Understanding energy savings in a crisis: The role of prices and non-monetary factors

Sophie M. Behr[*,†] Till Köveker[*,†] Merve Kücük[‡,§]

February 28, 2025

Abstract

Russia's invasion of Ukraine was accompanied by a significant reduction of its gas supply to Europe, causing sharp energy price surges and prompting governments to respond with public appeals and programs aimed at reducing consumption. This paper investigates the effects of price increases and non-monetary factors, such as public appeals and saving programs, on residential energy savings during the crisis. Using a unique building-level dataset on residential energy consumption and prices in Germany, we identify price-driven savings and energy price elasticities with a DiD-PSM approach. By comparing buildings which faced price increases to buildings with constant prices we can isolate price-driven savings from contemporaneous non-monetary effects. Our findings reveal that while increased prices led to moderate short-run energy savings, the majority of observed savings were driven by non-monetary factors. Consequently, we identify a relatively low short-run energy price elasticity of demand of -0.07. Going beyond average effect estimation, we use two machine learning methods to calculate building-level price-driven and non-price-driven savings and analyze their variation with socio-economic characteristics using census data.

[*] DIW Berlin
[†] Technical University Berlin
[‡] CMCC Foundation - Euro-Mediterranean Center on Climate Change
[§] Ca'Foscari University of Venice

1 Introduction

The Russian invasion of Ukraine caused an energy crisis that strongly affected European countries with high dependence on Russian natural gas imports. Interruption of Russian gas supply led to increasing gas prices with knock-on effects on other energy prices. Discussions about potential gas shortages in the winter of 2022/23 and fear of resulting power cuts, industrial production disruptions and job losses led to the formulation of saving targets by the European Union (EU) and several individual countries (European Council, 2022). Germany, Europe's largest economy and biggest importer of Russian gas until the crisis (Bachmann et al., 2022; U.S. Energy Information Administration, 2022), announced an ambitious gas saving target of 20%. German energy consumers were urged to save energy through numerous public calls by the government and various interest groups (More and Carrel, 2022). For example, the German minister for economic affairs, Robert Habeck, made a public appeal to *"[...] Join the effort! By saving energy, we can reduce Germany's dependence on Russian imports and help protect the climate."* (Bundesministerium für Wirtschaft und Klimaschutz, 2022; own translation). These appeals were backed with a portfolio of government programs, which included informing households about untapped energy saving potentials and their energy prices (The Press and Information Office of the Federal Government, 2022; Umweltbundesamt, 2022), mandatory optimization of buildings' heating system operation (Federal Ministry of Economic Affairs and Climate Action, 2022a), mandatory limits on heating public buildings (Goldenberg, 2022) and support measures to shield households from the financial burden of exploding energy prices (Federal Ministry of Economic Affairs and Climate Action, 2022b). A challenge for the design of these emergency measures was a lack of experience and resulting uncertainty about how much savings would result from higher prices and how much savings could be achieved through other non-price interventions.

This paper answers the question to which extent energy savings in the crisis were driven by higher prices (price-driven savings) and non-monetary factors (non-price-driven savings), such as public appeals. Disentangling these two types of savings allows us to accurately calculate the short-term price elasticity of energy demand in the energy price crisis. Since the government programs and public appeals were largely directed at households, our focus is on residential heat energy savings. We use building-level heat energy billing data of more than 140,000 two- and multi-family buildings from an energy billing service provider in Germany. Our empirical strategy includes three separate methods which rely on the variation of energy prices across different buildings. The variation in prices is due to households' different tariff types (mostly one- or two-year contracts), the timing of their contract renewal, and the contract position of their suppliers.[1]

In the first part of our analysis, we estimate the price-driven heat energy savings using a difference-in-differences framework with propensity score matching (DiD-PSM). Buildings that experienced an increase in energy prices in 2022 constitute the treatment group and buildings with constant prices the control group. Crucially, we estimate short-run heat energy price

[1] Energy suppliers who have purchased their energy through long-term contracts are less likely to default or to increase their prices than energy suppliers that had to buy their energy on the spot market in 2022.

elasticities during the crisis using the same DiD-PSM framework to carefully account for non-price-driven energy savings. Thus, we avoid that the estimated energy price elasticities are inflated by the contemporaneous appeals and programs to save energy.

Secondly, we employ double machine learning (DML) to flexibly analyze socio-economic heterogeneities in price-driven energy savings. To this end, we merge the building-level energy dataset with socio-economic characteristics obtained from the German microcensus at the zip-code level. Using the DML procedure we then estimate building-level treatment effects and use non-parametric kernel regressions to explore how they vary with socio-economic variables.

In the third part of our analysis, we again rely on machine learning techniques, this time to estimate non-price-driven savings and their heterogeneity. We first train a Lasso-regression model to predict energy consumption on pre-crisis data. Then, we use the model's predictions to obtain building-level counterfactual energy consumption for 2022 had the energy crisis not happened. The difference between observed and counterfactual energy consumption yields non-price-driven savings for buildings where prices did not rise during the crisis. Using these building-level estimates of non-price-driven energy savings, we can again use non-parametric kernel regressions to explore heterogeneities with respect to socio-economic variables.

The price elasticity of energy demand is an important parameter for assessing potential demand responses to crises and fundamental for energy and climate policy. The short-run energy price elasticity in times of energy price crises is the central parameter to understand and correctly anticipate price-driven savings.[2] It informs governments how to balance the use of price-based measures and non-price interventions, such as energy saving programs and campaigns, to achieve short-term demand reductions. There is a consensus that in the short-term energy demand is relatively price-inelastic during normal times with moderate prices (see Brons et al., 2008; Espey and Espey, 2004; Havranek et al., 2012; Labandeira et al., 2017). Yet, there is a lack of evidence on the short-term elasticity in times of crisis when prices surge beyond typical levels. The recent crisis thus offers an opportunity to study the demand elasticity at high prices and its potential non-linearity.

There has been an increasing interest in the demand response to the recent energy crisis (Dertwinkel-Kalt et al., 2024; Jamissen et al., 2024; Roth and Schmidt, 2023; Ruhnau et al., 2023). Roth and Schmidt (2023) find that savings during the energy crisis were not only due to mild temperatures. They calculate that savings exceed reductions in energy demand that would be expected from the warmer climate in 2022 alone. Relying on a field experiment with an energy provider, Dertwinkel-Kalt et al. (2024) also find that gas savings in Germany during the energy crisis were significant. However, economic incentives and prices had little impact due to limited consumer understanding, suggesting that societal engagement drove energy saving efforts. Ruhnau et al. (2023) find significant gas savings across small consumers, industry, and power stations, with variations in timing and magnitude linked to key wartime developments. They

[2] Long-run energy price elasticities, on the other hand, shape energy market design, are relevant for estimating emission reductions from carbon taxes and thus, to set carbon prices at levels that are consistent with Greenhouse-gas reduction targets. They also determine the incidences of, for example, energy and carbon taxes on consumers and producers and are therefore important to design equitable emission reduction strategies (see Trotta et al., 2022; Xiang and Lawley, 2019).

also estimate energy price elasticities but acknowledge that their elasticities may be inflated by "public attention and ethical considerations" of households to save energy (Ruhnau et al., 2023, p.624). These non-price-driven energy conservation motives may however have been particularly important during the recent energy crisis. To our knowledge, Jamissen et al. (2024) is the only other study to control for public attention in the crisis in their estimation of a gas price elasticity using an autoregressive distributed lag model. In contrast to our study, their analysis relies on aggregate data and can therefore only exploit across-time variation in prices and consumption to identify the effect of prices on natural gas demand. In spite of their different method and the use of aggregate data, they find a gas price elasticity, which is of similar magnitude as our elasticity estimate for gas-heated buildings. Importantly, they also find a much smaller gas price elasticity than parts of the literature which do not account for non-monetary factors in the crisis.

For a different energy price crisis in California, Reiss and White (2008) show that public calls to save energy were effective in achieving substantial and rapid demand reductions. Related, Ito et al. (2018) find that moral persuasion by firms or governments can be an effective short-term strategy to change households' consumption behavior. This indicates that not accounting for crisis-linked non-monetary motivations in estimation strategies may lead to an overestimation of the price elasticity of energy by falsely attributing non-price-driven savings to price-induced savings. Therefore, when analyzing energy price elasticities during energy crises, it is crucial to account for non-price motivations to conserve energy prompted by such events, which our data and identification strategy allows us to do.

We contribute to the literature in three major ways. First, we use unique building-level data that enables us to apply a DiD-PSM approach to identify short-term price-driven savings while comprehensively controlling for contemporaneous non-price-driven savings. Being able to isolate price-driven savings from non-price driven savings enables us to more precisely estimate the short-term price elasticity of heat energy demand during crises by avoiding that our estimated elasticities are inflated by contemporaneous non-price driven energy savings. Our findings help understand heat energy price elasticities during energy crises with large price increases.[3] We find that buildings with price increases saved on average 2.2 percentage points (pp.) more than buildings with constant prices. The associated arc-elasticity of energy demand is -0.07. Price-driven savings increase with the magnitude of the price increase. For buildings with low price increases of up to 25% we find small and statistically insignificant energy savings. However, buildings with larger prices increases of 25% - 50% and >50% reduced their heat energy consumption by statistically significant 2.2% and 4.4%, respectively. The arc-elasticities of energy demand are constant across price-increases - all three groups exhibit a price elasticity -0.07 (though insignificant for the group with low price increases). Other studies on the recent energy crisis find a wide range of energy price elasticities. Ruhnau et al. (2023) find relatively high gas price elasticities

[3] Several existing studies investigate price elasticities of residential heat energy demand during times-as-usual (see Hanemann et al., 2013; Ó Broin et al., 2015; Schulte and Heindl, 2017; Trotta et al., 2022) and identify an elasticity measure in the range of -0.18 to -0.64 for a number of countries and types of heating fuel. Labandeira et al. (2017) estimate in the most recent meta-analysis on energy price elasticities based on 428 papers produced between 1990 and 2016 providing 966 short-term price elasticities, the short-term energy price elasticity is around -0.22.

between between -0.16 and -0.27 without controlling for non-price motivations. Controlling for attention to the crisis using Google searches, Jamissen et al. (2024) find a gas price elasticity of -0.04. In spite of the above mentioned methodological differences, our estimated elasticity of -0.06 for gas-heated buildings is quite close to their estimated elasticity.[4] This underlines the importance of comprehensively controlling for non-price motivations as we do in our empirical strategy. While Jamissen et al. (2024) partly control non-price factors by accounting for public attention for the energy crisis, our control-group-based approach comprehensively controls for all non-price factors including governmental energy saving programs and regulations.

Second, we add to the literature of heterogeneous residential energy savings by applying DML to investigate how building-level price-driven savings vary with socio-economic characteristics that potentially affect the observed short-term response to a crisis. To this end, we merge the building-level energy dataset to administrative microdata on socio-economic characteristics from the German microcensus. Overall, we do not find strong heterogeneities of price-driven savings. We observe that price-induced savings increase with average net income per person, which contradicts other studies finding less price-responsiveness among high-income households (see Rubin and Auffhammer, 2024; Schmitz and Madlener, 2020). A possible reason for this finding could be a higher savings potential of richer households due to higher baseline energy consumption. Age, years of schooling, and unemployment do not vary with price-induced savings in a meaningful way.

Finally, we train a machine learning model to predict energy consumption in 2022 for a counterfactual scenario had the crisis not occurred. We estimate non-price-driven savings by comparing the counterfactual energy consumption in a non-crisis scenario with actual energy consumption for buildings where prices remained constant. This allows us to calculate the average magnitude of non-price-driven savings during the crisis. Moreover, as these estimates of non-price-driven savings are at the building-level, we can also explore how they vary with socio-economic characteristics. By proposing and implementing a methodology to estimate non-price-driven savings we add to the literature on energy demand response during energy crises. We find that across buildings with constant prices, non-price-induced savings are on average 8.5 % of a building's energy consumption in 2021. This underscores the importance of non-price factors to achieve energy savings in the recent crisis. Similar to price-driven savings, we do not observe strong heterogeneities of non-price-driven savings. A cautious interpretation of non-price savings shows that they first rise with the unemployment rate before falling again at very high unemployment rates. They do not, however, appear to vary with income, age and years of schooling.

The remainder of this paper is organized as follows: Section 2 provides insight into the background of this study by discussing the historical background of the gas trade relationship between Germany and Russia and how the Russian invasion of Ukraine in February 2022 has altered this relationship. This section also offers an overview of residential energy consumption and energy billing in Germany. Section 3 presents the data and descriptive statistics. Section 4 discusses our empirical strategy and explains the three methods that we use to estimate price-

[4] For buildings with district heating, we find a price elasticity of demand of -0.17.

driven and non-price-driven energy savings. In section 5 we present and discuss the results, before concluding in section 6.

2 Background

2.1 The Energy Crisis and its Implications for Germany

In 2021, before the Russian invasion of Ukraine, Russia was the main fossil fuel supplier of the European Union with shares of 26% petroleum oil and 44% natural gas imports (Wettengel, 2024). However, not all European countries were equally import dependent on fossil fuel imports from Russia. As of 2021, Germany was importing 98% of its oil and 95% of natural gas out of which 34% and 55% were originating from Russia, respectively (Bachmann et al., 2022; Wettengel, 2024).[5] However, this situation substantially changed with the Russian invasion of Ukraine in February 2022.

Already in the months preceding the invasion, Russia began restricting its gas supply to Germany. After the start of the war, Germany announced plans to reduce its dependence on Russian gas by 2024. With the EU embargos on Russian coal and oil, Germany also drastically reduced imports of these two energy carriers from Russia (Wettengel, 2024). Russia preempted Germany's plans of reducing its dependence on imports of Russian gas by gradually cutting the flow of gas to Germany until a complete halt in late summer 2022, briefly before the explosions of the Nord Stream pipelines (ibid.).

Following the invasion and the reduction of gas supply from Russia, gas wholesale prices in Germany temporarily increased to about four times the pre-war levels observed in 2021 (Bundesnetzagentur, 2023). The rising wholesale prices were ultimately also reflected in higher end-user prices. At their peak in the beginning of September 2022, households faced a gas price of around 40 cents/kWh when concluding new contracts compared to approximately 6 cents/kWh before the crisis (VERIVOX, 2025). These high prices and the supply interruption in late summer 2022 lead to intense public discussions and worries about the possibility of gas shortages in the winter 2022/23.

In response to this situation, the German government called upon all economic actors to save energy wherever and whenever it is possible while simultaneously trying to secure other sources for natural gas imports (More and Carrel, 2022). The resulting efforts to save proved successful, as the German economy reduced overall gas consumption by approximately 20% between July 2022 and March 2023 compared to previous years, with industrial gas consumption dropping by 26% and household consumption decreasing by 17% (Moll et al., 2023).

[5] Germany had 50-year long natural gas trade relationship with Russia which was founded on an agreement made in 1958 under which West Germany was providing pipes for the Druzhba pipeline ("Friendship Pipeline") in return for Russian natural gas. This pipeline, becoming operational in 1964, was the world's longest oil pipeline linking Russia with much of eastern Europe (Sullivan, 2022).

2.2 Residential Energy Consumption and Billing in Germany

The strong increase of residential energy prices in Germany posed substantial burdens on households (Kröger et al., 2023). At the end of 2022, the German government introduced one-off relief payments to alleviate these burdens to some extent. The support measure consisted in covering the monthly installments of residential gas and district heat customers for December 2022. Support measures for households using other heat energy carriers such as oil and wood pellets were designed differently. Since we can only exactly calculate the amount of support for gas and district heat, we choose to limit analysis to these two energy carriers (see section 3 for further details). In 2023 the German government introduced further interventions including a price cap on heat energy to further shield households from the high heat energy prices. As this price-cap strongly altered households' perception of marginal prices (Dertwinkel-Kalt et al., 2024) and thus households' reaction to energy prices, we decide to limit our analysis to the years up to 2022.

In Germany, heat energy is billed annually with consumers paying monthly or quarterly installments. The installments are calculated based on households' historical consumption. The objective of the monthly installments is to spread the cost of energy evenly throughout the year, and can be adjusted at any point during the billing cycle. At the end of the yearly billing cycle, the actual energy consumption is reconciled against the estimated energy consumption on which the advance payments are based. Additional payments or back payments are made to account for any overpayments or underpayments ensuring that consumers are billed accurately according to their actual consumption. Due to the annual billing cycle prices are not very salient and most households might not be aware of their exact heat energy price (Dertwinkel-Kalt et al., 2024). However, irregular price changes during a running contract are more visible to residential consumers as energy suppliers are legally obliged to inform them in writing with an advance notice of at least four weeks (Bundesnetzagentur, 2024, § 41 EnWG Energielieferverträge mit Letztverbrauchern Energiewirtschaftsgesetz). This advance notice must clearly outline the nature and reason for the change, allowing consumers ample time to assess its impact and, if necessary, switch energy providers or renegotiate their contract.

In some buildings the residents have a direct contract with the energy provider. In other - particularly in multi-apartment - buildings the landlord or building management company concludes a central energy contract for the entire building with the energy provider. In the latter buildings, the information flow of price changes to energy consuming residents may be less immediate than in the former. The reason is that for buildings with a central energy contract, the energy provider's contractual partner is the landlord or the building management company. In the occurrence of a price change, the energy provider is then obliged to inform the landlord or building management company. However, they do not have any legal obligation to immediately inform the energy consuming residents of the price change and may only do so at the end of the annual billing period. Nevertheless, their best interest would be of quickly informing the residents to avoid having to advance the price difference until the end of the billing period, to prevent conflicts with residents about high supplementary payments and to allow residents to react to the risen prices by adapting their heating behavior.

3 Data

3.1 Data Sources

Data on Heat Energy Consumption Our main dataset comprises residential annual heat energy billing data of 141,575 German multi-family buildings[6] from 2017 until 2022 by *ista*, one of the largest heat energy bill providers in Germany. The dataset contains information on energy consumption, energy expenditures, the billing periods, energy carrier, total living space, number of apartments in each building, retrofits[7], hot water consumption and the buildings' zip-code.[8] Based on the heat energy consumption and expenditures, we calculate kWh prices for each year and building.[9] As the buildings in our sample all have central heat energy contracts, all apartment units face the same kWh price, meaning that the calculated kWh price does not constitute the average but the actual heat energy price of each household in the building. During the energy crisis, the German government adopted a relief package to lift the energy price burden on households and firms. Part of this package were the so-called *December relief payments* which were announced in November 2022 for gas and district heat users. These payments meant that the state paid for the December heat energy costs of residents, i.e. one twelfth of the yearly payments. We can thus easily account for the December relief payments in our price calculation for buildings with as and district heat. The amount of the payments were independent of actual energy consumption to remain energy saving incentive-compatible. There were also relief payments for other energy carriers such as oil or wood pellets. In contrast to gas and district heating relief payments, these were not automatically paid out to households. Households had to actively apply for these relief payments. As we do not know which households received relief payments, we are therefore not able to calculate the prices they faced. This is the reason we solely focus on gas and district heating and exclude buildings with other heat energy carriers from the sample.[10]

Data on Climate Factors To adjust absolute heat energy consumption according to local temperature changes, we employ the publicly available zip-code level climate factors by the German Weather Service ("Deutscher Wetterdienst"). These climate factors are calculated for rolling 12-month periods as quotients of the mean annual degree days, which are calculated with reference to the time series of the Potsdam reference station, and the current annual degree days

[6] In 2022, approximately 62% of the German population lived in apartments and 12% lived in semi-detached housing, making the total share of the population living in multi-family houses 74% (Statista, 2023).

[7] We drop all households that have undergone thermal retrofitting procedures over the period of this study from our sample.

[8] Zip-codes in this paper refer to German postal codes ("Postleitzahlen").

[9] The heat energy expenditures that we observe from the heat energy bills are rounded to the full cent. Heat energy prices per kWh are, however, usually denominated at more granular intervals in most heat energy contracts. We calculate kWh prices by dividing the rounded observed expenditure by observed consumption. This results in prices with many decimal places. To ensure that we can identify if a price stayed constant, we round the kWh prices to 0.1 Euro cents. As our main empirical strategy relies on comparing buildings with price increases to buildings with constant prices, we drop the few buildings from our sample where prices decreased in 2022.

[10] Gas and district heat together cover 63% of German buildings (Destatis, 2018).

for the respective zip-code. In all parts of the analysis, we control for weather by including the climate factor up to the third polynomial as an explanatory variable.

Data on Home Office To account for potential changes in work and leisure time patterns due to the Covid-19 pandemic, we control for home office use. We rely on home office data at the zip-code level collected and made available by *infas360*. Specifically, we observe the number of days per week spent working from home on a scale from zero to five workdays for the period before, during, and after the pandemic at the zip-code level. [11] Our home office variable captures the mean of the number of days spent working from home, normalized by national home office trends from the federal statistical office of Germany, *Destatis*. Thus, the home office variable reflects the national home office average, but exploits variation at the sub-national level.

Data on Socio-Economic Characteristics We merge socio-economic data from the 2018 wave of the German microcensus with the data provided by *ista*. The microcensus is the biggest annually-repeated household survey in Germany and comprises information on the individuals living in the household as well as on the household characteristics. We use variables on age, income, unemployment, years of schooling, and share of people receiving social benefits. We use these variables to analyze how energy savings vary with socio-economic characteristics. Information on households' location is anonymized to a certain degree extend and follows a different spatial categorization than the zip-codes. In contrast to the heat energy billing data, the microcensus contains regional identifiers for the municipal identification number (AGS), not the zip-code. In a first step, we calculate AGS-level averages of the socio-economic variables of interest. In a second step, we merge the AGS-level dataset with the energy billing data using zip-code-to-AGS correspondence tables while carefully adjusting the AGS-level means when there is no one-to-one correspondence.[12]

3.2 Descriptive Statistics

Figure 1 presents the developments in heat energy prices and heat energy consumption from 2017 until 2022. The strong crisis-related energy price hike is visible in Panel (a) and the decline in energy consumption per square meter during the crisis is given in Panel (b).

Table 1 presents the descriptive statistics for the control and treatment group for the pre-treatment period as well as the year 2022, the year of the crisis. The top panel of the table shows building-level variables. The middle and lower panels show the climate factor and home office variable, as well as the socio-economic variables from the microcensus. For the treatment group, prices rose strongly between 2021 and 2022 with an average of 44.5%. Overall, weather-unadjusted heat energy consumption in both the treatment and the control groups fell by 16.0% in the same year. 5,254 buildings did not experience a change in their kWh price, meaning that the control group makes up a share of 3.7% of our sample in the year 2022.

[11] For the post-pandemic time frame, we observe not the realized but instead the planned number of home office days by the employer.

[12] Note that we use the word "zip-code" to refer to that community/neighborhood level also after merging the data.

Figure 1: Heat energy price and consumption between 2017-2022.

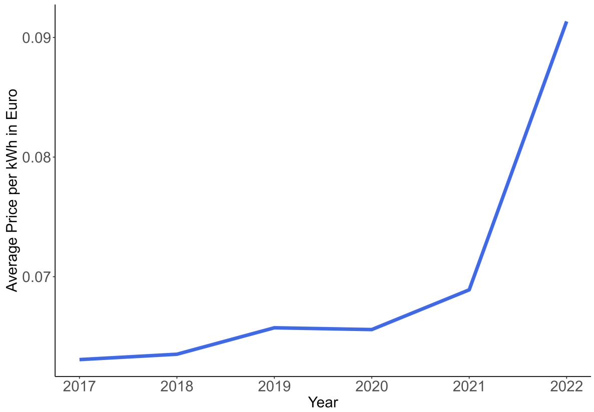

(a) Heat energy price in Euro per kWh between 2017-2022.

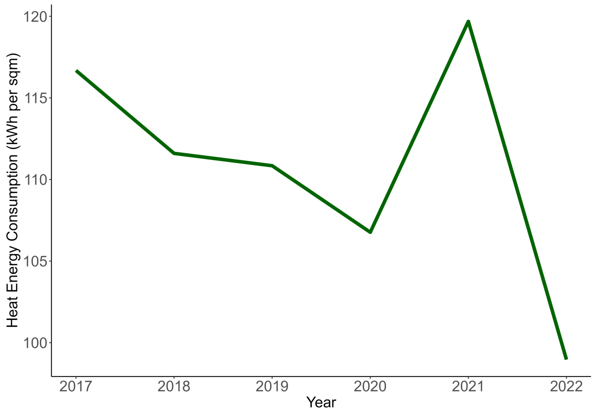

(b) Heat energy consumption in kWh per square meter between 2017-2022.

Table 1: Overview of descriptive statistics

	Control group		Treatment group	
	2017-2021	2022	2017-2021	2022
ista data				
Price (Cent/kWh)	5.71 (1.34)	5.94 (1.29)	6.54 (2.22)	9.88 (4.09)
Price change (%)	2.51 (12.39)	0.06 (0.72)	3.59 (14.35)	44.49 (43.52)
Heat consumption (kWh)	110.95 (41.69)	98.65 (38.57)	112.66 (45.76)	99.12 (42.22)
Heat consumption change (%)	3.60 (23.58)	-16.02 (22.55)	3.55 (27.00)	-15.99 (22.19)
Size apartment (sqm)	77.29 (28.18)	77.43 (27.97)	78.01 (28.37)	78.26 (28.41)
Number apartments	11.30 (10.40)	11.24 (10.37)	10.12 (10.82)	10.02 (10.78)
Share of gas (%)	94.41 (22.98)	94.44 (22.91)	86.13 (34.56)	86.43 (34.25)
German weather service DWD				
Climate factor	1.14 (0.10)	1.20 (0.09)	1.14 (0.10)	1.20 (0.09)
infas360				
Home office (%)	14.88 (10.50)	23.33 (8.52)	14.09 (10.06)	22.57 (8.03)
microcensus				
Year of birth	1972.79 (3.26)	1972.77 (3.26)	1972.72 (3.20)	1972.73 (3.19)
Social benefits (%)	3.98 (2.95)	3.98 (2.97)	4.12 (3.00)	4.11 (3.00)
Household income (€)	2867.85 (489.49)	2873.09 (496.86)	2867.94 (468.43)	2873.25 (472.17)
Years schooling	12.57 (0.71)	12.57 (0.72)	12.54 (0.73)	12.54 (0.73)
Unemployment (%)	3.17 (2.35)	3.16 (2.36)	3.22 (2.37)	3.21 (2.36)
Observations	22,921	5,254	588,425	136,321

Notes: The table reports the means and standard deviations (in parentheses) of the main variables used in the analysis for the unmatched sample.

4 Methodology and Empirical Strategy

We rely on several methods to estimate price-driven savings, the associated energy price elasticities and non-price driven savings during the energy crisis. At that time, price hikes coincided with several non-price factors that may have led to energy savings, such as public appeals and various government programs and regulations aimed at reducing energy consumption. The simultaneous occurrence of price hikes and non-price factors poses an empirical identification challenge to isolate price-driven from non-price driven savings. In appendix A.1 we illustrate this identification challenge using a simple theoretical model that features a household's decision on how much energy to consume depending on prices and public appeals to save energy. The model shows that when energy prices and public appeals are correlated, non-price driven savings can lead to an omitted variable bias in the estimation of price-driven savings and vice versa. In the energy crisis, there was a clear temporal correlation of price hikes and the rise of non-price factors for energy savings, such as public appeals. Hence, we design our empirical strategy to carefully isolate price-driven from non-price driven savings.

We apply a DiD-PSM approach to estimate price-driven savings and energy price elasticities in the crisis using households with constant prices as the control group. The DiD-PSM approach allows us to explore how price-driven savings vary by energy carrier and treatment intensity using coarse subgroup analyses. Moreover, we are interested in more granular and detailed heterogeneities of price-driven savings depending on socio-economic characteristics. To this end, we apply Double Machine Learning, again using buildings where prices did not rise in 2022 as the control group, which allows us to analyze how price-driven savings continuously vary with socio-economic characteristics. Finally, we estimate non-price driven savings that are attributable to factors such as savings programs, appeals and moral motivations to save. We predict counterfactual energy consumption had the crisis not happened, using a lasso-model trained on the pre-crisis period. The difference of observed energy consumption and predicted counterfactual energy consumption in 2022 for the buildings where prices did not rise in 2022 gives us an estimate of the non-price savings.[13] Using these building-level estimates, we can compute average non-price driven savings and explore how non-price driven savings vary with socio-economic characteristics.

4.1 Estimation of Price-Driven Savings & Energy Price Elasticities

4.1.1 Average Effects

Method 1: Difference-in-Differences with Propensity Score Matching To estimate price-driven energy savings we apply a DiD-PSM approach. We consider a building as treated if its energy price increased between 2021 and 2022. The control group consists of buildings for which prices stayed constant during this period. To increase the comparability of the treatment and the control group and to improve the plausibility of the parallel trends assumption we apply 1:1 nearest neighbor propensity score matching with replacement. In the results section, we show that our results are robust to a host of robustness checks including 5:1 matching. We discuss the validity of our control group, show evidence for parallel trends and matching diagnostics in the following subsection.

[13] For the buildings where prices rose this quantity would correspond to the sum of non-price and price-driven savings

We first estimate the propensity score in a logistic regression by regressing the treatment group indicator on a set of matching variables (the number of apartments in the building, the average apartment size, the average energy price in the pre-treatment period, the climate factor, and the home office rate) and on year and energy carrier fixed effects. As suggested by Austin (2011b) we match on the logit of the propensity score and use a caliper width of 0.2 standard deviations which is considered to be optimal by Austin (2011a). We furthermore enforce exact matching by year and energy carrier to avoid that an observation is matched to another observation from a different year or to a different type of energy carrier. Finally, we ensure common support by discarding observations that fall outside of the support of the treatment group's propensity score distribution.

Using the matched sample, we then estimate average price-driven savings in the treatment group with a DiD model. As we are in a setup with common treatment timing and only one post-treatment period, we can implement the DiD approach using a simple two-way fixed effects regression model (equation 1):

$$ln(e_{it}) = \lambda_i + \mu_t + \delta D_{it} + \boldsymbol{x'_{it}\beta} + \epsilon_{it} \tag{1}$$

In the estimation, each observation is weighted by its matching weight $w_{i,t}$. Matching weights take the value of one for treated observations. For the control group, they correspond to the number of times that a control observation was used as a match. $ln(e_{it})$ is the logarithm of residential energy consumption in building i during year t. D_{it} is a binary treatment indicator that takes the value of one for treated buildings in the year 2022. $\boldsymbol{x'_{it}}$ denotes the set of control variables which are the climate factor up to its third polynomial and the home office rate. γ_i and μ_t capture building and time fixed effects, ϵ_{it} is the error term. Our coefficient of interest is δ which indicates the ATT or the average price-driven energy savings in 2022 of buildings that were exposed to a price increase during the energy crisis.

To allow for potential non-linearities in the enery demand function we estimate price-driven savings by treatment intensity. To this end, we categorize buildings that experienced a price increase in 2022 into the following three treatment groups: those with price increases of <25%, between 25% - 50%, and >50%. The control group remains the same, i.e. group of buildings with constant prices in 2022. We repeat the above matching procedure by separately matching each one of the three treatment groups to the control group. For all three of the matched samples we apply the DiD-PSM described above to estimate the average price-driven savings for each price-increase level. Additionally, we estimate energy carrier specific effects, by analogously repeating the described matching procedure and the DiD-PSM estimation using the sample of gas-heated buildings only and using the sample of district-heated buildings only.

Identification, Validity of the Control Group and Matching Diagnostics The key identifying assumption of our DiD-PSM estimation is that the trend in average energy consumption would have been the same between the treatment group and the matched control group in a hypothetical crisis-scenario in which the treatment group is not exposed to price

hikes[14]. In other words, the parallel trends assumption implies that non-price-driven savings are on average equal between the two matched groups during the crisis year 2022.
While the assumption is untestable, we provide evidence in support of parallel trends with the event study in Figure 2. The figure shows the year-specific treatment coefficients that we obtain by estimating the following event-study specification on the matched sample:

$$ln(e_{it}) = \lambda_i + \mu_t + \sum_{\tau=2019}^{2021} \gamma_\tau D_{i\tau} + \delta D_{it} + \boldsymbol{x}'_{it}\boldsymbol{\beta} + \epsilon_{it} \tag{2}$$

The event study in Figure 2 clearly shows that on average the trends in energy consumption between the matched treatment and the control group were not significantly different before the energy crisis.

Figure 2: Event study supporting parallel trends

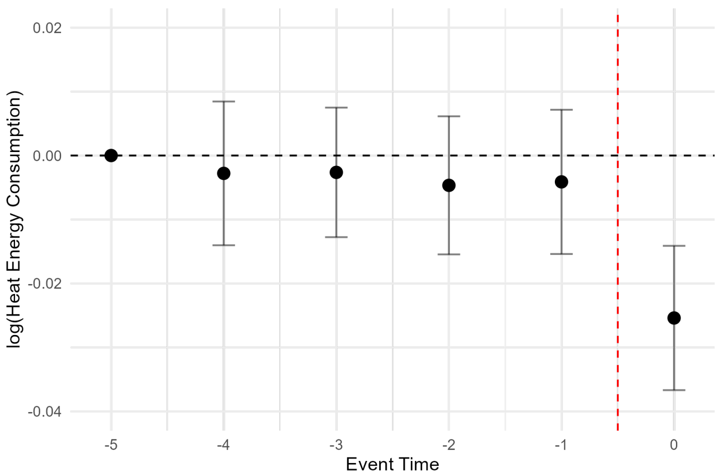

The figure shows the event study coefficients from estimating equation 2 on the matched sample. The x-axis displays the time to treatment in years, where 0 corresponds to the energy crisis in 2022. The displayed event study coefficients indicate the estimated year-specific treatment effects and placebo treatment effects. Vertical bars indicate 95% percent confidence intervals that are calculated using standard errors clustered at the building-level.

It is very unlikely that there may have been any form of active sorting into or, in our case, rather out of treatment at the building level. Landlords or building managers will not have strategically opted for shorter or longer contract periods to avoid having to renew the contracts in the middle of the crisis, as the crisis was not foreseeable when concluding the contracts.

[14] Besides parallel trends identification in a DiD also requires no anticipation and the stable unit treatment value assumption (SUTVA). In our setting, no anticipation most likely holds, as households probably did not anticipate the energy crisis. We nevertheless perform a robustness check where we exclude the year 2021 from our sample to address any concerns about possible anticipation (Section 5). For SUTVA to hold, we need to assume that of building i's energy consumption does not depend on the treatment status of building $j \neq i$. In other words, there should be no spillovers from the treatment status of one building on the energy consumption of other buildings. If our units of analysis were apartments, one may be concerned about spillover effects as it is well known that heating of one apartment has effects on the temperature in adjacent apartments. However, our analysis is at the building-level, which is why such concerns do not apply in our setting and we are confident that SUTVA holds.

However, prices could also increase if the contract period didn't end in the middle of the crisis; either because the energy provider went bankrupt or because they had to increase prices due to increased energy provision costs. Whether the energy provider went bankrupt or faced increased energy provision costs that they had to pass-on to consumers largely depended on their wholesale energy-provision contracts. It is unlikely that when concluding the energy contracts, landlords or building managers may have taken into consideration whether an energy provider buys its energy through short-term or long-term wholesale energy-provision contracts. Thus, we do not expect any building-level unobservables to be correlated with treatment assignment, because there was no active sorting into or out of treatment.

Nevertheless, treatment assignment may not have been entirely random. If in some areas energy providers relied more on long-term wholesale contracts than in others, then we may have an uneven distribution of treated and control buildings across Germany. This does not threaten the validity of the parallel trends assumption, as long as areas with higher and lower shares of treated buildings do not have different energy consumption trends. The event study results in Figure 2 indicate that such diverging energy consumption trends are most likely not present. To be on the safe side, we nevertheless additionally perform a robustness check where we do exact matching by "Raumordnungsregion" (hereon referred to as "spatial-planning region" or "region"), which is a geographic area defined for spatial planning purposes in Germany and lies between NUTS-2 and NUTS-3 regions.[15] This addresses the potential issue of diverging energy consumption trends between spatial-planning regions with higher and lower shares of treated buildings and additionally controls for any differences in region-level unobservables between treated and controls.

Even though any active sorting into or out of treatment seems unlikely, we choose to apply a DiD-PSM approach in order to balance out any observable differences between treatment and control group.

Figure 3 displays the distributions of the logit of the propensity score of treated and control buildings before and after matching. Additionally, Figure A.1 shows the covariate balance before and after matching. There is good overlap between the distributions of the treated and the control group before the matching (Figure 3 left panel) and most matching variables were already relatively balanced before the matching (Figure A.1). After matching, the two distributions are very similar, suggesting that the matching was successful in further improving the comparability of the treatment and the control group (Figure 3 Panel (b)). This is confirmed by Figure A.1 which shows that the matching procedure reduced the standardized mean difference between the treatment and the control group for all matching variables. The matching reduced the standardized mean difference below 0.1, which is the conventional threshold for good covariate balance (Austin, 2009), for all variables except for the average price in the pre-period. In the latter case, the matching nevertheless succeeded in reducing the standardized mean difference below 0.2. In the unmatched sample, the variance ratios of all variables were already between the conventional thresholds of 0.5 and 2 (**zhang_balance_2019**). Still, the matching further improved variance ratios.

[15] There are in total 96 spatial-planning regions in Germany.

Figure 3: Distributions of the Logit of the Propensity Score by Treatment Group

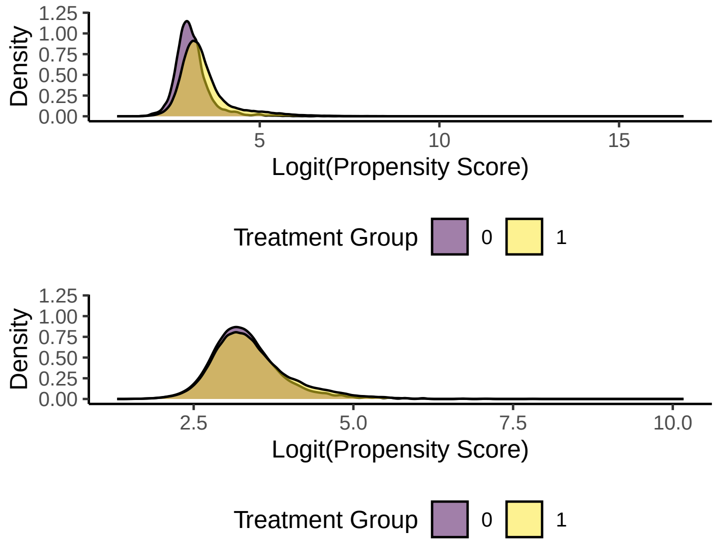

The figure shows the distribution of the logit of the propensity score by treatment group before (top panel) and after (bottom panel) matching.

Calculation of Elasticities Besides estimating the price-driven savings during the energy crisis, we are also interested in calculating the corresponding energy price elasticities. When percentage changes in prices or percentage changes in demand are large, it is common to use the arc formula for the price elasticity of demand Feehan (2018). The arc formula takes the percentage changes of price and demand relative to the midpoints of the intervals marked by the values before and after the demand and price changes:

$$\eta = \frac{\frac{e_{2022} - e_{2021}}{\frac{e_{2022} + e_{2021}}{2}}}{\frac{p_{2022} - p_{2021}}{\frac{p_{2022} + p_{2021}}{2}}} \tag{3}$$

To estimate the price-elasticity of energy demand during the crisis in 2022 we slightly adjust our main estimation equation 1. Instead of using the logarithm of energy consumption as the dependent variable, we now use a building's energy consumption normalized by the building's average energy consumption in 2021 and 2022:

$$\frac{e_{it}}{\frac{e_{i,2021} + e_{i,2022}}{2}} = \lambda_i + \mu_t + \delta D_{it} + \boldsymbol{x}_{it}' \boldsymbol{\beta} + \epsilon_{it} \tag{4}$$

In equation 4, δ stands for the average price-driven energy savings of the treatment group as a percentage of the average energy consumption in 2021 and 2022 which corresponds to the denominator of equation 3. We obtain $\hat{\eta}$, the estimate of the price elasticity of energy demand, by dividing $\hat{\delta}$ by the average price increase of the treatment group in 2022, normalized by the treatment group's average price in 2021 and 2022.

4.1.2 Heterogeneous effects

Method 2: Double Machine Learning We use double machine learning (DML) proposed by Chernozhukov et al. (2018) which allows us to explore heterogeneities in price-driven savings in a more flexible way than it would be possible with our previous DiD-PSM approach. In particular, the DML framework allows us to model non-parametric effect heterogeneity so that we can obtain a detailed picture of how non-price savings vary with socio-economic variables.

In our application of DML we closely follow the implementation described in Knaus (2022). As double machine learning models have mainly been developed for cross-sectional data, we only use the year 2022 for our analysis. Our dependent variable is Δe_i, the change in energy consumption compared to 2021 in kWh. The treatment variable D_i is a binary variable taking the value of one if a building's energy price increased during the crisis and zero if the price stayed constant. We use a rich set of control variables including the change in the climate factor up to the third polynomial, the one-year lag of the climate factor, the average energy price before the energy crisis, the change in the homeoffice rate, several building characteristics (number of apartments per building, average apartment size, energy carrier) and a set of socio-economic variables (average birthyear, share of social benefit recipients, unemployment rate, average net income and average years of schooling).
We first flexibly predict building-specific treatment probabilities, $\hat{\phi}(X_i)$, using a Lasso-model and five-fold cross-fitting. Then we predict the building-specific treatment-specific outcome (i.e. the change in energy consumption), $\hat{\pi}(d, X_i)$, again using a Lasso-model and five-fold cross-fitting. In a third step we plug the predicted $\hat{\phi}(X_i)$ and $\hat{\pi}(d, X_i)$ into the following formula for the doubly robust score of the potential outcome:

$$\hat{\Gamma}_i(d, X_i) = \hat{\pi}(d, X_i) + \frac{J(d)(\Delta e_i - \hat{\pi}(d, X_i))}{\hat{\phi}(X_i)} \tag{5}$$

$d \in \{0, 1\}$ takes the value of 0 to indicate the untreated state and 1 to indicate the treated state and the variable $J(d) = \mathbb{1}(D_i = d)$ indicates whether a building i's actual treatment status corresponds to d. For each building i, we thus obtain two doubly robust potential outcome scores: one potential outcome score under treatment $\hat{\Gamma}_i(1, X_i)$ and one potential outcome without treatment $\hat{\Gamma}_i(0, X_i)$. We then take the difference of the two doubly robust potential outcome scores to obtain a pseudo-outcome for each building:

$$\hat{\Delta}_i = \hat{\Gamma}_i(1, X_i) - \hat{\Gamma}_i(0, X_i) \tag{6}$$

The pseudo-outcome $\hat{\Delta}_i$ can be interpreted as a building-specific treatment effect. In a final step, we can therefore run a non-parametric kernel regression of $\hat{\Delta}_i$ on socio-economic variables to explore effect heterogeneity.

Identification DML identification relies on an unconfoundedness or conditional independence assumption Knaus (2022). This differs from the identifying assumption in our previous DiD-PSM approach which relies on parallel trends. Conditional independence is more difficult to defend than parallel trends. As discussed in section 4.1.1, it is very unlikely that buildings could have actively sorted into or out of treatment. We are therefore confident that there are no relevant unobserved building-level differences between the treatment and the control group. The double machine learning that flexibly controls for a rich set of covariates further reduces the likelihood of any remaining omitted variable bias.

4.2 Estimation of Non-Price-Driven Savings

Method 3: Lasso Prediction of Counterfactual Energy Consumption Our estimation of price-driven savings during the energy crisis rely on using a control group of buildings that were not exposed to price hikes to construct a counterfactual. Such natural control group is not available for the estimation of non-price-driven savings as all German households were exposed to the crisis. We therefore need to rely on a different method to estimate the counterfactual heat energy consumption in 2022 for a hypothetical scenario without the energy price crisis. We train a Lasso-model with ten-fold cross-validation on the years before the energy crisis[16]. The dependent variable is a building's energy consumption e_{it} and the explanatory variables are buildings' one-year lagged energy consumption, the local homeoffice rate and climate factors up to the third polynomial. We do not use further building-level characteristics as explanatory variables because the building's condition should be implicitly accounted for by controlling for the previous year's energy consumption. We then insert the 2022 climate factors, homeoffice rate and the energy consumption of 2021 into the Lasso-model trained on the pre-period. This provides an estimate of the counterfactual energy consumption in 2022, $\hat{e}_{i,2022}$, i.e. for a counterfactual scenario where the the energy price crisis had not happened.

$\widehat{\Delta e}_{i,2022}$ is the difference between the counterfactual energy consumption, $\hat{e}_{i,2022}$, and the observed energy consumption, $e_{i,2022}$. $\widehat{\Delta e}_{i,2022}$ has different interpretations for buildings with constant prices in the crisis and for buildings where prices increased. For buildings with constant prices we denote this difference as $\widehat{\Delta e}^{0}_{i,2022}$. This object is an estimate for building-level non-price driven savings (see discussion of identification below). For buildings where prices increased, the difference $\widehat{\Delta e}^{1}_{i,2022}$ needs to be interpreted as the sum of non-price driven savings and price-driven savings and is therefore not of much interest, as the two cannot be disentangled.

By taking the average of $\widehat{\Delta e}^{0}_{i,2022}$, we obtain an estimate of average non-price-driven savings for buildings with constant prices. Similarly to the final step of the DML procedure, we can now also regress $\widehat{\Delta e}^{0}_{i,2022}$ on socio-economic variables using non-parametric kernel regressions to explore heterogeneities in non-price savings during the energy crisis.

Identification Identification of the non-price-driven savings relies on the argument that short-term changes in energy consumption can only come from a limited set of factors. Besides the crisis-related non-price factors such as public appeals and saving programs, these factors

[16] As one of the explanatory variables is the buildings' lagged energy consumption, we can only use the four pre-treatment years 2018-2021 as compared to the DiD-PSM approach where have five pre-treatment years from 2017-2021.

are retrofits, changes in the annual temperature, reactions to price-changes and changes in homeoffice patterns[17]. Other factors such as socio-economic changes or changes in living habits are more sticky and will therefore only have a negligible effect on year-on-year changes in energy consumption.

As we exclude buildings that have been retrofitted from our sample and only focus on buildings that had a constant price from 2021 to 2022, changes in energy consumption in 2022 can only come from changes in temperature, changed homeoffice rates or the multiple non-price-factors related to the energy crisis such as appeals to save and energy savings programs. Since our machine learning model flexibly accounts for climate factors and for the homeoffice rate, the difference between the predicted counterfactual energy consumption and the observed actual energy consumption should correspond to the crisis-related non-price-driven energy savings.

5 Results and Discussion

5.1 Price-driven Savings and Heat Energy Price Elasticities

5.1.1 Average Price-Driven Savings

We estimate average price-driven energy savings using a DiD-PSM approach, as outlined in section 4.1.1. We find that energy savings in buildings where heat energy prices increased were on average 2.2 pp. larger than in control buildings where prices stayed constant (see Table 2 Column (a)). To corroborate our estimated price-driven savings, we perform a number of robustness checks, which all yield similar results. We present them at the end of this section.

Germany achieved the EU savings target of 15% in the energy crisis (Bundesnetzagentur, n.d.), which is also reflected in our sample, where we observe total heat energy savings of 16.0%. However, our results also show that price-driven savings of 2.2% only made a minor contribution. More important factors were non-price factors such as public appeals and savings programs (for a more detailed discussion and estimation results for non-price-driven savings please refer to section 5.2) and the weather. This is relevant for energy policy in times of crisis as it implies that sufficient short-term savings to avoid energy scarcity are unlikely to be achieved through high energy prices alone. This is particularly the case in countries with high information frictions in the residential energy market such as Germany (see section 2). If information frictions were lower, price-driven savings might have been higher. In 2022, Germany introduced regulation to promote the uptake of smart meters in the residential energy sector which may contribute to making households more price-responsive in future crises.

The estimated price-driven savings correspond to a short-term arc-elasticity of demand of -0.07, which we compute as described in section 4.1.1. This elasticity is smaller than the price elasticity found by other studies during the crisis that do not control for non-price-driven savings (Ruhnau et al., 2023). This underlines the importance of adequately controlling for non-price

[17] Changes in homeoffice patterns are particularly relevant when comparing the year 2021 which had several Covid-lockdowns and stay-at-home regulations with the year 2022 when Covid-restrictions have been largely waived

factors such as public appeals and attention, savings programs and regulations when estimating price-driven savings and energy price elasticities during times of energy crises. Not doing so might inflate estimated price elasticities. The results by Jamissen et al. (2024), who also control for one of the non-price factors (public attention), point in a similar direction, i.e. they find a much lower elasticity than Ruhnau et al. (2023). Similar to price-driven savings, the low short-run energy price elasticity is partly due to the information frictions in Germany's residential heat energy market. Another reason for our relatively low elasticity estimate is that we only estimate a short-run elasticity. Short-run energy price elasticities tend to be lower than long-run price elasticities in general because som price-induced adjustments such as energy efficiency investments take time to be implemented.

A point to note regarding the interpretation of the price-driven savings as well as the energy price elasticities is that they only reflect to what extent households reacted to actual building-level price changes. If households reacted to a general discourse about high energy prices to which the control group was equally exposed, then this effect will not be reflected in the estimated price-driven savings and elasticities.

With the objective of understanding if the low price responsiveness prevails across all levels of price increases, we estimate price-driven savings for three subgroups: buildings with a price increase of <25%, 25% - 50%, and >50% (as described in more detail in Section 4.1.1). Columns (b) - (d) of Table 2 and Figure 4 present the DiD-PSM results for the three subgroups. For buildings with a <25% price increase, we do not observe a statistically significant treatment effect. Buildings with price increases of 25% - 50% saved 2.2 pp. more compared to their untreated counterparts. For buildings where prices increased by >50% we observe the highest energy savings relative to the control group of 4.4 pp. The event study plots for for DID-PSM by subgroup indicate that parallel trends also hold for different levels of price-increase (see Appendix Figure A.3). Even though we find that households are not particularly price-responsive in the short run, we nevertheless see a clear pattern of larger savings in buildings with higher price increases.

However, these different price-driven savings of the three subgroups all translate to an energy-price elasticity of -0.07 (see Appendix Figure A.2a) meaning that the price elasticity of demand is quite constant with increasing prices.[18]

A subsequent subsample analysis shows that the average price-driven savings in buildings with district heating were 5.1% which is notably higher than the average price-driven savings of 2.1% in gas-heated buildings (see Figure 5). Again, the event study plots for these DiD-PSM subsample analyses indicates that parallel trends hold (see Appendix Figure A.4). This is also reflected in a higher price elasticity of demand for buildings with district heating. While we find a price elasticity of gas demand of -0.06, the elasticity for heat energy demand in buildings with district heating is almost three times as large with -0.17. There are several possible reasons for

[18] Note that the standard errors of the treatment group with price increases of up to 25% is comparatively large. The reason is the denominator in the arc elasticity calculation reflecting the average price increase between 2021 and 2022, meaning that $\hat{\delta}$ in equation 4 is simply divided by a smaller number than in the other two treatment groups with larger price increases.

Table 2: DiD-PSM: Full sample and by treatment intensity.

Dependent Variable:	log(Heat Energy Consumption)			
	(a)	(b)	(c)	(d)
Model:	Main DiD	%Δp<=25%	25%<%Δp<=50%	%Δp>50%
Variables				
Treated	−0.022***	−0.006	−0.022***	−0.044***
	(0.005)	(0.004)	(0.006)	(0.004)
Control variables	Yes	Yes	Yes	Yes
Fixed-effects				
Building	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	746,056	312,363	224,127	261,337
R^2	0.881	0.886	0.885	0.874

Control variables comprise homeoffice and the first to third polynomial of the climate factor.
Standard-errors are clustered at the building level and are given in parentheses.
** p <0.05, ** p <0.01, *** p <0.001.*

this stark difference. Gas-heated buildings have individual contracts with prices depending on a number of contract-specific factors such as the energy provider as well as the timing of entering into the contract. This means that two buildings next to each other might be paying completely different gas prices. In contrast, district heating, as its name reveals, usually covers an entire district where a single price is applicable throughout. This implies that the local newspapers might pick up on energy price increases supplied with district heating. While gas price increases were emphasized heavily in the national news, reports on increased district heating prices could have alerted households of the increased local prices and thereby reinforced the response to the energy crisis.

We perform a battery of robustness checks to confirm the validity of our results. In the first robustness check we exclude the year 2021. Gazprom started to reduce the supply of natural gas already in 2021, which caused increases in wholesale prices, although to a lesser extent than the price hikes observed in the second half of 2022 (Ruhnau et al., 2023). Due to our data structure, where prices are yearly average prices, this may bias the estimated price-driven savings toward zero: Assume, for example, that a building had to renew its contract just before the beginning of the heating period in the second half of 2021 with a price increase of 10%. Then only half of its annual consumption will be billed with the new higher price and the price increase that we observe in the annual data will be roughly 5%. In 2022, the full year will be billed with the new higher price. Compared to the average price of 2021, this will again correspond to a price increase of about 5%. Consequently, this building would be categorized as treated with a price increase in 2022 even though the price increase already happened in the second half of 2021. To

Figure 4: Price-driven savings by treatment group

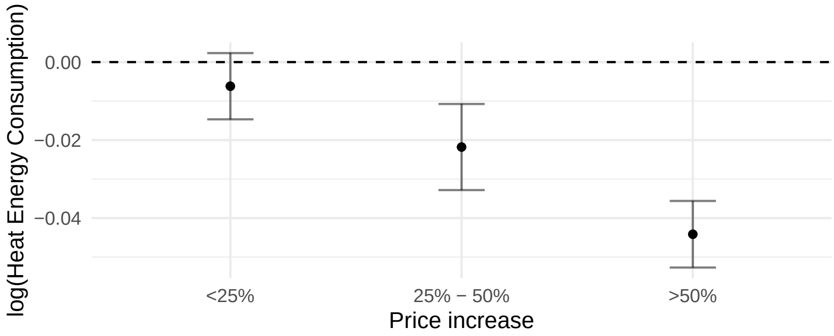

The figure shows estimates of price-driven energy savings during the crisis by treatment group. The coefficients indicate the estimated treatment effects on log energy consumption of being exposed to a price increase during the crisis. Vertical bars indicate 95% percent confidence intervals that are calculated using standard errors clustered at the building-level.

Figure 5: Price-driven savings by carrier

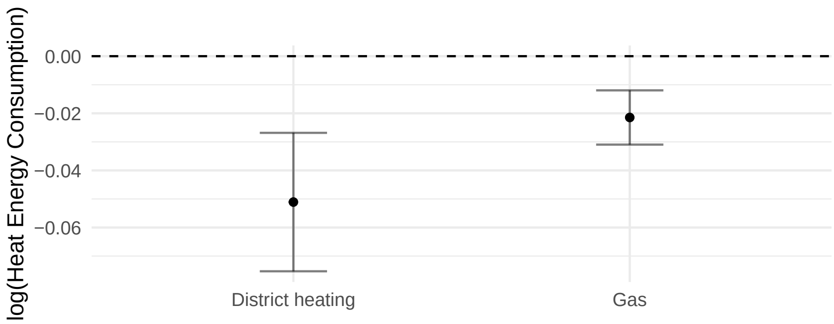

The figure shows estimates of price-driven energy savings during the crisis by energy carrier. The coefficients indicate the estimated treatment effects on log energy consumption of being exposed to a price increase during the crisis. Vertical bars indicate 95% percent confidence intervals that are calculated using standard errors clustered at the building-level.

avoid this potential issue, we drop the year 2021 in a first robustness check. Our estimated non-price-driven savings (Column (a) of Table A.2 in Appendix A.4) are not significantly different from the ones in our main specification, albeit smaller. This, however, is in line with our main result that price-driven savings only played a minor role in the overall savings during the crisis. In a second robustness check, we additionally enforce exact matches at the level of spatial-planning regions. This robustness check serves two purposes: First, it avoids that buildings from former East-Germany are matched to buildings from former West-Germany. Due to the historical differences, households in East- and West-German households may on average have different attitudes towards Russia, the War in Ukraine and in general different levels of trust in the media and government (Braun & Trüdinger, 2023). They may therefore have different non-price motives to save energy. Second, as discussed in the methodology section, some spatial-planning regions may have higher shares of treated buildings than others due to potentially different wholesale contract positions of the regional energy providers. By matching on spatial-planning regions, we thus also avoid that any region-level unobservables may bias our results. Column

(d) of Table A.2 in Appendix A.4 shows that the price-driven savings with matching at the level of spatial-planning region are, as in the first robustness check, not significantly different from our main specification. This, again, corroborates the result that price-induced savings did not constitute the lion's share of savings.

Third, we winsorize the energy consumption variable at the 5th and 95th percentile to address variation in energy consumption due to short-term changes in occupancies such as temporal vacancies due to vacation. The results for this robustness check are given in Column (c) of Table A.2 in Appendix A.4. They are very similar in magnitude to our main specification and again not significantly different from the price-driven savings that we estimate in our main specification.

Fourth, we perform 5:1 matching instead of 1:1 nearest neighbor matching. The results are shown in Column (b) of Table A.2 in Appendix A.4 and are again very similar to the price-driven savings from the main specification.

Finally, Figure A.5 in Appendix A.4 shows the event-study plots for all four robustness specifications. We do not observe significant pre-treatment coefficients in any robustness specification which is a further strong indication of the plausibility of the parallel trends assumption. Overall, our findings on price-driven savings appear to be robust against a number of potential concerns regarding the timing of the price increases, region-specific effects, as well as possible alternative methodological choices.

5.1.2 Socio-Economic Heterogeneities of Price-Driven Savings

To explore how price-driven energy savings vary with socio-economic characteristics, we regress the building-level pseudo outcomes $\hat{\Delta}_i$, that we estimated using the DML procedure, on socio-economic characteristics using a non-parametric local polynomial regressions. Figure 6 shows the estimated relationships between price-driven savings and four socio-economic variables: net income per person, the unemployment rate, age and years of schooling. The reason we include the unemployment rate as a variable of interest in this part of the analysis is that unemployed people do not pay for their own heating and therefore may have different incentives to save energy than those who do pay their own bills.

Price-driven savings seem to increase with the average net income per person (Figure 6 Panel (a)). This surprising observation may be due to the fact that richer households consume more energy per person which leads to a higher savings potential. However, the difference in price-driven savings between buildings with lower income vs. buildings with higher income appears not to be significant and should therefore be interpreted with caution.

The other three socio-economic variables of interest do not exhibit any clear visible relationship with estimated price-driven savings (Figure 6 Panel (b)-(d)).

Figure 6: Heterogeneity of price-driven savings

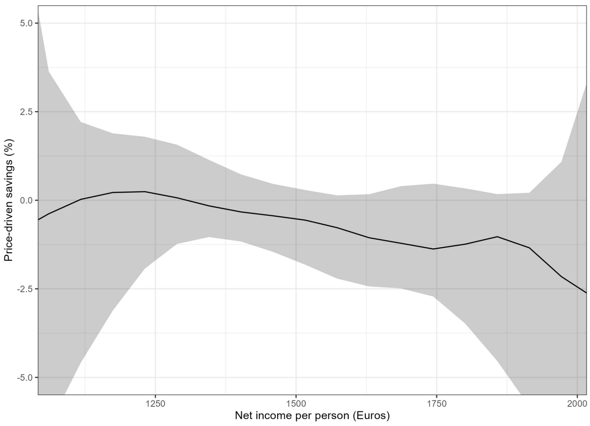

(a) Net income per person

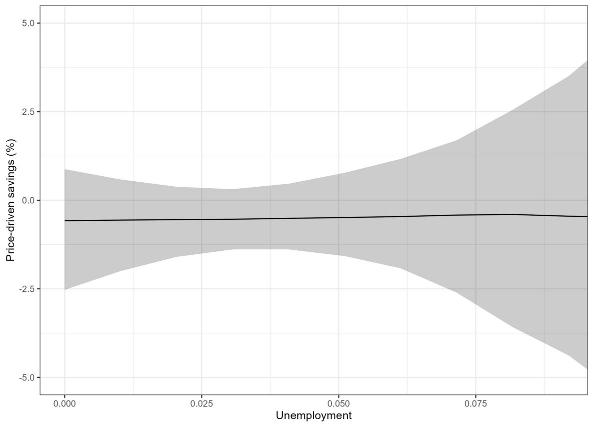

(b) Unemployment

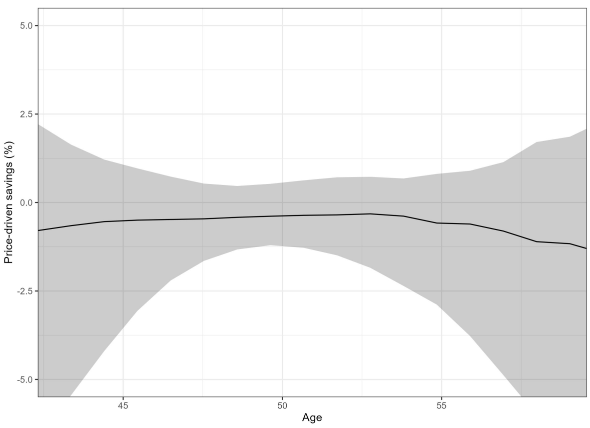

(c) Age

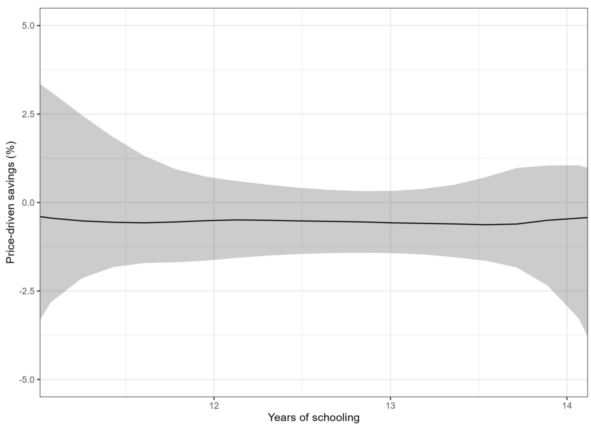

(d) Years of schooling

The figure shows how price-driven savings vary with different socio-economic characteristics. The black line is an estimated non-parametric local polynomial regression function that we obtain by regressing DML-pseudo outcomes $\hat{\Delta}_i$ on socio-economic characteristics. The shaded areas indicate 95% confidence intervals. Note that the confidence intervals only reflect the uncertainty in fitting the non-parametric regression on the estimate pseudo outcomes $\hat{\Delta}_i$. It does not include the uncertainty from the first step of estimating the building-level $\hat{\Delta}_i$ using DML.

5.2 Non-price-driven savings

5.2.1 Average non-price driven savings

To estimate non-price-induced savings we proceed as described in Section 4.2: We first predict counterfactual energy consumption in 2022 had the crisis not happened. For buildings with constant prices the difference between observed and predicted counterfactual energy consumption constitutes non-price-induced savings. Across buildings with constant prices in the crisis, the average building-level non-price-driven savings are on average 8.5 % of a building's energy consumption in 2021. In comparison to price-driven savings, the saving's effect of non-price factors was therefore almost four times as large. This underscores the importance of non-price factors to trigger short-term savings in the recent crisis.

5.2.2 Socio-Economic Heterogeneities of Non-Price-Driven Savings

We regress estimated non-price-driven savings on different socio-economic variables, again using non-parametric local polynomial regressions. Figure 7 shows how non-price driven energy savings vary with several socio-economic characteristics. Non-price-driven savings do not seem to vary with income, age and years of schooling (see Figure 7 Panel (a), (c), (d)).

Non-price-driven savings first increase with the unemployment rate and then appear to decrease again at very high unemployment rates (Figure 7 Panel (b)). However, this relationship should be interpreted with caution due to the large confidence intervals of the fitted regression function at higher unemployment rates.

Figure 7: Heterogeneity of non-price-driven savings

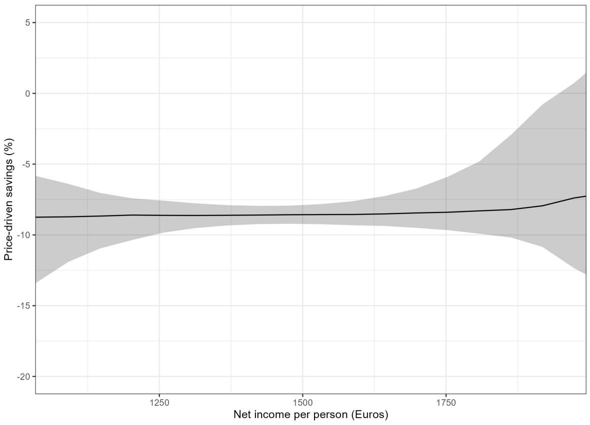

(a) Net income per person

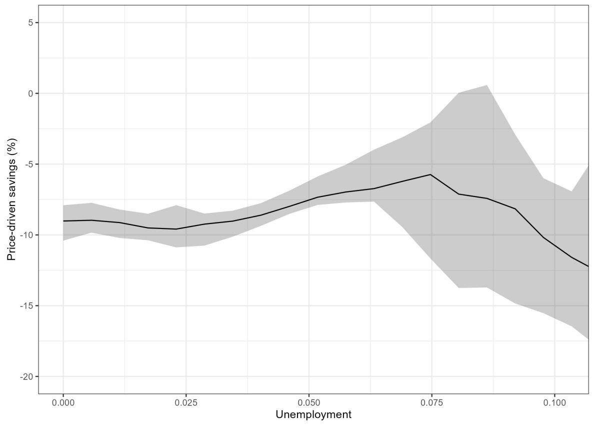

(b) Unemployment

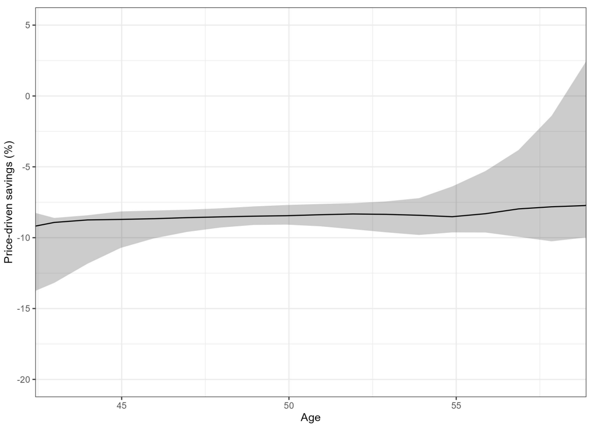

(c) Age

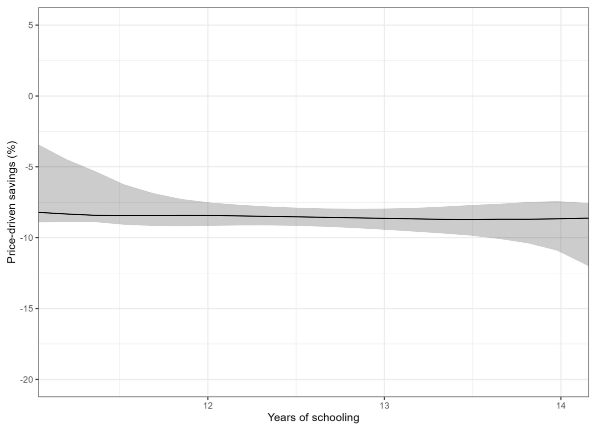

(d) Years of schooling

The figure shows how non-price-driven savings vary with different socio-economic characteristics. The black line is an estimated non-parametric local polynomial regression function that we obtain by regressing estimated non-price-driven savings on socio-economic characteristics. The shaded areas indicate 95% confidence intervals. Note that the confidence intervals only reflect the uncertainty in fitting the non-parametric regression on the estimated non-price-driven savings. It does not include the uncertainty from the first step of estimating the building-level non-price-driven savings.

6 Conclusion

In this paper, we examine the effects of the energy crisis in 2022 and the contemporaneous appeals and programs to save energy on energy consumption in Germany. We apply three different methods - DiD-PSM and two machine learning based approaches - to estimate price-driven energy savings, energy price elasticities and non-price-driven energy savings during the crisis. We also analyze how both types of energy savings vary with various socio-economic characteristics. Our analysis is based on a unique dataset of residential building-level heat energy prices and consumption that we combine with administrative data on socio-economic characteristics from the German microcensus.

Our findings confirm that rising energy prices contributed to a reduction in residential heat energy demand. However, we find that price-driven savings were, on average, relatively low compared to total heat energy savings during the crisis, even after controlling for the higher temperatures in 2022. The low price-driven savings translate to a low short-term energy price elasticity during the crisis of -0.07. The flip-side of the modest price-driven response is that the majority of observed savings were driven by non-price factors highlighting the significant role of public appeals and energy saving programs in achieving the observed short-term energy-savings.

Our analysis carries several methodological and policy implications. Methodologically, we propose and implement an approach for estimating energy price elasticities that prevents the overestimation of elasticity measures by comprehensively controlling for contemporaneous non-price-driven responses using an adequate control group. Moreover, we use double machine learning to analyse heterogeneities in non-price-driven savings and apply another machine learning based approach to estimate non-price-driven savings. On the policy side, our findings offer valuable insights into emergency interventions and the design of public policy during periods of energy crises to achieve short-term energy savings to avoid shortages.

Crucially, our results suggest that simply allowing energy prices to escalate in times of energy shortages is insufficient for achieving the necessary short-run energy demand reductions in the residential sector. This is particularly the case for countries such as Germany, where the residential heat energy sector is marked by strong information frictions. In the recent crisis, public appeals and energy saving programs have played a much larger role than price increases in achieving short-term heat energy savings in Germany's residential sector.

References

Austin, P. C. (2009). Balance diagnostics for comparing the distribution of baseline covariates between treatment groups in propensity-score matched samples. *Statistics in Medicine*, *28*(25), 3083–3107.

Austin, P. C. (2011a). Optimal caliper widths for propensity-score matching when estimating differences in means and differences in proportions in observational studies. *Pharmaceutical Statistics*, *10*(2), 150–161.

Austin, P. C. (2011b). An Introduction to Propensity Score Methods for Reducing the Effects of Confounding in Observational Studies. *Multivariate Behavioral Research*, *46*(3), 399–424.

Bachmann, R., Baqaee, D., Bayer, C., Kuhn, M., Löschel, A., Moll, B., Peichl, A., Pittel, K., & Schularick, M. (2022). *What if? The Economic Effects for Germany of a Stop of Energy Imports from Russia* (Research Report No. 028). ECONtribute Policy Brief.

Braun, D., & Trüdinger, E.-M. (2023). Communal and Exchange-Based Trust in Germany Thirty Years After Reunification: Convergence or Still an East–West Divide? [Publisher: Routledge]. *German Politics*.

Brons, M., Nijkamp, P., Pels, E., & Rietveld, P. (2008). A meta-analysis of the price elasticity of gasoline demand. A SUR approach. *Energy Economics*, *30*(5), 2105–2122.

Bundesministerium für Wirtschaft und Klimaschutz. (2022). Breites Bündnis ruft zum Energiesparen auf.

Bundesnetzagentur. (n.d.). Rückblick: Gasversorgung im Jahr 2022.

Bundesnetzagentur. (2023). Bundesnetzagentur veröffentlicht Zahlen zur Gasversorgung 2022.

Bundesnetzagentur. (2024). Bundesnetzagentur - Preise und Abschläge.

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., & Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, *21*(1), C1–C68.

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

Destatis. (2018). Dwellings, by type of heating and energy.

Espey, J. A., & Espey, M. (2004). Turning on the Lights: A Meta-Analysis of Residential Electricity Demand Elasticities. *Journal of Agricultural and Applied Economics*, *36*(1), 65–81.

European Council. (2022). Council adopts regulation on reducing gas demand by 15% this winter.

Federal Ministry of Economic Affairs and Climate Action. (2022a). Ordinances on energy saving.

Federal Ministry of Economic Affairs and Climate Action. (2022b, December). Bundestag adopts energy price caps – important relief for consumers.

Feehan, J. P. (2018). The long-run price elasticity of residential demand for electricity: Results from a natural experiment. *Utilities Policy*, *51*, 12–17.

Goldenberg, R. (2022). Germany implements energy-saving rules – DW – 09/01/2022. *Deutsche Welle.*

Hanemann, M., Labandeira, X., Labeaga, J., & Lopez-Otero, X. (2013). *Energy Demand for Heating: Short Run and Long Run.*

Havranek, T., Irsova, Z., & Janda, K. (2012). Demand for gasoline is more price-inelastic than commonly thought. *Energy Economics*, *34*(1), 201–207.

Ito, K., Ida, T., & Tanaka, M. (2018). Moral Suasion and Economic Incentives: Field Experimental Evidence from Energy Demand. *American Economic Journal: Economic Policy*, *10*(1), 240–267.

Jamissen, D., Vatne, J., Holz, F., & Neumann, A. (2024). The price elasticity of natural gas demand of small consumers in Germany during the energy crisis 2022. *Energy Efficiency*, *17*(8), 98.

Knaus, M. C. (2022). Double machine learning-based programme evaluation under unconfoundedness. *The Econometrics Journal*, *25*(3), 602–627.

Kröger, M., Longmuir, M., Neuhoff, K., & Schütze, F. (2023). The price of natural gas dependency: Price shocks, inequality, and public policy. *Energy Policy*, *175*, 113472.

Labandeira, X., Labeaga, J. M., & López-Otero, X. (2017). A meta-analysis on the price elasticity of energy demand. *Energy Policy*, *102*, 549–568.

Moll, B., Schularick, M., & Zachmann, G. (2023). The Power of Substitution: The Great German Gas Debate in Retrospect.

More, R., & Carrel, P. (2022). Germans told to unite over energy crisis amid EU divisions. *Reuters.*

Ó Broin, E., Nässén, J., & Johnsson, F. (2015). The influence of price and non-price effects on demand for heating in the EU residential sector. *Energy*, *81*, 146–158.

Perino, G. (2015). Climate Campaigns, Cap and Trade, and Carbon Leakage: Why Trying to Reduce Your Carbon Footprint Can Harm the Climate. *Journal of the Association of Environmental and Resource Economists*, *2*(3), 469–495.

Reiss, P. C., & White, M. W. (2008). What changes energy consumption? Prices and public pressures. *The RAND Journal of Economics*, *39*(3), 636–663.

Roth, A., & Schmidt, F. (2023). Not only a mild winter: German consumers change their behavior to save natural gas. *Joule*, *7*(6), 1081–1086.

Rubin, E., & Auffhammer, M. (2024). Quantifying Heterogeneity in the Price Elasticity of Residential Natural Gas [Publisher: The University of Chicago Press]. *Journal of the Association of Environmental and Resource Economists*, *11*(2), 319–357.

Ruhnau, O., Stiewe, C., Muessel, J., & Hirth, L. (2023). Natural gas savings in Germany during the 2022 energy crisis. *Nature Energy*, *8*(6), 621–628.

Schmitz, H., & Madlener, R. (2020). Heterogeneity in price responsiveness for residential space heating in Germany. *Empirical Economics*, *59*(5), 2255–2281.

Schulte, I., & Heindl, P. (2017). Price and income elasticities of residential energy demand in Germany. *Energy Policy*, *102*, 512–528.

Statista. (2023). Housing conditions by dwelling type in Germany 2022.

Sullivan, A. (2022). Russian gas in Germany: A 50-year relationship – DW – 03/09/2022.

The Press and Information Office of the Federal Government. (2022). Federal Cabinet: Additional energy-saving measures approved.

Trotta, G., Hansen, A. R., & Sommer, S. (2022). The price elasticity of residential district heating demand: New evidence from a dynamic panel approach. *Energy Economics*, *112*, 106163.

Umweltbundesamt. (2022). Economizing during the energy crisis.

U.S. Energy Information Administration. (2022). Europe is a key destination for Russia's energy exports.

VERIVOX. (2025). Gaskrise in Deutschland bis 2023.

Wettengel, J. (2024). Germany, EU remain heavily dependent on imported fossil fuels.

Xiang, D., & Lawley, C. (2019). The impact of British Columbia's carbon tax on residential natural gas consumption. *Energy Economics*, *80*, 206–218.

A Appendix

A.1 Theoretical Framework

We first present a simple theoretical model of energy consumption that describes a household's utility maximization problem which consists in allocating its budget between consuming energy and a consuming a composite consumption good. The particular feature of the model is that it allows for a specific disutility from energy consumption in 2022 reflecting the non-price driven motivations to save energy. In this model, a household derives utility from consuming energy e and a composite consumption good x. The household's utility from consuming energy also depends on non-monetary motivations to save energy. These non-monetary motivations can be influenced by public campaigns or government appeals for saving energy. Following Perino (2015), we model this using a utility function $u(e, x, m)$, where m is a parameter representing the effect of campaigns or appeals to save energy. A change in m changes the marginal rate of substitution between e and x. The price of energy is p and the composite consumption good is the numeraire. The household has income w that it can spend on e and x. For ease of exposition we again follow Perino (2015) and use a constant elasticity of substitution (CES) utility function:

$$u(e, x, m) = \left[\frac{1}{1+m}e^r + \frac{m}{1+m}x^r\right]^{1/r} \tag{7}$$

An increase in m reduces the weight in the utility function on e and increases the weight on x. An increase in m thus represents a campaign or appeal with the message to reduce energy consumption. Maximizing $u(e, x, m)$ subject to the budget constraint $pe + x = w$ yields the optimal energy consumption

$$e^* = \frac{p^{-\sigma}w}{m^\sigma + p^{1-\sigma}} \tag{8}$$

Taking the differential of e^* gives

$$de^* = \underbrace{-\frac{\sigma m(mp)^{\sigma-1}+1}{(m^\sigma + p^{1-\sigma})^2}}_{\frac{\partial e^*}{\partial p}} dp \underbrace{-(m^\sigma p^{1-\sigma}))^2 \sigma m^{\sigma-1}}_{\frac{\partial e^*}{\partial m}} dm \tag{9}$$

First note that both partial derivatives are negative because $\sigma, m, p \geq 0$ and $\sigma, m \geq 0$.

$$\frac{\partial e^*}{\partial p} = -\frac{\sigma m(mp)^{\sigma-1}+1}{(m^\sigma + p^{1-\sigma})^2} \leq 0 \tag{10}$$

$$\frac{\partial e^*}{\partial m} = -(m^\sigma p^{1-\sigma}))^2 \sigma m^{\sigma-1} \leq 0 \tag{11}$$

Intuitively, this means that the optimal consumption of energy e^* decreases with rising price p as well as with stronger public campaign m. Dividing both sides of the differential of e^* in equation 9 by dp gives us the total derivative of energy consumption with respect to the price of energy

$$\frac{de^*}{dp} = \underbrace{-\frac{\sigma m(mp)^{\sigma-1}+1}{(m^\sigma + p^{1-\sigma})^2}}_{\frac{\partial e^*}{\partial p}} \underbrace{-(m^\sigma p^{1-\sigma)})^2 \sigma m^{\sigma-1}}_{\frac{\partial e^*}{\partial m}} \frac{dm}{dp} \tag{12}$$

We easily see that the total derivative $\frac{de^*}{dp}$ only corresponds to the partial derivative $\frac{\partial e^*}{\partial p}$ if m is unrelated to p, i.e. if $\frac{dm}{dp} = 0$. However, $\frac{dm}{dp}$ was larger than zero during the 2022 energy price crisis when the price hikes were met with calls from several public actors to reduce energy consumption (i.e. an increase in m). During the crisis, the total derivative of $\frac{de^*}{dp}$ therefore corresponded to the combined partial effects of $\frac{\partial e^*}{\partial p}$ and $\frac{\partial e^*}{\partial m}$ (times $\frac{dm}{dp}$).

Consequently, regressing e^* on p without explicitly modeling m would yield an estimate of the effect of p on e^* which is biased by the effect of m on e^*. We therefore explicitly model non-monetary motivations in our empirical strategy to be able to properly isolate the effect of prices on energy consumption.

Similarly, equation 13 shows that the total derivative $\frac{de^*}{dm}$ only corresponds to the partial derivative $\frac{\partial e^*}{\partial m}$ if $\frac{dp}{dm} = 0$. As mentioned, this was not the case in the crisis. Hence, when estimating the effect of public appeals on energy consumption, or - more generally - estimating non-price-driven savings, one needs to carefully control for the contemporaneous effect of price hikes. To ensure that our estimated non-price-driven savings are not biased by the price hikes, we only use buildings where prices stayed constant during the crisis for the estimation of non-price driven energy savings (see Section 4.2).

$$\frac{de^*}{dm} = \underbrace{-\frac{\sigma m(mp)^{\sigma-1}+1}{(m^\sigma + p^{1-\sigma})^2}}_{\frac{\partial e^*}{\partial p}} \frac{dp}{dm} \underbrace{-(m^\sigma p^{1-\sigma)})^2 \sigma m^{\sigma-1}}_{\frac{\partial e^*}{\partial m}} \tag{13}$$

The model's results provide several empirical predictions. First, it trivially predicts that rising energy prices should lead to reduced energy consumption. Second, it projects that households reduce their energy consumption in 2022 compared to 2021 even if their prices stayed constant and that these savings vary with the strength of the non-monetary savings motives. And finally it shows that in order to isolate the partial effect of energy prices on energy consumption one needs to adequately control for non-monetary savings.

A.2 Further Matching Diagnostics

Figure A.1: Loveplot of the covariate balance

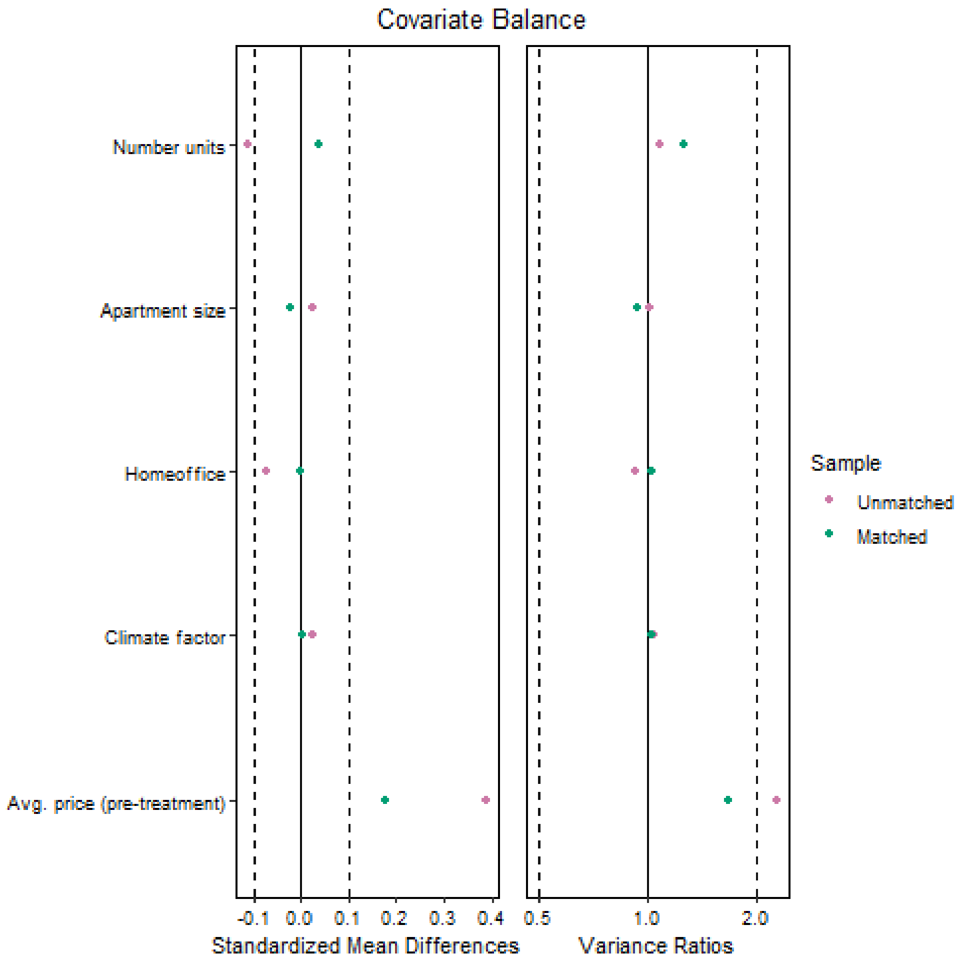

The figure shows standardized mean differences and variance ratios between the treatment and the control group before and after matching.

A.3 Results by Treatment Group and by Energy Carrier

Table A.1: DiD-PSM by type of carrier.

Dependent Variable:	log(Heat Energy Consumption)	
Model:	Gas	District Heat
Variables		
	Gas	District Heat
Treated	−0.021***	−0.051***
	(0.005)	(0.012)
Controls variables	Yes	Yes
Fixed-effects		
id	Yes	Yes
Year	Yes	Yes
Fit statistics		
Observations	647,462	98,594
R^2	0.867	0.911
R^2 Adj.	0.836	0.890
R^2 Within	0.011	0.016

Control variables comprise homeoffice and first to third polynomial of the climate factor.
Standard-errors are clustered at the building level and are given in parentheses.
** p <0.05, ** p <0.01, *** p <0.001*

Figure A.2: Elasticity by treatment group and carrier.

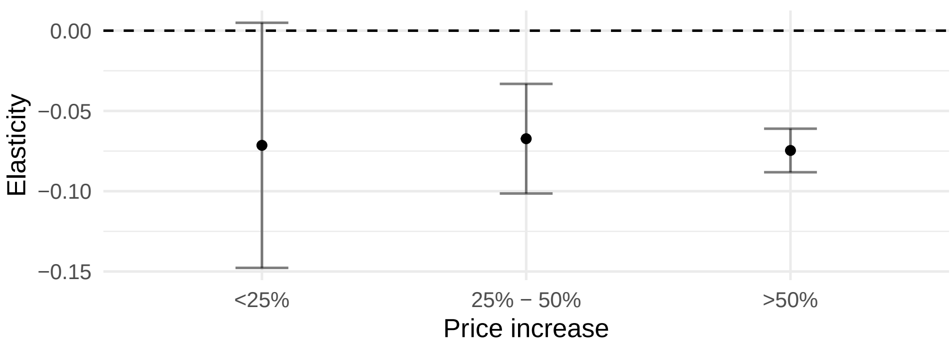

(a) Elasticity by treatment group.

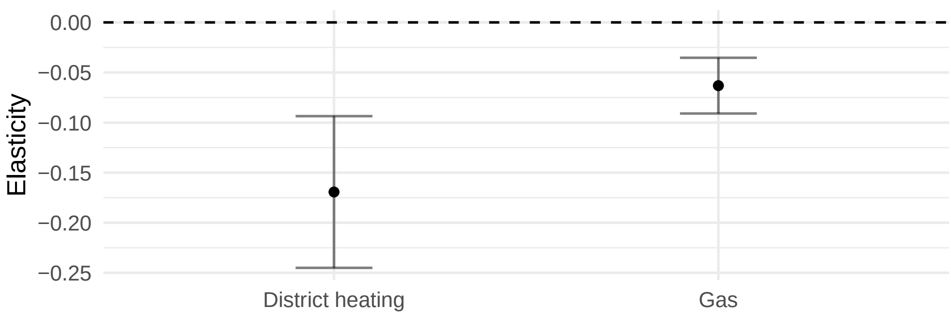

(b) Elasticity by energy carrier.

The figure shows estimates of the short-term energy price elasticity during the crisis by treatment group (Panel (a)) and by energy carrier (Panel (b)). Vertical bars indicate 95% percent confidence intervals that are calculated using standard errors clustered at the building-level.

Figure A.3: Event studies by treatment group.

(a) Event study for buildings with price increases of <25%

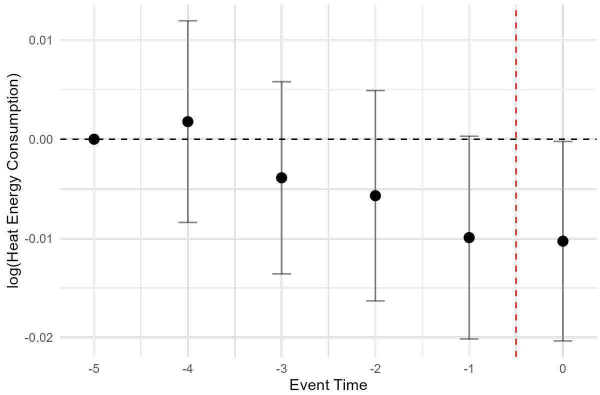

(b) Event study for buildings with price increases of 25% - 50%

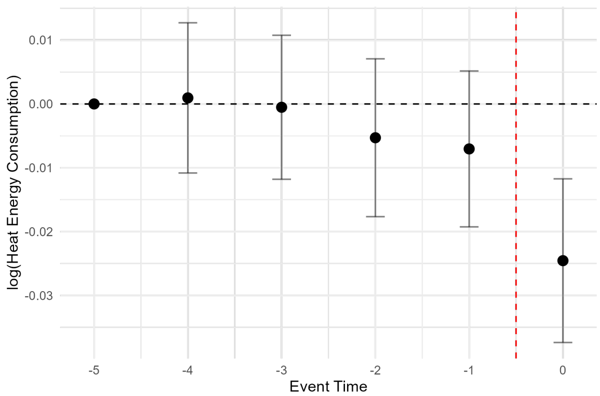

(c) Event study for buildings with price increases of >50%

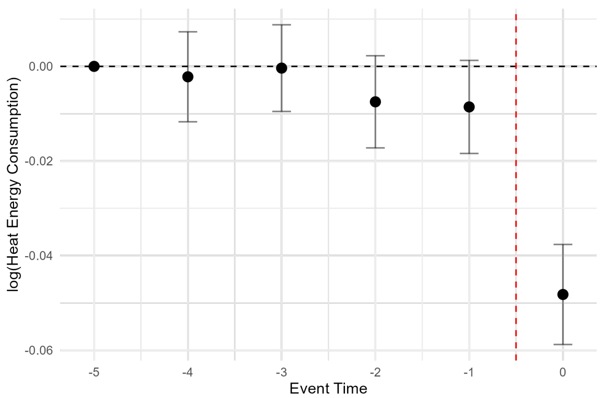

The figure shows the event study coefficients from estimating equation 2 on the matched sample of control buildings and treated buildings with price increases of <25% (Panel (a)), with price increases of 25% - 50% (Panel (b)) and with price increases of >50% (Panel (c)). The x-axis displays the time to treatment in years, where 0 corresponds to the energy crisis in 2022. The displayed event study coefficients indicate the estimated year-specific treatment effects and placebo treatment effects. Vertical bars indicate 95% percent confidence intervals that are calculated using standard errors clustered at the building-level.

Figure A.4: Event studies by carrier.

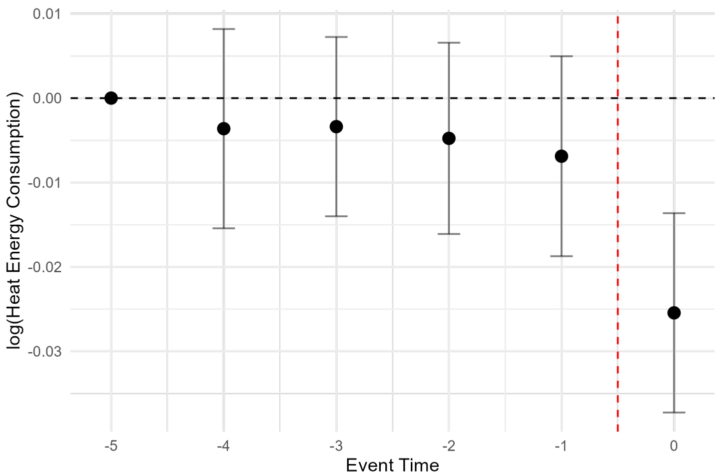

(a) Event study for gas.

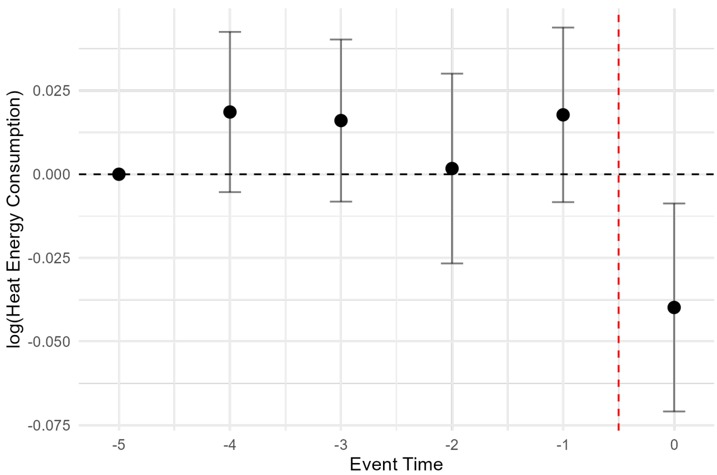

(b) Event study for district heat.

The figure shows the event study coefficients from estimating equation 2 on the matched sample of buildings heated with gas (Panel (a)) and buildings with district heating (Panel (b)) and with price increases of >50% (Panel (c)). The x-axis displays the time to treatment in years, where 0 corresponds to the energy crisis in 2022. The displayed event study coefficients indicate the estimated year-specific treatment effects and placebo treatment effects. Vertical bars indicate 95% percent confidence intervals that are calculated using standard errors clustered at the building-level.

A.4 Robustness

Figure A.5: Robustness: Event study plots of different DiD-PSM specifications

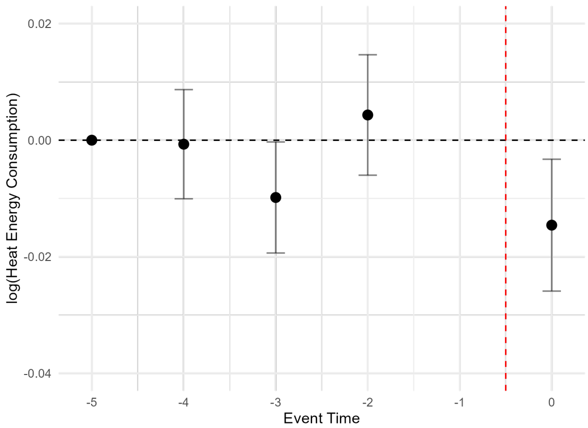

(a) Excluding the year 2021.

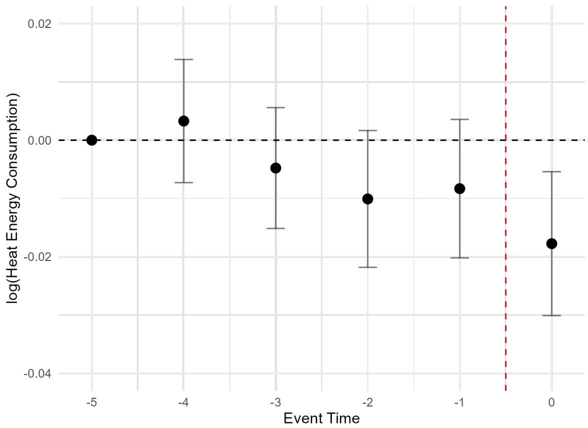

(b) Exact matching on the regional level.

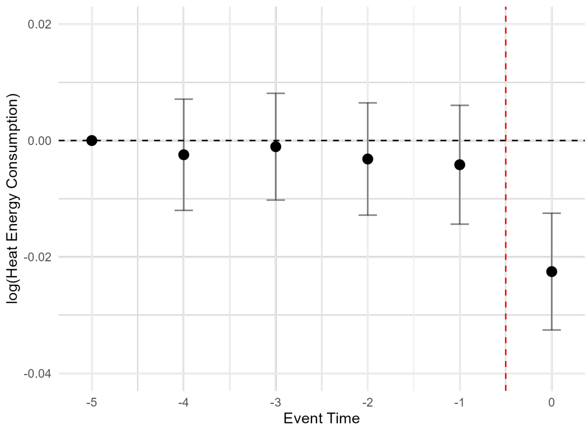

(c) Outcome variable winsorized.

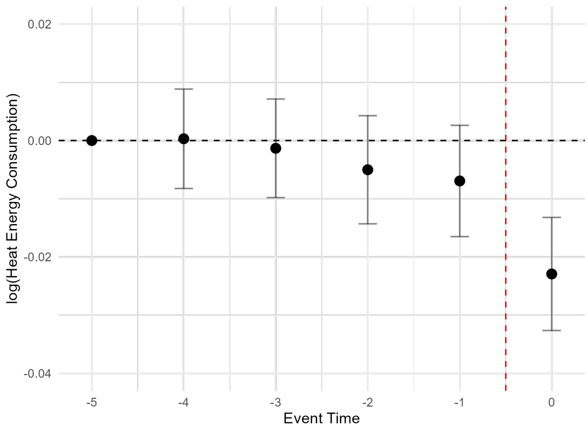

(d) Five to one matching.

The figure shows the event study coefficients from various DiD-PSM robustness checks. Panel (a) shows the main specification excluding the year 2021. Panel (b) shows the event study when doing an exact match on the regional level. Panel (c) shows the event study when winsorizing the outcome variable. Panel (d) shows using 5:1 matching. The x-axis displays the time to treatment in years, where 0 corresponds to the energy crisis in 2022. The displayed event study coefficients indicate the estimated year-specific treatment effects and placebo treatment effects. Vertical bars indicate 95% percent confidence intervals that are calculated using standard errors clustered at the building-level.

Table A.2: Robustness check: different DiD-PSM specifications.

Dependent Variable:	log(Heat Energy Consumption)			
Model:	Excluding 2021	5:1 Matching	Winsorized	Regional matching
Variables				
Treated	-0.013***	-0.020***	-0.020***	-0.013***
	(0.005)	(0.004)	(0.004)	(0.0046)
Control variables	Yes	Yes	Yes	Yes
Fixed-effects				
id	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Fit statistics				
Observations	597,136	746,362	746,056	636,618
R^2	0.883	0.881	0.888	0.876
Within R^2	0.011	0.012	0.013	0.011

Control variables comprise homeoffice and the first to third polynomial of the climate factor
Standard-errors are clustered at the building level and are given in parentheses.
** p <0.05, ** p <0.01, *** p <0.001.*