when expressed in levels. The cost shock lowers real wages, which theoretically has an ambiguous effect on labor supply, depending on the relative strengths of the income and substitution effects (Borjas, 2010). Our finding of a positive labor supply response suggests that the income effect outweighs the substitution effect. In contrast, Table 1 shows no effect on the share of working individuals per household, indicating that the energy crisis does not cause entry or exit into the labor force.[14]

Figure 3c unpacks the positive intensive margin labor earnings response and shows that the treatment effects are driven by the low-income households increasing their labor supply. For the lowest income quintile, the increase is about six times the size of the average effect, but we find no anticipation effects. The estimated effects decline with income, with point estimates turning negative for the highest-income group. Households that are liquidity constrained show a pronounced response on the intensive margin, and also a response on the extensive margin (Columns 3–4, Panel B, Table 1). Despite a larger intensive-margin response among low-income households, the effect in absolute size is smaller. On average, 15% of all households in the lowest income quintile participate in the labor market with average earnings €2,060 per worker, or €310 euros per household member average. In contrast, the second quintile has 29% worker participation with €2,680 per worker, or €780 per household member.

Figure 4 combines both intensive and extensive margin labor market responses by estimating equation (1) separately using indicator $\mathbb{1}[Y > y]$ as the dependent variable. We repeat the estimation for varying thresholds y of average earnings per household member, Y. This approach allows us to assess whether the energy crisis increases the probability that household earnings exceed specific thresholds, from €0 to €5,000 with €250 intervals. The method captures both extensive and intensive margins by including households with zero labor income (Chen and Roth, 2024). We find a small negative effect for the first group ($\mathbb{1}[Y > 0]$), but not for other thresholds, suggesting that the negative effect is driven by households with very low earnings. This may reflect low-income households falling into basic social assistance or facing wage garnishments due to payment defaults, thereby reducing their disposable income from labor. We find a positive treatment effect on the probability of households earning between €1,750 and €2,250 per month. This aligns with the positive intensive margin labor supply response observed in the lowest-income quintile. We find no statistically or economically significant effect at income levels above €3,000 per month, consistent with Figure 3c, which shows no intensive margin response for the highest income quintile, and only small effects for the third and the fourth quintiles.

[14]In the Appendix Table B.2 we address possible sample selection by constructing Lee-style bounds within a difference-in-differences framework, following Lee (2009) and Rathnayake et al. (2024). For the full population the bounds include zero, whereas for the working-age subsample both bounds are strictly positive. In both cases, the point estimate lies within the bounds.

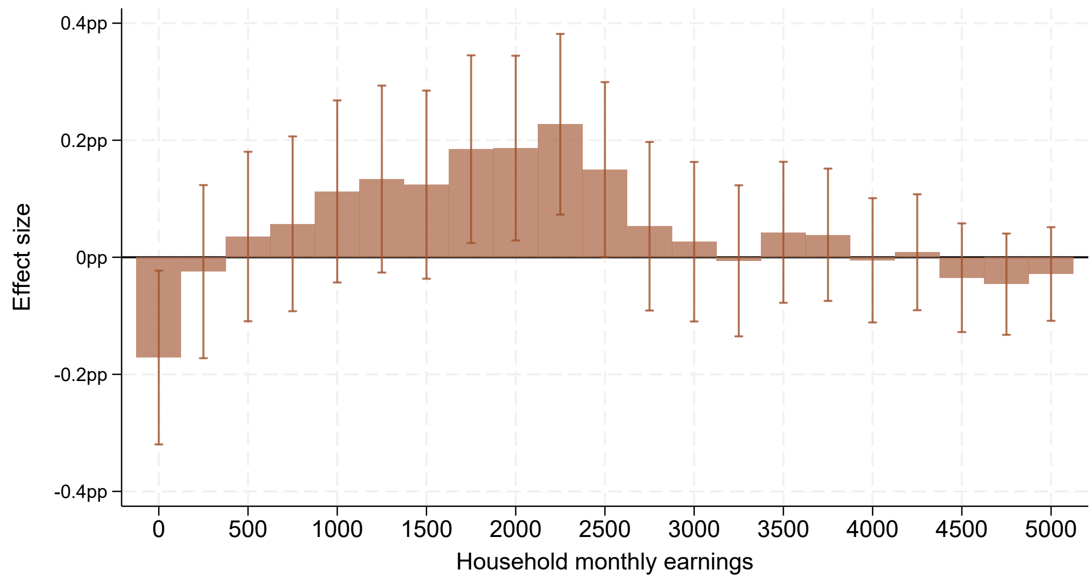

Figure 4: Labor supply, combined intensive and extensive margin responses

Notes: The figure estimates equation (1) where the dependent variable is $\mathbb{1}(Y > y)$; probability that the average earnings per household member, Y is more than a given threshold y. Following Chen and Roth (2024) we repeat the estimation for different values of $y = 0, 250, ..., 5000$, as shown on the x-axis (21 separate regressions). All regressions control for the anticipation effect (not shown), and include household-cohort (ih) and cohort-month (ht) fixed effects. Error bars indicate 95% confidence intervals and standard errors are clustered at the household level.

3.4 Residual consumption and savings

Figure 2d shows the impact on the residual term, capturing consumption, savings decisions, and other margins of household adjustment. This effect reflects the net impact of the cost shock after accounting for other adjustment margins – electricity use, income, and payment defaults. The effect appears after contract expiration and is alleviated in the spring of 2023, when prices fell and the government distributed electricity support payments. The effect is smaller for households with longer anticipation periods, explained by their greater reduction in electricity use and somewhat smaller price increases. In column 5 of Table 1 we find an average reduction of 2.8% in residual consumption and savings due to higher electricity prices.

Figure 3d depicts the heterogeneous effects of the energy crisis on residual consumption and savings. The effect is largest for the lowest-income groups, which already had lower baseline consumption levels. For each additional income quintile, the effect is 0.9 percentage points smaller in absolute value. This finding can be explained by two main factors. First, low-income households spend, on average, a larger share of their income on electricity. Second, low-income households reduced their electricity use less in response to higher prices. The positive earnings response on the intensive margin is partly offset by the unemployed, pensioners, and other households without labor income that are not found to adjust their earnings. The negative effect on residual consumption and savings falls most heavily on households that use more electricity than the median, cutting back the most. Indebted households cut back less than others, so the effect is smaller for them. Finally, the size of the negative effect declines as the anticipation period grows longer.

4 Conclusions

This paper provides causal evidence on how households adjust to a large and salient cost-of-living shock, the European energy crisis. We have traced these adjustments across four key margins – energy consumption, labor earnings, financial distress, and residual consumption and savings – and documented substantial heterogeneity across the income distribution. The interpretation of the quantitative effects should be understood in the context of our study period. On the one hand, the energy crisis was highly salient, with price awareness far greater than in normal times. On the other hand, the crisis was short-lived, and overall responses to sustained policy-driven changes in commodity prices, such as those created by climate policies, are likely to be larger. It remains an open empirical question how households would respond to other commodity price shocks, such as gasoline and food. Nevertheless, the broader patterns in household responses to the European energy crisis reveal economically meaningful mechanisms with direct implications for both distributional analysis and policy design in other contexts as well.

Differences across households in their ability to reduce energy use, adjust earnings, and smooth other consumption or savings are central to determining the optimal composition of price- and transfer-based relief policies. The first general insight is that heterogeneity in households' ability to adapt calls for targeted interventions rather than across-the-board subsidies. Second, the finding that low-income households' electricity use is less responsive implies that price-based interventions targeted at low-income groups distort choices only minimally, and thereby cause relatively small efficiency losses. Third, anticipation periods and gradual policy phase-ins matter, as they allow forward-looking households time to adjust before prices rise. In our setting, however, anticipation was insufficient to fully offset the adverse effects manifested through increased payment defaults. Finally, labor supply adjustments are concentrated among low- and middle-income workers and are primarily driven by income effects. Relief policies are therefore likely to directly influence labor supply responses at the lower end of the income distribution. The evidence presented here provides a first step. Future work should examine how substitution and income effects interact across multiple goods and over time, and how households respond to other types of cost shocks.

References

Adams-Prassl, A., K. Huttunen, E. Nix, and N. Zhang (2024). Violence against women at work. *The Quarterly Journal of Economics 139*(2), 937–991.

Ahlvik, L., M. Liski, and M. Mäkimattila (2024). Pigovian income taxation. *CESifo Working Paper* (11174).

Alberini, A., O. Khymych, and M. Ščasnỳ (2020). Responsiveness to energy price changes when salience is high: Residential natural gas demand in Ukraine. *Energy Policy 144*, 111534.

Atkinson, A. B. and J. E. Stiglitz (1976). The design of tax structure: direct versus indirect taxation. *Journal of Public Economics 6*(1-2), 55–75.

Baker, S. R. and C. Yannelis (2017). Income changes and consumption: Evidence from the 2013 federal government shutdown. *Review of Economic Dynamics 23*, 99–124.

Barreca, A., R. J. Park, and P. Stainier (2022). High temperatures and electricity disconnections for low-income homes in California. *Nature Energy 7*(11), 1052–1064.

Behr, S. M., T. Köveker, and M. Kücük (2025). Understanding energy savings in a crisis: The role of prices and non-monetary factors. *DIW Discussion Papers*.

Boehm, J., E. Fize, and X. Jaravel (2025). Five facts about MPCs: Evidence from a randomized experiment. *American Economic Review 115*(1), 1–42.

Borjas, G. J. (2010). *Labor economics*. McGraw-Hill/Irwin Boston.

Burger, S. P., C. R. Knittel, I. J. Pérez-Arriaga, I. Schneider, and F. Vom Scheidt (2020). The efficiency and distributional effects of alternative residential electricity rate designs. *The Energy Journal 41*(1), 199–240.

Byrne, D. P., L. A. Martin, and J. S. Nah (2022). Price discrimination by negotiation: A field experiment in retail electricity. *The Quarterly Journal of Economics 137*(4), 2499–2537.

Campbell, J. R. and Z. Hercowitz (2019). Liquidity constraints of the middle class. *American Economic Journal: Economic Policy 11*(3), 130–155.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics 134*(3), 1405–1454.

Cesarini, D., E. Lindqvist, M. J. Notowidigdo, and R. Östling (2017). The effect of wealth on individual and household labor supply: evidence from Swedish lotteries. *American Economic Review 107*(12), 3917–3946.

Chen, J. and J. Roth (2024). Logs with zeros? Some problems and solutions. *The Quarterly Journal of Economics 139*(2), 891–936.

Cronin, J. A., D. Fullerton, and S. Sexton (2019). Vertical and horizontal redistributions from a carbon tax and rebate. *Journal of the Association of Environmental and Resource Economists 6*(S1), S169–S208.

Çürük, M., A. Martens, K. Vringer, and D. van Soest (2025). The price elasticity of energy demand and the implications for climate policy. *Working paper*.

Dertwinkel-Kalt, M., C. Feldhaus, A. Ockenfels, and M. Sutter (2024). Household reduction of gas consumption in the energy crisis is not explained by individual economic incentives. *Proceedings of the National Academy of Sciences 121*(48), e2411740121.

Deryugina, T., A. MacKay, and J. Reif (2020). The long-run dynamics of electricity demand: Evidence from municipal aggregation. *American Economic Journal: Applied Economics 12*(1), 86–114.

Doligalski, P., G. P. Dworczak, M. Akbarpour, and G. S. D. Kominers (2025). Optimal redistribution via income taxation and market design.

Doremus, J. M., I. Jacqz, and S. Johnston (2022). Sweating the energy bill: Extreme weather, poor households, and the energy spending gap. *Journal of Environmental Economics and Management 112*, 102609.

Douenne, T. (2020). The vertical and horizontal distributive effects of energy taxes: A case study of a French policy. *The Energy Journal 41*(3), 231–254.

Fabra, N., C. Leblanc, and M. Souza (2025). Unpacking the distributional implications of the energy crisis: Lessons from the Iberian electricity market. *CESifo Working Paper* (12093).

Fagereng, A., M. B. Holm, and G. J. Natvik (2021). MPC heterogeneity and household balance sheets. *American Economic Journal: Macroeconomics 13*(4), 1–54.

Ferey, A., B. B. Lockwood, and D. Taubinsky (2024). Sufficient statistics for nonlinear tax systems with general across-income heterogeneity. *American Economic Review 114*(10), 3206–3249.

Fetzer, T., L. Gazze, and M. Bishop (2024). Distributional and climate implications of policy responses to energy price shocks. *Economic Policy 39*(120), 711–756.

Finnish Energy Authority (2021). National report on the state electricity and gas markets in Finland to the agency for the cooperation of energy regulators and to the European Commission.

Fischer, C. and W. A. Pizer (2019). Horizontal equity effects in energy regulation. *Journal of the Association of Environmental and Resource Economists 6*(S1), S209–S237.

Gelman, M., Y. Gorodnichenko, S. Kariv, D. Koustas, M. D. Shapiro, D. Silverman, and S. Tadelis (2023). The response of consumer spending to changes in gasoline prices. *American Economic Journal: Macroeconomics 15*(2), 129–160.

Gelman, M., S. Kariv, M. D. Shapiro, D. Silverman, and S. Tadelis (2020). How individuals respond to a liquidity shock: Evidence from the 2013 government shutdown. *Journal of Public Economics 189*, 103917.

Goldschmidt, D. and J. F. Schmieder (2017). The rise of domestic outsourcing and the evolution of the German wage structure. *The Quarterly Journal of Economics 132*(3), 1165–1217.

Golosov, M., M. Graber, M. Mogstad, and D. Novgorodsky (2024). How Americans respond to idiosyncratic and exogenous changes in household wealth and unearned income. *The Quarterly Journal of Economics 139*(2), 1321–1395.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics 225*(2), 254–277.

Grainger, C. A. and C. D. Kolstad (2010). Who pays a price on carbon? *Environmental and Resource Economics 46*, 359–376.

Gravert, C. (2024). From intent to inertia: Experimental evidence from the retail electricity market. Technical report, CEBI Working Paper Series.

Hahn, R. W. and R. D. Metcalfe (2021). Efficiency and equity impacts of energy subsidies. *American Economic Review 111*(5), 1658–1688.

Hellwig, C. and N. Werquin (2025, June). Using consumption data to derive optimal income and capital tax rates. Working Paper 2022-03, Federal Reserve Bank of Chicago. Originally published January 2022; revised June 16, 2025.

Ito, K. (2014). Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing. *American Economic Review 104*(2), 537–563.

Johnson, D. S., J. A. Parker, and N. S. Souleles (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review 96*(5), 1589–1610.

Kaplan, G., G. L. Violante, and J. Weidner (2014). The wealthy hand-to-mouth. *Brookings Papers on Economic Activity*.

Lee, D. S. (2009). Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *The Review of Economic Studies 76*(3), 1071–1102.

Levell, P., M. O'Connell, and K. Smith (2025). The welfare effects of price shocks and household relief packages: Evidence from an energy crisis. *Working paper*.

Levinson, A. and E. Silva (2022). The electric gini: Income redistribution through energy prices. *American Economic Journal: Economic Policy 14*(2), 341–365.

Malani, A. and J. Reif (2015). Interpreting pre-trends as anticipation: Impact on estimated treatment effects from tort reform. *Journal of Public Economics 124*, 1–17.

Motiva (2015). Sähkölämmityksen tehostamisohjelma Elvari. Website. Available at https://www.motiva.fi/koti_ja_asuminen/energiatehokas_pientalo/energiatehokas_sahkolammitys/sahkolammityksen_tehostamisohjelma_elvari.

Parker, J. A., N. S. Souleles, D. S. Johnson, and R. McClelland (2013). Consumer spending and the economic stimulus payments of 2008. *American Economic Review 103*(6), 2530–2553.

Picchio, M., S. Suetens, and J. C. van Ours (2018). Labour supply effects of winning a lottery. *The Economic Journal 128*(611), 1700–1729.

Pizer, W. A. and S. Sexton (2019). The distributional impacts of energy taxes. *Review of Environmental Economics and Policy*.

Prest, B. C. (2020). Peaking interest: How awareness drives the effectiveness of time-of-use electricity pricing. *Journal of the Association of Environmental and Resource Economists 7*(1), 103–143.

Rathnayake, G., A. Negi, O. Bartalotti, and X. Zhao (2024). Difference-in-differences with sample selection. *arXiv preprint arXiv:2411.09221*.

Reguant, M., N. Fabra, and J. Wang (2025). The distributional impacts of time-varying electricity pricing: A novel approach to estimating household income. *Conditionally accepted at the Journal of the European Economic Association*.

Roth, J., P. H. Sant'Anna, A. Bilinski, and J. Poe (2023). What's trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics 235*(2), 2218–2244.

Rubin, E. and M. Auffhammer (2024). Quantifying heterogeneity in the price elasticity of residential natural gas. *Journal of the Association of Environmental and Resource Economists 11*(2), 319–357.

Steen, F., S. A. Ulsaker, and T. A. Aursland (2021). Unemployment shocks, cyclical prices and shopping behavior. NHH Dept. of Economics Discussion Paper 03/2021.