Firms in the energy crisis: evidence from 2021-22 *

Matteo Alpino[†a], Luca Citino[‡a], and Annalisa Frigo[§a]

[a]Bank of Italy

April 2024

Abstract

We study the causal impact of energy price shocks on Italian industrial firms during the 2021-22 energy crisis. For identification we exploit the staggered expiration of fixed-price energy contracts, which generates large increases in the retail price of energy paid by firms. Upon contract expiration, firms do not reduce electricity demand and reduce gas demand only in the second half of 2022. In that period, the estimated price elasticity of gas demand is -1.1 on average, but it is much lower for gas-intensive firms. Additional evidence suggests that energy demand reductions are accommodated via both input substitution and output reduction.

(*JEL classification*: Q41)

* A previous version of this paper which analyzed only data on 2021 was circulated as Bank of Italy Occasional paper under the title "The effects of the 2021 energy crisis on medium-sized and large industrial firms: evidence from Italy" (Alpino et al., 2023). We are extremely thankful to Andrea Petrella for help in designing the survey questions and to Fabrizio Balassone for early encouragement to work on this topic. We thank Antonio Accetturo, Fabio Bertolotti, Marco Bottone, Federico Cingano, Guido de Blasio, Domenico Depalo, Alberto Felettigh, Sara Formai, Rossella Greco, Andrea Lamorgese, Andrea Linarello, Francesca Lotti, Giovanni Marin, Enrico Sette, Francesco Vona and seminar and conference participants at the Bank of Italy, the ESCB Cluster 2 workshop, EIEF, the EEA conference, the EMEE workshop, Frisch Center, the IAERE conference, the CEPR-NBB Workshop on Net-Zero, Oslo Metropolitan University and Roma Tre University. Matteo Alpino is grateful to EIEF for hosting him while working on this paper. All errors are our own. The opinions expressed in this paper do not necessarily reflect those of the Bank of Italy.

[†]Contact: matteo.alpino@bancaditalia.it

[‡]Contact: luca.citino@bancaditalia.it

[§]Contact: annalisa.frigo@bancaditalia.it

1 Introduction

How do firms cope with a large and sudden upsurge in energy prices? Until recently, this question had little practical relevance, but the 2021-2022 energy crisis has changed this perspective, sparking serious concerns with regard to its consequences on the European economy.[1] Understanding how such episodes affect firms and how they react is key to inform policymakers on the most efficient way to manage future energy crises. Beyond its relevance for emergency periods, addressing this question carries important implications for the green transition, given the steadily increasing prices of fossil fuels along the path towards achieving Net Zero.

In this paper, we answer the question above by providing fresh evidence from the 2021-22 European energy crisis, when the wholesale price of natural gas rose from around 30 euros per megawatt hour (MWh) to 300, triggering a similar increase in the price of electricity.[2] Existing works on this episode have relied on timely but aggregate time-series data (Ruhnau et al., 2022; Corsello et al., 2023; Alessandri and Gazzani, 2023; Moll et al., 2023), while papers using micro data on previous time periods (von Graevenitz and Rottner, 2022; Fontagné et al., 2023; Gerster and Lamp, 2023) often leverage comparably smaller energy price shocks and thus may not be readily applicable when trying to understand how firms navigate severe energy crises. On the contrary, we use survey micro data on manufacturing firms covering the period 2021-22, which allow us to identify large energy price shocks to firms and their effects on their energy demand, input substitution, capacity utilization, and final output prices.[3]

For identification, we exploit the fact that wholesale energy prices are transmitted only gradually and partially to the actual energy prices paid by firms. Our novel empirical strategy is based on the staggered expiration of fixed-price energy contracts at the firm level, which we proxy using ad-hoc survey questions. We use a staggered difference-in-differences approach to compare firms that experience an idiosyncratic increase in their energy costs with those that will experience the same event in the future, or that never will during our observation window. In order to estimate the effect, we mainly rely on the imputation estimator proposed

[1] For example, a joint statement by European industrial energy consumers in March 2022 reads: *"The events [ed. the invasion of Ukraine] have further precipitated Europe in a profound energy crisis that compromises the future of Europe's industrial base and the independence of its economy."*.

[2] The design of the wholesale electricity market is such that the price is set by the marginal producer, which is a gas-fired power plants in most cases.

[3] The survey is the *Indagine sulle imprese industriali e dei servizi* (Inquiry into investments of industrial and services firms), an annual survey conducted by the Bank of Italy since 1984. Over the years, it has been used to address a number of research questions (Pozzi and Schivardi, 2016; Rodano et al., 2016; Guiso and Parigi, 1999; Bond et al., 2015; Schivardi et al., 2021).

by Borusyak et al. (2021), but our results are robust to alternative difference-in-differences estimators (De Chaisemartin and dHaultfoeuille, 2020; Callaway and SantAnna, 2021; Sun and Abraham, 2021) and to the synthetic difference-in-differences approach (Arkhangelsky et al., 2021).

We show that the expiration of a fixed-price contract generates unanticipated and persistent increases in the retail price of electricity and natural gas at the firm level, even after accounting for government policies that partially mitigated price increases. For the average firm, retail prices increase up to 47% in the case of electricity and 29% in the case of natural gas. This result is relevant in and of itself, because it highlights the importance of fixed-price inventory contracts, routinely used by firms for purchasing several inputs (Kumar and Wesselbaum, 2024), in the transmission of macro shocks.

In response to higher retail energy prices, firms do not change their demand for electricity. Instead, they reduce their demand of natural gas, but *only* in the second half of 2022, irrespective of the cohort of treatment. Pooling effects across cohorts, we find that in that period the average firm reduces gas demand by 34%. We argue that this delayed response in natural gas adjustment is because of the pessimistic expectations about the continuation of the crisis emerging in the summer months of 2022, when the spot price reached unprecedented levels (more than 300 euros per MWh) and futures markets were forecasting that the gas price would stay at very high levels (around 200 euros per MWh) at least until mid-2023. These patterns of behaviour are consistent with a model of adjustment costs (Pindyck and Rotemberg, 1983; Atkeson and Kehoe, 1999), where (perceived) temporary shocks lead to inaction (as in second half of 2021 and the first of 2022), while (perceived) permanent ones prompt action (like in the second semester 2022).

We conduct an extensive heterogeneity analysis with random forests and show that the drop in gas consumption in the second half of 2022 is larger (-41%) among firms who declare that this input is *not* essential in their production process, and smaller (-28%) for those who declare that gas is an essential input (the difference is significant at the 10% level). Interestingly, firms declaring that gas is essential constitute a sizable group (more than 50%), spread over many economic sectors, and include almost all gas intensive firms (so called *gasivore*). Among the latter group, made of a thousand firms accounting for 20 per cent of national natural gas consumption, the downward adjustment in the second half of 2022 is even smaller (-8%).

We then proceed to calculate price elasticities of electricity and gas demand by rescaling the quantity effects by the price effects estimated above. As for electricity, this parameter is zero on average and not different across different types of firms. While this appears in sharp contrast with other estimates from recent studies on the pre-crisis years (e.g. between -0.4 and -0.6 on average in Marin and Vona (2021); Fontagné et al. (2023); von Graevenitz and Rottner (2022)), recent estimates using large shocks on large electricity consumers find smaller elasticities (between -0.09 and -0.2 in Gerster and Lamp (2023)). Moreover, the studies above also find smaller elasticities for larger shocks and in more recent periods, consistently with our results.

As for natural gas we find an average demand elasticity equal to -1.1, higher compared to the older literature (Labandeira et al., 2017), but broadly in line with the recent findings of Fontagné et al. (2023) (between -0.9 and -1.2). In addition, we find substantial heterogeneity in the gas elasticity across firms: -0.5 for firms for which natural gas is an essential input and -2.5 for other firms; -0.03 for gas intensive firms, -1.3 for the other firms. Our elasticity estimates have direct policy implications. Among other things, we can use these numbers to estimate to what extent price-distorting support measures used during the crisis have increased natural gas demand. Similarly to Deryugina et al. (2020), we use a simple and standard tax incidence formula to study how the equilibrium consumption of natural gas changes after the introduction of a per-unit quantity subsidy. We show that the heterogeneity in demand elasticities that we uncover bear sizable implications for the equilibrium consumption response, provided that the supply elasticity is not too low.

In the second part of the paper, we investigate the potential mechanisms behind these patterns of adjustment in energy demand. On the one hand, firms that adjust gas consumption downwards can substitute it with other inputs and/or reduce their output. On the other hand, firms that do not reduce their energy consumption can shift their cost increases to the prices of their final goods and/or reduce their profit margins.

In order to look at input substitution, we rely on administrative data on plants subject to the European Emission Trading System (EU ETS), which record physical quantity use of all energy and non-energy fossil fuels at an annual frequency starting from 2018. We find that, upon expiration of fixed-price contracts, EU ETS plants decrease gas demand and increase the use of other fossil fuels. On net, our point estimates indicate that this substitution accounts for at most half of the energy content that is lost with decreased gas consumption, though these effects are rather imprecise. Substitution is incomplete at best, as treated firms are unable or unwilling to completely offset the natural gas drop by using other fossil fuels.

In order to look at the physical output produced by firms, their profit margins, and prices of their final goods, we turn to a longer yearly panel of our survey covering the 2018-2022 period. Here we exploit quantitative survey questions on plant capacity utilization, the growth rate of final good prices and a categorical variable indicating profit levels. Upon expiration of a fixed-price contract, in 2022 plant capacity utilization decreases by less than 2 p.p. (not significant) from an average baseline of 80%. Also, we detect no significant heterogeneity across different types of firms. Conversely, being exposed to the energy shock reduces the probability to have a positive profit margin by 10 percentage points in 2022. This effect is sizable as in our sample about 80% of firms declare they make profits. When looking at the growth rate of final prices, in 2022 we find that the shock leads to a decrease of 2.7 percentage points, not significant. The average price *increase* in our sample was 11% in the same year and no firms *decreased* their price. The empirical pattern in our data is consistent with a model of *strategic complementarities*, where rivals' shocks matter for one's price adjustment, or a model of price taking, where firms' final output prices do not depend on idiosyncratic cost shocks, but just market-wide cost shocks (Duprez and Magerman, 2018; Amiti et al., 2019; Muehlegger and Sweeney, 2022). As for prices, some interesting patterns emerge among gas intensive firms: upon exposure to the shock, they do increase significantly their final good prices. This indicates that pass-through was a successful strategy for this type of firms which did not cut gas consumption by much; this results is in line with recent evidence in Lafrogne-Joussier et al. (2023).

Our paper contributes to different strands of the literature. First, we contribute to the literature on the effects of the 2021-2022 energy crisis on firms. In this respect, recent macro studies have either looked at the impact crisis on inflation and on output using models or aggregate data (Alessandri and Gazzani, 2023; Bachmann et al., 2022; Moll et al., 2023). In this paper we provide the first firm-level evidence on the impact of the ongoing energy crisis on industrial firms' input demand, together with other relevant response margins, highlighting the role of treatment effect heterogeneity, also in the time dimension.

Second, we contribute to the literature estimating natural gas and electricity demand elasticities for firms. While there is more credible evidence for households (Reiss and White, 2005; Jessoe and Rapson, 2014; Auffhammer and Rubin, 2018; Hahn and Metcalfe, 2021), estimates for industrial firms are more limited and traditionally used instrumental variables relying on sector or time-series variation.[4] Our work complements and is very much related to very recent pre-crisis contributions trying to leverage across-firm variation in energy prices to estimate demand elasticities (Marin and Vona, 2021; Fontagné et al., 2023; von Graevenitz

[4]Previously used instruments for the price of energy in a demand equation include: price paid by the household or industry sector only (Burke and Abayasekara, 2018; Csereklyei, 2020), domestic natural gas reserves and distance weighted reserves in other countries (Burke and Yang, 2016), lagged prices (Graf and Wozabal, 2013), the spot price of Brent crude oil (Davis and Muehlegger, 2010), weather shocks (Hausman and Kellogg, 2015), wholesale prices Faiella et al. (2022) among others.

and Rottner, 2022; Gerster and Lamp, 2023). In our work we propose a new and credible identification strategy based on the availability and staggered expiration of fixed-price contracts and other hedging instruments already in place *before* the crisis, which slow down the transmission of wholesale to retail energy prices for some firms, more than others. This prevents our effects to be confounded by aggregate contemporaneous shocks.

Third, our results speak to the potential effects of higher energy prices during the energy transition. How firms react to energy input price shocks is related to how they react when facing carbon pricing schemes (Martin et al., 2014; Cui et al., 2018; Colmer et al., 2023; Martinsson et al., 2024).

Finally, the paper is related to the literature on the pass-through of cost shocks (Ganapati et al., 2020; Amiti et al., 2019; Muehlegger and Sweeney, 2022; Lafrogne-Joussier et al., 2023). While we do not estimate a pass-through rate due to data limitations, we show that during the crisis firms did not differentially increase their prices depending on their own idiosyncratic shocks, with the exception of gas intensive firms.

The paper is structured as follows. In Section 3 we describe the data, how we validate survey answers against administrative sources and national accounts, together with the measurement of key variables for our analysis. In Section 2 we provide background on the energy crisis in Italy. In Section 5 we describe in detail our identification strategy. In Section 6 we show the main results on how firms' average unitary energy costs react to the expiration of fixed-price energy contracts, and its consequences for energy demand. In Section 7 we show how firms react not only by changing their energy demand, but also by setting new prices, changing the utilization rate of their plants and substituting away from more expensive energy inputs. In Section 8 we illustrate the policy implications of our results with some simple incidence calculations and in Section 9 we conclude.

2 Background: 2021-22 energy crisis in Italy

One year before the beginning of the energy crisis (in 2020), Italy was a net importer of natural gas: imports accounted for 93% of gross inland consumption. Of these imports, Russia accounted for 43%, making it a key supplier. Italy was also a net importer of electricity, and natural gas accounted for approximately half of domestic power generation. The high reliance on natural gas in electricity production coupled with the marginal price system at work in the day-ahead power market implies that shocks to the wholesale price of natural gas almost completely pass-through to the wholesale price of electricity.[5] Figure 2.1 plots the evolution of the wholesale price of natural gas in Italy. The price was rather stable at low levels until mid-2021 (around 30 euro per MWh). After that, it slowly started to rise above historical levels. The first major upswing occurred in the fall of 2021, when the price went above 100 euros per MWh; the second took place in December of the same year, when it almost reached 200 euros. After a temporary drop, the price rose up to around 250 in February when Russia invaded Ukraine, but rapidly decreased thereafter and stabilized below 100 until the summer of 2022. At that point, the price climbed again quite rapidly, reaching a historical peak at over 300 euros in late August 2022. Before the end of the year, the price dropped to much lower levels, and then up again, before a final descent to 70.

Figure 2.1: Wholesale gas price (PSV) at the daily frequency

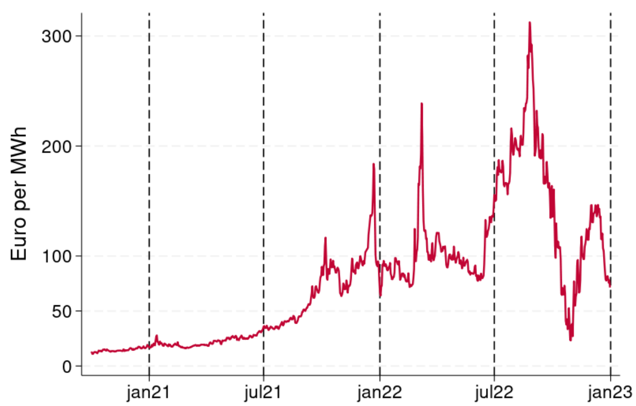

Source: GME. *Note*: The figure shows the spot price of natural gas traded on the Italian trading point.

3 Data

The main data source used in this paper is the *Indagine sulle imprese industriali e dei servizi* (Inquiry into investments of industrial and services firms; henceforth, Invind), an annual survey conducted by the Bank of Italy since 1984 and representative of industrial and services firms with at least 20 employees. The Bank conducts the survey between February and May of every year t and contains information on standard firm-level variables such as sales, profits, employment, costs, capacity utilization, actual and expected own price changes and actual and expected investment in year $t-1$. Invind data have been used before in the literature to address a number of research questions[6] and it is also routinely used by the Bank of Italy

[5]The two time series of gas and electricity prices are indeed highly correlated (98% at the daily frequency).
[6]These include the impact of productivity and demand shocks on firms' growth (Pozzi and Schivardi, 2016), bankruptcy law and bank financing (Rodano et al., 2016), the determinants of investment demand (Guiso and

to provide timely evidence on sales and investment dynamics as well as other issues in its official reports.

For the purpose of this study, at the end of 2021 we designed an *ad hoc* section on the 2021 energy crisis, administered in the spring of 2022 (henceforth, 2021 wave) only to industrial firms with 50 employees or more. The following year, at the end of 2022, we designed a new survey wave that was administered to firms in the spring of 2023 (henceforth, 2022 wave). We restrict our sample to firms that use energy as an *input*, that is we drop NACE sectors 19 (manufacture of coke and refined petroleum products) and 35 (electricity, gas, steam and air conditioning supply).

The main advantage of this survey is that we could gather timely information on firms' energy expenditures and consumption, together with hedging strategies against the crisis, which we could directly link to other more standard firm-level variables. This is particularly useful as no other representative firm-level data source on 2021-2022 is available at the moment for research purposes.

3.a Structure of the energy survey section

The 2021 wave contained nine quantitative questions, while the 2022 wave contained twelve quantitative questions. We report all of them exactly as they appeared to firms in Appendix A. The results of the analysis based on wave 2021 alone were carried out during 2022 and reported in a Bank of Italy working paper (Alpino et al., 2023). Questions in the second wave were designed after conducting such analysis.

In both waves, the survey asks firms to indicate both expenditures (in thousands of €) and physical quantities (in MWh and standard cubic meters) for electricity and natural gas purchases separately, both during the first and second semester of the previous year (2021 and 2022).[7] Dividing expenditures by physical quantities at the half-yearly frequency allow us to construct firm-level average unitary costs (retail energy prices, henceforth) for electricity and natural gas, separately for each semester of 2021 and 2022.

Importantly for our identification strategy, in both waves we ask firms questions about their fixed price contracts or equivalent hedging tools that firms were endowed with *before* the start of the crisis (more precisely at the beginning of 2021). We explain in detail how we use these variables to build our treatment variable in Section 3.e.

Parigi, 1999; Bond et al., 2015), mechanisms behind agglomeration economies (Andini et al., 2013) and the role of management practices during the Covid-19 pandemic (Schivardi et al., 2021)

[7] This purposefully excludes self-production of energy by firms.

In the 2022 wave we also collect additional information on subsidies received under the tax credit scheme implemented by the Italian government in 2022 to mitigate the impact of the energy crisis on firms. This is the single largest policy implemented in Italy to cushion firms against higher prices and the only one that would not be directly visible in the energy bills that we observe in the survey (such as cuts in VAT and administrative and environmental fees). In Section 6 we show that our treatment variable is still strongly predictive of changes in average unitary costs, *net* of any government transfers, eliminating concerns that our treatment may not have enough power.[8]

3.b Other administrative data sources

We supplement the survey data with other confidential administrative information, which we match through firms' unique tax identification numbers. We gather information on whether firms have at least one plant subject to the EU Emissions Trading System (EU ETS). For these factories we have detailed input use by fuel at the yearly frequency from the Italian Institute for Environmental Protection and Research (ISPRA). We use these data on fuel consumption both to validate the gas consumption measures in our survey, and to study the substitutability away from gas towards other inputs. In addition, we use micro-level energy consumption data from the Fund for Energy and Environmental Services (CSEA) on firms which are eligible for energy subsidies, because of their high electricity or gas intensity and levels of consumption (*energivore* firms).[9] For electricity intensive firms we observe electricity consumption (in MWh) at the monthly frequency since 2018.[10] For gas intensive firms, we observe natural gas consumption (in standard cubic meters) at the monthly frequency in 2019, 2021 and 2022.[11] These companies, slightly less than 4,000 in Italy, belong to the right tail of the energy-intensity distribution.[12] Similarly to the ISPRA data, we also use the data

[8]Estimates included in the Bank of Italy Annual Report for 2022 indicate that the Italian tax credit helped industrial firms to reduce their average unitary costs for electricity and natural gas by 13 and 18%, respectively.

[9]The registry is publicly available on the website of the Fund for energy and environmental services (portale elettrivori, Cassa per i servizi energetici e ambientali, CSEA).

[10]The subsidies have been in place for several years and grant a permanent discount on the component of the electricity price that is earmarked to finance subsidies for renewable energy generation (*oneri di sistema A3*SOS*). This component was completely lifted for all firms starting from January 2022, while in the last three quarters of 2021 it was lifted for low-voltage consumers (e.g. households and small firms).

[11]The concept of gas intensive firms was introduced during the crisis to identify firms more exposed to the energy shock. Since 2023, these firms enjoy a permanent discount on the component of the natural gas price that is earmarked to finance policies aimed at reaching emission reductions to meet EU goals (*componente* RE^{IG} and RE^{TIG}).

[12]In order to qualify for the subsidy scheme, firms must consume at least 1 GWh of electricity or gas per year. In addition they must belong to a specific set of 4-digit NACE industrial sectors defined by the EU regulation on State Aid; for a sub-set of these sectors there is the additional requirement of having the ratio between energy expenditure and value added or sales higher than a threshold.

from CSEA to validate our survey measures.

Finally, we use additional information from the Italian National Institute of Statistics (Istat) on the energy intensity of the Italian industrial sectors at the level of 2-digit NACE industries and Eurostat data on average retail prices and consumption for industrial consumers by consumption bracket. These data are useful to validate our survey measures, as succinctly described in the next subsection. All of the validation analyses are reported in Appendix C.

3.c Validation of survey answers

Given that respondents might not be familiar with physical units of energy, we validate our measurement by verifying whether self-reported quantities of gas and electricity or the respective expenditures take plausible values. Intuitively, excessively high or low figures suggest that respondents get the order of magnitude wrong, e.g. kWh instead of MWh. To this end, we rely on two benchmarks: (i) Eurostat data on average unitary prices for non-household consumers; and (ii) Invind data on the ratio between total energy costs and turnover.

In 22% (21%) of our observations on electricity (natural gas) we detect a systematic mistake in the units of measurement that can be adjusted by re-scaling the values in the first wave of the survey. The next year, such shares become 22.5% and 18% respectively. Whenever we cannot reconcile the replies with plausible units of measurement, we adopt a precautionary approach and disregarded such observations from the estimation samples; this is the case for 6.1% (6.3%) of the observations on electricity (natural gas) consumption in the 2021 wave and 3.8% (6.2%) of the observations in the 2022 wave. In general, the comparisons with quantities consumed as recorded by administrative sources (for energy intensive or EU ETS firms) lend credibility to the accuracy of our raw data and of our adjustment algorithm. Only 2.8% (4.5%) of the electricity sample and 2.1% (2.5%) of the natural gas sample in 2021 (in 2022) were conservatively dropped due to differences larger than 35% in absolute value between survey and administrative records. We report more detailed information on our data validation algorithm in Appendix C.

3.d Non-response bias

Invind is a yearly business survey administered by the Bank of Italy since the 1980s, routinely used in the Bank's official reports and that has been used before in economic research to answer a variety of questions.[13] The Bank interviews the same set of firms every year, adjusting for exit. In order to maintain survey representativeness of the target population in each year, *unit* non-response is taken care of using a standardized raking post-stratification procedure. This amounts to adjusting survey weights ex-post such that marginal distributions of key variables are equal to those in the population (Bank of Italy, 2017).

Even absent issues related to unit non response, we need to consider that not all Invind respondents provide valid replies to our *ad hoc* energy sections, generating an *item* non response issue. In order to mitigate concerns that this generates bias in Appendix D we use an inverse-probability-weighting strategy (Wooldridge et al., 2002; Stantcheva, 2022) and show that our results remain unaltered even after assigning larger weights to firms having attributes that make them more likely not to respond to the energy sections.

3.e Measurement of key variables

In this section we illustrate how we use answers to our survey questions to build cohort-of-treatment dummies that we use in our staggered difference-in-differences design.

In the 2021 wave of the survey, we ask: *"At the beginning of 2021, did your firm own any instrument that protected it, wholly or partly, from the energy price increases over the second half of the year?"*. This comes with four possible replies: (a) No (b) Yes, fixed-price contracts (c) Yes, financial derivatives (d) Yes, other instrument. Given that the vast majority of protected firms use fixed-price contracts, we collapse all "Yes" answers and construct a dummy variable called I_i^{2021} (=1 if protected).

Two things are worth noticing. First, the formulation specifies that the question refers to contracts already in place *at the beginning of 2021*, a time when the markets did not foresee the upcoming crisis.[14] This is key to ensure that the question does not pick up firms that subscribe fixed price contracts as an endogenous response to the crisis. Second, notice that this question conflates both protection from increases in electricity and in natural gas prices, due to space constraints in the survey. The fact that we cannot measure these separately can introduce measurement error. Reassuringly, this variable is strongly predictive of changes in average unitary cost of both electricity and natural gas between the first and second semester of 2021.[15]

[13]See (Pozzi and Schivardi, 2016; Rodano et al., 2016; Guiso and Parigi, 1999; Bond et al., 2015; Schivardi et al., 2021)

[14]As late as March 2021, future markets were expecting the TTF price of natural gas next October to be in line with the average value for the same month in the previous five years (around 15 euro). By July, the expectation had climbed to 36 euro, still well below the final realized price of 76.

[15]The correlation is stronger for electricity. This is consistent with evidence from the Bank of Italy Survey

In the 2022 wave of the survey, this time separately for electricity and natural gas, we ask: *"In 2022, did your firm have instruments (for example fixed-price contracts or derivatives) to protect itself, even partially, from rises in the prices of electricity (natural gas)?"*, with two possible replies: (a) Yes (b) No. Similarly to before, for a given input $j = \{\text{electricity, natural gas}\}$, we construct a dummy called $I_i^{j,2022}$ taking value 1 if the firm had any instruments. Furthermore, separately for electricity and natural gas, we ask: *"If yes, how many months did this protection last in 2022?"*, with open answer. We call this variable m_i^j, again for a given input $j = \{\text{electricity, natural gas}\}$.

Restricting our attention to firms present in both waves, we combine answers from both questions to build an input-specific treatment cohort variable E_i^j, indicating the semester (h) when the firm is first exposed to higher prices for a given input $j = \{\text{electricity, natural gas}\}$ (0 if it is never exposed in the observation window). The variable is constructed as follows:

$$E_i^j = \begin{cases} 2021h2, & \text{if } I^{2021} = 0 \text{ and } I^{j,2022} = 0 \\ 2022h1, & \text{if } I^{2021} = 1 \text{ and } I^{j,2022} = 0 \\ 2022h2, & \text{if } I^{2021} = 1 \text{ and } I^{j,2022} = 1 \text{ and } m_i^j = 6 \\ 0, & \text{if } I^{2021} = 1 \text{ and } I^{j,2022} = 1 \text{ and } m_i^j = 12 \end{cases} \tag{1}$$

The interpretation of this variable goes as follows. We consider all firms to be protected in $2021h1$, as the crisis has not yet started. Firms with $I^{2021} = 0$ and $I^{j,2022} = 0$ are immediately exposed ($E_i = 2021h2$) because they were not protected nor in 2021 nor in 2022. We call these firms the "early treated". Firms with $I^{2021} = 1$ and $I^{j,2022} = 0$ are first exposed in the first semester 2022 ($E_i = 2022h1$). We call these firms the "mid treated". Firms with $I^{2021} = 1$ and $m^j = 6$ are "late treated": protected during 2021 and during 6 months in 2022. We assume these are the first 6 months of the year and set $E_i = 2022h1$. In other words, we assume that at the beginning of 2021 the firm was holding a fixed-price contract that eventually expired in mid-2022. Note that this is the only sensible assumption because it was almost impossible to purchase insurance against higher energy prices after Russia invaded Ukraine, as no retailer was supplying fixed price contracts. Nevertheless, we will test this assumption in our event-study analysis by checking that price of j increase in line with our assumed timing. Finally, firms with $E_i = 0$ are protected during 2021 ($I^{2021} = 1$) and during 12 months in 2022 ($m^j = 12$). They constitute our "pure control group". In total, our electricity sample

on Inflation and Growth Expectations, where this information was asked separately. Evidence therein shows that almost all firms with a fixed price contract for gas have the same type of contract also for electricity, while among those firms with fixed price for electricity, two thirds have the same type of contract for gas. More details are provided in Alpino et al. (2023).

is made of 413 firms, while our gas sample is made of 308 firms.[16]. Note that in principle the formulation of our questions does not exclude the possibility that the initial contracts expire at the end of 2021 and that the firm immediately buys a new one. In Appendix F we use a synthetic diff-in-diff design to show that this is not the case. In other words we check that price of j for the "mid treated" increases relative to the "pure control group" starting in 2022h1. Thanks to these variables we construct event-time ($K_i = t - E_i$) dummies (specific to electricity or natural gas) that we use in the event-study specifications described below.

Through this sample selection procedure we are dropping firms that are protected for only *part* of a given semester in 2022. These are firms that have $1 \leq m_i \leq 5$ or $7 \leq m_i \leq 11$, of which there are 31 in the case of gas and 20 in the case of electricity. As it is visible in equation 1, this sample selection does not affect the composition of the "early treated" and "mid treated", but would affect the composition of the "late treated" and the "pure control group". Despite the small number of "late treated", we prefer to exclude these firms, as this creates a better alignment between the time dimension of the treatment variable and of the outcome variable. If firms are fully protected in their last semester of protection, the treatment becomes binary. Otherwise, the treatment would be continuous in the last semester of protection, making it harder to interpret the effects.

There is another group of firms that we exclude: (ii) firms that were not insured in 2021 ($I^{2021} = 0$) but were protected during 2022 ($I^{2022} = 1$). These are 14 firms in the gas sample and 32 firms in the electricity sample. This is because we see this form of protection as potentially endogenous to potential outcomes. Firms who choose to insure themselves after receiving the shock in 2021 may do so in anticipation of their treatment effects, leading to bias. Instead, in all of our regressions we condition on firms being insured at the beginning of 2021, when market participants were not foreseeing the upcoming crisis. Thus these firms were exposed sooner rather than later to the energy shock because of the timing of expiration of their contracts, and not because some of them decided to take action in anticipation of the crisis. We further discuss the validity of our identification assumptions in Section 5.b.

For some of our analyses and specifically those on input substitution and firm performance, we must rely on a yearly panel. In that setting we change the definition of our cohorts of treatment, again to better align the time dimension of our treatment and our outcome variables. When we collapse the treatment at the yearly frequency, many more firms see their electricity contracts expire when their gas contracts expire. Since we do not have a lot of separate variation for the two inputs, we construct a joint yearly treatment cohort E_i^Y in the following way:

[16] The share of firms in each cohort, together with other summary statistics can be found in Table 4.1

$$E_i^Y = \begin{cases} 2021, & \text{if } I^{2021} = 0 \text{ and } I^{2022,ele} = 0 \text{ and } I^{2022,gas} = 0 \\ 2022, & \text{if } I^{2021} = 1 \text{ and } I^{2022,ele} = 0 \text{ and } I^{2022,gas} = 0 \\ 0, & \text{if } I^{2021} = 1 \text{ and } I^{2022,ele} = 1 \text{ and } I^{2022,gas} = 1 \text{ and } m_i^{ele} = 12 \text{ and } m_i^{gas} = 12 \end{cases} \quad (2)$$

The definition of the cohorts follows very closely that for the half-yearly frequency in equation 1. We end up with a greater sample of 837 firms. The number of firms is substantially higher than in the half-yearly setting. This is because we also use firms that responded to our price-protection questions but then do not respond to the energy expenditure and consumption questions. As we show in in Table A.3 in Appendix J the sample is not too different from the half-yearly one in terms of observable characteristics, a point on which we return below in the results section.

3.f Outcome variables on firm performance

When looking at the impact of higher retail prices on firm-level outcomes other than energy demand, we focus on three distinct outcomes that we draw from the Invind survey: the growth rate in final goods prices[17]; capacity utilization and a Likert categorical variable on profit margins. Here we describe in detail how they are constructed.

As for the growth rate in final good prices, we rely on a recurring quantitative open question asking to report the "Average annual percentage change in selling prices of goods and services" between year t and year $t-1$.

As for the measurement of output responses, firms are asked about the capacity utilization of their plants, defined as the percentage ratio between actual production and maximum possible output.

On profit margins, we rely on a recurring question asking: "Please describe the firm's operating result for 2022?", that has five qualitative options: "1 = large profit; 2 = small profit; 3 =broad balance; 4 = small loss; 5 =large loss.". Based on this variable we construct a dummy taking value one when the firm operates on "large profit" or "small profit"

[17]The level of final good prices is not available in the Invind survey.

4 Summary statistics

In Table 4.1 we report summary statistics for our energy demand estimation samples, i.e. an electricity sample and a natural gas sample. Characteristics are measured before the start of the crisis, that is in the first semester of 2021. As from the Invind sampling design, included firms belong to the industrial sector and have at least 50 employees. All statistics are weighted by survey weights. At the bottom we report the number of firms. The number of observations can be retrieved by multiplying the number of firms by four (the number of semesters).

Overall, the electricity and the natural gas sample have similar characteristics. As for the sectoral distribution, more than half of the firms are in the metalworking industry, while the remaining half is more or less evenly split across the other sectors. Non-metallic minerals and water & waste remain residual categories in both samples. More than 80% of firms are located in the North of the country, while only between 5 and 7% of the sample is in the South or Islands. The unitary retail price of electricity is 16 euro cents per KWh for the electricity sample and the average firm consumes around 6000 MWh of electricity in a semester. As for gas, the average firm in the gas sample pays 10 euros per GJ of natural gas, while it consumes 63 million standard cubic meters of gas. Both in the electricity and gas sample, around a third of firms are energy intensive according to the Italian legislation definition[18], while around 5% of them have plants subject to the EU-ETS. Around half of firms in the gas sample declare that gas is an essential input in their production process. Cohorts of treatment are not evenly split. Only 14 or 18% in the electricity and gas sample, respectively, belong to the pure control group. The biggest groups are the early treated (exposed in the second half of 2021) and the mid treated (exposed in the first half of 2022). The late treated (exposed in the second half of 2022) is a residual category, accounting for 1% and 3% of the respective samples. Finally, as noticed at the beginning, firms in our sample are relatively large, with yearly revenues in the order of 80-90 million euros and more than 200 employees on average (although these distributions are very skewed).

[18] See *Ministerial Decree* n.541 of 2021.

Table 4.1: Summary statistics for the electricity and gas samples

	Electricity sample	Gas sample
	(1)	(2)
Variables	mean	mean
Sectoral composition		
Food and beverages	8%	6%
Textiles & apparel	13%	11%
Chem, pharma, rubber	12%	15%
Non-metallic minerals	4%	4%
Metalworking industry	51%	51%
Wood, paper, furniture	10%	11%
Water & waste	3%	3%
Macroarea		
North-West	40%	43%
North-East	39%	40%
Center	14%	13%
South or Islands	7%	5%
Energy-related variables		
Price of electricity (euro/KWh)	0,16	
Price of natural gas (euro/GJ)		10,41
Quantity of electricity (GWh)	6,161	
Quantity of natural gas (mil. smc)		63,406
Energy-intensive firm (0/1)	30%	29%
Subject to the EU-ETS	4%	5%
Gas is an indispensable input* (0/1)		54%
Cohorts of treatment		
Pure control	18%	14%
Early treated	44%	45%
Mid treated	35%	40%
Late treated	3%	1%
Other firm-level information		
Sales (million euro)	86,26	97,76
Labour force	204,2	224,9
Number of observations	413	308

Note: Invind data. The table reports summary statistics for the energy demand analyses used in Section 6. Characteristics are measured in the first semester of 2021, at baseline (thus the number of observations corresponds to the number of firms). *The variable "Gas is an essential input" is taken from the Business Outlook survey of the Bank of Italy and it refers to the beginning of 2022.

5 Empirical strategy

5.a Main specification

In this Section we describe how we isolate the causal effect of higher electricity and natural gas prices on firms' respective input demands. We consider the following econometric model:

$$\log y_{ijt} = \alpha_i + \gamma_t + \sum_{k=0}^{3} \beta_{ijk} \cdot \mathbf{1}(t - E_i^j = k) + \epsilon_{ijt}, \tag{3}$$

where y_{ijt} can either be the quantity q or retail price p of energy type j (electricity or natural gas) for firm i in semester t. α_i is a set of firm-level fixed effects that capture time-invariant differences in outcomes between firms. γ_t is a set of calendar time fixed effects that capture unobservable common trends across different cohorts of firms. $\mathbf{1}(t - E_i^j = k)$ are dummies that capture time relative to the contract expiration date E_i^j (event time, henceforth).[19] The coefficients β_{ijk} are potentially heterogeneous treatment effects of contract expiration on firm i at horizon k after the contract expiration.

Note that E_i^j is specific to each energy type j. We do not include E_i^j for both types in the same equation because they are highly correlated and lead to multicollinearity. The high correlation is partly due to data limitations (the insurance question is not fuel specific in 2021; see Section 3.e) and partly due to institutional features: data for 2022 shows that it is very common to have the same type of contract for both fuels.[20] Thus, when using quantity as outcome, our design has two limitations. First, it does not allow to estimate cross elasticities. We see this as a minor limitation because in our context the crisis induced an increase to the wholesale price of both natural gas and electricity.[21] As such, there was limited scope for substituting one fuel with the other to reduce costs. In addition, substitution between these two energy sources is technically difficult in the short run. Second, our design effectively identifies the increase of the price of both electricity and natural gas; thus, even in the absence of cross-substitution, quantity of one input might be affected by changes in the price of the other input via a (negative) scale effect. This is not a concern in our setting. As we show in the results section, electricity does not respond to higher prices, eliminating the concern that part of the gas response could be driven by an electricity-induced scale effect. A similar argument can be made for gas: since electricity does not respond at all, we detect no gas-induced scale

[19] See Section 3.e on how we construct this variable.

[20] The correlation between $I_i^{electricity}$ and I_i^{gas} is 0.59, between $m_i^{electricity}$ and m_i^{gas} is 0.72.

[21] The design of the wholesale electricity market implies that the price of electricity is equal to the bid of the marginal power producer, which is a gas-fired power plant in most hours of the year.

effect on electricity.

We estimate equation 3 on a balanced panel of firms observed during all the four semesters of 2021-2022 using the "imputation estimator" of Borusyak et al. (2021). This estimator is consistent for the average treatment effect on the treated (ATT) under standard parallel trends and no-anticipation assumptions, which we discuss below. The estimator works in three steps. First, it uses untreated (i.e. never-treated or not-yet-treated) observations of each firm to estimate the following model: $\log y_{ijt} = \alpha_i + \gamma_t + u_{ijt}$. Under the parallel trend assumption this gives us an estimate of each firm's counterfactual outcome absent treatment, i.e. $\widehat{\log y_{ijt}(0)}$. Then, for each treated period of each eventually treated firm it computes a treatment effect (τ_{ijt}) as the difference between the observed and the counterfactual outcome, i.e. $\tau_{ijt} = \log y_{ijt} - \widehat{\log y_{ijt}(0)}$. Lastly, it aggregates up individual treatment effects (τ_{ijt}) using weights of choice, depending on the precise estimand one is interested in. In all of our specifications we use survey sampling weights. In order to study treatment effect heterogeneity more systematically, we compute (survey-weighted-)averages over subsets of firms. Standard errors are clustered at the firm level to avoid known serial correlation issues (Bertrand et al., 2004). Finally, as suggested by Borusyak et al. (2021), we estimate pre-trend coefficients by a separate regression of the outcome on time fixed effects, firm fixed effects and dummies for periods before treatment. This is estimated using only untreated observations for each firm. As a consequence, in each event-study graph pre and post treatment coefficients are from different regressions; pre-treatment coefficients are relative to the first untreated period ($k = -3$), while post-treatment coefficients are relative to the average of the pre-treatment periods. Despite our preferences for the imputation procedure, Appendix J presents results of our baseline specifications obtained with other estimators, including De Chaisemartin and dHaultfoeuille (2020), Callaway and SantAnna (2021), Sun and Abraham (2021), and plain OLS.

5.b Validity of the difference-in-differences design

As it is standard in difference-in-differences designs, we make two assumptions. First, we assume that firms' potential log-outcomes absent treatment evolve according to the following simple additive model: $\log y_{it}(0) = \alpha_i + \gamma_t + u_{it}$, a parallel trend assumption. Second, we assume that observed log-outcomes are equal to untreated potential outcomes before treatment i.e. $\log y_{ik} = \log y_{ik}(0)$, a no-anticipation assumption.

As for the parallel trend assumption, we have good reason to believe that it holds in our design. This is because the expiration dates for fixed-price contracts are predetermined and not influenced by the crisis or firms' endogenous responses to it. In defining our treatment variable, we always condition on firms being already insured at the beginning of 2021, when the *futures* market was forecasting a stable and low price of natural gas.[22] Whether a firm's fixed-price contract expires sooner rather than later is only a result of when the contract was last *signed*, a matter of luck more than anything else. While we cannot directly test the parallel trend assumption, we always provide tests for parallel pre-trends and find no evidence thereof. As a robustness exercise, we also estimate synthetic diff-in-diff specifications (Arkhangelsky et al., 2021), which make pre-trends parallel by construction. Our results are confirmed qualitatively and quantitatively, indicating that the estimates obtained with the Borusyak et al. (2021) estimator are not driven by pre-treatment trends.

One may be concerned that the "early treated" cohort ($E_i = 2021h2$) makes an exception to the logic above. Indeed this group may include two sub-categories of firms, which we cannot separate in the data. The first sub-category consists of firms that had fixed-price contracts at the beginning of 2021 that expired in the second half of 2021. These firms subscribed fixed-price contracts but simply were unlucky with the timing. The second sub-category consists of firms that were under variable-price contracts and thus were immediately exposed to higher prices. This latter group of firms might not have had protection because they expected to perform particularly well in the event of an energy crisis. One may be concerned that this part of early treated firms is fundamentally different and that the other cohorts do not represent a good counterfactual for what would have happened to them, absent the treatment. Given that for this group of early-treated we only observe one period before treatment, we cannot perform a pre-trend test for them. To address this issue, in Appendix J we present cohort-specific estimates and show that the results are not driven by any specific cohort. Furthermore, notice that the "early treated" cohort never serves as a control group for the other cohorts.

As for the no-anticipation assumption, we always present pre-trend tests and verify that treatment effects do not materialize before contract expiration. Also, the fact that treatment effects are not significantly different across cohorts signals that anticipation is likely not an issue in this setting.

[22] An analysis of energy futures suggests that market participants did not anticipate any surge in gas prices until at least May 2021. For example, the contract expiring at the end of 2021 which eventually closed at 107 euro/MWh was trading at 18 at the beginning of February, at 20 at the beginning of April, at 24 at the beginning of May, at 41 at the beginning of August.

6 Effects on energy prices and energy demand

Figure 6.1 reports our baseline results. All plots report point estimates and associated 95% confidence intervals for average causal effects τ_k, with $k = \{0, 1, 2\}$ and pre-trend coefficients from a separate regression for $k = \{-2, -1\}$ (omitted category is $k = -3$). In the top two panels the outcome is the log of retail price of electricity and natural gas. Correspondingly, in the two bottom panels the outcome is the log of physical quantities of electricity and gas purchased by the firm. Since many of the τ_k coefficients are relatively large, in the text we describe the magnitude of the effects by commenting on $e^{\tau_k} - 1$, which transforms log changes in exact percentage changes.

Figure 6.1: The effect of the expiration of a fixed-price contract on average prices and quantities of energy inputs at the firm level

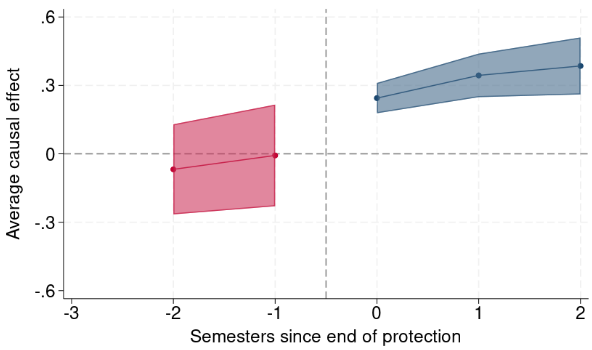

(a) Average costs of electricity

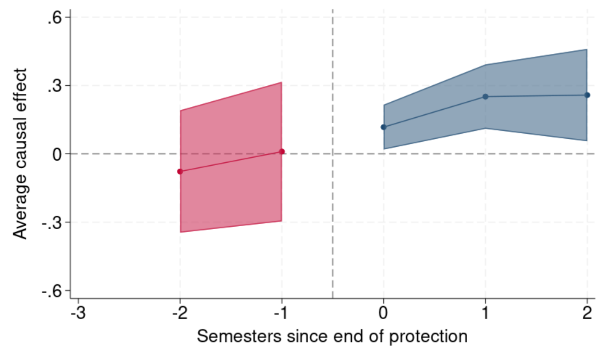

(b) Average costs of natural gas

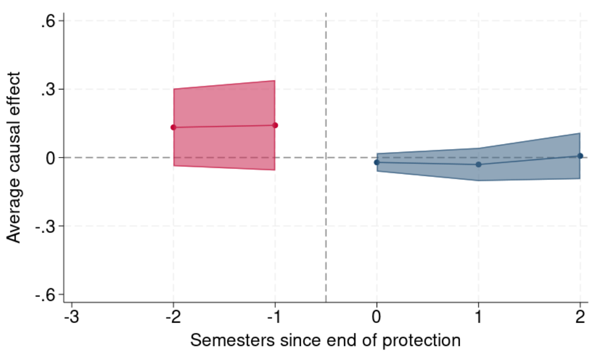

(c) Quantity of electricity

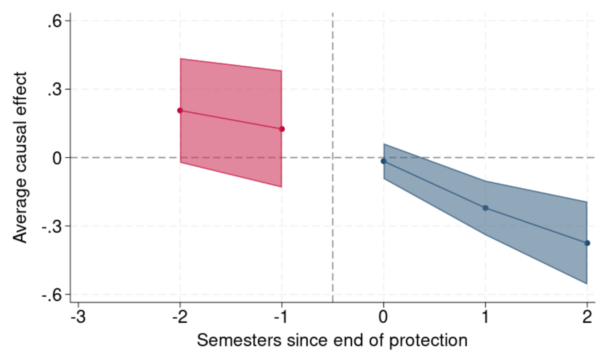

(d) Quantity of natural gas

Note: The figures show average causal effects of the expiration of a fixed-price contract on the average cost of electricity and natural gas (panels (a) and (b)) and the corresponding demanded quantities (panels (c) and (d)). Outcome variables are always in logs. Average causal effects before and after the treatment are estimated in two separate regressions, using the "imputation" estimator by Borusyak et al. (2021), as described in Section 5. Confidence intervals are at the 95% level.

6.a Effects on energy prices

In panel (a) we see that the retail price for electricity significantly increases when the fixed-price contract expires, while there is no pre-trending of causal effects before the treatment actually takes place.[23] One year and a half after expiration, the increase is 47% relative to the baseline, quite precisely estimated. We find a similar pattern in panel (b), when looking at the retail price for natural gas. Again, there is no evidence of a pre-trend[24] and estimates are precise. One year and a half after expiration, the increase is equal to 29% compared to the baseline. Cohort-specific estimates, reported in Appendix J, show that the results are not driven by any specific cohort, and that magnitudes are similar across them. Synthetic diff-in-diff estimates (Arkhangelsky et al., 2021), reported in Appendix F, show that the timing and magnitude of the results hold for every cohort in isolation when using only the "pure control group" as a comparison. From an econometric standpoint, this evidence provides support to the validity of our design and of the coding of our treatment variable.[25] From an economic standpoint, our findings highlight the importance of fixed-price inventory contracts, routinely used by firms for a variety of inputs (Kumar and Wesselbaum, 2024), in the transmission of macro shocks.

6.b Effects on energy demand

In panel (c) we study the (log-)quantity of (purchased) electricity. Despite the large price increase, this outcome does not respond to the treatment. Coefficients are positive, close to zero and confidence intervals are tight. We interpret these as precisely estimated zeros. We do not find evidence of treatment effect heterogeneity[26].

In panel (d) we investigate what happens to the (log-)quantity of natural gas. Here we detect a different pattern. In the first treatment period, the point estimate is virtually zero. In the second, it becomes negative, and more so in the third. The coefficient corresponding to the last period implies that natural gas consumption decreases around 32% compared to a

[23]The p-value of the pre-trend test is 0.13.

[24]The p-value of the pre-trend test is 0.39.

[25]These results also underline that commonly used shift-share identification strategies (Linn, 2008; Ganapati et al., 2020; Marin and Vona, 2021), which combine nation-wide swings in energy prices and cross-sectional variation in energy intensity, may suffer from measurement error, as noted by Lafrogne-Joussier et al. (2023). The implicit assumption behind these strategies is that all firms are exposed to the same fluctuation in energy prices at the same time. Our results cast doubt on this assumption by showing that firms face largely different prices depending on their contractual arrangement and the exact timing of contract expiration.

[26]The estimated effects are very similar between firms that cover part of their electricity consumption by generating their own power from renewable sources, and other firms. Furthermore, they are also very similar between electricity intensive and other firms; for the former group the results are confirmed by a companion analysis of monthly electricity consumption from administrative sources (see Appendix I).

counterfactual with no price increase. By visually inspecting the graph, we can see some evidence of a pre-existing downward trend in the case of gas, which may confound at least part of the effect.[27] We probe the validity of this result with a synthetic diff-in-diff exercise à la Arkhangelsky et al. (2021), which matches explicitly on the pre-trends. Results, reported in Appendix F, confirm a similarly large decline.

What explain this delayed response in gas consumption? It could be due to both cohort and calendar factors. On one hand, earlier treated cohorts might have more time to adjust. On the other hand, the crisis worsened over time after Russia invaded Ukraine, thus not all semesters are alike. One of the advantages of our staggered design is that we can disentangle whether the dynamics of the effect are driven by cohort or calendar factors.

In the lower panel of Figure 6.2 we display average causal effects separately by cohort. In the second half of 2021, only the "early treated" (black triangles) had already experienced the energy price shock, but the treatment effect on the quantity of gas is zero. In the first semester of 2022, the effect is again zero both for the "early treated", and for the "mid treated" (red circles), which experienced the shock for the first time in that period. Finally, all cohorts display a large negative effect in the second half of 2022, which corresponds to the first period since the shock for the "late treated" (blue diamonds), to the second for the "mid treated", and to the third for the "early treated".

Overall, we find that the negative effects in gas consumption are exclusively driven by what happens in the second half of 2022, while time elapsed since contract expiration does not matter, as all treated cohorts in a given calendar time period have similar coefficients. The exact same dynamics emerge in the synthetic diff-in-diff exercise, reported in Appendix F.

Why do firms adjust exclusively in the second half of 2022? First, we check whether the magnitude of the energy price shock was very different across calendar periods or, in other words whether the quantity reductions are different, but the demand elasticity is constant. This does not seem to be the case. The upper panel of Figure 6.2 shows that, despite the treatment effect on the price is slightly increasing over calendar time, this cannot explain the heterogeneity in quantity adjustments.[28] This conclusion will be confirmed in Section E where we combine our estimates to derive the price elasticity of demand.

[27]The p-value of the pre-trend test is 0.05.

[28]The fact that the effect on the firm-level energy price is very similar in the first and second half of 2022 is only apparently inconsistent with the evolution of the wholesale price depicted in Figure 2.1. In fact, while it is true that the wholesale price reaches its peak in the third quarter of 2022, in the last quarter it falls back to the same levels observed in the first semester of 2022. Since consumption is typically lower in the summer months, the consumption-weighted average wholesale price in the last semester of 2022 is probably not much higher than in the previous semester.

Figure 6.2: Natural gas: heterogeneous effects by cohort

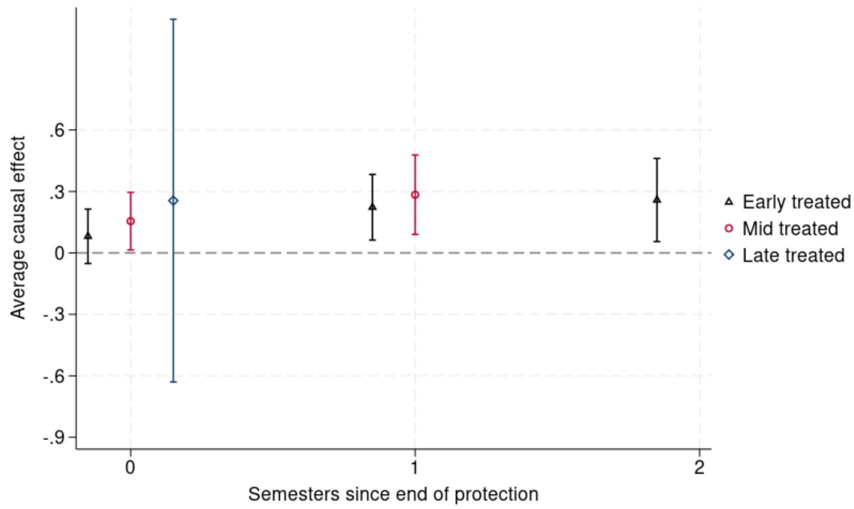

(a) Average cost

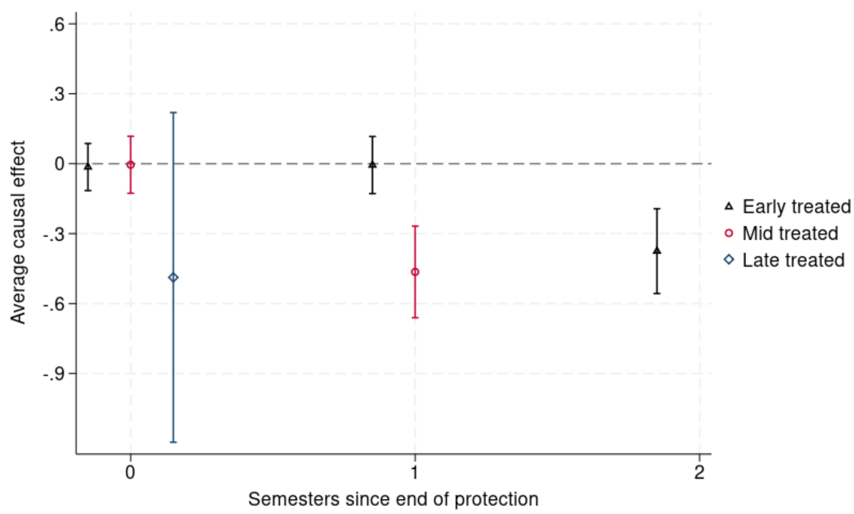

(b) Quantity

Note: the estimates in the upper and lower panel are from the same regression as in panel (b) and (d) of Figure 6.1 respectively; here they are reported by each cohort separately. Confidence intervals at the 95% level.

Next, we check whether the expectation on the duration of the energy crisis was somewhat different in the second half of 2022. Indeed, the 2022 summer was the period of highest market pessimism about future evolutions of the crisis: the spot price reached its highest peak (see Figure 2.1), and the *futures* market was pointing towards a long lasting crisis (see Figure B.1). A story consistent with our results goes as follows: firms were playing a "wait and see" strategy in 2021h2, when markets were expecting a short-lived crisis, and again in 2022h1, when the wholesale gas price shot up after the invasion of Ukraine but rapidly decreased afterwards. The strategy changed in the summer 2022 amid fears that Europe could be in short supply of natural gas in the forthcoming winter: presented with expectations pointing towards a wholesale gas around 200 euro per MWh until mid-2023, many business leaders decided to take action. The evidence is thus consistent with a "putty-putty" model with adjustment cost (Pindyck and Rotemberg, 1983; Atkeson and Kehoe, 1999).

6.c Heterogeneity in demand response

Pooling across cohorts, the average effect on gas consumption in the second semester of 2022 is large (-35%) but confidence intervals remain wide (between -44% and -25% at the 95% level).[29] A possible explanation has to do with treatment effect heterogeneity: it may be that only some types of treated firms are able or willing to scale down energy demand, while others cannot adjust or do not find it convenient. Assuming Cobb-Douglas production function, economic theory would suggest that energy intensive firms should scale down energy demand the most when its price rises. However, the production function of energy intensive firms might be better approximated by a Leontief function, in which case energy intensity would be associated with a smaller elasticity.

To explore this issue, we test whether treatment effects are heterogeneous along the following covariates: a (self-reported) dummy for whether gas is an essential input in production[30], which can be thought as a proxy for Leontief production function; a gas intensive dummy[31]; sector dummies; an EU ETS dummy; and employment in 2021.

To investigate treatment effect heterogeneity in a credible way, we turn to machine learning (ML) techniques, which are becoming increasingly popular for this aim in causal analysis and in particular in the context of randomized control trials (Haaland and Roth, 2020; Allcott et al., 2020; Alpino et al., 2022). To the best of our knowledge, we are the first to apply these tools in the context of staggered difference-in-differences. The estimator by Borusyak et al. (2021) is particularly well suited for this, as it provides a treatment effect estimate for each treated observation. We can thus use it as an outcome variable and use random forests

[29]Magnitude and precision are virtually unaffected when including firm-semester fixed effects, which control for firm-specific seasonality. This robustness mechanically excludes the "early treated", for which the second semester is not observed both before and after treatment.

[30]The yes/no question comes from a different Bank of Italy survey (*Business Outlook Survey of Industrial and Service Firms*, Sondtel) run in fall 2022 and reads: "At the beginning of 2022, was gas an essential input for your firms manufacturing process?". Essential inputs are defined as follows: Inputs are essential when given the plants and machinery installed and used in the manufacturing process the total or partial lack thereof would make it impossible to produce the good in the short term.

[31]The dummy is based on the official definition by the Italian legislation. Gas intensive firms must have annual consumption above 1 GWh, belong to certain industrial sectors, and have gas intensity (measured as expenditure for gas natural on sales and/or value added) above certain thresholds.

to find its best predictors, in the spirit of Athey and Imbens (2016) and Wager and Athey (2018). The advantages of this approach relative to a more traditional heterogeneity analysis are twofold: a) the quest for heterogeneity is more efficient and also explores non-linear combinations of the available covariates; b) the procedure is less prone to bias arising from multiple hypotheses testing.

The ML analysis reveals three key predictors of treatment effect heterogeneity: the dummy for whether natural gas is essential in the production process; the dummy for gas intensive firms; the dummy for the EU ETS; the dummy for the food sector (see Appendix G for details about the methodology and the results). Starting from the first dimension of heterogeneity, in the second half of 2022 the average treatment effect is equal to -41% among firms for which gas is *not* essential, and to -28% for other businesses.[32] The result that firms for which gas is essential reduce its consumption less when its price increases is not surprising, but it is important for at least two reasons. First, this is a large group, accounting for approximately 40% of treated observations. Second, it includes all gas intensive firms.

Coming to the other dimensions of heterogeneity, the average effect is zero among gas intensive firms, slightly positive in the food industry, and -19% among EU ETS firms. However, it is difficult to draw definitive conclusions about the magnitude of these effects from the Invind survey, considering the small sample size (respectively 22, 21 and 22 treated firms), and that almost all of them also declare gas to be an essential input. In order to estimate the effect among gas intensive firms more credibly, we turn to administrative data from the Fund for Energy and Environmental Services (CSEA), where we observe gas consumption at the monthly frequency. After matching this data with the fixed-price contract questions from Invind, we run an higher-frequency version of our staggered diff-in-diff analysis on more than one hundred gas intensive firms. Results, reported at length in Appendix H, support the following findings: a) in the second half of 2022, the effect on gas consumption among gas intensive firms is -8%, marginally not significant at conventional levels; b) the effect is zero in the previous periods.

6.d Price elasticity of energy demand

The price elasticity of energy demand is an important parameter which affects quantitatively the equilibrium responses derived in economic models. In our setting we can compute it by combining our estimates of the effect of fixed-price contract expirations on (log) prices and on (log) quantities. Since we only use within-firm variation, our estimate can be described as

[32]Both estimates are significant at the 99% level. The p-value of the test that the two effects are the same is 0.09.

a "micro" elasticity. This behavioral response does not take into account the fact that part of the input substitution process derives from factor reallocation across firms (and sectors), as opposed to within firms (Bachmann et al., 2022). However, note that credibly identified *micro* parameters are a useful disciplining device in macro models when used as target moments (Nakamura and Steinsson, 2018).

In practice, we follow an IV-LATE approach, where we scale the average treatment effect estimated in the quantity equation by the average treatment effect estimated in the price equation.[33] Under standard IV-LATE assumptions, we can construct an estimate of the elasticity θ as follows:

$$\theta = \frac{e^{\hat{\tau}^q} - 1}{e^{\hat{\tau}^p} - 1} \tag{4}$$

where $\hat{\tau}^q$ and $\hat{\tau}^p$ are the estimates of, respectively, the average treatment effect on the log quantity and on the log price of energy (Angrist et al., 2000). We construct standard errors using the delta method (see Appendix E for details).

Figure 6.3 reports estimates by calendar period for the entire sample and for some selected sub-samples. We do not report estimates by cohort because we find no heterogeneity along this dimension.

As for electricity, the elasticity is always zero. At first glance, this result appears in contrast with recent findings related to the pre-crisis years. For example, both Marin and Vona (2021); von Graevenitz and Rottner (2022) and Fontagné et al. (2023) estimate an elasticity between -0.4 and -0.5 using French and German data from the periods 1997-2015, 2009-2017 and 1996-2019, respectively. These studies however use comparably smaller shocks. Instead, Gerster and Lamp (2023) finds electricity elasticities between -0.09 and -0.2 using a large discontinuity in electricity prices ($\approx$ 30%) faced by very large German firms. Moreover, also von Graevenitz and Rottner (2022); Fontagné et al. (2023) find that this elasticity is smaller with larger shocks and in later periods, consistently with our findings.

As for natural gas, our estimated average elasticity is equal to zero in the second half of 2021 and in the first of 2022, and equal to -1.1 in the second semester of 2022.[34] This value is higher

[33] An alternative approach could be to regress the individual treatment effects from the price equation on the individual treatment effect from the quantity equation following Deryugina et al. (2020). However, this strategy is unfeasible in our application because we construct prices as the ratio between expenditures (in monetary terms) and physical quantity (see section 3). Thus any measurement error in quantity will translate in a measurement error of opposite sign in prices, thus yielding a negative correlation by construction. This issue is reminiscent of the "division bias" discussed by Borjas (1980) in the context of using the ratio of earnings and hours as a proxy for wage in the regression of hours on wages.

[34] In the case of natural gas, confidence intervals are particularly wide in the first semester 2021. This is due to the fact that the first-stage is not very strong in that period because the insurance question was not fuel-specific

Figure 6.3: Price elasticity of demand by calendar period

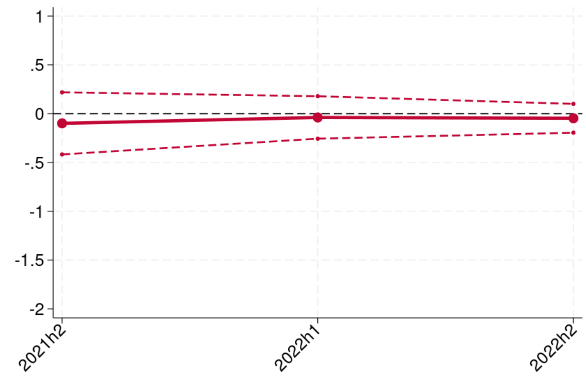

(a) Electricity

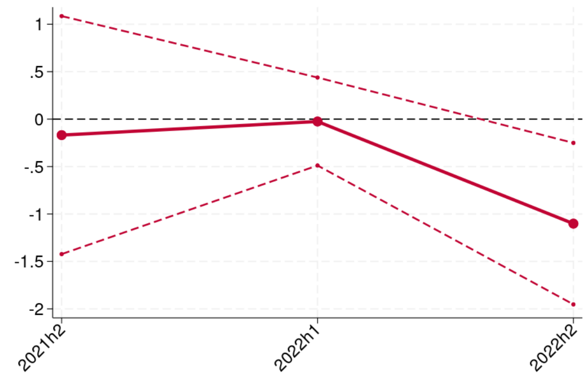

(b) Natural gas

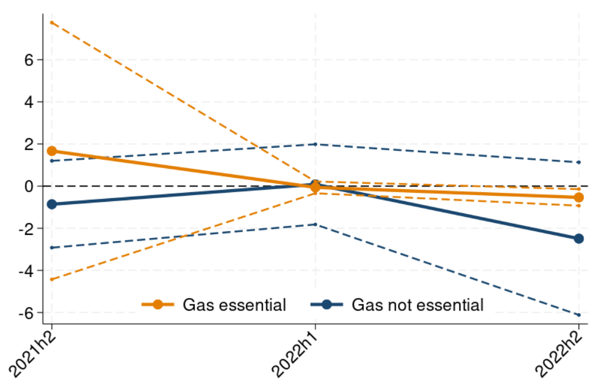

(c) Natural gas: essential or not

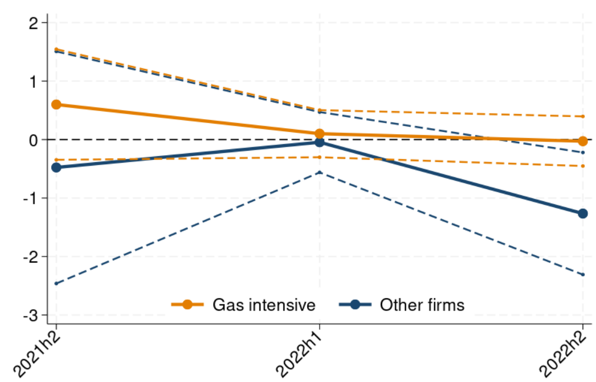

(d) Natural gas: intensive or not

Note: the elasticity is computed as $\frac{e^{\hat{\tau}^q}-1}{e^{\hat{\tau}^p}-1}$ where $\hat{\tau}^q$ and $\hat{\tau}^p$ are the estimates of, respectively, the average treatment effect on the log quantity and on the log price. Estimates in panel (a) combine effects from panels (a) and (c) of Figure 6.1; estimates in panel (b), (c) and (d) combine effects from panels (b) and (d). The upper panels plot average elasticities; the lower panel elasticities for selected sub-samples. Standard errors are constructed using the delta method. Confidence intervals are at the 95% level.

compared to the older literature (Labandeira et al., 2017), but broadly in line with the recent findings of Fontagné et al. (2023) (between -0.9 and -1.2). Furthermore, we find that in the second semester 2022, gas elasticity is smaller in absolute value for firms for which gas is essential (-0.5) and for gas intensive firms (-0.03) relative to their complement groups (-2.5 and -1.3). Although these exact values must be taken with caution due to the imprecision of the estimates, they suggest relevant heterogeneity in gas elasticity across-firms. These estimates are directly policy relevant. In Section 8 we use these estimates to compute by how much natural gas equilibrium quantities change after the introduction of a per-unity

in the 2021 wave of our survey. For more details on this see Alpino et al. (2023).

quantity subsidies.

6.e Input substitution

In this section we study whether the negative effect on gas consumption in the second half of 2022 was partly compensated by a larger use of other fossil fuels. To this end, we turn to administrative data on plants subject to EU ETS, for which we observe consumption of each fossil fuel separately, although at the annual frequency. We use firm identifiers to match this data to the Invind survey, which we use to define treatment cohorts. For this yearly design we must collapse our original half-yearly treatment into a yearly treatment. Thus firms end up being divided into three cohorts only: those that are treated in 2021, those that are treated in 2022 and those that are never treated (a pure control group). We provided details on how we construct these cohorts in Section 3.e. While on average natural gas is the main source of fossil energy for factories in the EU ETS, they also consume half as much solid fuels (e.g. coal). The use of other liquid (e.g. kerosene) or gas (e.g. LPG) fuels is much more limited.

In this section, we rely on an annual panel of plants which begins in 2018 and use outcomes in levels to facilitate the comparison across fuels. Our test for input substitution is as follows. First, we use consumption of natural gas as an outcome variable to confirm that our results in the previous section extend to the EU ETS sub-sample, where 90% of the plants belong to firms that declare gas to be an essential input. Panel (a) of Figure 6.4 shows that the effect on gas consumption is indeed negative, building up over time. There is no evidence of a pre trend.[35] The effect is equal to -26 terajoules (Tj) in 2021 and -89 Tj in 2022, respectively -4% and -14% relative to the 2018-2020 average[36], in line with the evidence provided in the previous sections. Second, we use total fossil energy consumption as an outcome variable, namely the sum of natural gas plus all other fuels (e.g. coal, LPG, kerosene, lignite, gasoline, etc.). If plants are able to completely substitute natural gas with other fuels, average treatment effects would be exactly equal to zero. On the contrary, if plants cannot substitute gas with other inputs, average treatment effects should be the same as when using natural gas as outcome. Results, presented in the panel (b) of the same figure, lie between these two extremes. The profile of the event-study is similar to the previous one, but attenuated in magnitude. The effects are equal to +27 Tj in 2021 and -56 Tj in 2022; confidence intervals are wider. This evidence suggests that input substitution is incomplete at best, as treated firms reduce their total fossil energy consumption to a greater extent, relative to control firms. In other words,

[35] The p-value of the pre-trend test is 0.47.

[36] The second effect is statistically significant at the 95% level. Results are quantitatively very similar if we estimate the model in logs.

treated firms are unable or unwilling to completely offset the drop in natural gas by using other fossil fuels. To further explore this issue, we study one by one the consumption of

Figure 6.4: Input substitution test among EU ETS plants

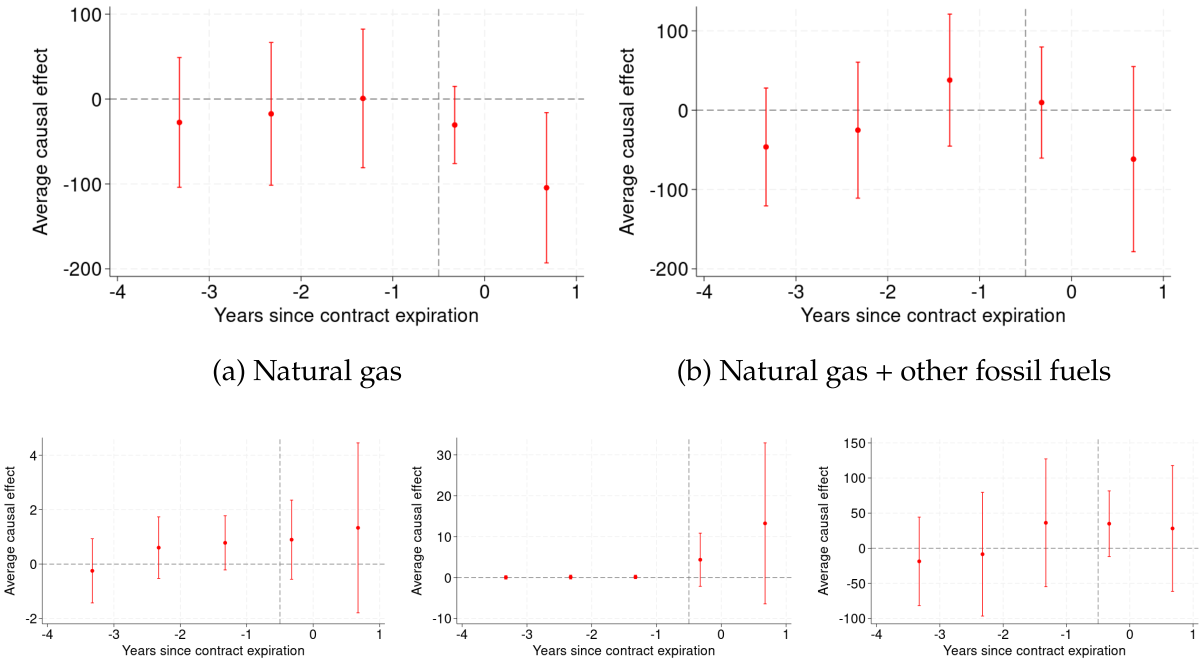

(a) Natural gas

(b) Natural gas + other fossil fuels

(c) Other fossil gas fuels

(d) Other liquid fossil fuels

(e) Other solid fossil fuels

Note: The figures show average causal effects of the the expiration of a fixed-price contract on different outcomes in levels. The outcome is reported underneath each event-study. Average causal effects are estimated on the set of firms belonging to both the Invind and the EU ETS sample. Average causal effects before and after the treatment are estimated in two separate regressions, using the "imputation" estimator by Borusyak et al. (2021), as described in Section 5. Confidence intervals are at the 95% level.

different types of fuels, to test whether some of these are more of a substitute in the recent crisis. The event studies (reported at the bottom of Figure 6.4) show that consumption of all types of fossil fuel (other than natural gas) increases more in the treatment group than in the control group, but confirms the idea that substitution was incomplete at best. In the case of other gases, the increase is quantitatively negligible. For liquid fuels the effect is larger (6 Tj in 2021 and 10 in 2022, both marginally insignificant a the 90% level), but still small compared to the drop in natural gas (one tenth). In the case of solid fuels, the effect is large, but it arises one period in advance, and it is larger in 2021 than in 2022, casting doubts on the validity of the identifying assumptions.[37] Overall, even though we can not rule out that some firms managed to substitute natural gas with other fuels, it is safe to conclude that input substitution was not the main explanation behind the drop in natural gas consumption

[37]The p-value of the pre-trend test is 0.02.

identified in this work, at least in case of plants subject to the EU ETS. Note that these enterprises are different from the average firm and in particular from those for which natural gas is not essential, which is the group that reduce the most its gas consumption, according to our results.

Even when restricting the attention to EU ETS firms, one limitation of this exercise is that in our data we only observe fossil fuels, while substitution might occur via other fuels (e.g. hydrogen or electricity). As for electricity, as already noted elsewhere in this paper, its price was increasing in tandem with natural gas during the recent crisis. As a consequence it would have made little economic sense to substitute one with the other. As for hydrogen, despite recent efforts to promote it in industry at the EU level, its use is still residual in Italian industry.[38] Finally, our analysis does not exclude the possibility of substituting energy by importing energy intensive intermediate goods from outside Europe, a channel emphasized by Moll et al. (2023). Unfortunately, at the moment we lack data to test this hypothesis.

7 Effects on output prices, production and profit margins

When firms cannot substitute away from more expensive inputs, they have two other options at their disposal. On the one hand they could reduce the quantity of output they produce; on the other hand they could pass on some of the cost increases to consumers via higher prices. The combination of lower output, higher prices, together with input substitution, will be reflected in their profit margins.

In this section we study these margins of adjustment by exploiting a longer yearly panel of firms that answered the Invind survey over the 2018-2022 period. Since not all firms are interviewed in all years, and do not answer to all questions, we are forced to use an unbalanced panel of firms. As detailed in Section 3.e and similarly to Section 6.e, we collapse our original half-yearly treatment into a yearly treatment. Thus firms end up being divided into three cohorts only: those that are treated in 2021 (corresponding to the "early treated" in the previous analysis), those that are treated in 2022 (corresponding to the "mid treated") and those that are never treated (a pure control group). Aside from this, the regression model is exactly the same as in Section 5. We consider three outcomes at the firm level that are available in the Invind survey: the yearly growth rate in the prices of goods sold[39]; the average degree of plant capacity utilization, defined as a percentage of the maximum

[38]According to the *Italian National Strategy for Hydrogen* the share of hydrogen in final energy consumption is at 2% in the country and less than 1% in “hard to abate” industries like chemicals or refineries.

[39]The level of prices is not available in the Invind survey

output attainable with the given capital[40]; and a dummy for whether the profit margin is strictly positive. The yearly sample we use in this analysis is similar to the half-yearly one used in previous sections. We detect no appreciable difference both in terms of sectoral and geographical composition. Firms in the yearly sample have a lower probability of declaring that gas is an essential input (37% vs 54% in the half-yearly sample) and of being energy intensive firms, according to the Italian legislation (21% vs 30% in the half-yearly sample). For completeness, in Table A.3 of Appendix J we report summary statistics.

Event-studies are depicted in Figure 7.1 while estimates by calendar year are reported in Table 7.1. Table 7.2 reports heterogeneous treatment effects for year 2022 according to the same categories used in the previous sections.

Figure 7.1: The effect of energy price-protection lifting on prices, capacity utilization and profit

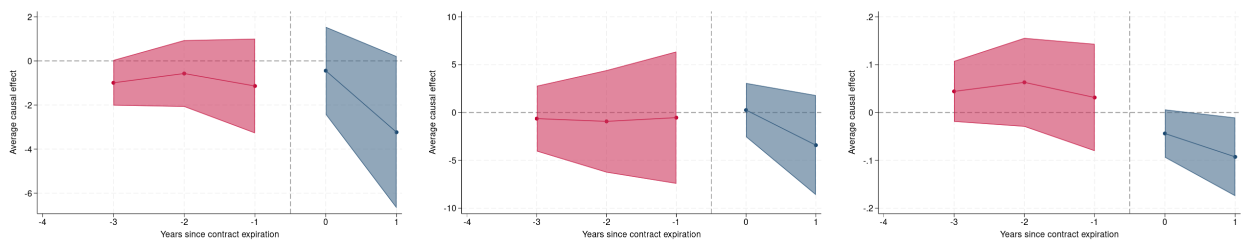

(a) Change in price (%) (b) Capacity utilization (%) (c) Positive profit (0/1)

Note: The figures show average causal effects of the expiration of a fixed-price contract on different outcomes, reported underneath each event-study. Samples are different across the three panels. Average causal effects before and after the treatment are estimated in two separate regressions, using the "imputation" estimator by Borusyak et al. (2021), as described in Section 5. Confidence intervals are at the 95% level.

Starting with the first outcome, we find that when treated firms see their fixed-price contract expire, they increase their final price less than the control group. The effect is equal to -2.7 percentage point, not significant at conventional levels, and it is entirely driven by what happens in 2022. Note that in that year, the average price increase was around 11 per cent, and no firms decreased its price. So our results suggests that, amid generalized upward price pressure, firms exposed to energy price shock did not increase their price more than other businesses. How to rationalize this finding? First, there is mounting evidence that firms' price updating decisions do not depend much on idiosyncratic cost shocks, but rather on rivals shocks (via strategic complementarities) and on market-wide cost shocks (Amiti et al., 2019; Muehlegger and Sweeney, 2022; Duprez and Magerman, 2018). If these models of price formation are a good representation of reality, our reduced form approach is not

[40]The survey records also changes in capacity relative to the previous year; we adjust our outcome to account for this to get a good measure of physical output.

appropriate to gauge the the effect of the energy crisis on inflation, because it can only identify partial equilibrium effects. Second, even if idiosyncratic shocks matter, identifying their effect for the average firm is quite challenging because the share of energy cost on firms' variable cost is very low on average Alpino et al. (2022). Indeed, our heterogeneity analysis reveals that when focusing on firms for which energy represents a larger fraction of their costs (gas intensive and EU ETS firms), the treatment effect is positive, large and significant. In this respect, our findings are in line with Lafrogne-Joussier et al. (2023), who estimate the pass-through of energy shocks on producer prices in the recent crisis. They find that, due the relatively small share of energy in firms' variable costs and despite substantial pass-through of positive shocks, the recent energy price surge only moderately impacted manufacturing inflation.

Table 7.1: Average treatment effects by calendar year

	Annual price change (%)	Capacity utilization (%)	Profit>0 (0/1)
	(1)	(2)	(3)
2021	0.58	0.02	-0.00
	(1.20)	(1.75)	(0.03)
2022	-2.73	-1.89	-0.10***
	(1.66)	(2.52)	(0.04)
N	2887	3655	3464

Note: the table reports point estimates and standard errors of the treatment effects by calendar year from the same regressions as in Figure 7.1.

Coming to the second outcome, we estimate that upon contract expiration treated firms reduce capacity utilization by less than 2 percentage points in 2022. The drop is quite small, considering that capacity utilization is on average around 80 per cent in the sample years. We do not find evidence of treatment effect heterogeneity. How to rationalize the the limited responsiveness of physical output with the large drop in input identified in the previous sections? We believe there are at least three non mutually exclusive explanations. First, natural gas only drops in the second semester, while capacity utilization refers to the average over the year. Second, our measure of capacity utilization might understate changes over time. Lacking hard data to answer this question, some respondents might be tempted to provide the same figure as in the previous year. Third, firms might be able to substitute natural gas might other non-energy inputs, and keeping output unchanged. Unfortunately, our data does not allow us to discriminate between these possibilities.

Finally, we find that being exposed to the energy shock reduce the probability to have a positive profit margin by 10 percentage points in 2022. The effect is sizable, considering that over our sample period approximately 80 per cent of firms declare to have positive profits. The drop is widespread with the exception of the EU ETS firms, for which the effect is zero; note that this is the group for which we estimated the largest positive effect on the growth of final prices.

Table 7.2: Heterogeneous treatment effects in 2022 by type of firm

(a) Annual price change (%)

	Gas essential	Gas intensive	EU ETS	Electricity intensive
	(1)	(2)	(3)	(4)
No	-4.47**	-3.34**	-3.44**	-4.21**
	(1.75)	(1.68)	(1.67)	(1.68)
Yes	1.81	9.09**	18.29***	2.68
	(1.90)	(3.89)	(5.06)	(2.15)
P-value equality test	0.00	0.00	0.00	0.00

(b) Capacity utilization (%)

	Gas essential	Gas intensive	EU ETS	Electricity intensive
	(1)	(2)	(3)	(4)
No	-2.05	-1.93	-1.88	-1.90
	(2.64)	(2.53)	(2.53)	(2.60)
Yes	-1.06	-1.25	-2.21	-1.85
	(2.98)	(4.27)	(3.68)	(2.94)
P-value equality test	0.66	0.85	0.91	0.98

(c) Profit>0 (0/1)

	Gas essential	Gas intensive	EU ETS	Electricity intensive
	(1)	(2)	(3)	(4)
No	-0.11***	-0.10***	-0.10***	-0.10**
	(0.04)	(0.04)	(0.04)	(0.04)
Yes	-0.12**	-0.09	0.00	-0.11*
	(0.05)	(0.09)	(0.07)	(0.06)
P-value equality test	0.96	0.95	0.12	0.82

Note: the table reports point estimates and standard errors of the treatment effects in 2022 by different types of firms from the same regressions as in Figure 7.1. The outcome variable is reported above each panel. The row "No" refers to the firms for which the dummy on top of each column is switched off; the row "Yes" refers to the firms for which the dummy on top of each column is switched on. The last row reports the p-value of the test that the two treatment effect in the column are equal to each other.

8 Policy implications

Our results bear policy implications for the design of support measures aimed at firms during energy crises. Such tools were popular in the EU during the recent 2021-22 crisis. According to recent estimates by Bruegel, since March 2022 EU governments allocated around €670 billion to help businesses face energy shocks (McWilliams et al., 2024). In spite of their popularity, several economists were concerned that these policies would exacerbate the crisis if designed in such a way to reduce the marginal price of energy faced by firms. In fact, this would amount to subsidize the demand for energy thus reducing the incentive to save it at a time of very short supply (Gros, 2022; Signorini, 2022).

Unintended upward pressure on energy demand could be avoided if policies were targeted to firms with a demand elasticity close to zero. In the paper we show that firms' price elasticity of demand was close to zero for electricity, while demand was somewhat more responsive in the case of natural gas, although only in the second half of 2022. Our heterogeneity analysis highlights an important distinction in the case of gas. The price-responsiveness is almost entirely explained by firms for which gas is not an essential input. To the contrary, the remaining group, which also includes almost all gas-intensive firms, displays an elasticity close to zero.

In order to avoid or mitigate unintended demand responses, it follows that support measures should target firms for which gas is an essential input or, alternatively, gas intensive firms, which is an observable characteristic. In the Italian context, the subsidies were initially targeted to energy intensive firms for both electricity and natural gas consumption, but they were expanded to all firms later in the crisis. Our results suggest that such expansion in case of natural gas consumption might have caused some upward pressure on demand.

To get a sense of the quantitative magnitude of these policy implications in the case of natural gas, we follow Deryugina et al. (2020) in presenting some simple calculations calibrated to the Italian context. Between 2015 and 2019, industrial gas consumption in the second semester of the year was equal to 6,700 million of Standard Cubic Meters (SCM) on average (source: SNAM, the network operator); the average retail price in the business sector was equal to 1.35 euro per SCM in the second half of 2022 (source: Eurostat).[41] Assuming for simplicity that the market for industrial gas deliveries is perfectly competitive, the introduction of a per-unit subsidy s on the price P would result in an increase of the equilibrium quantity Q

[41] Eurostat computes this price as the average cost per unit of energy net of taxes. As such, this price already include government aid in the form of reduced taxes and fees, but does not include the subsidy which took the form of a tax credit.

which depends on the (absolute value) of the demand ϵ_S and supply ϵ_D elasticities:

$$\frac{\partial Q}{\partial s} = \frac{\epsilon_S \epsilon_D}{\epsilon_S + \epsilon_D} \frac{Q}{P} = \frac{1}{1/\epsilon_S + 1/\epsilon_D} \frac{Q}{P}. \tag{5}$$

Figure 8.1: Effect of a subsidy on equilibrium gas consumption

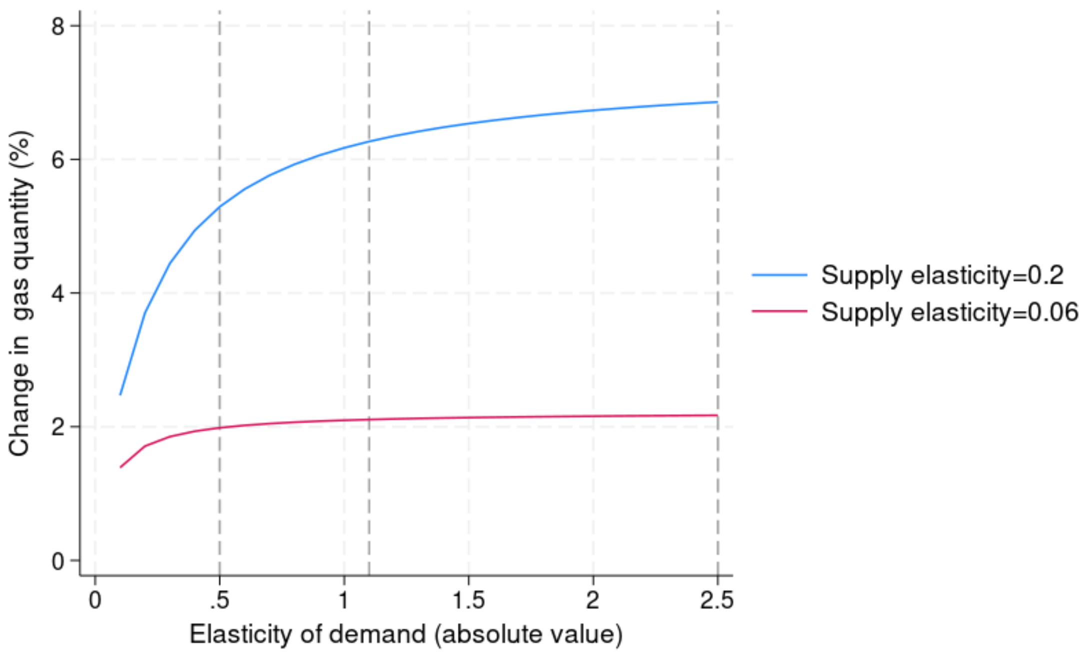

Note: The figure shows the percentage increase in the equilibrium quantity of gas purchases induced by a 50 cent subsidy for different values of demand and supply elasticities. It is calculated according to the formula in (5) and scaled by the baseline quantity; Q is set equal to 6,700 million of Standard Cubic Meters (SCM) and P to 1.35 euro per SCM. Vertical dotted lines are drawn in correspondence of some values of the elasticity of demand estimated in the paper: 0.5 (firms for which gas is essential), 1.1 (average) and 2.5 (firms for which gas is not essential.

Supply of gas is usually considered quite inelastic in the short term, but there is scarce evidence regarding its magnitude. As a first exercise, we assume a very low supply elasticity. In particular, we follow Albrizio et al. (2022) and use the value $\epsilon_S = 0.06$, as estimated by Krichene (2002). In this case, if one were to use our average estimates of the demand elasticity ($\epsilon_D = -1.1$), a 50 cent price subsidy would increase the equilibrium quantity by 141 million of SCM (2%) in the second half of 2022. Such effect would be only marginally lower using the value for firms for which gas is essential ($\epsilon_D = -0.5$) and marginally higher using the value for firms for which gas is not essential ($\epsilon_D = -2.5$) (red line in Figure 8.1). Intuitively, in this case the supply curve is so inelastic that producers can cash in a great share of the subsidy without increasing quantities much and irrespective of the demand elasticity.

As a second exercise, we consider a slightly higher value of the supply elasticity, which is a more sensible assumption in our context. Italy has a very developed natural gas infrastructure: three LNG terminals and five pipelines connecting it to several producers other than Russia (Azerbaijan, Algeria, Lybia and Northern Europe). In 2022, to the surprise of many commentators, Italy managed to secure new large gas delivery contracts from these countries; the annual increase in import from these suppliers (14.8 billion of SCM; +34%) almost offset the drop in imports from Russia (15 billion of SCM). These are large numbers, roughly equal to the annual consumption of the Italian industrial sector. Thus, in our second exercise we assume an higher supply elasticity, namely $\epsilon_S = 0.2$. In this case, how large is the demand elasticity becomes important (see blue line in Figure 8.1). When assuming $\epsilon_D = -0.5$, the subsidy-induced increase in quantity is equal to 350 million of SCM (+5%), while when assuming $\epsilon_D = -2.5$, the increase is equal to 460 million of SCM (+7%).[42]

9 Conclusions

We provide evidence on the effect of large energy price shocks on firms. We do so with a novel identification strategy based on the staggered expiration of fixed-price energy contracts, which expose some businesses to spikes in energy prices sooner rather than later. In our difference-in-differences design, we estimate that during the 2021-22 crisis average unitary costs for exposed firms increased up to 45% and 30% for electricity and gas respectively. Despite this sizeable shock, the demand adjustment is relatively small. Firms do not cut electricity consumption, and reduce natural gas only in the second half of 2022 (-35%). Furthermore, this drop is very heterogeneous across firms. In particular it is smaller for those declaring natural gas to be an essential in their production process, which include gas intensive firms that account for 20% of national consumption.

Additional evidence suggests that the drop in gas consumption was not fully compensated by substitution with other inputs, and that, as a consequence, output fell somewhat. Furthermore, we find that the idiosyncratic energy price cost shocks that we identify do not induce a significantly higher increase in the price of final output and decreases the probability of reporting positive profits, in a context of generalized upward price pressure.

[42] Notice that in presence of heterogeneous elasticities, it is not clear which one should be used in the calculation, as it depends on whether the marginal consumer is a firm for which gas is essential or not.

References

Albrizio, S., Bluedorn, J. C., Koch, C., Pescatori, A., and Stuermer, M. (2022). Market size and supply disruptions: Sharing the pain from a potential russian gas shut-off to the eu. *IMF Working Papers*, 2022(143).

Alessandri, P. and Gazzani, A. G. (2023). Natural gas and the macroeconomy: Not all energy shocks are alike. *Available at SSRN 4549079*.

Allcott, H., Braghieri, L., Eichmeyer, S., and Gentzkow, M. (2020). The welfare effects of social media. *American Economic Review*, 110(3):629–676.

Alpino, M., Citino, L., and Frigo, A. (2023). The effects of the 2021 energy crisis on medium and large industrial firms: evidence from italy. *Banca dItalia, Occasional working paper series n.776*.

Alpino, M., Hauge, K. E., Kotsadam, A., and Markussen, S. (2022). Effects of dialogue meetings on sickness absenceevidence from a large field experiment. *Journal of Health Economics*, 83:102615.

Amiti, M., Itskhoki, O., and Konings, J. (2019). International shocks, variable markups, and domestic prices. *The Review of Economic Studies*, 86(6):2356–2402.

Andini, M., De Blasio, G., Duranton, G., and Strange, W. C. (2013). Marshallian labour market pooling: Evidence from italy. *Regional Science and Urban Economics*, 43(6):1008–1022.

Angrist, J. D., Graddy, K., and Imbens, G. W. (2000). The interpretation of instrumental variables estimators in simultaneous equations models with an application to the demand for fish. *The Review of Economic Studies*, 67(3):499–527.

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., and Wager, S. (2021). Synthetic difference-in-differences. *American Economic Review*, 111(12):4088–4118.

Athey, S. and Imbens, G. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27):7353–7360.

Atkeson, A. and Kehoe, P. J. (1999). Models of energy use: Putty-putty versus putty-clay. *American Economic Review*, 89(4):1028–1043.

Auffhammer, M. and Rubin, E. (2018). Natural gas price elasticities and optimal cost recovery under consumer heterogeneity: Evidence from 300 million natural gas bills. Technical report, National Bureau of Economic Research.

Bachmann, R., Baqaee, D., Bayer, C., Kuhn, M., Löschel, A., Moll, B., Peichl, A., Pittel, K., Schularick, M., et al. (2022). What if? the economic effects for germany of a stop of energy imports from russia. *ECONtribute Policy Brief*, 28:2022.

Bank of Italy (2017). Invind, metodi e fonti: note metodologiche. Technical report, Bank of Italy.

Becker, S. O. and Ichino, A. (2002). Estimation of average treatment effects based on propensity scores. *The stata journal*, 2(4):358–377.

Bertrand, M., Duflo, E., and Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *The Quarterly journal of economics*, 119(1):249–275.

Bond, S. R., Rodano, G., and Serrano-Velarde, N. A. B. (2015). Investment dynamics in italy: Financing constraints, demand and uncertainty. *Bank of Italy Occasional Paper*, (283).

Borjas, G. J. (1980). The relationship between wages and weekly hours of work: The role of division bias. *The Journal of Human Resources*, 15(3):409–423.

Borusyak, K., Jaravel, X., and Spiess, J. (2021). Revisiting event study designs: Robust and efficient estimation. *arXiv preprint arXiv:2108.12419*.

Burke, P. J. and Abayasekara, A. (2018). The price elasticity of electricity demand in the united states: A three-dimensional analysis. *The Energy Journal*, 39(2).

Burke, P. J. and Yang, H. (2016). The price and income elasticities of natural gas demand: International evidence. *Energy Economics*, 59:466–474.

Callaway, B. and SantAnna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of econometrics*, 225(2):200–230.

Clarke, D., Pailañir, D., Athey, S., and Imbens, G. (2023). Synthetic difference in differences estimation. *arXiv preprint arXiv:2301.11859*.

Colmer, J., Martin, R., Muûls, M., and Wagner, U. J. (2023). Does pricing carbon mitigate climate change? firm-level evidence from the european union emissions trading scheme.

Corsello, F., Flaccadoro, M., and Villa, S. (2023). Quantity versus price dynamics: the role of energy and bottlenecks in the italian industrial sector. Technical report, Bank of Italy, Economic Research and International Relations Area.

Csereklyei, Z. (2020). Price and income elasticities of residential and industrial electricity demand in the european union. *Energy Policy*, 137:111079.

Cui, J., Zhang, J., and Zheng, Y. (2018). Carbon pricing induces innovation: evidence from china's regional carbon market pilots. In *AEA Papers and Proceedings*, volume 108, pages 453–457. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Davis, L. W. and Muehlegger, E. (2010). Do americans consume too little natural gas? an empirical test of marginal cost pricing. *The RAND Journal of Economics*, 41(4):791–810.

De Chaisemartin, C. and dHaultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–2996.

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

Duprez, C. and Magerman, G. (2018). Price updating in production networks. Technical report, NBB Working Paper.

Faiella, I., Lavecchia, L., Michelangeli, V., and Mistretta, A. (2022). A climate stress test on the financial vulnerability of italian households and firms. *Journal of Policy Modeling*, 44(2):396–417.

Fontagné, L., Martin, P., and Orefice, G. (2023). The many channels of firms adjustment to energy shocks: Evidence from france. Technical report, Technical Report, Mimeo.

Ganapati, S., Shapiro, J. S., and Walker, R. (2020). Energy cost pass-through in us manufacturing: Estimates and implications for carbon taxes. *American Economic Journal: Applied Economics*, 12(2):303–42.

Gerster, A. and Lamp, S. (2023). Energy tax exemptions and industrial production.

Glynn, A. N. and Quinn, K. M. (2010). An introduction to the augmented inverse propensity weighted estimator. *Political analysis*, 18(1):36–56.

Graf, C. and Wozabal, D. (2013). Measuring competitiveness of the epex spot market for electricity. *Energy Policy*, 62:948–958.

Gros, D. (2022). *Why gas price caps and consumer subsidies are both extremely costly and ultimately futile*. CEPS.

Guiso, L. and Parigi, G. (1999). Investment and demand uncertainty. *The Quarterly Journal of Economics*, 114(1):185–227.

Haaland, I. and Roth, C. (2020). Labor market concerns and support for immigration. *Journal of Public Economics*, 191:104256.

Hahn, R. W. and Metcalfe, R. D. (2021). Efficiency and equity impacts of energy subsidies. *American Economic Review*, 111(5):1658–88.

Hausman, C. and Kellogg, R. (2015). Welfare and distributional implications of shale gas. Technical report, National Bureau of Economic Research.

Jessoe, K. and Rapson, D. (2014). Knowledge is (less) power: Experimental evidence from residential energy use. *American Economic Review*, 104(4):1417–38.

Krichene, N. (2002). World crude oil and natural gas: a demand and supply model. *Energy economics*, 24(6):557–576.

Kumar, S. and Wesselbaum, D. (2024). Contracts and firms' inflation expectations. *Review of Economics and Statistics*, 106(1):246–255.

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

Lafrogne-Joussier, R., Martin, J., and Mejean, I. (2023). Cost pass-through and the rise of inflation. *Unpublished Manuscript*.

Linn, J. (2008). Energy prices and the adoption of energy-saving technology. *The Economic Journal*, 118(533):1986–2012.

Marin, G. and Vona, F. (2021). The impact of energy prices on socioeconomic and environmental performance: Evidence from french manufacturing establishments, 1997–2015. *European Economic Review*, 135:103739.

Martin, R., De Preux, L. B., and Wagner, U. J. (2014). The impact of a carbon tax on manufacturing: Evidence from microdata. *Journal of Public Economics*, 117:1–14.

Martinsson, G., Sajtos, L., Strömberg, P., and Thomann, C. (2024). The effect of carbon pricing on firm emissions: Evidence from the swedish co2 tax. *The Review of Financial Studies*, page hhad097.

McWilliams, B., Sgaravatti, G., Tagliapietra, S., and Zachmann, G. (2024). Europes under-the-radar industrial policy: intervention in electricity pricing. Technical report, Bruegel.

Moll, B., Schularick, M., and Zachmann, G. (2023). The power of substitution: The great german gas debate in retrospect. *Brookings Papers on Economic Activity*.

Muehlegger, E. and Sweeney, R. L. (2022). Pass-through of own and rival cost shocks: Evidence from the us fracking boom. *Review of Economics and Statistics*, 104(6):1361–1369.

Nakamura, E. and Steinsson, J. (2018). Identification in macroeconomics. *Journal of Economic Perspectives*, 32(3):59–86.

Pindyck, R. S. and Rotemberg, J. J. (1983). Dynamic factor demands and the effects of energy price shocks. *The American Economic Review*, 73(5):1066–1079.

Pozzi, A. and Schivardi, F. (2016). Demand or productivity: What determines firm growth? *The RAND Journal of Economics*, 47(3):608–630.

Reiss, P. C. and White, M. W. (2005). Household electricity demand, revisited. *The Review of Economic Studies*, 72(3):853–883.

Rodano, G., Serrano-Velarde, N., and Tarantino, E. (2016). Bankruptcy law and bank financing. *Journal of Financial Economics*, 120(2):363–382.

Ruhnau, O., Stiewe, C., Muessel, J., and Hirth, L. (2022). Gas demand in times of crisis. the response of german households and industry to the 2021/22 energy crisis.

Schivardi, F., Patnaik, M., Linarello, A., and Lamorgese, A. (2021). Management practices and resilience to shocks: Evidence from covid-19.

Signorini, L. F. (2022). Sustainable investment choices: emergencies and transition. Remarks by Senior Deputy Governor Luigi Federico Signorini at the Centesimus Annus Pro Pontifice Foundation, url =https://www.bancaditalia.it/pubblicazioni/interventi-direttorio/int-dir-2022/en-SIGNORINI-11-giugno-2022.pdf?language$_i d = 1,$.

Stantcheva, S. (2022). How to run surveys: A guide to creating your own identifying variation and revealing the invisible. Technical report, National Bureau of Economic Research.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.

von Graevenitz, K. and Rottner, E. (2022). Do manufacturing plants respond to exogenous changes in electricity prices? evidence from administrative micro-data. *Evidence From Administrative Micro-Data*, pages 22–038.

Wager, S. and Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242.

Wooldridge, J. M. et al. (2002). Inverse probability weighted m-estimators for sample selection, attrition, and stratification. *Portuguese economic journal*, 1(2):117–139.

Appendices

Appendix A Questionnaires

Figure A.1: Survey questions for the energy section

Rising energy prices

At the beginning of 2021, did your firm own any instruments that protected it, wholly or partly, from energy price increases over the second half of the year?	E11

1 No
2 Yes, fixed-price contracts
3 Yes, financial derivatives
4 Yes, other instrument

	In the first half of the 2021		In the second half of the 2021	
Please indicate, even approximately, the purchased quantity and the respective cost of the following products:	Purchased quantity	Total cost (thousands of euros)	Purchased quantity	Total cost (thousands of euros)
Electricity	E9A MWh	E7A €	E9B MWh	E7B €
Natural gas	E10A Scm	E8A €	E10B Scm	E8B €

(a) 2021 wave

Rising energy prices

In 2022, did your firm have instruments (for example fixed-price contracts or derivatives) to protect itself, even partially, from the rises in the prices ...				If yes, how many months did this protection last in 2022?
A ...of electricity?	(Yes/No)	E11A	⇨	E12A
B ...of gas?	(Yes/No)	E11B	⇨	E12B

Please indicate, even approximately, the amount of electricity and natural gas purchased and their costs (gross of any tax credit): *(put 0 if you didn't purchase any during the semester)*

	In the first half of the 2022		In the second half of the 2022	
	Purchased quantity	Total cost (thousands of euros)	Purchased quantity	Total cost (thousands of euros)
Electricity	E9A MWh	E7A €	E9B MWh	E7B €
Natural gas	E10A Scm	E8A €	E10B Scm	E8B €

(b) 2022 wave

Note: The figures displays the original questionnaires of the energy section of the Invind survey, both in the 2021 and 2022 wave.

Appendix B Background pictures

Figure B.1: Expectations on wholesale gas price (TTF) implied by futures

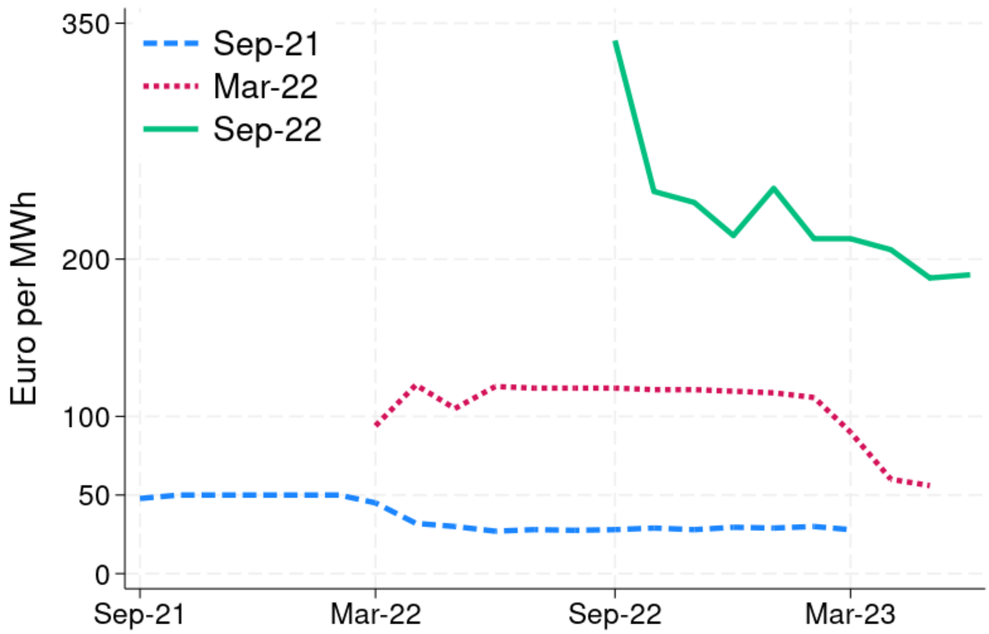

Note:The figure shows futures curve for the Title Transfer Facility (TTF) price at three different points in time: September 2021 (in blue), March 2022 (in red) and September 2022 (in green). The first point of each line is the spot price at that date.

Figure B.2: Fossil fuel energy mix for firms in the matched EU ETS - Invind sample

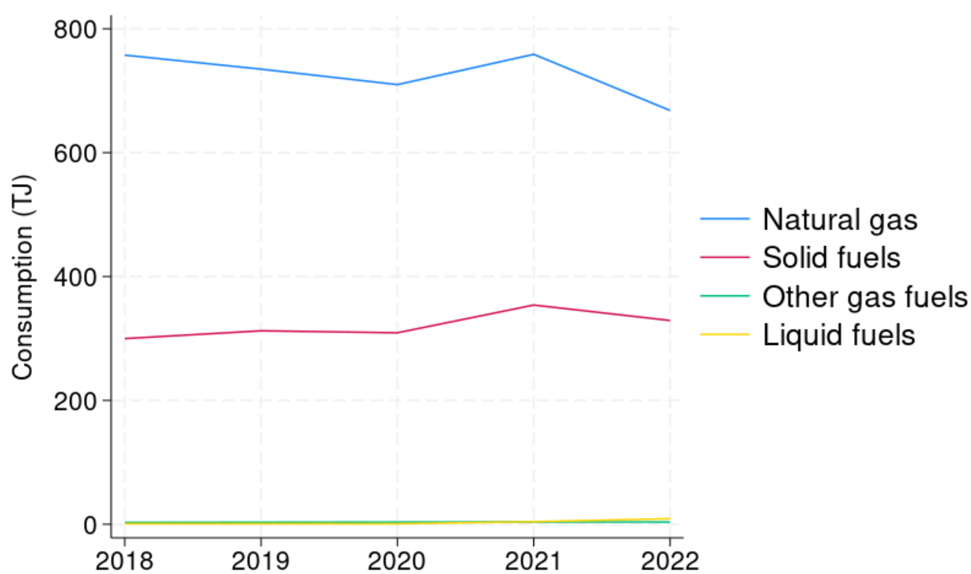

Note: The figure displays the average consumption of different fossil fuels (measured in Terajoules, TJ) for firms belonging to the matched Invind-EU ETS sample.

Appendix C Validation of survey answers

In this Appendix we detail the validation procedure implemented to check the quality of the Invind survey data. Considering that respondents might not be familiar with physical units of measurement, we verify whether quantities and costs of gas and electricity assume plausible values. To this end, we implement the algorithm described below.

First, we exclude from our sample of interest the firms that did not reply to all the energy-related questions. Note that consumed quantities must be strictly positive to be able to compute a valid retail price of energy inputs.

We then rely on two references to cross-check the plausibility of the Invind replies. In fact, given that respondents might not be familiar with physical units of energy, excessively high or low figures might indicate that respondents got the order of magnitude wrong, e.g. kWh instead of MWh. Hence, we recover the type of systematic mistake made for those replies taking on implausible values. To this end, we resort to two criteria based on year- and semester-specific parameters, and defined for electricity and natural gas separately.

We compute the average unitary price paid by firms for each semester and compare it with the corresponding average price recorded by Eurostat for the Italian market. The Invind and the Eurostat prices are constructed similarly, as they both include levies and taxes. However, Eurostat includes all non-household consumers, while we only have industrial firms with at least 50 employees. Therefore, we adopt a loose criterion and flag only those observations in which the unit price is not included in a price range defined as half the minimum price and double the maximum of the reference Eurostat statistics across consumption classes and semesters.[43] We employ a second criterion[44] based on the examination of the ratio between energy costs and turnover. We flag observations above and below the 99th and 1st percentile of the distribution[45], respectively. These correspond to cost-turnover ratios above 50% and below 0.1%, respectively. Combining the two criteria, we identify 6 error categories for the firm-level replies on electricity and 4 categories for the ones on natural gas (Table C.1). This exercise is performed for both semesters separately. In 18.7 and 2.8% of the electricity-related replies and 21 and 6.4% of the gas-related replies in the 2021 and 2022 wave respectively, we observe a consistent mistake across semesters and rescale the values accordingly. Only

[43] In light of the upward trend in prices over time, While the parameter of the maximum price is semester-specific, the lower-bound of the price distribution is considered constant over time and equal to half the minimum price reported by Eurostat in the first semester of 2021.

[44] Given that our unit prices in the Invind data could be out of the sensible range because of mistakes in filling in total expenditure (i.e the numerator) and/or the consumption quantity (i.e. the denominator), we need two criteria to reconcile implausible unitary prices with specific errors in the units of measurement.

[45] The percentiles are computed over the unweighted set of firms with strictly positive costs.

firms replying with (possibly rescaled) valid values in all 4 semesters are part of the final estimation samples.

Next, we examine the coherence of the stated quantities against administrative data collected by CSEA. This latter record is available only for the subsample of energy intensive firms (*energivore* and *gasivore*). Whenever the difference between the two figures differ by more than 35% in absolute value for at least one semester, we conservatively drop the firm from the estimation sample. The scatterplots of Figure C.1 are reassuring insofar as both the non-manipulated and the manipulated quantities lie very close to the 45-degree line, indicating that our data match administrative records pretty well and that our correction alghoritm works well.

As regards natural gas quantities, and limited to firms with plants subject to the EU-ETS, we can perform the same check also on the administrative records collected by ISPRA. In this exercise though, firms are excluded from the sample only in case the reported quantities are smaller than the ETS ones. On the contrary, larger values are compatible with multi-plant firms having lines of production not emitting CO2 emissions.

These two steps combined lead to 45 and 28 firms eliminated from the electricity and gas estimation sample, respectively.

Table C.1: Validation of data quality: units of measurement in quantities and expenditure

(1) Cost-share criterion	(2) Price-range criterion	(3) Expenditure	(4) Quantity	(5) Prevalence 2021	(6) 2022
Panel A: Natural gas					
✓	✓	000 €	SCM	70%	90%
✗ - upper tail	✗ - higher price (000-fold)	€	SCM	3%	0%
✓	✗ - higher price (000-fold)	000 €	000 SCM	18%	4.9%
✓	✗ - higher price (million-fold)	000 €	million SCM	0%	0.7%
✗ - lower tail	✗ - lower price	Million €	’000 SCM	0%	0.8%
Residual observations (dropped)				9%	3.6%
Total				100%	100%
Panel B: Electricity					
✓	✓	000 €	Mwh	71.7%	94.2%
✓	✗ - lower price	000 €	Kwh	14.3%	1.9%
✗	✗ - higher price	€	Mwh	2.7%	0.1%
✗	✓	€	Kwh	2.3%	0.1%
✓	✗ - higher price	000 €	Gwh	0%	0.7%
✗ - lower tail	✓	Million €	Gwh	0%	0.3%
✗ - lower tail	✗ - lower price	Million €	Mwh	0.1%	0%
✗ - lower tail	✗ - lower price	Million €	Twh	0.1%	0%
Residual observations (dropped)				7%	2.8%
Total				100%	100%

Note: The table presents the result of the data validation procedure. As respondents might be unfamiliar with physical units of measurement, we reviewed the plausibility of the expenditure and quantity replies, separately for gas (Panel A) and electricity (Panel B). Depending on whether unitary prices satisfy two reference criteria (Column 1 and 2), observations are sorted into mutually exclusive compilation mistakes. In more details, the two checks allow us to determine the univocal units of measurement used by the respondent (Column 3 and 4) compatible with the mistake category. This exercise is performed for both semesters. In case we observe a consistent mistake across semesters, we rescale the values with the goal of harmonising all observations in terms of thousands of euro for expenditure, and Mwh and SCM for purchased quantities of electricity and natural gas, respectively. We operate this correction in 18.7 and 2.8% of the electricity-related replies and 21 and 6.4% of the gas-related replies in the 2021 and 2022 wave, respectively (Column 5 and 6). The distributions are unweighted.

Figure C.1: Consistency with other data sources

(a) Electricity consumption in 2021

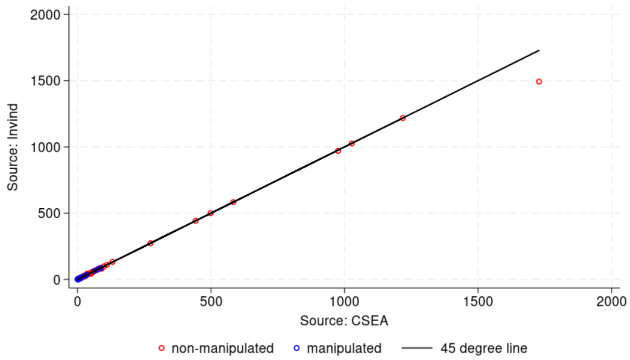

(b) Natural gas consumption in 2021

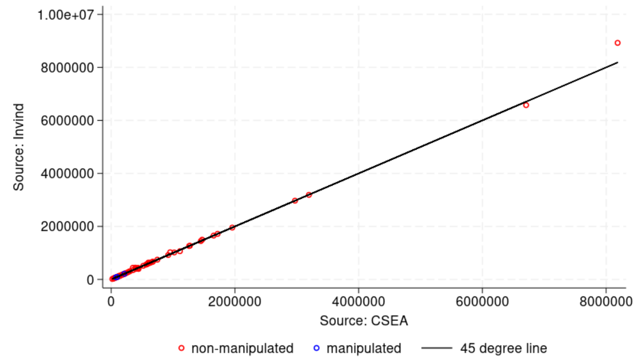

(c) Electricity consumption in 2022

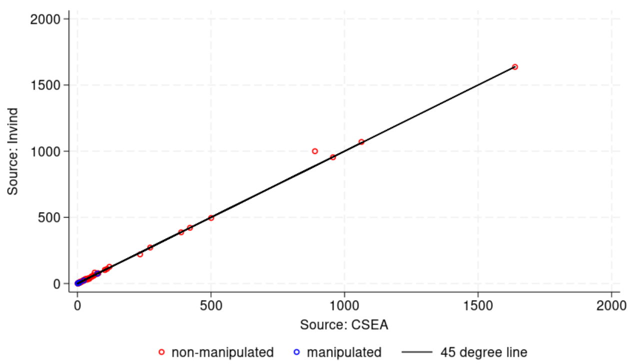

(d) Natural gas consumption in 2022

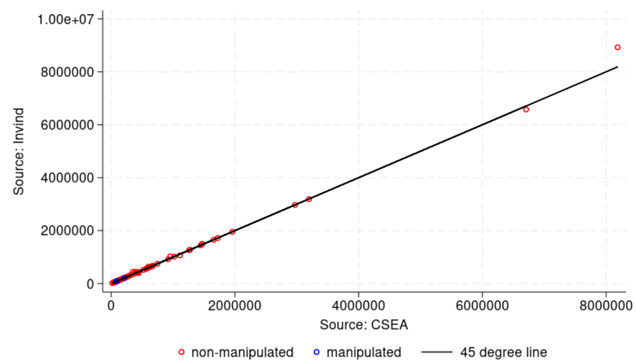

(e) Gas consumption in 2021

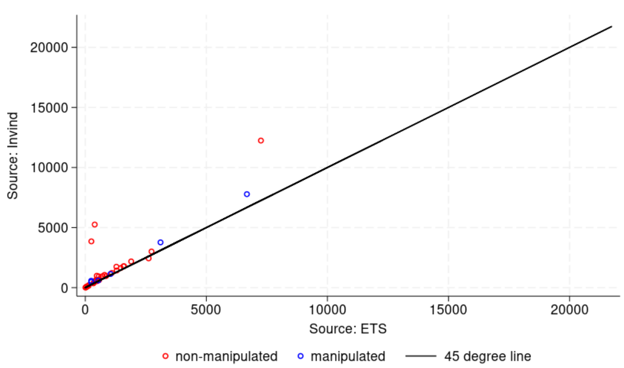

(f) Gas consumption in 2022

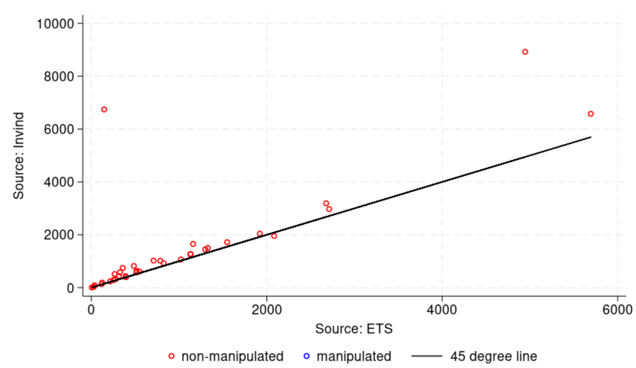

Note: The figures show the consistency between data sourced via the Invind survey and quantities of electricity and natural gas recorded in administrative data for a sub-sample of firms belonging to the *energivore* and *gasivore* lists (panel (a) and (b) for 2021, and panel (c) and (d) for 2022) and to the EU-ETS (panel (e) and (f)).

Appendix D Non-response bias

In this section we examine the robustness of our findings to a correction method called "inverse probability weighting"(Wooldridge et al., 2002; Stantcheva, 2022). This method is commonly used to address differential attrition by utilizing the relationships among observed covariates to re-weight the observed data to approximate the distribution in the full data set (Stantcheva, 2022; Glynn and Quinn, 2010). In practice, we run our baseline specification weighting observations by the inverse of the probability of being part of the respective estimation sample. The latter probability is obtained as the propensity score from estimating by logit equations 6 and 7, where X_i include covariates measured at baseline.[46].

$$\mathbf{1}(\text{Electricity sample}_i) = X_i'\beta^e + \varepsilon_i \tag{6}$$

$$\mathbf{1}(\text{Gas sample}_i) = X_i'\beta^g + \xi_i \tag{7}$$

Figure D.1 graphically indicates that for electricity and gas samples separately, the support of the propensity score overlaps between out of sample and in sample observations. We test and verify the balancing of covariates within bins (or "blocks") of the propensity score following Becker and Ichino (2002).

In Figure D.2, we compare our baseline results with those obtained by rerunning the same specification with inverse probability weighting. The two sets of results are remarkably similar, mitigating concerns about item non response biasing our results.

Figure D.1: Common support of propensity score

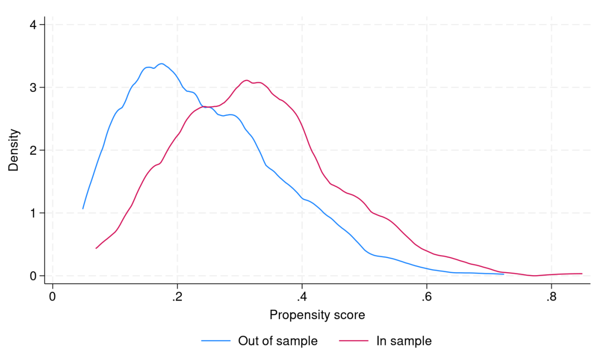

(a) Natural gas sample

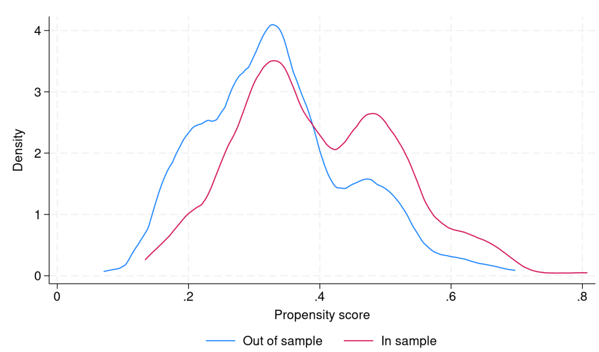

(b) Electricity sample

Note: The figures show the distribution of the propensity score of out of sample and in sample observations.

[46]We include total sales, total investment, dummies for size class, sector dummies, macroregion dummies, a dummy for being in the EU ETS, a dummy for being an electricity intensive firm, a dummy for being a gas intensive firm

Figure D.2: Inverse-probability-weighted estimates

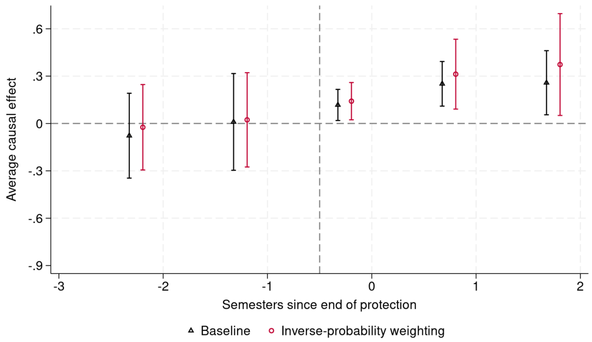

(a) Price of natural gas

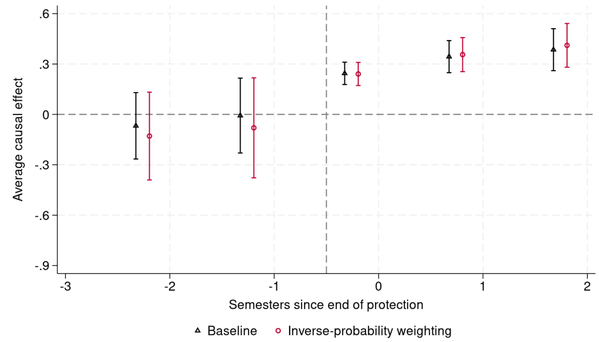

(b) Price of electricity

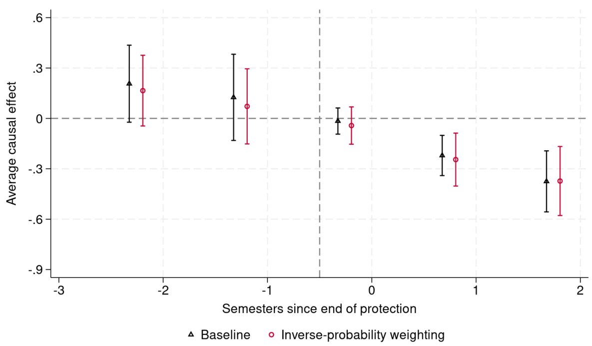

(c) Gas demand

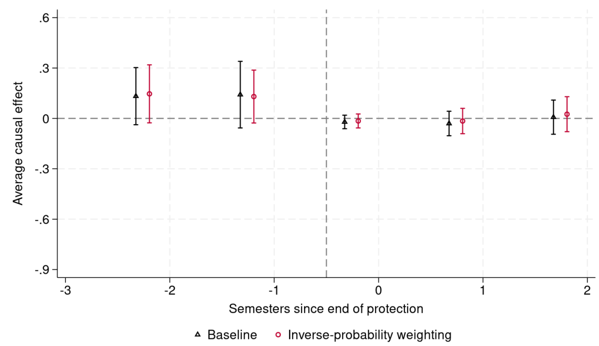

(d) Electricity demand

Note: The figures show average causal effects of the expiration of a fixed-price contract on the average costs of electricity and natural gas (panels (a) and (b)) and the corresponding demanded quantities (panels (c) and (d)). The charts compare our baseline results (in black) with those obtained by rerunning the same specification with inverse-probability weighting (Wooldridge et al., 2002; Stantcheva, 2022).

Appendix E Inference of elasticity

We derive the standard errors of the IV-style elasticity using the delta method. The elasticity of interest is:

$$\theta = \frac{e^{\hat{\tau}_t^q} - 1}{e^{\hat{\tau}_t^p} - 1}.$$

The corresponding vector of partial derivatives is: $\nabla_\theta = \begin{bmatrix} \frac{e^{\hat{\tau}_t^q}}{e^{\hat{\tau}_t^p}-1} \\ -\frac{(e^{\hat{\tau}_t^q}-1)e^{\hat{\tau}_t^p}}{(e^{\hat{\tau}_t^p}-1)^2} \end{bmatrix}$.

The variance of θ is given by:

$$\sigma_\theta^2 = \nabla_\theta^T \Sigma \nabla_\theta$$

where

$$\Sigma = \begin{bmatrix} \sigma^2_{\hat{\tau}_t^q} & \sigma_{\hat{\tau}_t^p \hat{\tau}_t^q} \\ \sigma_{\hat{\tau}_t^p \hat{\tau}_t^q} & \sigma^2_{\hat{\tau}_t^p} \end{bmatrix}$$

is the 2x2 variance covariance matrix obtained using the formula in Theorem 3 in Borusyak et al. (2021).

The variance of the elasticity is thus

$$\sigma_\theta^2 = \frac{e^{2\hat{\tau}_t^q}\sigma^2_{\hat{\tau}_t^q}(e^{\hat{\tau}_t^p} - 1)^2 + e^{2\hat{\tau}_t^p}\sigma^2_{\hat{\tau}_t^p}(e^{\hat{\tau}_t^q} - 1)^2}{(e^{\hat{\tau}_t^p} - 1)^4} - 2\frac{(e^{\hat{\tau}_t^q} - 1)e^{\hat{\tau}_t^p}e^{\hat{\tau}_t^q}\sigma_{\hat{\tau}_t^p\hat{\tau}_t^q}}{(e^{\hat{\tau}_t^p} - 1)^3}.$$

Appendix F Synthetic diff-in-diff

Motivated by the possibility of underlying pre-trends in the event study graphs presented in the main body of the paper, we probe the robustness of our results with an alternative design that explicitly matches on the path of pre-treatment outcomes: the synthetic difference-in-differences (SDID) estimator (Arkhangelsky et al., 2021).[47]

We follow the procedure outlined in Clarke et al. (2023) to implement the SDID method

[47]Note that this estimator requires a balanced sample and does not allow weights. When we replicate the staggered diff-in-diff analysis without weights the results are virtually unchanged.

in the staggered case and conduct valid bootstrap inference. Figure F.1 reports the event-study for electricity and in Figure F.2 for natural gas. We present the results for the three treatment cohorts separately. In all of the three cases, the donor pool comes from the pure control group. Naturally, the pre-trend matching uses more pre-treatment period for the "late treated" than for the "mid treated" and "early treated", as we only have four periods at disposal.

The estimates of the treatment effects are very similar to those estimated in the staggered diff-in-diff. This shows that our research design is robust to the choice of different estimation techniques that take pre-trends explicitly into account.

In addition, as in the staggered diff-in-diff, all of the three cohorts display very similar treatment effects. Thus, this exercise confirms that the results are not driven by any specific cohort. Finally, the synthetic diff-in-diff shows that the negative treatment effect on the quantity of gas is driven only by what happens in the second half of 2022.

Figure F.1: Synthetic diff-in-diff estimates of the effect of price-protection lifting on retail electricity price and quantities of electricity

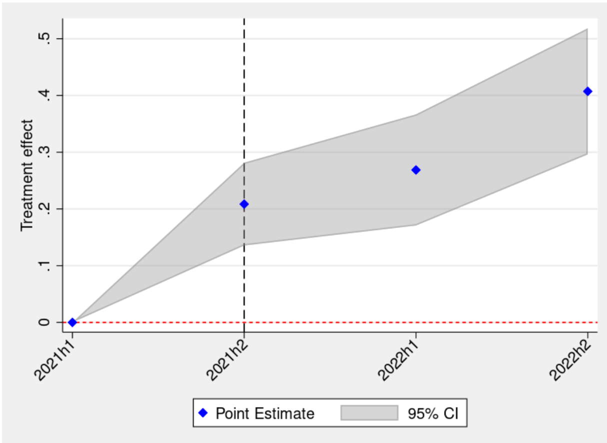

(a) Price of electricity for "early treated"

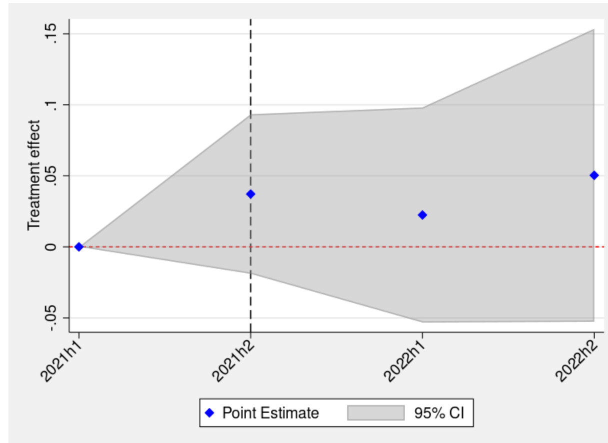

(b) Quantity of electricity for "early treated"

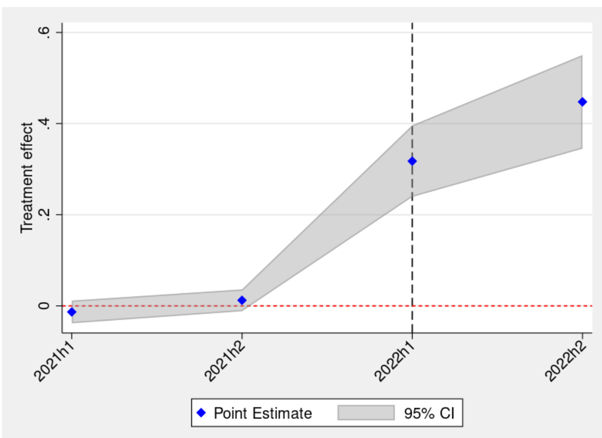

(c) Price of electricity for "mid treated"

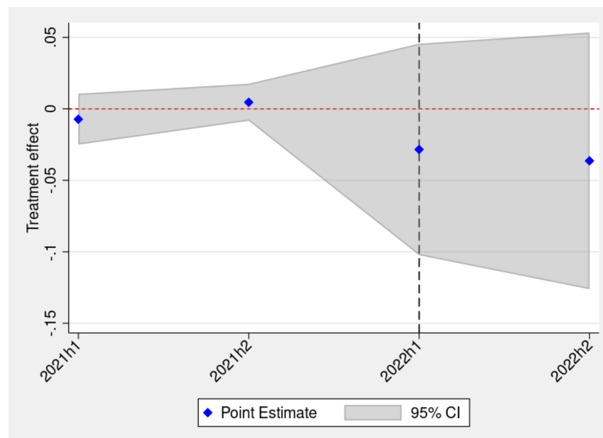

(d) Quantity of electricity for "mid treated"

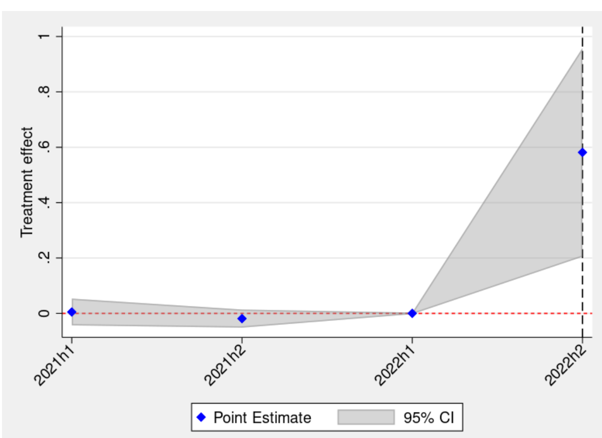

(e) Price of electricity for "late treated"

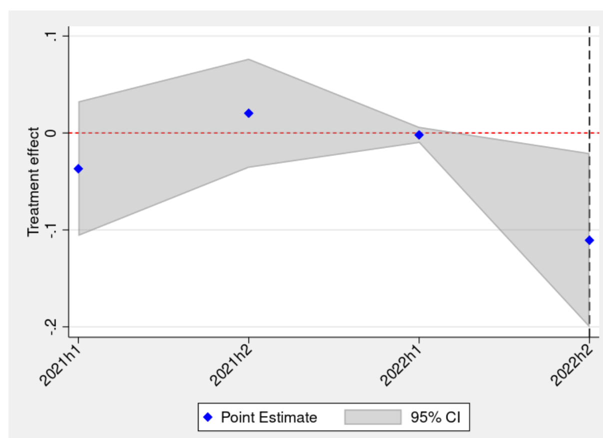

(f) Quantity of electricity for "late treated"

Note: The figure shows average causal effects of price-protection lifting according to the SDID method in the staggered case. The outcome variable is always in logs. Bootstrapped confidence intervals are at the 90% level.

Figure F.2: Synthetic diff-in-diff estimates of the effect of price-protection lifting on retail gas price and quantities of gas

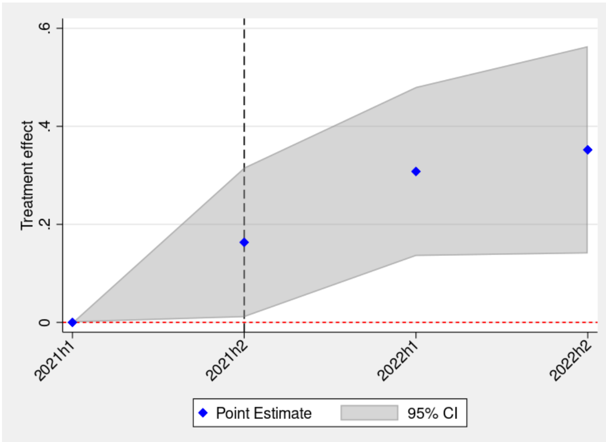

(a) Price of natural gas for "early treated"

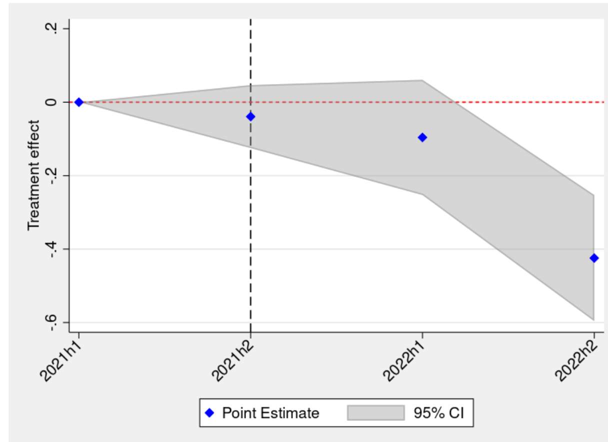

(b) Quantity of natural gas for "early treated"

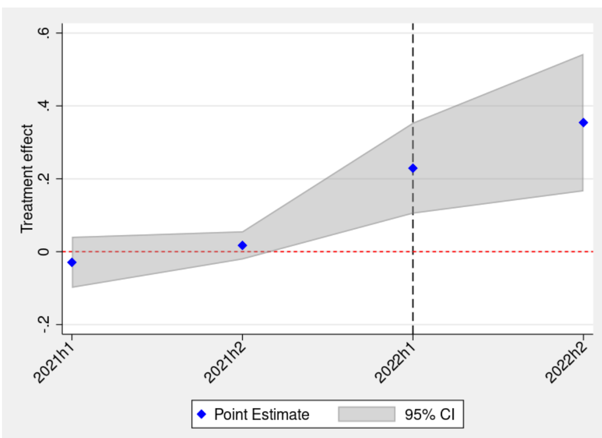

(c) Price of natural gas for "mid treated"

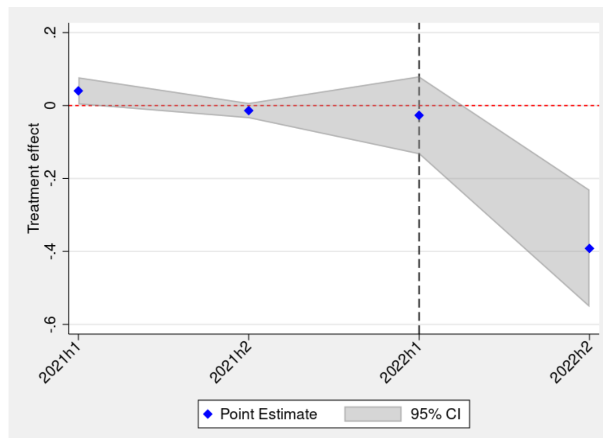

(d) Quantity of natural gas for "mid treated"

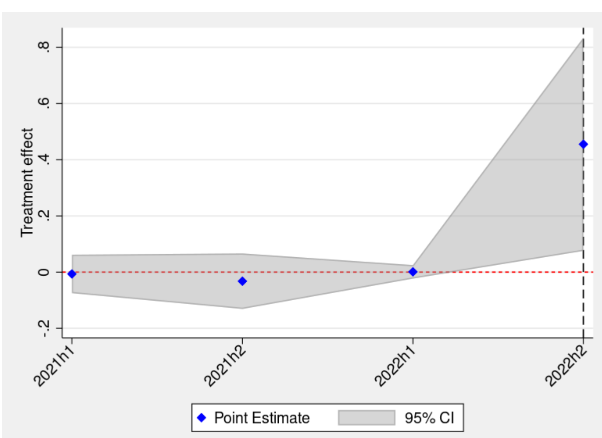

(e) Price of natural gas for "late treated"

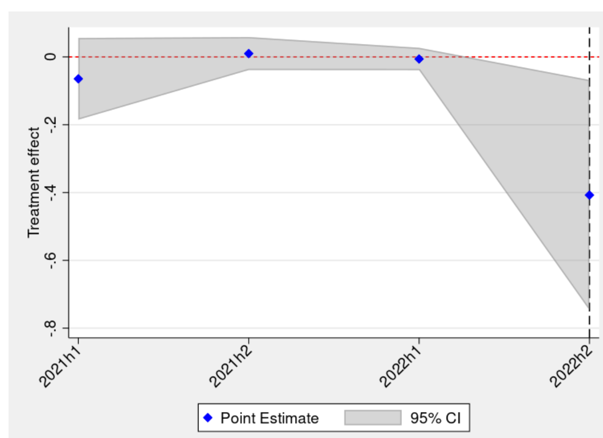

(f) Quantity of natural gas for "late treated"

Note: The figure shows average causal effects of price-protection lifting according to the SDID method in the staggered case. The outcome variable is always in logs. Bootstrapped confidence intervals are at the 90% level.

Appendix G Effect heterogeneity with machine learning

Our focus is on the effect on gas consumption in the second semester of 2022. We use random forests in the spirit of Athey and Imbens (2016) and Wager and Athey (2018). First, we split the treated observations in two random subsamples: a learn subsample (60 per cent of the overall sample) and a test subsample (40 per cent). Second, we build a forest of 5,000 trees using the learn subsample only. Each tree can pick only a random half of the considered covariates. Third, we use the forest to predict treatment effects out-of-sample in the test sample. Finally, we test whether machine learning (ML) predictions carry over to the test sample. The sample splitting approach ensures that overfitting does not drive our results.

Figure G.1 plots the distribution of individual treatment effects as estimated by the Borusyak et al. (2021) imputation method (in blue), as predicted by ML in sample (black) and out of sample (red). The point estimate is always very close to -0.45 across the three distributions.[48] The blue distribution is very dispersed. However note that this variation could be due to treatment effect heterogeneity (along observables or unobservables) or due to noise (Borusyak et al., 2021), as individual treatment effects contain the error term ϵ_{ijt} in equation 3. The distribution of in-sample-ML-predicted treatment effects (in black) is much less dispersed, but still displaying economically relevant heterogeneity and including values around zero for some observations.[49] The distribution of out-of-sample-ML-predicted treatment effects (in red) is similar to the black one, but even more compressed, with the difference plausibly due to overfitting in the in-sample predictions. Still, the red distribution has a support going from -1 and zero; this means that the forest predict that some treated firms would decrease gas consumption by as much as 60 per cent, while others would not change it at all.

In order to understand which observables predict heterogeneity in treatment effects, and in which direction, we first regress the in-sample-ML-predicted treatment effect on each covariate separately (see Table G.1). The forest predicts that firms having lower-than-average treatment effects (i.e. a small gas demand reduction) are concentrated: in the food sector; in the chemicals-pharmaceutical-rubber sector; firms declaring natural gas to be an essential input; those subject to the EU ETS; the gas intensive ones according to the Italian state aid regulation. Firms having higher-than-average treatment effects (i.e. a large gas demand reductions) are those in the wood and paper industry. The results are confirmed in the test sample when using as an outcome variable the out-of-sample-ML-predicted treatment

[48] Note that here we refer to the point estimates, and not their exponential transformation that we comment in most of the paper.

[49] This suggests that the extreme values in the blue distribution are probably due to noise, although we cannot rule out that they are due to treatment effect heterogeneity along unobservables.

Figure G.1: Distribution of treatment effects on gas demand in 2022h2

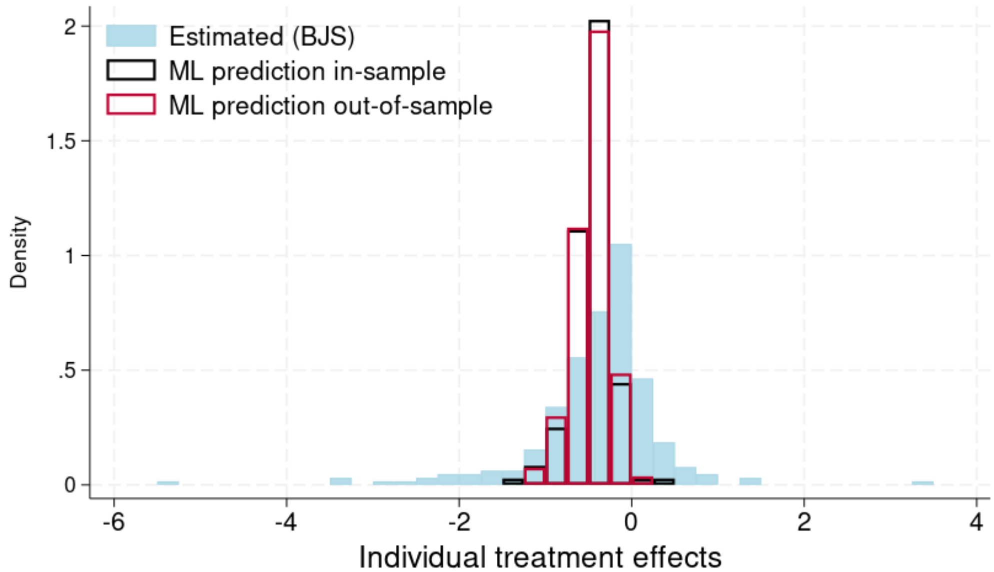

Note: The figure shows the distribution of individual treatment effects in 2022h2. The outcome is natural gas consumption.

effects (Table G.2). In order to test whether these results represent true heterogeneity and not a statistical fluke due to overfitting, we estimate the same regressions in the test subsample using as an outcome variable the estimates of the treatment effect obtained with the Borusyak et al. (2021) method. If treatment effect heterogeneity is real, we would expect to see the same signs and similar coefficients in these data, because the Borusyak et al. (2021) estimates in the test sample were never used to train the forest. Results are presented in Table G.2. For some covariates, signs are different and/or coefficients are greatly attenuated, but some of the predictions are confirmed out-of-sample, both in terms of sign and size of the coefficients. In particular, four (non mutually exclusive) groups of observations display lower gas adjustment than the average: firms in the food industry, firms declaring that gas is an essential input, firms in the EU ETS, and gas intensive firms. We take this as evidence that treatment effect heterogeneity exists in this context along these covariates.

Table G.1: Characterizing in-sample predictions of treatment effect heterogeneity

	in-sample ML predictions of treatment effects										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Food (0/1)	0.384***										
	(0.03)										
Textiles appareal (0/1)		-0.0847									
		(0.06)									
Chem., pharma., rubber (0/1)			0.0953**								
			(0.04)								
Non-metallic minerals (0/1)				-0.0825							
				(0.14)							
Metalworking (0/1)					-0.0324						
					(0.04)						
Wood, paper, furniture (0/1)						-0.195***					
						(0.06)					
Water, waste (0/1)							-0.109				
							(0.12)				
Nat. gas indispensable (0/1)								0.157***			
								(0.04)			
Employment (heads)									-0.0000121		
									(0.00)		
EU ETS (0/1)										0.135***	
										(0.05)	
Gas intensive (0/1)											0.169***
											(0.04)
R2	0.21	0.01	0.02	0.00	0.00	0.08	0.01	0.11	0.00	0.03	0.04
N	144	144	144	144	144	144	144	144	144	144	144

Note: OLS regressions in the learn sub-sample. The outcome variable is the treatment effect (on gas consumption in 2022h2) as predicted by ML. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table G.2: Characterizing out-of-sample predictions of treatment effect heterogeneity

	out-of-sample ML predictions of treatment effects					
	(1)	(2)	(3)	(4)	(5)	(6)
Food (0/1)	0.307***					
	(0.04)					
Chem., pharma., rubber (0/1)		0.0916*				
		(0.05)				
Wood, paper, furniture (0/1)			-0.205***			
			(0.06)			
Nat. gas indispensable (0/1)				0.153***		
				(0.04)		
EU ETS (0/1)					0.126***	
					(0.03)	
Gas intensive (0/1)						0.141***
						(0.04)
R2	0.16	0.03	0.08	0.12	0.02	0.03
N	107	107	107	100	107	107

Note: OLS regressions in the test sub-sample. The outcome variable is the treatment effect (on gas consumption in 2022h2) as predicted by ML. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table G.3: Testing treatment effect heterogeneity in the test sample

	treatment effects estimated using BJS					
	(1)	(2)	(3)	(4)	(5)	(6)
Food (0/1)	0.223*					
	(0.12)					
Chem., pharma., rubber (0/1)		-0.215				
		(0.14)				
Wood, paper, furniture (0/1)			0.0503			
			(0.22)			
Nat. gas indispensable (0/1)				0.341***		
				(0.13)		
EU ETS (0/1)					0.231***	
					(0.07)	
Gas intensive (0/1)						0.313***
						(0.11)
R2	0.01	0.02	0.00	0.08	0.01	0.02
N	107	107	107	100	107	107

Standard errors in parentheses

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

Note: OLS regressions in the test sub-sample. The outcome variable is the treatment effect (on gas consumption in 2022h2) as estimated by the imputation methods by Borusyak et al. (2021).

Appendix H Gas intensive firms

We analyze administrative data from CSEA where we observe monthly natural gas consumption by gas intensive firms, as defined by the Italian legislation, for years 2019, 2021 and 2022.[50] Gas consumption is measured at the gas meter level and transmitted to CSEA from the retailers; thus measurement error should be very small. We match our data to the Invind survey to obtain information on fixed-price contracts, and restrict ourselves to the period 2021-2022 for which we have more observations. Our final matched sample is a balanced panel of 126 firms. In order to avoid having cohorts with very few firms, we collapse the data at the quarterly frequency and define cohorts of treatment as follows:

$$E_i = \begin{cases} 2021q3, & \text{if } I^{2021} = 0 \text{ and } I^{gas,2022} = 0 \\ 2022q1, & \text{if } I^{2021} = 1 \text{ and } \text{ and } 0 \leq m_i^{gas} < 3 \\ 2022q2, & \text{if } I^{2021} = 1 \text{ and } \text{ and } 3 \leq m_i^{gas} < 6 \\ 2022q3, & \text{if } I^{2021} = 1 \text{ and } \text{ and } 6 \leq m_i^{gas} < 9 \\ 2022q4, & \text{if } I^{2021} = 1 \text{ and } \text{ and } 9 \leq m_i^{gas} < 12 \\ 0, & \text{if } I^{2021} = 1 \text{ and } m_i^{gas} = 12 \end{cases} \tag{8}$$

Going from the earliest treated to the latest treated, cohorts have the following number of firms: 38, 31, 4, 4 and 14. The pure control group includes 35 firms. The outcome variable is the log consumption of natural gas. Data displays strong seasonal patterns which are heterogeneous across firms. Thus we present results from two specification: our baseline (as in equation (3)) and an augmented version which includes firm-by-quarter fixed effects. In the latter model, the earliest treated cohort drop out because the last two quarters are not observed both before and after treatment. We use the Borusyak et al. (2021) estimator without any weights.

Figure A.1 presents the event-study graph from the two specifications, baseline in panel (a) and augmented in panel (b). There is no evidence of a pre-trend. Treatment effects are small and close to zero for the first four quarters. In the last two quarters, the baseline model shows a drop with very large standard errors, while the augmented model cannot estimate these two effects.

[50]See *Ministerial Decree* n.541 of 2021.

Figure A.1: Event-study estimates for the log quantity of gas among gas intensive firms

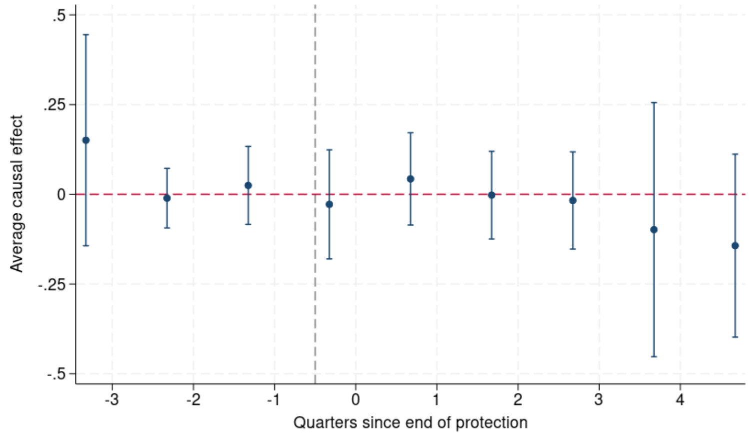

(a) Baseline model

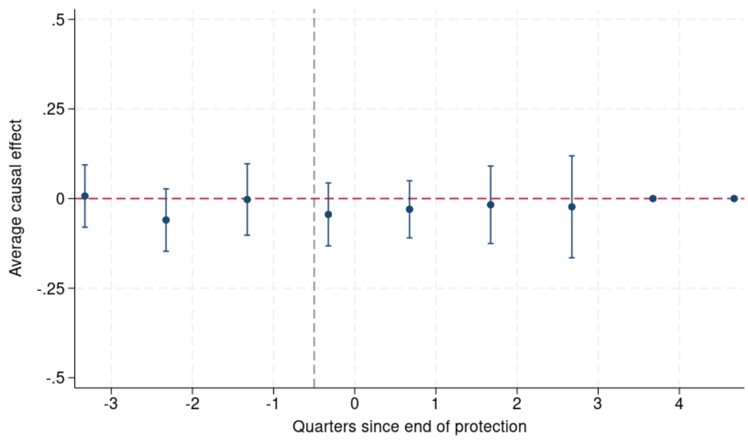

(b) Augmented model with firm-by-quarter fixed effects

Note: The figures show average causal effects of the expiration of a fixed-price contract on the log quantity of natural gas. Average causal effects before and after the treatment are estimated in two separate regressions, using the "imputation" estimator by Borusyak et al. (2021).

To test for the presence of the heterogeneity across calendar periods, we compute treatment effects by quarter and by semester. The results, reported in Table A.1, are qualitatively in line with the evidence from Invind presented in the main body of the paper. Treatment effects are close to zero until mid-2022; afterwards they are negative. In our augmented specification, the coefficients imply a 8% reduction in the second half of 2022, barely insignificant at the 90% level. In the Invind analysis, the treatment effect for this group in this period was very small and positive, but only 28 gas intensive firms were included in the sample; the effect was equal to - 25% for the remaining firms. Overall, we think that the evidence presented in this section corroborates two results of our main analysis: a) gas consumption drops, if anything, only in the second semester 2022; b) in that period, gas intensive firms reduce their gas consumption much less than other firms.

Table A.1: Treatment effects by calendar time among gas intensive firms

	(1)	(2)	(3)	(4)
ATT 21q3	0.01	0.00		
	(0.15)	(.)		
ATT 21q4	0.03	0.00		
	(0.10)	(.)		
ATT 22q1	0.04	0.01		
	(0.07)	(0.06)		
ATT 22q2	-0.05	0.00		
	(0.08)	(0.07)		
ATT 22q3	-0.05	-0.08		
	(0.11)	(0.05)		
ATT 22q4	-0.07	-0.09		
	(0.07)	(0.05)		
ATT 21h2			0.02	0.00
			(0.11)	(.)
ATT 22h1			-0.00	0.01
			(0.06)	(0.06)
ATT 22h2			-0.06	-0.08
			(0.07)	(0.04)
pvalue pre-trend	0.55	0.31	0.55	0.31
FirmXquarter FE	No	Yes	No	Yes
N	1008	854	1008	854

Note: the table presents the treatment effects of the expiration of a fixed-price contract in different calendar periods. Standard errors are reported in parentheses.

Appendix I Electricity intensive firms

We analyze administrative data from CSEA where we observe monthly electricity consumption by electricity intensive firms, as defined by the Italian legislation, for years 2018-2022.[51] Electricity consumption is measured at the meter level and transmitted to CSEA from the retailers; thus measurement error should be very small. We match our data to the Invind survey to obtain information on fixed-price protection, and restrict ourselves to the period 2020-2022. Our final matched sample is a balanced panel of 279 firms. In order to avoid having cohorts with very few firms, we collapse the data at the quarterly frequency and define cohorts of treatment as follows:

$$E_i = \begin{cases} 2021q3, & \text{if } I^{2021} = 0 \text{ and } I^{ele,2022} = 0 \\ 2022q1, & \text{if } I^{2021} = 1 \text{ and } \text{ and } 0 \leq m_i^{ele} < 3 \\ 2022q2, & \text{if } I^{2021} = 1 \text{ and } \text{ and } 3 \leq m_i^{ele} < 6 \\ 2022q3, & \text{if } I^{2021} = 1 \text{ and } \text{ and } 6 \leq m_i^{ele} < 9 \\ 2022q4, & \text{if } I^{2021} = 1 \text{ and } \text{ and } 9 \leq m_i^{ele} < 12 \\ 0, & \text{if } I^{2021} = 1 \text{ and } m_i^{ele} = 12 \end{cases} \tag{9}$$

Going from the earliest treated to the latest treated, cohorts have the following number of firms: 106, 99, 5, 5 and 19. The pure control group includes 45 firms. The outcome variable is the log consumption of electricity. We present results from our baseline, as in equation (3). We use the Borusyak et al. (2021) estimator without any weight. Figure A.1 presents the event-study graph. Treatment effects are very small and fluctuates around zero. Estimates by calendar periodo (not reported) are always close to zero.

[51]See *Ministerial Decree* n.541 of 2021.

Figure A.1: Event-study estimates for the log quantity of electricity among electricity intensive firms

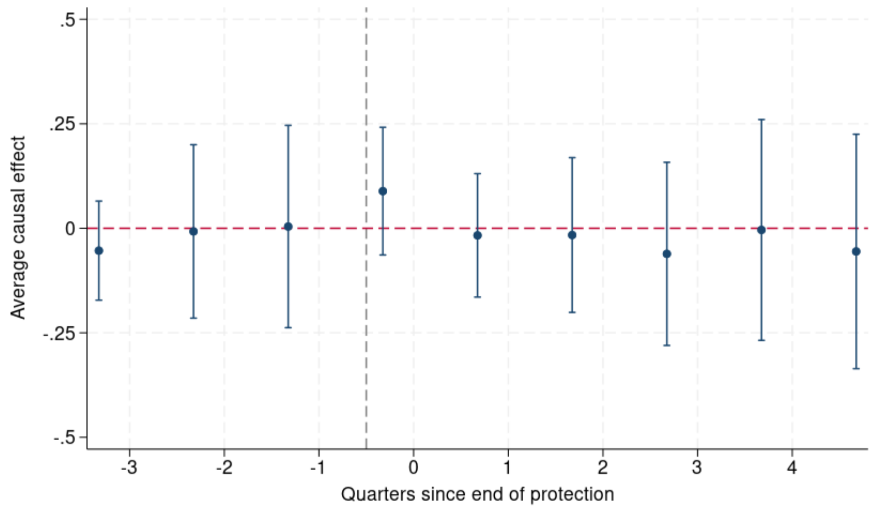

Note: The figures show average causal effects of the end of price protection on the log quantity of electricity. Average causal effects before and after the treatment are estimated in two separate regressions, using the "imputation" estimator by Borusyak et al. (2021). Confidence intervals at the 95% level.

Appendix J Additional results

Figure A.1: Baseline results with different diff-in-diff estimators

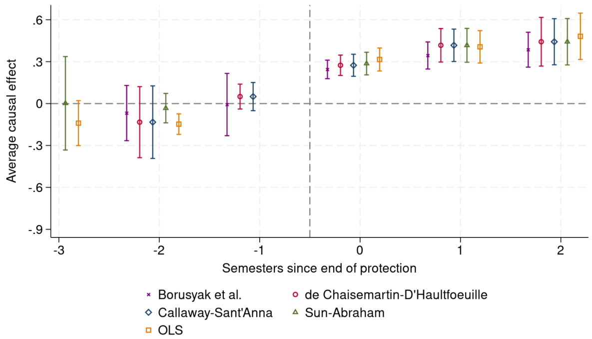

(a) Average costs of electricity

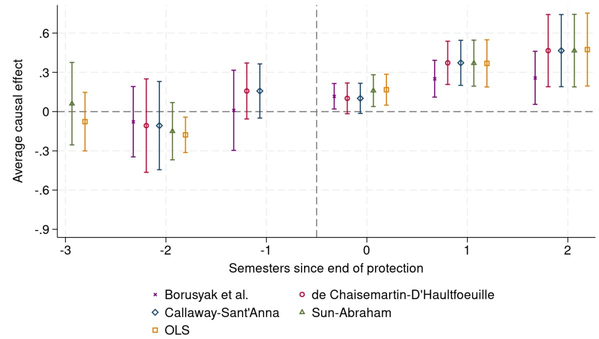

(b) Average costs of natural gas

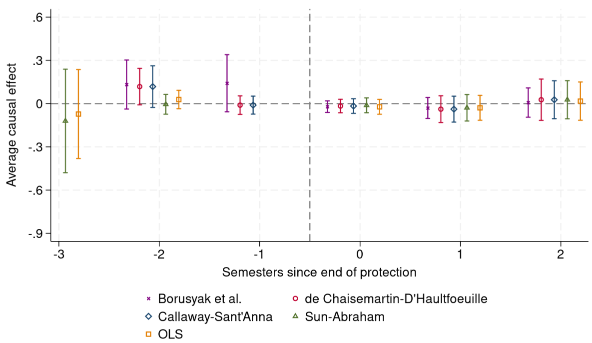

(c) Quantity of electricity

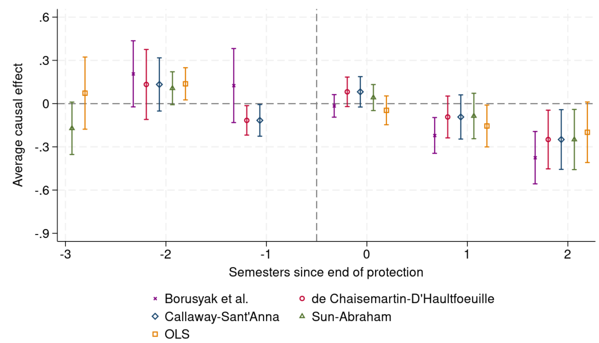

(d) Quantity of natural gas

Note: The figures show average causal effects of the end of price protection on the average costs of electricity and natural gas (panels (a) and (b)) and the corresponding demanded quantities (panels (c) and (d)). Each color corresponds to a different estimation procedure. Standard errors are clustered at the firm level. Confidence intervals are at the 95% level.

Figure A.2: Baseline results controlling for differential time trends

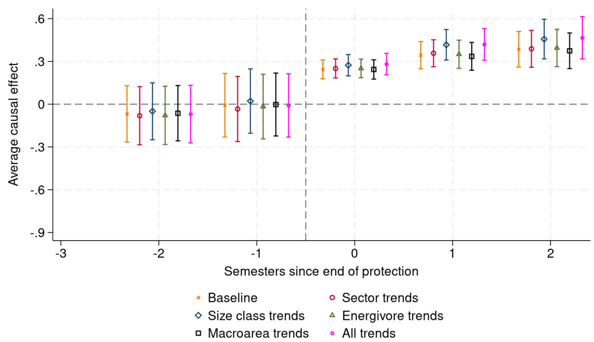

(a) Average costs of electricity

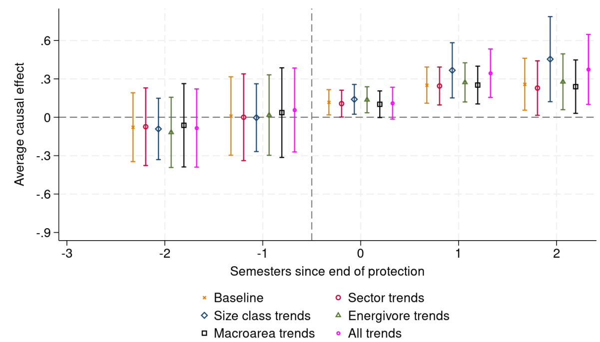

(b) Average costs of natural gas

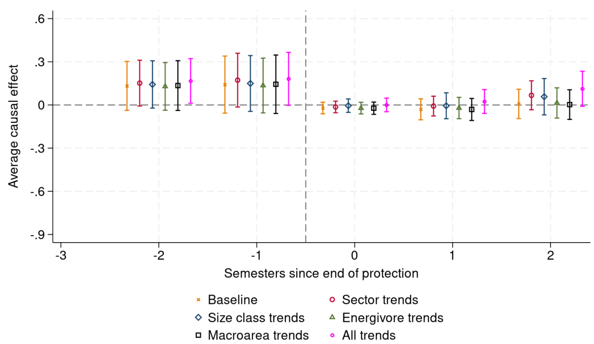

(c) Quantity of electricity

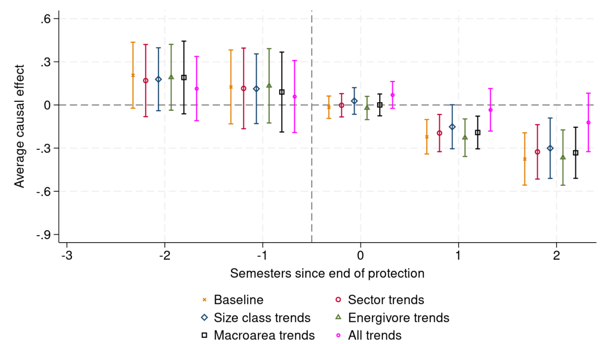

(d) Quantity of natural gas

Note: The figures show average causal effects of the end of price protection on the average costs of electricity and natural gas (panels (a) and (b)) and the corresponding demanded quantities (panels (c) and (d)). Average causal effects before and after the treatment are estimated in two separate regressions, using the "imputation" estimator by Borusyak et al. (2021), as described in Section 5. Each color corresponds to a different model where the baseline equation has been augmented with time fixed effects interacted with a different firm characteristic. Standard errors are clustered at the firm level. Confidence intervals are at the 95% level.

Figure A.3: Summary statistics for the yearly sample

	Yearly sample
	(1)
Variables	mean
Sectoral composition	
Food and beverages	10%
Textiles & apparel	7%
Chem, pharma, rubber	12%
Non-metallic minerals	3%
Metalworking industry	49%
Wood, paper, furniture	13%
Water & waste	6%
Macroarea	
North-West	41%
North-East	27%
Center	17%
South or Islands	14%
Firm-level outcomes	
Capacity utilization [0-100]	78%
%change in the price of output	2%
Firm with profit (0 if balance or loss	75%
Negative margin	12%
Energy-related variables	
Gas is an indispensable input* (0/1)	37%
Subject to EU ETS	5%
Energy intensive firm	21%
Cohorts of treatment	
Pure control	9%
Treated in 2021	58%
Treated in 2022	33%
Number of observations	595

Note: Invind data. The table reports summary statistics for the yearly sample used in Section 7. Characteristics are measured in 2019, at baseline. *The variable "Gas is an essential input" is taken from the Busines Outlook survey of the Bank of Italy and it refers to the beginning of 2022. The number of firms for "% change in the price of output" is 387 while the number of firms for the profit and loss variables is 542.