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

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

[a]Bank of Italy

May 2026

Abstract

We quantify how large, unexpected energy-price hikes affect Italian industrial firms during the 2021-22 energy crisis. For identification we exploit the staggered expiration of fixed-price energy contracts. Contract expiration raises firms' average per-unit cost of electricity and gas by 47 percent and 29 percent, respectively. Electricity demand does not respond, while gas consumption falls only in the second half of 2022, with substantial heterogeneity across firms. Gas-intensive firms, which account for 80% of industrial gas consumption, are almost perfectly inelastic, while other firms' price elasticity is -1.3. The estimated heterogeneity has important implications for policy design. A simple incidence framework shows that subsidies targeted to gas-intensive firms, combined with pay-for-reduction schemes for other firms, can focus resources to the most exposed consumers while achieving energy savings.

(*JEL classification*: Q41)

Keywords: Energy crisis, price-elasticity of energy demand

---

* 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 during part of this project. All errors are our own. The opinions expressed in this paper do not necessarily reflect those of Banca d'Italia.

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

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

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

1 Introduction

How do firms cope with a large and sudden surge in energy prices? Until recently, this question had little practical relevance. However, the 2021-2022 EU energy crisis changed that perspective, with natural gas prices skyrocketing tenfold – from 30 to 300 euros per megawatt-hour (MWh) – triggering an unprecedented spike also in electricity costs across Europe.

The energy shock elicited swift policy responses from EU governments, including energy bill subsidies, VAT reductions, and excise tax cuts to mitigate the increase in energy cost for firms – especially for energy-intensive ones. Yet, these support measures raised concerns among economists about diminishing incentives for energy conservation, potentially worsening the crisis (Signorini, 2022; Zettelmeyer et al., 2022).

The relevance of these concerns hinges on a key parameter: the price elasticity of energy demand (Gros, 2022). Despite its importance, there is limited micro-level evidence on this parameter during periods of large price volatility. Existing papers using microdata have typically focused on smaller, less disruptive energy price shocks from earlier periods (von Graevenitz and Rottner, 2022; Fontagné et al., 2024), limiting their applicability to understanding how firms navigate severe crises. In fact, a small-shock elasticity is a local derivative and extrapolates to crisis-sized price changes only under isoelastic demand. This may not hold true in reality: on the one hand, large shocks may be more salient, raising managerial attention and thereby triggering responses that are absent from normal times[1]; on the other hand, technical constraints, minimum throughput requirements and limited substitution possibilities may make demand progressively less elastic as quantities fall.

This paper examines firm-level energy demand and input substitution during the 2021-2022 energy crisis, leveraging a unique dataset that combines survey responses with detailed administrative data. For identification, we exploit the staggered expiration of fixed-price energy contracts. In the survey, firms report whether they had fixed-price contracts secured before the crisis and their duration, which we use in a staggered difference-in-differences framework. Specifically, we compare firms that lose fixed-price protection during the observation period to those whose protection will expire later or does not expire within the observation window.

To estimate effects, we rely on the imputation estimator of Borusyak et al. (2024). However, our results are robust to alternative staggered difference-in-difference estimators (De Chaisemartin and d'Haultfoeuille, 2020; Callaway and SantAnna, 2021; Sun and Abraham, 2021)

[1]For evidence on the role of salience in other settings see Chetty et al. (2009) and Jessoe and Rapson (2014).

and to the synthetic difference-in-differences method (Arkhangelsky et al., 2021), which considerably loosens the requirement of parallel trends between treated and control units and addresses concerns related to pre-trend testing (Roth, 2022). Results are also robust when controlling for non-parametric time trends that vary at the level of industry, energy-intensity, size and participation in the EU Emission Trading System (ETS).

Our analysis reveals that the expiration of fixed-price contracts leads to a sudden increase in the average unit costs of electricity and natural gas at the firm level. The effect sizes are economically meaningful: average unit costs rise up to 47% in the case of electricity and 29% in the case of natural gas, even after accounting for government policies that partially mitigated cost increases. These findings underscore the critical role of fixed-price contracts, widely used for various input purchases (Kumar and Wesselbaum, 2024), in the transmission of macroeconomic shocks.

In response to these substantial cost increases, firms exhibit markedly different patterns of adjustment across energy inputs. On average, firms make virtually no changes to their electricity consumption. In contrast, natural gas consumption declines sharply, but this reduction, estimated at 35%, occurs exclusively in the second half of 2022. The delayed adjustment aligns with firms increasingly pessimistic expectations during the summer of 2022, as spot gas prices surpassed €300 per MWh, and futures markets projected sustained high prices of around €200 per MWh through mid-2023. Combining these observed quantity adjustments and corresponding energy unit cost increases, we estimate the elasticity of natural gas demand to be -1.1 in the second half of 2022, while it remains effectively zero in the two previous semesters.

During the recent energy crisis, EU policymakers gave larger subsidies to energy-intensive firms (McWilliams et al., 2024). Providing an assessment on whether this type of targeting was more or less detrimental for energy conservation requires us to go beyond estimating elasticities for the "average" firm. Economic theory offers mixed predictions about whether elasticities would be higher for energy-intensive firms, as this depends on several parameters, including the elasticity of substitution across inputs (Cahuc et al., 2014). This leaves the question to be resolved empirically.

Using random forests to detect treatment effect heterogeneity, we show that gas-intensive companies reduced gas consumption significantly less than other firms, a result confirmed also when using confidential administrative data on monthly gas usage among these firms. Gas-intensive firms exhibit a very low elasticity of -0.03 in the second half of 2022, compared to -1.3 for other firms. This difference appears driven by a low elasticity of substitution between gas and other inputs among gas-intensive firms. Two pieces of evidence support this. First, survey responses indicate that gas-intensive producers almost always classify natural gas as an *essential* input. Second, plant-level data for EU-ETS facilities, which are large gas consumers, reveal little switching from this energy carrier to alternative fossil fuels.

The limited adjustment in terms of energy demand suggests that, on average, firms did not reduce output substantially in 2021-22. We provide complementary evidence that firms, and even more so energy-intensive ones, raised output prices substantially, passing a large share of the shock onto customers, in a macro environment characterized by strong post-pandemic recovery. To a lesser extent, firms, which on average entered the energy crisis with substantial liquidity buffers and good financial conditions (Banca d'Italia, 2021), absorbed the shock in their margins.

Our results carry significant implications for policy design. In the final part of the paper we present an incidence framework in partial equilibrium to study the effects of targeted and untargeted support measures on equilibrium prices, quantities and fiscal costs. We keep the model simple so that its key objects can be mapped directly to quantities we estimate in the data, in the spirit of Deryugina et al. (2020) and Hahn and Metcalfe (2021). For targeted policies, group-specific elasticities govern the direct demand response. But once equilibrium price adjustments are accounted for, both targeted and untargeted policies depend on a *quantity*-weighted elasticity across all firms, since *all* firms respond to the subsidy-induced price increase, irrespective of whether it is targeted or untargeted.[2] In our setting, gas consumption is highly concentrated among a small number of gas-intensive firms with near-zero demand elasticity, so the quantity-weighted elasticity is far below the unweighted average and determines how much support measures leak into higher wholesale prices rather than reaching firms as genuine relief.

Calibrating the model to Italy in 2022H1 and using our estimated elasticities, we first compare an untargeted per-unit subsidy of €0.5/Smc to an equivalent one targeted just to gas-intensive firms. The model predicts that the untargeted subsidy raises aggregate gas use by 5.3% and the gas price by 26.3%, largely due to the response of non-gas-intensive firms, costing €3.5 billion in total. The targeted subsidy instead shrinks the equilibrium price increase to just 2.2% and the quantity increase to 0.4%. The fiscal cost of the targeted subsidy is approximately one fifth of that of the untargeted one.

Next, we extend the analysis to pay-for-reduction schemes – which reward firms for verified consumption cuts relative to a given baseline.[3] These schemes raise the opportunity cost

[2]This incidence result is not new and can be found in Chetty (2009), with an application to income taxation.

[3]For example, Germany's *Gasminderungsausschreibungen* and Italy's gas demand reduction auctions. See

of consuming gas, and are isomorphic to a per-unit tax. Thus equilibrium responses on the energy price and quantities have the same magnitude, but opposite sign compared to those generated by subsidies. Finally, we show that policymakers willing to help the most vulnerable businesses, while minimizing equilibrium responses that exacerbate the crisis, can resort to a policy mix: subsidies for gas-intensive firms, and pay-for-reduction for the others. In our calibration, a €0.5/Smc subsidy to gas-intensive firms combined with an equal per-unit reward for reductions to other firms reduces equilibrium gas consumption by 4.4% and lowers prices by 21.9%. The policy mix has a fiscal cost of less than €3 billion, mostly transferred to gas-intensive firms.

We contribute to the literature on energy demand elasticities for firms.[4] From a substantive perspective, we make a unique contribution by estimating elasticities during a period of severe crisis. As explained above, small-shock elasticities extrapolate to large shocks only under isoelastic demand. One possibility is that technical constraints make energy demand progressively less elastic as quantity decreases. Consistently with this, studies on electricity demand that examine small price shocks tend to find larger elasticities (0.4-0.5) than those exploiting larger variations (0.1-0.2) (Marin and Vona, 2021; von Graevenitz and Rottner, 2022; Fontagné et al., 2024; Blonz, 2022; Gerster and Lamp, 2024). Using even larger shocks, we find that the elasticity can approach zero.

For natural gas, we provide two novel findings. First, elasticities are close to zero in 2021 and the first half of 2022, and then increase in the second half of the year, irrespective of time elapsed since contract expiration. This later phase of the crisis coincided with deteriorating expectations about its duration. The existing literature already established that long-run elasticities are higher (Deryugina et al., 2020; Labandeira et al., 2017; Sikl et al., 2025); our evidence suggests that firms hit by the shock react only when they stop to perceive it as temporary, making it rational to undertake costly adjustments to reduce gas consumption (Pindyck and Rotemberg, 1983; Atkeson and Kehoe, 1999).

Second, gas-intensive firms are less elastic compared to the average industrial firm, despite energy representing a larger share of their costs. We think that this result reflects low substitutability of gas with other inputs. Consistent with this interpretation, most gas-intensive firms declare gas to be an essential input in their production process, and additional evidence from EU ETS plants shows that lower gas use is only partly offset by switching to other fuels.

Sgaravatti et al. (2021) for a survey.

[4]While household demand elasticities are well-studied (Reiss and White, 2005; Jessoe and Rapson, 2014; Auffhammer and Rubin, 2018; Hahn and Metcalfe, 2021; Costa and Gerard, 2021; Byrne et al., 2021; Deryugina et al., 2020; Ahlvik et al., 2025), estimates for industrial firms are scarce.

From a methodological perspective, our study introduces a novel identification strategy based on the staggered expiration of firms' fixed-price contracts. This approach appears more credible than earlier studies, which exploit sectoral or time-series variation (Burke and Abayasekara, 2018; Csereklyei, 2020; Graf and Wozabal, 2013; Davis and Muehlegger, 2010; Hausman and Kellogg, 2015; Faiella et al., 2022). Closest to our design, Ahlvik et al. (2025) exploit quasi-random expiration dates of fixed-term electricity contracts to study household responses to the Finnish energy crisis.

We also contribute to the literature on the 2021-2022 energy crisis. Previous studies have mostly focused on inflation and output effects (Ruhnau et al., 2023; Alessandri and Gazzani, 2023; Moll et al., 2023; Lafrogne-Joussier et al., 2023; Corsello et al., 2023; Wehrhöfer, 2026). In this paper we provide the first firm-level evidence on industrial firms' input demand, highlighting the role of treatment effect heterogeneity across time, energy source and firms' type.[5] Our micro elasticities can also inform macroeconomic models on the crisis impact (Bachmann et al., 2022; Nakamura and Steinsson, 2018).

Finally, by combining firm-level elasticities with a simple incidence framework, we connect to the literature on energy subsidies and their cost (Davis, 2014; Hahn and Metcalfe, 2021).

2 Background and research design

In this section, we provide a brief account of the timeline of the 2021-22 energy crisis, separately for the wholesale and retail markets. Here we limit ourselves to those elements that are central to our research question and essential for understanding our research design, while we defer further details to Appendix A.

[5] For sectoral level estimates on the effects on gas consumption see Emiliozzi and Favero (2025).

2.a Wholesale market developments

Figure 2.1: Wholesale energy prices

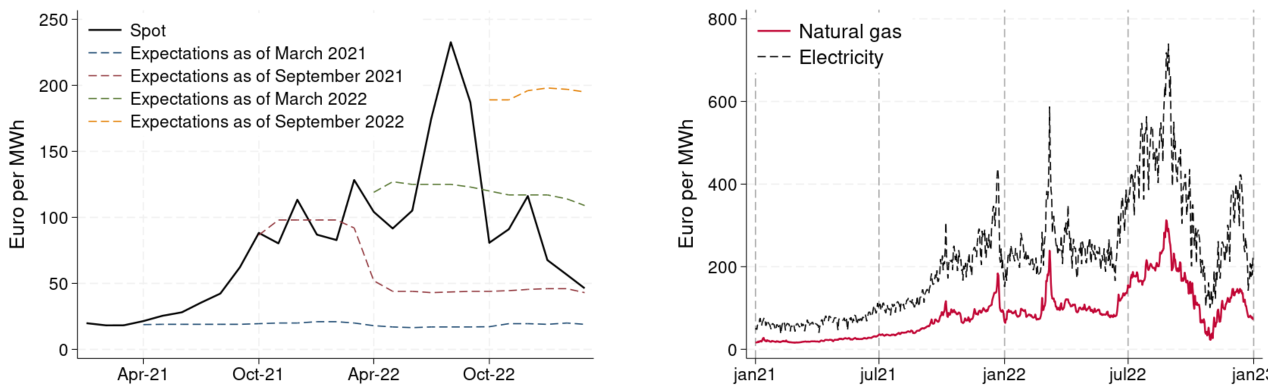

(a) Wholesale prices for natural gas in Europe (b) Wholesale energy prices in Italy

Source: Italian Power Exchange (GME) for the spot prices and our calculations on www.barchart.com data for the forward curves. *Note*: Panel (a) plots the spot price of natural gas in the Italian wholesale market together with forward curves for the Dutch TTF price observed at four points in time (March 2021, September 2021, March 2022, and September 2022), derived from futures prices. Forward curves are interpreted as market expectations of future spot prices. All series are monthly averages. Panel (b) shows the spot price of natural gas and electricity traded in the Italian day-ahead wholesale market at the daily frequency. The correlation between the two is 98%.

As late as March 2021, forward curves implied by natural gas futures contracts showed no indication of the impending energy crisis: market participants expected gas prices to remain close to their historical average (approximately 20 euro/MWh) throughout 2021 and 2022 (see the blue line in panel (a) in Figure 2.1).

The first signs of a potential energy crisis emerged around May 2021, when spot prices at the main European gas hubs began to gradually rise above historical levels[6]. The first major upswing occurred in the fall of 2021, when spot prices exceeded 100 euro/MWh. By September 2021, markets expected prices to remain at that level throughout the winter (red line in panel (a) in Figure 2.1). These forecasts proved overly optimistic: spot prices increased further following Russias invasion of Ukraine in late February 2022. Markets subsequently revised their expectations upward, with forward curves in March 2022 indicating gas prices above 100 euro/MWh through late 2023 (see green line in panel (a) in Figure 2.1).

These forecasts again proved too optimistic, as spot prices surged to unprecedented levels during the summer of 2022, exceeding 200 euro/MWh amid concerns that Europe would be

[6] The main price benchmark is the Dutch TTF. The Italian PSV is closely correlated due to high market integration, but usually sits at slightly higher levels.

unable to replenish gas storage ahead of the following winter. By September 2022, forward curves projected gas prices of around 200 euro/MWh well into 2023 (yellow line in panel (a) in Figure 2.1). This time, however, forecasts proved overly pessimistic, as spot prices began a steep decline, reaching approximately 50 per MWh by March 2023.

In Italy, the high reliance on natural gas in electricity production coupled with the marginal price system at work in the wholesale power market implies that shocks to the wholesale price of natural gas almost completely pass through to the wholesale price of electricity. The two prices are indeed almost perfectly correlated (Panel (b) in Figure 2.1).

2.b Retail market developments

Wholesale prices do not transmit immediately to retail prices. Firms purchase gas and electricity through retail contracts with durations ranging from 12 to 24 months, signed with energy retailers operating in a relatively liberalized and unconcentrated market (ARERA, 2022).[7] By law, energy purchased under these contracts cannot be resold[8].

Contracts available in the liberalized market can differ in terms of both contractual conditions and pricing structures. For the purposes of our analysis, however, the most relevant distinction is between fixed-price and variable-price contracts. Fixed-price contracts lock in the retail price for 12 to 24 months, whereas variable-price contracts are indexed to (lagged) developments in wholesale markets.

Firms on variable-price contracts experienced the energy shock with a lag of approximately one month. By contrast, firms on fixed-price contracts[9] signed before the onset of the crisis, certainly before July 2021, were fully shielded from the energy price shock until the predetermined expiration date of their contracts. In the extreme case, a firm that signed a 24-month fixed-price contract in March 2021 was able to avoid the 2021–22 energy crisis entirely. Firms whose fixed-price contracts expired during the crisis, however, were forced to replace them with new contracts, either fixed or variable, and therefore necessarily faced

[7] Small firms, which are outside the scope of our analysis, still had access to regulated tariffs; see Appendix 2 for details.

[8] This distinction is explicit in Directive (EU) 2019/944 and is reflected in standard retail contracting practice: retail supply contracts deliver electricity at the firm's metering point for on-site use and do not confer a right to resell that delivered energy on the wholesale market. Some firms may operate in the wholesale market to resell self-generated electricity, but this requires to be registered in the official register of wholesale market operators. This does not appear to be empirically relevant in our data: after matching our sample to the official list of participants in this register, only 4 firms in our sample are registered market operators. We provide more details on the institutional setting in which Italian industrial firms operate in Appendix A.c.

[9] According to the Regulatory Authority for Energy, Networks and Environment (ARERA), in 2021 45% of firms had fixed-price contracts for electricity and 44% for natural gas.

a large increase in energy costs.

Figure 2.2: Retail energy market developments in Italy

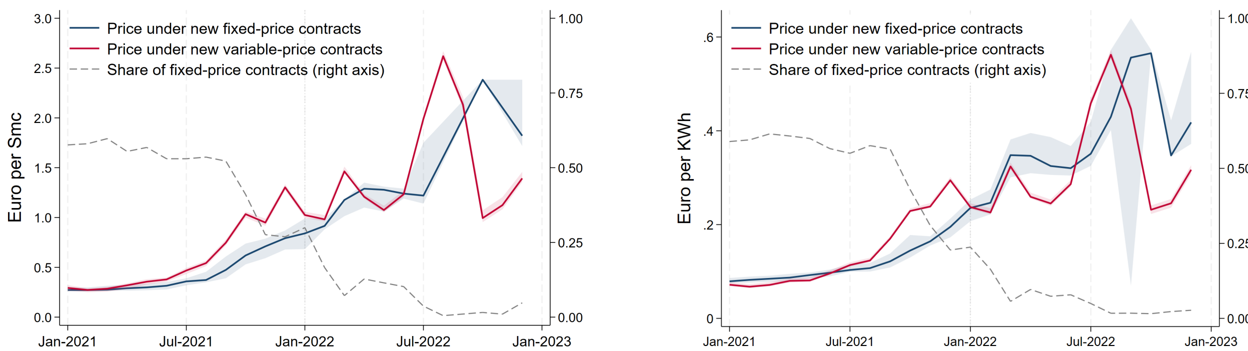

(a) Natural gas

(b) Electricity

Source: own elaborations on data from Portale Offerte, ARERA. *Note*: The figure reports information on retail energy contracts available to business consumers at different points in time; each observation corresponds to the last day of the month. The solid lines show the median per-unit energy price (net of taxes, system and network charges, and the fixed component of two-part tariffs) across available contracts, while the shaded area the corresponding interquartile range. The dashed line plots the share of contracts available for sale that feature fixed prices.

As the crisis unfolded, developments in wholesale markets were progressively transmitted to newly issued retail contracts. Prices for both fixed-price and variable-price contracts available on the retail market increased sharply beginning in June 2021. Firms signing new fixed-price contracts for gas in the fall of 2021 locked in prices already 113% (68% for electricity) above historical levels, while firms doing so in the first half of 2022 locked in prices 200% (165% for electricity) above historical levels. By mid-2022, expectations in wholesale markets had become so adverse that fixed-price contracts had almost disappeared from the retail market, leaving expensive and risky variable-price contracts as the only option for firms whose existing contracts expired.

To identify the effect of the energy shock on Italian firms, we exploit the predetermined staggered expiration dates of fixed-price contracts signed before the crisis. Crucially, once the crisis began, firms could no longer insure themselves in a way that would shield them from a sizable energy shock.

3 Data and Measurement

3.a The Invind survey

The primary data source is the Inquiry into investments of industrial and services firms (henceforth, Invind), an annual Bank of Italy survey, representative of industrial and services firms with at least 20 employees and routinely used for institutional purposes. The Bank surveys the same companies every year, adjusting for firm exit and addressing *unit* non-response through raking post-stratification that aligns survey weights with population distributions. Widely used in the literature, Invind has informed research on topics such as productivity shocks (Pozzi and Schivardi, 2016), bankruptcy law (Rodano et al., 2016), investment demand (Guiso and Parigi, 1999; Bond et al., 2015), agglomeration economies (Andini et al., 2013), and management practices during Covid-19 (Lamorgese et al., 2024).

A key advantage of Invind is its detailed, timely data on firms' energy expenditures, consumption, and hedging strategies in 2021 and 2022 – information that is not available in balance-sheet datasets.[10] To collect these data, we introduced an *ad hoc* energy section to the 2021 wave, administered in spring 2022, targeting industrial firms with 50 or more employees.[11] This section gathered semester-level electricity and gas expenditures (in €) and consumption (in megawatthour, MWh and standard cubic meters, smc) for the previous year, excluding self-produced energy (Section 3.c). Firms were also asked about fixed-price contracts or equivalent hedging tools in place at the beginning of 2021 – before the start of the crisis –, forming the basis of our treatment variable (Section 3.d). The following year's wave collected analogous data for 2022, adding information on tax credits under the energy cost relief program implemented by the Italian government in that year.[12] The response rate to the energy section is high (around 50%) and our results are robust to addressing selective item non-response with inverse-probability weighting (Appendix E).

In all our analyses, we focus on firms that use energy as an *input*, thus excluding NACE sectors 19 (manufacture of coke and refined petroleum products) and 35 (electricity, gas, steam and air conditioning supply).

[10] For background information on the retail energy market in Italy see Appendix 2.

[11] According to Istat's Business Register, in 2021 this group accounted for 49% of sectoral employment and 56% of sectoral turnover. Energy use is even more concentrated: administrative data from the Fund for Energy and Environmental Services (CSEA) show that energy-intensive firms with at least 50 employees alone account for 74% of industrial natural gas consumption and 40% of industrial electricity consumption.

[12] We report all the questions, including the tax-credit item, in Appendix C.

3.b Administrative data and other sources

For the subset of energy-intensive firms, we use confidential administrative records from the Fund for Energy and Environmental Services (CSEA), which we match to our survey through unique tax identifiers. This group includes nearly 3,800 firms, which collectively account for 80% of industrial natural gas consumption, 45% of industrial electricity consumption, and 17% of the sector's value added[13]. For electricity-intensive firms, CSEA data include monthly electricity consumption since 2018 and up to 2022. As for gas-intensive firms, we observe monthly gas consumption for the years 2019, 2021, and 2022.[14]

For the subset of firms owning plants subject to the EU Emissions Trading System (EU ETS), we gather additional confidential information from the Italian Institute for Environmental Protection and Research (ISPRA), including yearly plant-level fuel use, which we use to study input substitution. In Italy, the EU ETS includes around 1,000 plants belonging to 300 firms. In 2019 these facilities accounted for 63.5% of natural gas consumption in the industrial sector, while their share of value added was approximately 1.5%.

Additionally, we use the 2022 Bank of Italy's Business Outlook Survey of Industrial and Service Firms (Sondtel) to identify firms that declare natural gas to be an "essential" input; we consider this a proxy for low elasticity of substitution between gas and other inputs.

3.c Measurement of energy-related outcome variables

The first outcome variable of interest is the (log of) quantity consumed (of both electricity and natural gas separately). The surveys record semester-level electricity and gas consumption (in megawatthour, MWh and standard cubic meters, smc, excluding self-produced energy) in 2021 and 2022. We ensure self-reported data on energy consumption are plausible and aligned with administrative records, where available, applying a cleaning algorithm to correct systematic errors (e.g., kWh instead of MWh). Validation results, detailed in Appendix D, confirm the data reliability.

The second outcome variable is the (log of) average unit cost. While the quantity of energy consumed is directly observable, we construct energy unit costs by dividing overall expenditures (in €) for a given energy input in a semester by the corresponding physical

[13]Restricting to firms with at least 50 employees, energy-intensive firms account for 74% of industrial natural gas consumption and 40% of industrial electricity consumption.

[14]These firms qualified for energy bill discounts even before the crisis. Eligibility required annual energy consumption over 1 GWh and operation in specific 4-digit NACE sectors per EU State Aid rules. Some sectors faced additional criteria, related to energy intensity. Subsidies reduced price components like renewable energy fees, which were waived for smaller firms in late 2021 and all firms by January 2022.

quantity. Importantly, we make sure that our expenditure variables exclude all forms of aid provided by the government to cushion the effects of the energy crisis[15]. We use Eurostat data on average unitary cost by consumption bracket for industrial consumers to validate our survey-based measures (see Appendix D).

Retail energy contracts are usually in the form of two-part tariffs, thus the average unit cost departs from the per-unit marginal price by a wedge equal to the ratio of the fixed-charge and the quantity consumed. In Appendix B we use open data on the universe of retail energy contracts published by ARERA and show that for natural gas – the input for which we estimate demand response – the wedge is between 0.7 and 1% before the crisis. For electricity, the fixed-charge structure implies a somewhat larger wedge (between 8 and 10%). The implications for our estimates are discussed in Section 5.[16]

The elasticity with respect to the average unit cost is the policy-relevant parameter in our context: the support measures implemented during the 2021–22 energy crisis reduced billed unit costs by acting on both fixed and per-unit components of the energy bill simultaneously.[17] Furthermore, evidence from the household sector suggests that energy users respond to average rather than marginal prices under nonlinear schedules (Ito, 2014), making the average unit cost the behaviorally relevant price index for demand decisions.

3.d Treatment variables

In the 2021 wave of the survey, firms were asked:

> "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?"

[15]Starting from the second semester of 2021, the Italian government implemented support measures including system charge reductions and VAT cuts. These are automatically netted out from reported expenditures, as these are directly reflected in firms' energy bills. In 2022, the government also introduced tax credits covering a fraction of energy expenditures. These are not automatically netted out from firms' expenditures, but we observe them in the Invind data thanks to a specific question, so we are able to net them out. Therefore, we are able to construct an accurate measure of unit costs actually faced by firms, *net* of any government support. Concretely, our average unit cost constructed net of government support reads:

$$\widetilde{p}_{it} = (\text{Expenditure}_{it} - \text{Tax Credit}_{it}) \, / \, \text{Quantity}_{it}.$$

[16]These values represent an upper bound during our sample period. From 1 January 2022, ARERA set system charges to zero for all non-household consumers.

[17]In the Italian context, such policies included reductions of the following: a) system charges, that have both fixed and per-unit components; b) VAT, that applies to total bill, made by both fixed and per-unit components; c) excise duties, levied per unit of energy. Furthermore, the government granted tax credits calculated as a fraction of the energy component of the energy bill, again composed by both fixed and per-unit components.

This comes with four possible replies: (a) No (b) Yes, fixed-price contracts (c) Yes, financial derivatives (d) Yes, other instrument. Since most protected firms reported using fixed-price contracts, we combined all "Yes" responses into a dummy variable called I_i^{2021} (=1 if protected).[18]

Two things are worth noticing. First, contracts had to be in place *at the beginning of 2021*, prior to any anticipation of the crisis (see Section 2). Thus the question does not capture firms' endogenous responses to the crisis. Second – due to space constraints – the question conflates protection for electricity *and* gas, introducing potential measurement error. Using aggregate information from ARERA, we estimate that most firms had either fixed-price contracts for both energy carriers, or for none; this suggests that the extent of this misclassification is quantitatively not relevant.[19] Consistently, we find that this variable strongly predicts changes in average unit costs for both inputs between the first and second semester of 2021 (Alpino et al., 2023).

In the 2022 wave of the survey, the question was split by input type:

> "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)?"

The response options were: (a) Yes (b) No. As before, for a given input j = {electricity, natural gas}, we created a dummy called $I_i^{j,2022}$ taking value 1 if the firm i had any protection instruments for input j. Firms with protection were also asked:

> "If yes, how many months did this protection last in 2022?"

The open-ended responses were recorded as m_i^j for the respective input j.

[18] Results are virtually unchanged when excluding groups (c) and (d) from the analysis (10.2% of the electricity sample and 10.7% of the gas sample). We report them in Appendix I.

[19] In 2021, $P(E)$ = 45% of firms had fixed-price contracts for electricity and $P(G)$ = 44% had fixed-price contracts for gas. In our survey, $P(E \cup G)$ = 47% answered "yes" to E11 (fixed-price contract for at least one of the two). Basic set arithmetic implies that the vast majority of "yes" responses must reflect *joint* coverage, thus limiting the scope for measurement error in the treatment cohorts and suggesting that the treatment variable derived from question E11 is a reliable proxy for protection against both electricity and gas price shocks.Formally, using the inclusion-exclusion identity:

$$P(E \cup G) = P(E) + P(G) - P(E \cap G) = 0.45 + 0.44 - 0.47 = 0.42,$$

so 42% of firms are covered for *both* energy carriers. Given this result, it follows that only 5% of firms are estimated to be covered for only one carrier:

$$P(\text{only one}) = P(G \setminus E) + P(E \setminus G) = P(E) + P(G) - 2P(E \cap G) = 0.45 + 0.44 - 2 \cdot 0.42 = 0.05.$$

Restricting our attention to firms present in both waves[20], we constructed an input-specific treatment cohort variable E_i^j that identifies the semester h when a firm is first exposed to higher prices for input j (0 if it is never exposed in the observation window). The variable is constructed as follows:

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

All firms are untreated in $2021h1$ since the wholesale market crisis had not yet started. Firms with fixed-price contracts signed before the crisis remain untreated in later periods until their contracts expire.[21] We classify firms into four treatment cohorts based on E_i^j.

Firms with $I^{2021} = 0$ and $I^{j,2022} = 0$ are immediately exposed in 2021h2 ($E_i^j = 2021h2$; **early treated**). This cohort includes firms on variable-price contracts and those with fixed-price contracts expiring during 2021h2. In Section 5.b we confirm the absence of differential treatment effects for this group, mitigating concerns that its composition may affect the results.

Firms with $I^{2021} = 1$ and $I^{j,2022} = 0$ are exposed in 2022h1 ($E_i^j = 2022h1$) due to contract expiration at the end of 2021 (**mid treated**). By 2022, these firms face higher energy costs regardless of whether they purchase a new fixed-price or variable-price contract, as both became more expensive following the start of the crisis in late 2021.

Firms with $I^{2021} = 1$ and $m^j = 6$ are protected until mid-2022, when their fixed-price contracts expire (**late treated**). We assume these are the first 6 months of the year. In particular, we assume that at the beginning of 2021 the firm was holding a fixed-price contract signed at the beginning of 2021 that eventually expired in mid-2022. Thus, we set $E_i^j = 2022h2$ for these firms. We test the validity of this assumption in our event-study analysis by checking that price of input j increases in line with our assumed timing for this specific cohort (see Section

[20]Exit is essentially absent from our sample during the crisis period. Using the *Infocamere* administrative firm registry, which records the dates of all liquidatory procedures, voluntary liquidations, and definitive cessations for Italian firms, we checked every firm in our 2021 electricity and gas samples for exit-related events occurring during 2021-2022. We find zero liquidatory procedures, zero definitive cessations, one voluntary non-insolvency exit in both the electricity and gas sample, and three (one) non-liquidatory distress procedure openings in the electricity (gas) sample.

[21]Fixed-price contracts signed *after* the crisis began offer limited protection due to increased prices. Contracts signed early in 2022 may have provided savings over variable pricing signed at the same moment but a much higher cost than pre-crisis contracts (see Section 2).

5.b).

Firms with $E_i^j = 0$ are protected throughout 2021 and 2022 (**pure control group**). Note that the formulation of our question in 2022 does not ensure that the pure control group are protected under the exact same contract signed at the beginning of 2021, for which they declared to be protected in the first survey. In principle, they might have purchased a new fixed-price contract in January 2022, but in the results section we show that this is not consistent with our data. In particular the price of energy input j for the "mid treated" increases relative to the "pure control group" starting in 2022h1, ruling out the possibility that both cohorts had contracts expiring at the end of 2021.

The electricity and gas samples comprise 413 and 308 firms, respectively.[22] We excluded two subgroups of firms: (a) firms protected for only part of a semester in 2022 ($1 \leq m_i \leq 5$ or $7 \leq m_i \leq 11$), which ensures binary treatment status during protection periods and helps interpretability. This subgroup consists of 31 and 20 firms in the case of electricity and gas, respectively; (b) firms unprotected in 2021 ($I^{2021} = 0$) but protected in 2022 ($I^{2022} = 1$), as this protection may reflect endogenous responses to the shock. These are 14 and 32 firms in the electricity and gas sample, respectively.

In Appendix G we report balancing tables across different treatment cohorts. We highlight some significant level differences across groups. One concern with this could be that firms with higher energy intensity or in particular sectors would select into longer contracts – and this type of selection is related to potential underlying *trends* in outcomes. To avoid any concerns, in Appendix H we show that results are virtually unchanged when controlling for several covariate-specific non-parametric time trends. In Section 5.b we show that results are also robust when using the synthetic difference-in-differences estimator, which explicitly matches on pre-trends.

Extensions with confidential administrative data Depending on the structure of our additional data sources, we adapt the definition of the E_i^j variable. As for electricity and gas-intensive firms, we observe *monthly* electricity and natural gas consumption, respectively. We can thus define more fine-grained cohorts. We match our data to the Invind survey to obtain information on fixed-price contracts. In order to avoid having cohorts with very few firms, we collapse the data at the quarterly frequency and re-define cohorts of treatment as

[22] In the case of electricity, going from the early treated to the pure control group, the number of firms in each cohort is 176, 150, 11, 76. As for gas, the same figures are 128, 124, 7, 49. Summary statistics can be found in Appendix G.

follows:

$$E_i^j = \begin{cases} 2021q3, & \text{if } I^{2021} = 0 \text{ and } I^{j,2022} = 0 \\ 2022q1, & \text{if } I^{2021} = 1 \text{ and } 0 \leq m_i^j < 3 \\ 2022q2, & \text{if } I^{2021} = 1 \text{ and } 3 \leq m_i^j < 6 \\ 2022q3, & \text{if } I^{2021} = 1 \text{ and } 6 \leq m_i^j < 9 \\ 2022q4, & \text{if } I^{2021} = 1 \text{ and } 9 \leq m_i^j < 12 \\ 0, & \text{if } I^{2021} = 1 \text{ and } m_i^j = 12 \end{cases} \tag{2}$$

For the electricity-intensive case, the balanced matched sample includes 279 firms[23] over the period 2020-2022. For the gas-intensive case, our balanced matched sample includes 126 firms[24] in 2021 and 2022.

As for firms subject to the EU ETS, we have yearly information on fuel use that we match to the Invind survey. We aggregate treatment cohorts into a yearly variable for gas only, E_i^{gas}, again aligning the time dimension of treatment and outcome variables:

$$E_i^{gas} = \begin{cases} 2021, & \text{if } I^{2021} = 0 \text{ and } I^{2022,gas} = 0 \\ 2022, & \text{if } I^{2021} = 1 \text{ and } I^{2022,gas} = 0 \\ 0, & \text{if } I^{2021} = 1 \text{ and } I^{2022,gas} = 1 \text{ and } m_i^{gas} = 12 \end{cases} \tag{3}$$

The matched EU ETS-Invind sample includes 107 plants (66 firms).[25]

4 Empirical strategy

We estimate the effect of higher electricity and gas prices on firms' input demands using models of the type:

$$\log y_{ijt} = \alpha_i + \gamma_t + \sum_{k=0}^{3} \beta_{ijk} \cdot \mathbf{1}(t - E_i^j = k) + \epsilon_{ijt}, \tag{4}$$

where y_{ijt} is the quantity (q) or effective unit cost (p) of energy input j for firm i in semester t. Variable p is constructed from the ratio of expenditures and quantities, net of any government

[23]Going from the earliest treated to the latest treated, cohorts have the following number of firms: 106, 99, 5, 5 and 19

[24]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.

[25]The 2021 (2022) cohort is made of 53 (20) plants. The pure control group has 34 plants.

subsidies, as described in detail in Section 3. Firm (α_i) and calendar-time (γ_t) fixed effects control for time-invariant firm-specific characteristics and common time-varying shocks, respectively. Event-time dummies, $\mathbf{1}(t - E_i^j = k)$, capture time relative to the contract expiration date E_i^j (event time, henceforth).[26] We are interested in potentially heterogeneous treatment effects (β_{ijk}) of contract expiration on p and q.

We estimate separate reduced-form models for $\log(p)$ and $\log(q)$.[27] The elasticities we report are Wald-type ratios of the estimated average treatment effects on $\log(q)$ and $\log(p)$; they should be interpreted as the percent change in quantities per percent change in average unit costs induced by contract expiration.

We estimate separate models for electricity and natural gas to avoid multicollinearity stemming from the high correlation in expiration dates across inputs. This could be problematic if gas and electricity effects reinforce each other, leading to omitted variable bias. This concern arises in two cases: substitution between the inputs and scale effects. However, in this setting, neither issue is problematic. Substitution was minimal during the crisis, as both gas and electricity prices rose sharply. Regarding scale effects, as shown in Section 5, electricity demand does not respond to higher prices, ruling out electricity-induced scale effects on gas. Similarly, the lack of response in electricity demand rules out gas-induced scale effects on electricity.

We estimate each model on a balanced panel of firms over 2021-2022 using the Borusyak et al. (2024) (BJS) "imputation estimator", which identifies the ATT under standard parallel trends and no-anticipation assumptions. We cluster standard errors at the firm level to avoid known serial correlation issues (Bertrand et al., 2004). We use sample weights included in the Invind survey to ensure representativeness of our estimates. Pre-treatment coefficients are estimated in a *separate* regression that only uses untreated firm-year observations. Separate estimation of pre-treatment coefficients is a deliberate feature of the BJS imputation estimator. It serves three purposes: (i) it separates validation of the parallel trends assumption from estimation given that assumption; (ii) it improves efficiency, since all untreated observations contribute to the imputation stage; and (iii) it removes the correlation between treatment effect estimates and pre-trend coefficients that otherwise introduces bias when researchers condition on a pre-trend test passing (Roth, 2022).

Parallel trends assumption. Parallel trends are plausible because the expiration dates for fixed-price contracts are predetermined and not influenced by the crisis or firms' endogenous

[26] See Section 3.d on how we construct this variable.

[27] We avoid regressing quantities on prices directly, due to potential division bias (Borjas, 1980).

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 (Figure 2.1). Thus, we only exploit variation in when a pre-existing contract – one the firm already held at the beginning of 2021, before the price surge – expired. In this way, the expiration date and protection status in 2022 were determined at contract signature, which predates the crisis period we study. Firms cannot retroactively extend an already-signed contract once wholesale prices rise (see Section 2). Pre-trend tests, performed on a separate regression that uses only untreated observations as in Borusyak et al. (2024), consistently support this assumption. The resulting joint F-test is a diagnostic for the parallel trends assumption and is uninformative about the magnitude or significance of post-treatment coefficients.

One concern involves the "early-treated" cohort ($E_i^j = 2022h1$), which includes both firms with expiring fixed-price contracts and those with variable-price contracts. The latter set of firms may have had different characteristics, such as weaker preferences for price protection or different expectations about future demand. In Section 5.b we show that effects are similar across cohorts, limiting the validity of this concern. Additionally, early-treated firms never serve as controls for other cohorts.

Another concern is selection into longer contracts based on unobserved preferences for price certainty. Since firm-level fixed effects control for level differences, this only threatens validity if these preferences correlate with abrupt changes in outcomes specifically at contract expiration. This is unlikely, and the absence of pre-trends (see Section 5) supports this assumption.

No-anticipation assumption. We assess the plausibility of no-anticipation directly in Section 5, where we show that quantities do not systematically jump before the treatment date. Cohort-specific estimates reported in Section 5.b provide further reassurance: significant gas reductions emerge uniformly in 2022H2 across all cohorts, rather than accumulating earlier among firms with longer pre-expiration exposure to crisis conditions. If anticipation were driving results, later-treated cohorts would show smaller post-treatment responses having already adjusted; we observe the opposite pattern.

5 Results

Figure 5.1 reports our baseline results. Following Borusyak et al. (2024), 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 effective unit cost of electricity and gas. Correspondingly, in the two bottom panels the outcome is the log of physical quantities of electricity and gas purchased by the firm. In the text we describe the magnitude of the effects by commenting on $e^{\tau_k} - 1$.

Figure 5.1: The effect of the expiration of a fixed-price contract on average prices and quantities of energy inputs at the firm level.

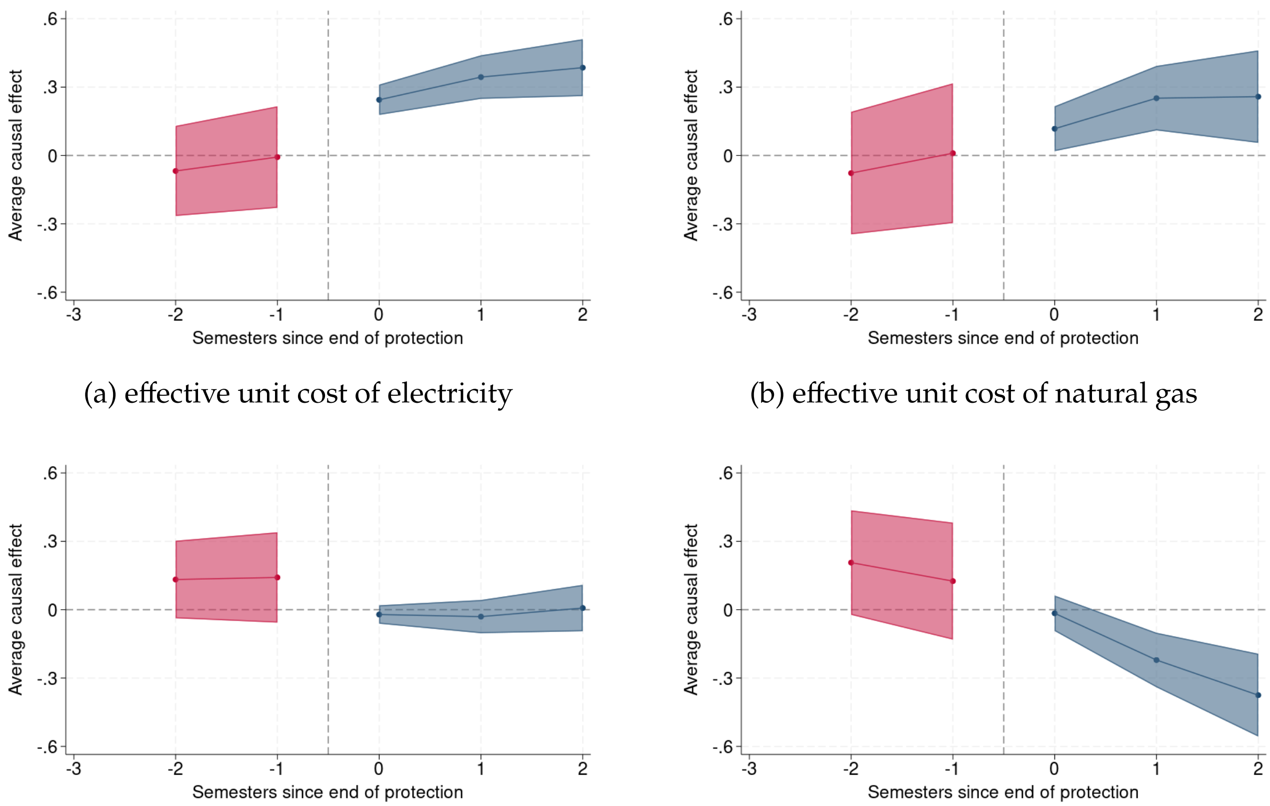

(a) effective unit cost of electricity

(b) effective unit cost 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 unit costs 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. (2024), as described in Section 4. All estimates are weighted by survey weights. Standard errors are clustered at the firm level. Confidence intervals are at the 95% level.

In panel (a) we observe a significant increase in the effective unit cost of electricity following the expiration of fixed-price contracts, with no evidence of pre-trending effects before the treatment occurs. A year and a half after expiration, the price increase is 47% relative to the baseline. All three post-expiration coefficients are precisely estimated and strongly significant: $\hat{\tau}_0 = 0.244$ ($p < 0.001$), $\hat{\tau}_1 = 0.342$ ($p < 0.001$), and $\hat{\tau}_2 = 0.384$ ($p < 0.001$), corresponding to price increases of 28%, 41%, and 47% in levels. The pre-trend joint test does not reject parallel trends (F= 2.02, $p = 0.13$). A similar pattern is seen in panel (b) for the effective unit cost of natural gas. One and a half years after expiration, the increase in gas price is 29% relative to the baseline. The post-expiration coefficients are $\hat{\tau}_0 = 0.116$ ($p = 0.021$), $\hat{\tau}_1 = 0.250$ ($p = 0.001$), and $\hat{\tau}_2 = 0.256$ ($p = 0.013$), corresponding to increases of 12%, 28%, and 29% in levels. The pre-trend joint test does not reject parallel trends (F= 0.94, $p = 0.39$).[28]

In panel (c) we analyze the (log-)quantity of electricity purchased. Despite the large price increase, this outcome does not respond to the treatment. Coefficients are positive, close to zero, and confidence intervals are tight.[29]

Our results differ from most of the existing studies: although the literature has found the demand for electricity to be quite inelastic in the short-term (Labandeira et al., 2017; Sikl et al., 2025), the response is typically not zero[30]. We offer two explanations. First, for non iso-elastic demand functions, quantity responses to a large price change – like those studied in our paper – do not necessarily correspond to responses to small price shock, like those studied in most previous work (e.g. Fontagné et al. (2024); von Graevenitz and Rottner (2022)). For example, if electricity demand has a minimum-throughput floor (machines that cannot be ramped below a threshold without irreversible damage), local elasticity shrinks toward zero as price rises and the floor becomes binding. Second, technical accounts suggest that energy-carrier substitution is especially capital-embodied when electricity is involved: while switching between combustible carriers can sometimes be accommodated within existing equipment through minimal retrofits, substituting electricity for combustible fuelsor vice versatypically changes the end-use conversion technology itself and requires process-specific replacement or integration of new equipment (Wei et al., 2019; Mallapragada et al., 2023; Leicher et al., 2024). Such capital-embodied adjustments are unlikely to be an important margin of adjustment over the short horizon studied in our paper, consistent with a broad investment literature showing that capital adjustment is costly, lumpy, and often subject to fixed/non-convex adjustment costs (Doms and Dunne, 1998; Caballero and Engel, 1999;

[28]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).

[29]This result also resolves the measurement question raised in Section 3. As documented in Appendix B, the fixed-charge component of Italian electricity tariffs implies a non-negligible wedge between our average per-unit cost of electricity and its per-unit marginal price. A zero quantity response renders this distinction immaterial: regardless of the size of the wedge, $\Delta Q = 0$ implies that the elasticity with respect to average unit cost of gas or electricity and the elasticity with respect to the marginal price are both zero.

[30]One exception is the recent study by Scarazzato (2025), which finds that end-user electricity demand – including industrial demand – is perfectly inelastic in the short run.

Cooper and Haltiwanger, 2006). Reducing electricity should then come at the cost of lower output, something for which we do not find evidence

In panel (d), we examine the (log-)quantity of gas and detect a different pattern. In the first treatment period, the point estimate is almost zero, but it turns negative in the second period and even more so in the third. The coefficient corresponding to the last period is $\hat{\tau}_2 = -0.372$ ($p < 0.001$), which implies that gas consumption decreases by 32% compared to the counterfactual with no price increase.

The joint pre-trend test for the quantity of gas leads us to reject the null of no pre-trends (F= 4.74, $p = 0.01$), driven by an upward pre-trend in gas consumption among units treated later. We address this concern in Section 5.b, where we use a synthetic difference-in-differences estimator that explicitly matches pre-treatment consumption paths. This yields quantitatively identical results. In Appendix H we further show that our results are robust to controlling for covariate-specific trends.

5.a Heterogeneity of gas effects across calendar periods

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. On one hand, earlier treated cohorts might have had more time to adjust. On the other hand, the crisis worsened over time after Russia invaded Ukraine, potentially leading to changes in salience of the shock.

We find negative effects on gas consumption to be entirely driven by 2022h2, with zero effects in earlier periods across cohorts (see Appendix I). This pattern is not due to differences in the magnitude of the energy shock across calendar periods for a constant elasticity, as the treatment effect on prices is relatively stable over time. Pooling across cohorts, the ATT on gas consumption in 2022h2 is -35%, with 95% confidence interval between -44% and -25%. We compute price-elasticities by scaling average quantity effects by average price effects (see Appendix N for details on how we compute the elasticity). For gas, it is equal to -1.1 in 2022h2, and very close to zero in previous semesters.

The increasing gas-demand elasticity, *irrespectively* of time-since-treatment, could reflect changing expectations about the energy crisis. Early on, when the crisis was expected to be short-lived, firms likely employed a "wait-and-see" strategy. In summer 2022, market pessimism peaked: spot prices hit record highs, futures markets projected a prolonged crisis and business confidence reached its trough (Figure 2.1 and Appendix Figure A.1). Fears of winter gas shortages and prolonged high prices in 2022h2 likely prompted many firms to act, consistent with "putty-putty" models where adjustment costs shape responses under uncertainty (Pindyck and Rotemberg, 1983; Atkeson and Kehoe, 1999).

One might worry that the 2022H2 response reflects macro shocks correlated with treatment timing rather than contract expiration per se. The covariate-specific trend robustness in Appendix H addresses this directly: results are unchanged when we allow sector, size, ETS, and energy-intensity groups to follow arbitrary time paths. We discuss this further in Section 5.b.

One might alternatively attribute the absence of quantity responses in earlier periods to firms anticipating government bailouts. The Italian government did implement support measures from 2021H2 onward, including system charge reductions, VAT cuts on gas, and eventually tax credits for energy-intensive firms , with their scale expanding substantially over time and peaking in 2022H2. However, this interpretation is at odds with our price results. As shown in Appendix I, contract expiration triggered a substantial increase in effective unit costs already in 2021H2, precisely because our price measure is net of all support received (as described in Section 3). If bailout anticipation were suppressing quantity responses, we would expect this to matter most in 2022H2, when the effective cost shock was largest and government support most generous, yet that is exactly the period when we *do* observe a large gas demand reduction. The timing is therefore more consistent with firms responding to realized out-of-pocket costs and to changing expectations about crisis duration than with bailout anticipation.

5.b Robustness: synthetic diff-in-diff

To further validate our findings, we also implement a synthetic diff-in-diff (SDID) exercise (Arkhangelsky et al., 2021), following the procedure in Clarke et al. (2024) for staggered treatment and using bootstrap for inference[31]. This method explicitly matches on pre-trends and addresses concerns related to pre-trend testing (Roth, 2022).

Figure 5.2 reports event-study graphs for average unit costs of electricity and natural gas using the synthetic diff-in-diff methodology, and Figure 5.3 reports the corresponding results for log quantities. We present the results for the three treatment cohorts separately, with the "donor pool" always drawn from the pure control group. Naturally, the pre-trend matching can use more pre-treatment periods for the "late treated" than for the "mid treated" and "early treated", as we only have four periods at our disposal.

Across all cohorts, panels (a)–(f) of Figure 5.2 show that the retail price increases *exactly* at the time of expiration of the respective fixed-price contract. This reassures us that our coding

[31] The SDID estimator does not support survey weights.

of the cohort-of-treatment variable is correct. Moreover, the ATT estimates in Figure 5.3 are qualitatively and quantitatively similar to our baseline results, which shows robustness to a relaxation of the parallel trend assumption. As in the staggered diff-in-diff, all three cohorts display very similar treatment effects, confirming results are not driven by a specific cohort.

Finally, Figure 5.3 corroborates that the negative treatment effect on the quantity of gas is driven only by what happens in 2022h2, while the time elapsed since contract expiration does not matter, as all treated cohorts in a given calendar time period have similar coefficients.

In panel (f) of Figure 5.3 we observe a statistically significant decrease in the quantity of electricity for the late treated. This cohort is composed of a smaller number of firms, which is why this does not show up in the average effect that we estimated with the Borusyak et al. (2024) estimator.

Figure 5.2: Synthetic diff-in-diff estimates of the effect of the expiration of a fixed-price contract on log net unit cost of electricity (left column) and natural gas (right column), by treatment cohort.

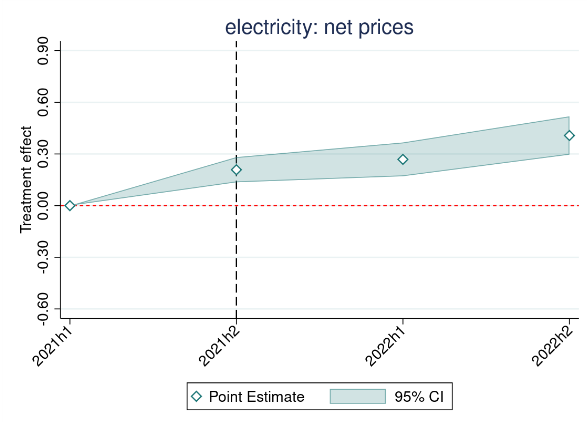

(a) Electricity; early treated

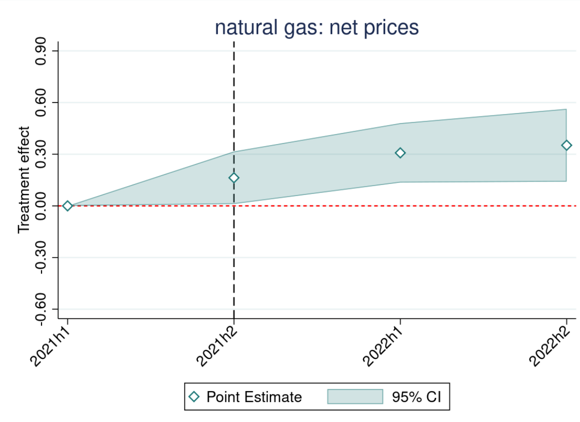

(b) Natural gas; early treated

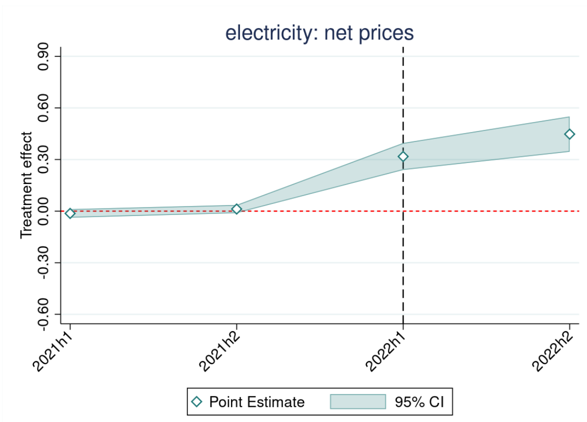

(c) Electricity; mid treated

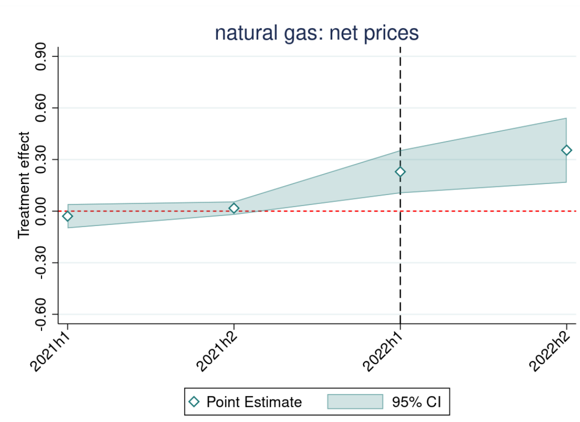

(d) Natural gas; mid treated

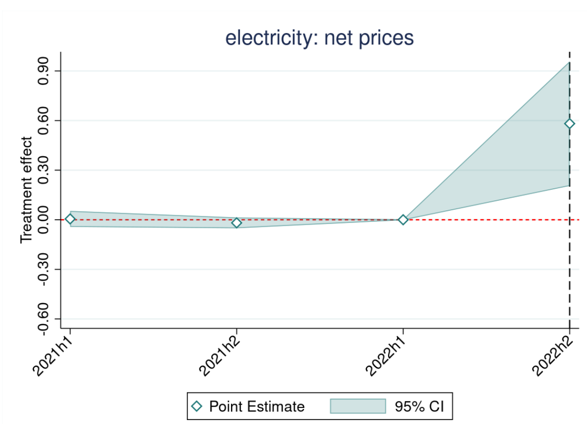

(e) Electricity; late treated

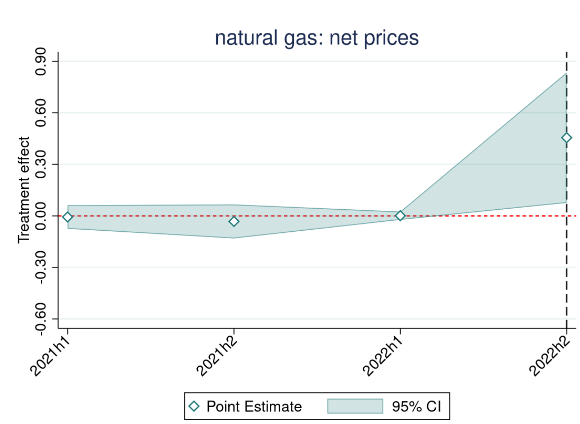

(f) Natural gas; late treated

Note: The figure shows average causal effects of the expiration of a fixed-price contract on log net unit costs, estimated with the synthetic difference-in-difference method in the staggered case (Arkhangelsky et al., 2021; Clarke et al., 2024). The SDID estimator does not support survey weights. Bootstrapped confidence intervals are at the 95% level.

Figure 5.3: Synthetic diff-in-diff estimates of the effect of the expiration of a fixed-price contract on log quantities of electricity (left column) and natural gas (right column), by treatment cohort.

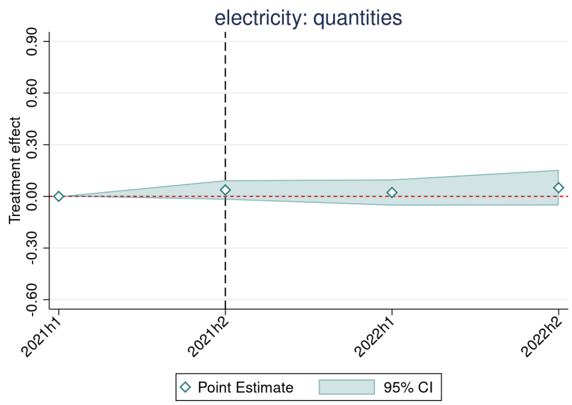

(a) Electricity; early treated

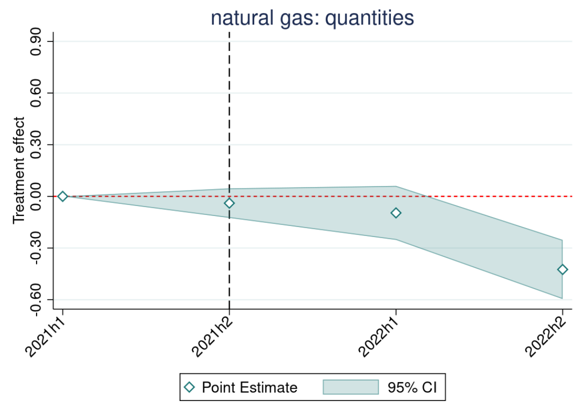

(b) Natural gas; early treated

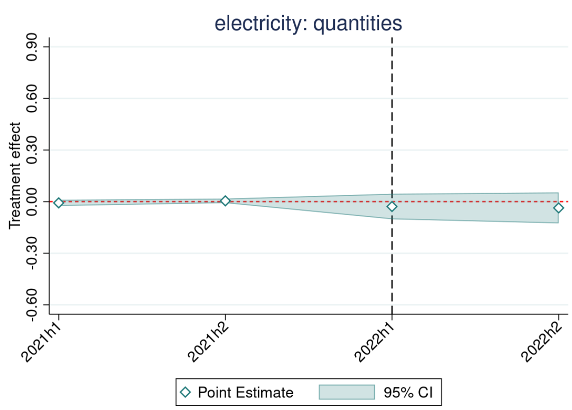

(c) Electricity; mid treated

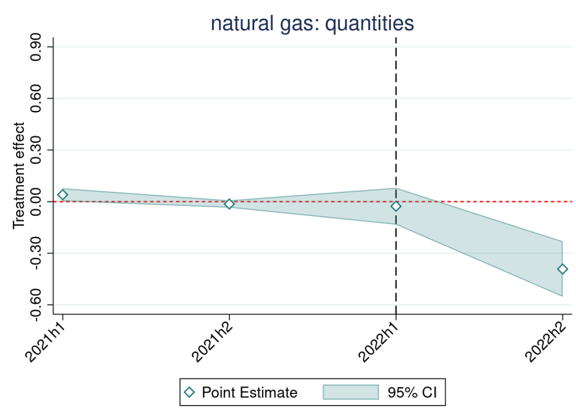

(d) Natural gas; mid treated

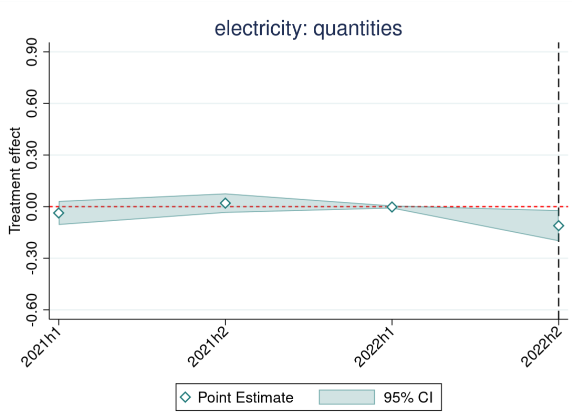

(e) Electricity; late treated

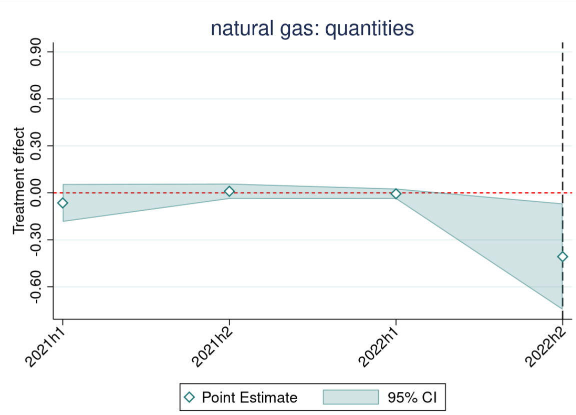

(f) Natural gas; late treated

Note: The figure shows average causal effects of the expiration of a fixed-price contract on log purchased quantities, estimated with the synthetic difference-in-difference method in the staggered case (Arkhangelsky et al., 2021; Clarke et al., 2024). Bootstrapped confidence intervals are at the 95% level.

5.c Other robustness tests

Confounders due to macro shocks. The second half of 2022 coincides with major macroeconomic and policy shocks: Russias supply cutoff, elevated EU-ETS carbon prices, the Nord Stream sabotage, and EU demand-reduction mandates (Regulation 2022/1369). These could affect gas demand through non-price channels and identification would be compromised if responsiveness to these shocks were correlated with treatment timing. In order to mitigate these concerns, Figure H.1 in Appendix H, shows that results are robust to the inclusion of sector, size, energy-intensity, and ETS-specific non-parametric *trends*.

Confounders due to seasonality. Finally, given that we observe few periods across two years, one may wonder whether results are driven by seasonality. We address this by re-estimating the baseline specification with firm × semester fixed effects, which absorb any *firm-specific* seasonal pattern (e.g. a firm that systematically uses more energy in winter). The τ_0 and τ_1 coefficients are virtually unchanged in magnitude and precision relative to the baseline (Appendix Figure H.2). One qualification applies: because the panel spans only four semesters (2021H1-2022H2), early-treated firms (E_i^j = 2021H2) are treated in every second-semester period in the sample and therefore have no untreated H2 observation from which the BJS estimator can impute their firm × H2 fixed effect. These firms are consequently excluded at H2 horizons (τ_0 and τ_2), though they contribute normally to τ_1. In addition, τ_2 is suppressed for all cohorts as it would require a 2023H1 observation outside the window. The stability of the remaining estimates across both specifications confirms that seasonal patterns do not drive our results.

Robustness to alternative estimators. Appendix I confirms robustness using alternative diff-in-diff estimators, including De Chaisemartin and d'Haultfoeuille (2020), Callaway and SantAnna (2021), Sun and Abraham (2021), and OLS.

5.d Heterogeneity in gas demand response across firms in 2022h2

The ATT in 2022h2 for gas consumption might hide substantial heterogeneity, due to firm characteristics, sectoral specificities, and particularly the large variation in energy intensity across firms. Energy-intensive firms were the earliest and largest government aid recipients during the 2021-22 crisis (McWilliams et al., 2024), but whether their demand is more or less elastic than other companies is theoretically ambiguous, as it depends on several parameters, including the elasticity of substitution between inputs (Cahuc et al., 2014).[32]

[32]We provide formulae in Appendix L.

To investigate treatment effect heterogeneity in a credible way, we turn to machine learning (ML) techniques, which have become popular for this aim in causal analysis in the context of randomized control trials (Haaland and Roth, 2020; Allcott et al., 2020; Alpino et al., 2022). We are among the first to apply these tools in the context of staggered difference-in-differences (Hatamyar et al., 2023). In our view, the estimator by Borusyak et al. (2024) 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* to find its best predictors, in the spirit of Athey and Imbens (2016) and Wager and Athey (2018).[33] The advantages of this approach relative to a more traditional heterogeneity analysis are twofold: first, the quest for heterogeneity is more efficient and also explores non-linear combinations of the available covariates; second, the procedure is less prone to bias arising from multiple hypotheses testing.

The random forest algorithm highlights some key predictors of heterogeneity: a self-reported dummy variable indicating whether gas is an essential input in production;[34] a dummy for being a gas-intensive firm according to the Italian legislation; and the EU ETS dummy. Interestingly, the ML algorithm does not detect a large amount of between-industry heterogeneity, with only the food sector dummy being selected as a predictor. Firms in this industry overwhelmingly declare gas to be an essential input.

Next, we compute heterogeneous treatment effects with our baseline Borusyak et al. (2024) estimator in the sub-samples selected by the ML algorithm, computed as subsample averages of the individual-level BJS treatment effects; no machine learning enters this estimation step. In the second half of 2022, the ATT is equal to -41% among firms for which gas is *not* considered essential, and to -28% among those for which it *is* essential. Both estimates are significant at the 99% level and their difference is significant at the 10% level. The associated price elasticities of demand are -2.5 and -0.5, respectively, which suggests that the gas-essential dummy can be considered as proxy for low elasticity of substitution, as the formulation of the survey question suggests (see Appendix N for details on how we compute elasticities).

Interestingly, we find the ATT, and consequently the demand elasticity, to be near zero among gas-intensive firms. Since all gas-intensive firms classify gas as an essential input, their lower demand elasticity appears directly linked to the reduced substitutability of gas

[33] We provide more details and extensive results in Appendix M.

[34] The dummy variable is derived from the Bank of Italys *Business Outlook Survey of Industrial and Service Firms* (Sondtel) conducted in fall 2022. The survey posed the question: "At the beginning of 2022, was gas an essential input for your firms manufacturing process?". In the survey instructions, essential inputs are defined as those whose absence -given the existing plants and machinery- would make production impossible in the short term.

with other inputs. The implied elasticities are $\varepsilon_D = 0.03$ for gas-intensive firms and $\varepsilon_D = 1.3$ for other firms; we use these two values in the policy analysis of Section 7, as the gas-intensive classification is the administrative tag actually employed by policymakers for subsidy targeting.[35] Note that the heterogeneity unveiled in this section cannot be attributed to sectoral dynamics, or else the ML algorithm would have selected several sector dummies as key predictors. Furthermore, gas-intensive firms are spread throughout all industries except water and waste management.

5.e Energy intensive firms

Previous analyses indicate that energy-intensive firms exhibit inelastic demand for both gas and electricity. Given the policy significance of this group, we conduct a complementary analysis using confidential administrative data on their energy consumption, which is exclusively available for these firms. Figure 5.4 presents event-study graphs for the evolution of gas and electricity consumption around the expiration of fixed-price contracts, at a quarterly frequency. For electricity-intensive firms, the impact of losing fixed-price protection on the (log-)quantity of electricity consumed remains negligible across all treated periods. This finding closely mirrors the patterns observed for the average industrial firm, as detailed in the first part of Section 5.

In the case of gas-intensive firms, the effects on gas consumption are minimal during the first four quarters but become negative - albeit not statistically significant at conventional levels - in the final two quarters. When the estimates are aggregated by calendar time, a noticeable negative impact on gas consumption emerges only toward the end of 2022, consistently with the baseline results. Aggregating by semester, the results indicate an ATT of -6% in the second half of 2022, not statistically different from zero at conventional levels. This stands in stark contrast to the -35% ATT observed for the average industrial firm in the baseline analysis.

[35]Throughout the paper, our price outcome is the effective unit cost net of government support, as defined above. All conclusions carry through if one uses gross-of-subsidy unit costs instead: the implied gas demand elasticity for the average firm changes from -1.1 to -1.0, for non-gas-intensive firms from -1.3 to -1.2, and for gas-intensive firms from -0.03 to -0.02.

Figure 5.4: Event-study estimates for the (log-)quantity of energy among energy intensive firms

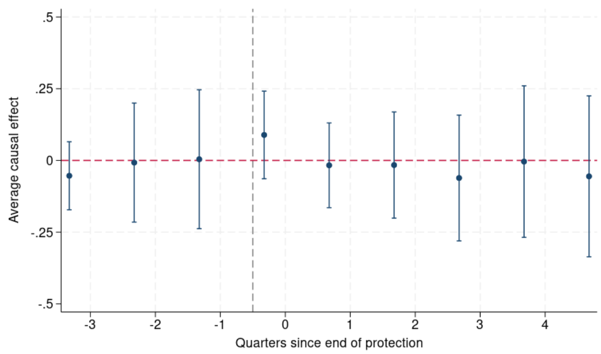

(a) Electricity

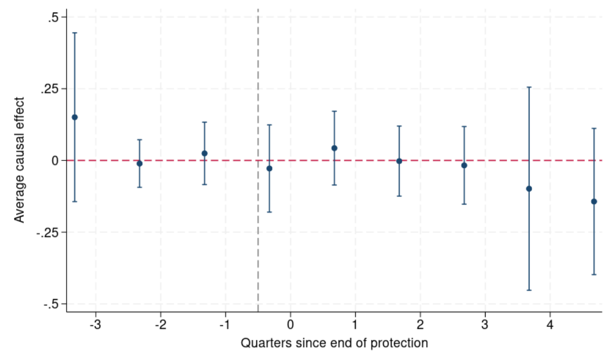

(b) Natural gas

Note: The figures show average causal effects of the expiration of a fixed-price contract on the (log-)quantity of electricity for electricity intensive firms (panel a) and on the (log-)quantity of natural gas for gas intensive firms (panel b). The outcomes are measured with administrative data from CSEA. Average causal effects before and after the treatment are estimated in two separate regressions, using the "imputation" estimator by Borusyak et al. (2024). Standard errors are clustered at the firm level. Confidence intervals are at the 95% level.

Pass-through to final good prices The limited quantity response just documented for energy-intensive firms does not imply that they bore the entire cost shock internally. A plausible alternative adjustment channel is output-price pass-through. Using complementary information from the Invind survey, we report cohort averages of year-on-year changes in selling prices in Appendix Figure A.4. Two things are noticeable in the descriptive evidence. First, all cohorts raise prices markedly and almost synchronously in 2021 and even more in 2022, in line with the general inflationary episode. This is consistent with recent 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[36] (Duprez and Magerman, 2018; Amiti et al., 2019; Muehlegger and Sweeney, 2022). Second, energy-intensive firms increased their prices substantially more than other industrial firms, indicating that price setting constituted an important margin of adjustment for them. By increasing output prices at rates higher than other firms, they were able to accommodate higher energy costs, thereby reconciling their near-zero demand elasticity with continued financial viability during the crisis.

[36]The near-synchronous price movement across cohorts documented above also explains why a causal event-study design analogous to equation (4) would be uninformative for estimating the pass-through.

5.f Input substitution

In this section, we explore whether the decline in gas consumption during 2022h2 was partly offset by increased reliance on other fossil fuels. The analysis is restricted to firms subject to the EU ETS for which we observe, from administrative sources, energy consumption across all fossil fuels[37].

While natural gas is typically the primary fossil fuel consumed by EU ETS plants – serving as an essential input for 90% of these firms – these plants also utilize solid fuels (e.g., coal) to a significant extent. The use of liquid (e.g., kerosene) or gas fuels (e.g., LPG) is much less common.

Figure 5.5: Input substitution test among EU ETS plants

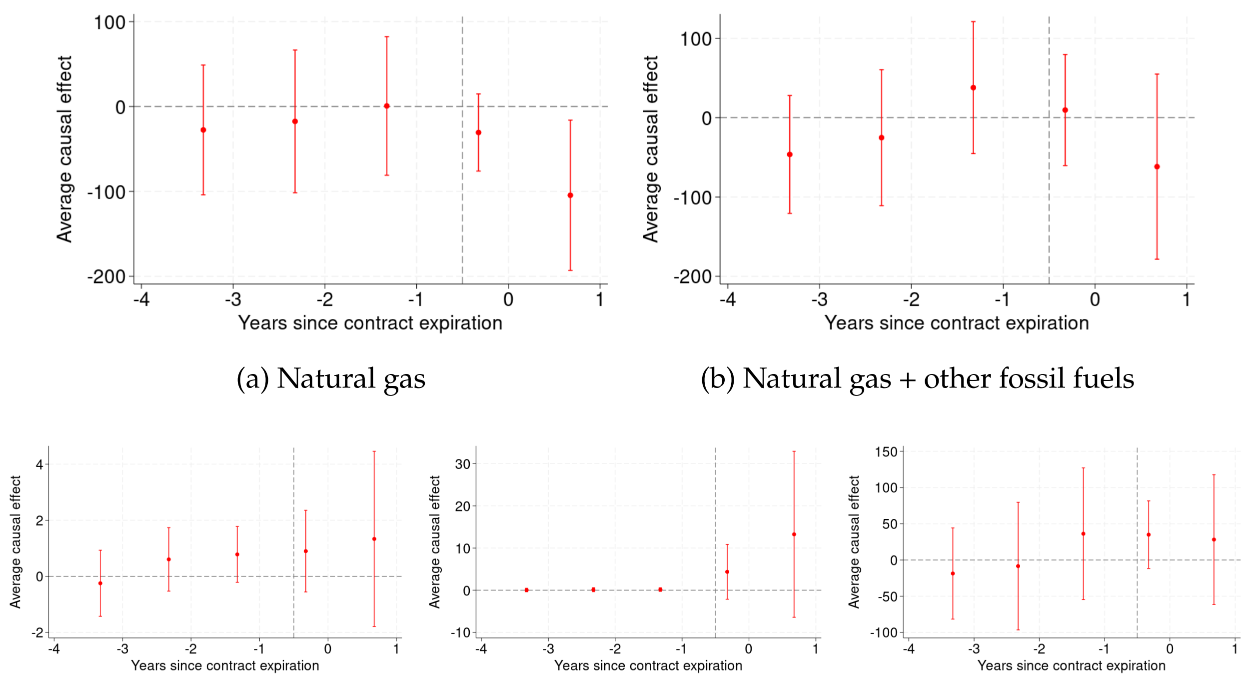

(a) Natural gas

(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 expiration of a fixed-price contract on different outcomes (in Terajoules, TJ) 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. (2024). All estimates are weighted by survey weights. Standard errors are clustered at the firm level. Confidence intervals are at the 95% level.

[37]We do not observe electricity consumption. In principle, firms with unexpired electricity contracts facing lower electricity tariffs could shift some production load from gas-fired to electric processes. However, the majority of firms in Italy were either protected against both electricity and gas, or against neither (see Section 3.d). Furthermore, the tight link between Italian electricity and gas wholesale prices excludes any cost advantage of electrification, unless self-generation from renewables is employed.

Panel (a) of Figure 5.5 illustrates a negative effect on gas consumption over time, with no evidence of a pre-trend (p-value: 0.47). The ATT estimates are -26 terajoules (Tj) in 2021 and -89 Tj in 2022, corresponding to declines of 4% and 14%, respectively, relative to the 2018-2020 average. We next examine total fossil energy consumption, defined as the combined use of gas and other fuels, including biofuels. If firms completely substituted gas with other fuels, the ATT would equal zero, while no substitution would produce effects equivalent to those for gas alone. Panel (b) shows that point estimates fall between these two extremes, with ATT estimates of +27 Tj in 2021 and -56 Tj in 2022. The event-study profile is similar but attenuated, with wider confidence intervals. These findings indicate incomplete substitution, with treated firms reducing total fossil energy consumption more than control firms.

Panels (c)-(e) of Figure 5.5 display changes in the consumption of fuels (including biofuels) other than natural gas, classified by physical state: gaseous, liquid, and solid fuels. Results show slight increases in the use of all other fuels among treated firms, but substitution remains incomplete. For gaseous fuels, the increase is negligible. For liquid fuels, the effects are modest (6 Tj in 2021 and 10 Tj in 2022) and marginally insignificant at the 10% level. Solid fuels show a larger increase, but the effect arises earlier and is more pronounced in 2021 than in 2022, raising doubts about the validity of identifying assumptions for this particular outcome (pre-trend test p-value: 0.02). While some substitution from gas to other fossil fuels appears to have occurred, it is incomplete and does not fully explain the observed reduction in gas consumption among EU ETS plants. This is consistent with gas-intensive firms declaring that gas is an essential input.

Additional descriptive evidence on substitution of gas The Invind survey included a direct qualitative question on gas substitution in 2022. Only 6% of firms in our sample declared to have substituted gas with other inputs (fossil fuels, electricity, or other), and the share was lower among firms for which gas was essential (5.5%), EU ETS firms (2%), and gas-intensive firms (3%). Electricity was the preferred option on average, but less so for these three groups (see Appendix Figure A.3).

6 Discussion of the empirical results

Our estimates point to substantial short-run rigidity in industrial energy demand, with important differences across energy carriers, firm types, and time horizons. In this section, we review the main empirical findings and discuss their interpretation.

Energy demand less elastic than in the literature Most existing micro-level estimates are based on smaller, normal-time price variation (Fontagné et al., 2023; von Graevenitz and Rottner, 2022). Our design, by contrast, identifies an arc-elasticity over the crisis-relevant price range. When energy demand is not isoelastic, local elasticities estimated around normal prices need not extrapolate to very large shocks, particularly once firms approach technological constraints or operational minima.

Inelastic demand for electricity Our results suggest that electricity demand is even more rigid than previously documented (Labandeira et al., 2017; Sikl et al., 2025). Technical accounts emphasize that adjusting electricity use often requires changing capital goods, rather than simply switching energy inputs for given capital (Wei et al., 2019; Mallapragada et al., 2023; Leicher et al., 2024); over our short horizon, such adjustment is limited by costly and lumpy capital investment (Doms and Dunne, 1998; Caballero and Engel, 1999; Cooper and Haltiwanger, 2006).

Delayed gas response Gas demand adjusts to the shock only in 2022H2, and the effect is common to all cohorts regardless of when they lose protection. Our interpretation is that firms initially followed a "wait-and-see" strategy, as financial markets expected the crisis would be short-lived. As expectations turned very pessimistic in 2022H2, many firms likely decided to act, consistent with "putty-putty" models where adjustment costs shape responses under uncertainty (Pindyck and Rotemberg, 1983; Atkeson and Kehoe, 1999).

Energy-intensive firms have less elastic energy demand One might expect energy-intensive firms to display *higher* demand elasticity: by definition they have a larger energy cost share, and the resulting scale effect should raise the own-price elasticity of energy demand (Cahuc et al., 2014). Our results run counter to this prior: gas-intensive firms exhibit near-zero elasticity, while non-gas-intensive firms are substantially more responsive. The difference must therefore be driven by lower substitutability of gas in the energy-intensive production processes. Consistent with this interpretation, most gas-intensive firms declare gas to be an essential input, and additional evidence from EU ETS plants shows that lower gas use is only partly offset by switching to other fuels.

How did firms bear the cost of the shock? The very limited adjustment in terms of energy demand suggests that, on average, firms did not reduce output substantially in 2021-22, and therefore absorbed the cost shock. This raises the question of how that cost was ultimately borne.

Year 2021 was characterized by a strong post-pandemic demand recovery, and the Italian corporate sector entered the energy crisis with sizable liquidity buffers and good financial conditions (Banca d'Italia, 2021). In this macroeconomic environment, it is plausible that many firms chose to absorb the shock through nominal and financial margins, by raising output prices, compressing profit margins, or both.

Figure 6.1: Firm strategies in response to energy price increases in 2022

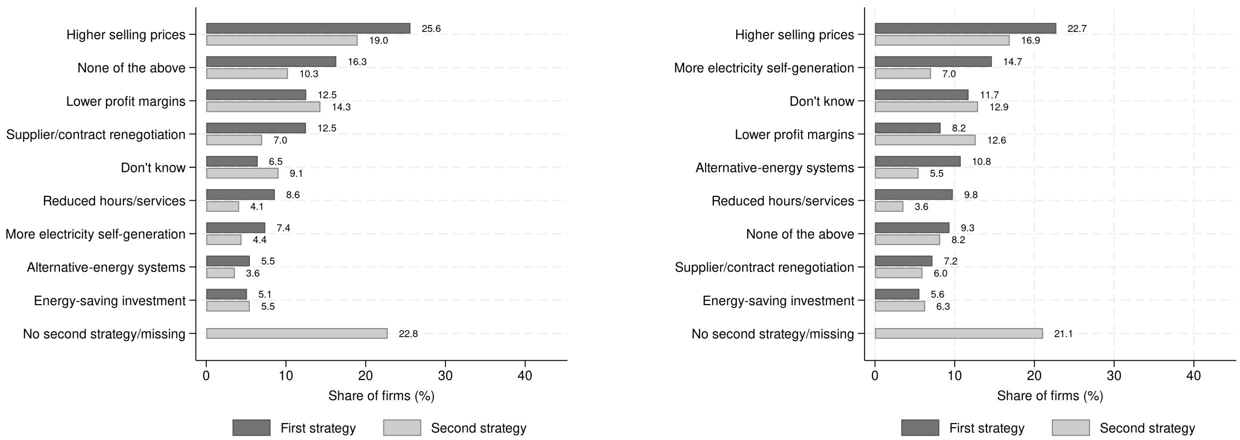

(a) Adopted in the first nine months of 2022 (b) Planned for the following six months

Notes: Bars report weighted percentages of firms using survey weights. In both panels, first and second strategies are computed on the same analysis population: firms with a valid first-strategy response. The residual category for the second strategy includes firms reporting no second strategy, missing second strategy, or a value outside the valid response range.

Our evidence is consistent with this interpretation. As noted in Section 5.e, firms in our sample increased output prices significantly in 2021 (8%) and even more in 2022 (12%) (see Appendix Figure A.4).[38] Qualitative questions in the 2022 *Business Outlook Survey of Industrial and Service Firms* (Sondtel) reinforce this picture: in the first nine months of 2022, by far the most common strategy reported by manufacturing firms in response to higher energy costs was to raise selling prices, followed by a reduction in profit margins (panel a of Figure 6.1a). By contrast, only a relatively small share reported reducing plant operating time or curtailing production, even partially[39]. Investment-based responses, such as installing energy-saving

[38]The price adjustment is broadly the same across treatment cohorts, consistent with ample evidence on strategic complementarities in price updating (Duprez and Magerman, 2018; Amiti et al., 2019; Muehlegger and Sweeney, 2022).

[39]As reported in the *Sondtel* survey, a small share of firms decided to change energy suppliers and/or renegotiate contracts in the first nine months of 2022. Note that this strategy can cushion only partially, if at all, from the shock, as prices in the retail market rose substantially already in fall 2021 (see Section 2 for more details). Furthermore, note that this has not implications for the validity of our empirical design, because we classify firms as protected only if they hold a fixed-price contracts signed at the *beginning* of 2021, i.e. before the crisis (see Section 3).

equipment, expanding self-generation, or adapting plants to alternative energy sources, were also uncommon at that stage, though they became notably more frequent when firms were asked about their plans for the coming months (panel b). This pattern suggests that technological adjustment was feasible but slow.

7 Policy implications for the design of support measures

We now turn to the implications of our results for the design of support measures for firms during energy crises. While EU governments widely implemented such policies in 2021-22 (Sgaravatti et al., 2021; McWilliams et al., 2024), some economists warned that these interventions could weaken incentives to conserve gas at a time of scarcity, thus exacerbating the crisis (Gros, 2022; Signorini, 2022). We present a simple incidence framework that relates policy interventions to equilibrium prices, quantities and fiscal costs. We keep the framework intentionally simple so that each object maps directly to what we can estimate in the data, similarly to Deryugina et al. (2020) and Hahn and Metcalfe (2021). The model captures heterogeneity in elasticity across gas-intensive and non gas-intensive firms, as well as their different share in overall baseline gas consumption, thus allowing us to study targeted policies.

7.a A simple incidence framework

The supply side is characterized by an upward-sloping isoelastic short-run supply function $Q^S(P) = BP^\eta$, which responds to a producer price P, and depends on the price elasticity of supply $\eta > 0$.[40] On the demand side, we assume final users can be of two types (indexed by g): gas-intensive and non-gas-intensive firms $g \in \{\mathrm{GI}, \mathrm{NGI}\}$. Each group is characterized by a downward-sloping isoelastic short-run demand function, which responds to a type-specific consumer price p_g, so that $Q^D(p_g) = A_g p_g^{-\varepsilon_g}$, where $\varepsilon_g > 0$ is group g's price elasticity of demand.

The producer price P and group-specific consumer prices p_g are linked by the following reduced-form rule:

$$p_g \;=\; \alpha_g + P, \qquad \alpha_g \in \mathbb{R}. \tag{5}$$

Here α_g captures g-specific pre-existing additive wedges between P and p_g. These allow us

[40] This reflects the fact that expanding gas output or imports requires drawing on increasingly costly sources, so that higher prices P are needed to elicit greater quantities.

to match the well-known empirical fact that large energy-intensive gas consumers pay lower energy prices (Davis et al., 2013).[41] Market clearing requires that

$$Q^S(P) = \sum_{g \in \{\text{GI,NGI}\}} Q_g^D(p_g). \tag{6}$$

We now consider the introduction of two alternative policies, indexed by the group g for which they apply to: (i) a per-unit subsidy $s_g > 0$ to gas consumption, and (ii) a per-unit reward $r_g > 0$ for reducing gas consumption relative to a historical baseline ("pay-for-reduction").[42] Setting $s_g = s \ \forall \ g$ and $r_g = r \ \forall \ g$ lets us study the case where the same policy is made available for all firms.

The per-unit subsidy and the reward both enter equation 5 additively, but with opposite sign. Concretely, equation 5 becomes

$$p_g = \alpha_g + P - s_g + r_g.$$

While the subsidy decreases the price, the reward acts like a per-unit tax on gas, since consuming one additional unit of gas costs the firm p_g, plus the foregone reward r_g. In this way, the equilibrium effects of introducing subsidies or rewards are the mirror images of one another.

We start by considering the introduction of a constant per-unit subsidy to all firms. Implicit differentiation of equation (6) allows us to derive equilibrium effects on prices and quantities:

$$\frac{dP}{ds} = \frac{\overbrace{\sum_g \varepsilon_g \frac{Q_g}{p_g}}^{\text{direct demand effect}}}{\underbrace{\frac{\eta Q}{P}}_{\text{equilibrium supply effect}} + \underbrace{\sum_g \varepsilon_g \frac{Q_g}{p_g}}_{\text{equilibrium demand effect}}}; \qquad \frac{dQ}{ds} = \frac{\eta Q}{P} \cdot \frac{dP}{ds}, \tag{7}$$

which is a standard application of the competitive incidence formula to a market with

[41] For example, energy-intensive firms are often granted discounts on levies and taxes by governments under EU State Aid rules (Gerster and Lamp, 2024).

[42] In Appendix O, we also consider two more policies implemented in some EU countries: (iii) lump-sum transfers; and (iv) retail price caps. Equilibrium responses under (iii) are obviously zero, while under (iv) are qualitatively similar to those generated by subsidies, but they are larger.

heterogeneous demand groups.[43] A key result is that the relevant parameter for evaluating this generalized per-unit subsidy is not the simple average demand elasticity across firms, but a Q_g/p_g-weighted elasticity. When prices are similar across groups, as they are in our calibration, this is well-approximated by a quantity-weighted elasticity. The relevance of quantity-weighted elasticities has been noted in other policy contexts (e.g., Chetty (2009) in income taxation), but it acquires first-order quantitative importance for energy subsidies, because consumption is highly concentrated among a small number of firms with relatively inelastic demand.

For a subsidy targeted to a generic group g, we have:

$$\frac{dP}{ds_g} = \frac{\varepsilon_g \dfrac{Q_g}{p_g}}{\dfrac{\eta Q}{P} + \sum_t \varepsilon_t \dfrac{Q_t}{p_t}}; \qquad \frac{dQ}{ds_g} = \frac{\eta Q}{P} \cdot \frac{dP}{ds_g}. \tag{8}$$

In this case the direct effect is given by the response of group g only. The equilibrium effects at the denominator are the same as in equation 7, as both groups now pay higher prices because of the subsidy. As for the per-unit reward, we notice that first-order effects on the equilibrium are identical to the case of the subsidies, but with opposite sign. All derivations are identical. The framework can also accommodate the introduction of the two policy instruments at the same time. Totally differentiating market clearing around the no-policy equilibrium yields

$$\frac{\eta Q}{P}\, dP = \sum_g \varepsilon_g \frac{Q_g}{p_g} (ds_g - dr_g - dP).$$

Defining $M_g \equiv \varepsilon_g Q_g/p_g$, the producer-price response is

$$dP = \frac{\sum_g M_g(ds_g - dr_g)}{\frac{\eta Q}{P} + \sum_g M_g}, \tag{9}$$

and aggregate quantity changes by

$$dQ = \frac{\eta Q}{P}\, dP. \tag{10}$$

[43]The direct demand effect – via ds, keeping P fixed – at the numerator is weighted average of group-specific elasticities, with weights depending on group-level consumed quantity and retail price. The terms at the denominator rescale this effect by equilibrium changes in supply and demand occurring via P. In case of homogeneous demand elasticities and no price wedges the formula nests the more familiar conditions $dP/ds = \epsilon_D/(\epsilon_S + \epsilon_D) \leq 1$ and $dQ/ds = \epsilon_S\, \epsilon_D/(\epsilon_S + \epsilon_D)$.

Equations 9 and 10 nest the uniform subsidy, targeted subsidy, and pay-for-reduction cases considered above. They also imply that, to first order, the effects of multiple instruments are additive, with subsidies and rewards entering with opposite signs.

Fiscal costs of subsidies depend on both the mechanical cost of subsidizing baseline consumption and the behavioral response induced by the subsidy. For a subsidy targeted to group g, the fiscal cost is

$$C(s_g) = \underbrace{s_g \, Q_{0,g}}_{\text{mechanical cost}} + \underbrace{s_g^2 \frac{dQ_g}{ds_g}}_{\text{behavioral fiscal externality}}, \tag{11}$$

where $Q_{0,g}$ denotes baseline consumption for the eligible group g. The mechanical term is large when the targeted group consumes a large amount of gas. The behavioral term captures the induced change in consumption by the eligible group.

The government's fiscal outlay differs from the subsidy case. Under a subsidy, the government pays s per unit consumed; under the reward, it pays r per unit *saved*. Denoting the pre-policy quantity of the targeted group by $Q_{0,g}$, the fiscal cost of introducing a small reward $dr_g = r_g - 0$ is:

$$C(r_g) = r_g \cdot \left(Q_{0,g} - Q_g(r_g)\right) \approx r_g^2 \cdot \left(-\frac{dQ_g}{dr_g}\right), \tag{12}$$

where the approximation follows from a first-order Taylor expansion. Two features are worth noting. First, the fiscal base is the quantity *saved*, not the quantity *consumed*, so fiscal outlays are proportional to the behavioral response and are zero if demand is perfectly inelastic. This is the mirror image of the per-unit subsidy, where mechanical outlays (on baseline consumption) dominate when demand is inelastic. Second, because the reward lowers rather than raises the wholesale price, untargeted groups benefit indirectly through lower input costs.

For a policy mix combining subsidies and rewards, the fiscal cost is computed from the group-specific quantity changes induced by the simultaneous policy, rather than by adding the fiscal costs of the two policies considered separately. While equilibrium price and aggregate quantity effects are additive to first order, behavioral fiscal costs are not: each instrument changes the common producer price and therefore affects the fiscal base of the other instrument. Below we show that these interactions are quantitatively small in our calibration.

7.b Policy simulations

We use the model to study the introduction of targeted vs. untargeted gas policies, in a context where gas-intensive firms have relatively inelastic demand, but account for the majority of industrial gas consumption. Because our empirical analysis was conducted using Italian data, our calibration employs data from this country on gas consumption and prices. Our results do not hinge on the specific data employed, since the analysis is centered on relative comparisons between targeted and untargeted policies. We use the demand elasticities implied by our estimates for both gas-intensive ($\varepsilon_{\text{GI}} = 0.03$) and not gas-intensive ($\varepsilon_{\text{NGI}} = 1.3$) firms, and a supply elasticity equal to 0.2 from Abuin (2024) as in Harstad and Holtsmark (2024). We provide all details on the calibration in Table 7.1.[44]

Table 7.1: Calibration of model parameters

Description	Parameter	Value	Source
Demand elasticity, non gas-intensive (NGI) firms	ε_{NGI}	1.3	Own estimate
Demand elasticity, gas-intensive (GI) firms	ε_{GI}	0.03	Own estimate
Supply elasticity	η	0.20	Abuin (2024)
Consumer price for NGI (€ /Smc)	p_{NGI}	1.143	Eurostat
Consumer price for GI (€ /Smc)	p_{GI}	1.142	Eurostat
Producer price (€ /Smc)	P	1.093	GME
Industrial gas consumption (Smc)	Q	6.70×10^9	SNAM
NGI share of consumption	Q_{NGI}/Q	0.2	CSEA
GI share of consumption	Q_{GI}/Q	0.8	CSEA

Notes: Price and quantity data refer to the first semester of 2022, i.e. the baseline with respect to our policy experiment, which is meant to be conducted in the second semester of 2022.

Table 7.2 summarizes equilibrium and fiscal effects of two alternative policies: a subsidy equal to 0.5 per Smc of gas consumed (in panel a), and a reward of the same amount per Smc of gas saved (panel b). Each policy is simulated under three designs: uniform for all firms, targeted to gas-intensive (GI) firms, and targeted to non-gas-intensive (NGI) firms.

Under a uniform subsidy (panel a), aggregate gas use increases by 5.3% and the gas price received by producers (P) rises by 26.3%, reflecting the combination of inelastic short-run supply and a sizable outward shift in demand. In contrast, targeting the subsidy to GI firms generates minimal equilibrium distortions: prices increase by only 2.2% and quantities by

[44] It is straightforward to allow a fraction of firms to be covered by fixed-price contracts within each group. This modifies the denominator of equation 7: firms with fixed-price contracts F face $dp_{gF}/dP = 0$ and therefore drop from the demand-side equilibrium response, amplifying pass-through to the producer price. Restricting subsidies to exposed firms reduces both the direct demand impulse (numerator) and the equilibrium dampening (denominator); quantitatively, the policy conclusions are robust to this extension.

0.4%. As a consequence, the fiscal cost is composed almost exclusively of mechanical outlays, which are however quite significant (2.6 billion) due to the high GI baseline consumption.

The comparison of these two designs suggests that, by targeting gas-intensive firms rather than providing uniform subsidies, policymakers can support the most exposed group, while minimizing upward price pressure, and reducing the fiscal cost by almost one quarter.

It is instructive to see that, despite much smaller baseline consumption (and thus substantially lower mechanical outlays of €659 million), NGI targeting yields large equilibrium responses: prices rise by 24.1% and quantities by 4.8%. This policy has similar equilibrium effects compared to the uniform design, as in both cases NGI firms drive the behavioral response. Fiscal costs are the lowest among the three designs, but with a much higher behavioural share.

Summing up, subsidizing the elastic margin amplifies scarcity pressures and raises producer prices, whereas subsidizing the inelastic, high-exposure group primarily reallocates resources with comparatively small effects on aggregate gas use.

Table 7.2: Results of policy simulations

Instrument	**Target**	$\Delta P/P$	$\Delta Q/Q$	**Fiscal cost** (€ bn)		
		(%)	(%)	**Mechanical**	**Behavioral**	**Total**
Panel (a)						
Per-unit subsidy	Uniform	+26.3	+5.3	3.35	0.18	3.53
	GI only	+2.2	+0.4	2.68	0.03	2.71
	NGI only	+24.1	+4.8	0.67	0.18	0.85
Panel (b)						
Pay-for-reductions	Uniform	−26.3	−5.3	0.00	0.18	0.18
	GI only	−2.2	−0.4	0.00	0.03	0.03
	NGI only	−24.1	−4.8	0.00	0.18	0.18
Panel (c)						
Mix: Subsidy to GI + reward to NGI		−21.9	−4.4	2.68	0.25	2.93

Notes: The table displays the first-order price and quantity effects of introducing alternative instruments: a per-unit subsidy, a per-unit reward for gas savings, and a combined policy mix. Each instrument is equal to €0.50/Smc. The policy mix combines a subsidy targeted to gas-intensive (GI) firms with a pay-for-reductions scheme targeted to non-gas-intensive (NGI) firms. Numbers are computed using equations 7, 8, 11, and 12. Fiscal costs are in € billions.

Panel (b) shows that a pay-for-reductions scheme is the mirror image of the subsidy policy in terms of equilibrium effects on prices and quantities, inducing behavioral response of the same size, but opposite sign. The reward scheme is less expensive for the government because the fiscal base is the energy saved, not the energy consumed, thus transferring fewer resources to the corporate sector.

Taken together, results from panels (a) and (b) in Table 7.2 suggest that policymakers willing to help the most vulnerable businesses, while minimizing equilibrium responses that exacerbate the crisis, can resort to a policy mix: subsidies for gas-intensive firms, and pay-for-reduction for the others. This combination allows the government to exploit the most elastic segment of the market to attain energy savings at little cost, while transferring significant resources to the less elastic segment, with little equilibrium consequences.

Panel (c) reports simulated effects of this combined policy. In our calibration, a €0.5/Smc subsidy to gas-intensive firms combined with a €0.5/Smc reward for reductions by non-gas-intensive firms reduces equilibrium gas consumption by 4.4% and lowers the producer price by 21.9%. The total fiscal cost is €2.93 billion, mostly reflecting the mechanical cost of transfers to gas-intensive firms. The policy mix therefore achieves conservation through the elastic segment of demand while directing financial support toward the high-exposure, inelastic firms.

8 Conclusions

What determines how industrial firms respond to large, sudden energy price increases, and how should policymakers design support measures in the face of such shocks? These questions, once largely academic, have become central to economic policy. The 2021-22 energy crisis sent natural gas prices to historically unprecedented levels and forced governments across the EU to intervene at substantial fiscal cost. With geopolitical tensions over energy supply persisting, and European industry's structural dependence on imported gas far from resolved, these questions will not remain confined to the recent past.

This paper provides the first firm-level causal evidence on industrial energy demand during a period of severe price stress. We exploit the staggered expiration of fixed-price energy contracts held by Italian industrial firms before the onset of the crisis, using a research design that compares firms as they lose price protection at different points in time. Our findings reveal a pattern of substantial short-run rigidity, with important differences between energy carriers, firm types, and time horizons. For electricity, we find no quantity response despite large cost increases. For natural gas, consumption falls meaningfully, but only in the second half of 2022. The timing aligns with a shift in market expectations during the summer of that year, when futures markets began projecting prolonged scarcity.

A central finding concerns treatment effect heterogeneity. Using machine learning to detect heterogeneity, an application that is novel in this literature, we show that gas-intensive firms, which account for most of industrial gas consumption, have a price elasticity of demand close to zero, while other firms are substantially more responsive.

These findings have direct implications for the design of energy support measures. We develop a partial equilibrium incidence framework that incorporates demand heterogeneity and general equilibrium price effects, and use it to compare targeted subsidies, untargeted subsidies, and pay-for-reduction schemes. A key result is that the relevant sufficient statistic for evaluating any policy is a quantity-weighted demand elasticity, not the simple average across firms. Because gas-intensive firms are both large consumers and highly inelastic, this quantity-weighted elasticity is far below the unweighted mean - with major consequences for subsidy pass-through and fiscal cost. The optimal policy mix follows directly: subsidies should be targeted to inelastic, gas-intensive users - minimizing equilibrium distortions while supporting the most exposed firms - while pay-for-reduction schemes should be directed at the elastic segment, where conservation incentives are most effective and fiscally cheap.

More generally, our findings underline the importance of incorporating heterogeneity in adjustment margins into both microeconomic policy evaluation and macroeconomic models of energy shocks. As recent geopolitical developments renew concerns about the resilience of industrial production to disruptions in global energy markets, understanding which firms can adjust and along which margins remains central for the design of efficient stabilization policies. A natural extension of this paper is the longer-run response, as firms invest in more energy-efficient capital, switch processes, or substitute toward alternative carriers. Understanding whether the 2021-22 crisis ultimately triggered durable structural change in industrial energy use, and whether current geopolitical risks are accelerating the energy transition by making the option value of self-generation and fuel diversification more salient, remains an open and important question.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process During the preparation of this work the authors used Claude, ChatGPT and Refine in order to check the internal coherence of the draft. After using this tool service, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

Abuin, C. (2024). Power decarbonization in a global energy market: The climate effect of us lng exports. Technical report, Working Paper.

Ahlvik, L., Kaariaho, T., Liski, M., and Vehviläinen, I. (2025). *Household-level responses to the european energy crisis*. JSTOR.

Alessandri, P. and Gazzani, A. G. (2023). Natural gas and the macroeconomy: Not all energy shocks are alike. *Bank of Italy Temi di Discussione (Working Paper) No*, 1428.

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 d'Italia, 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.

ARERA (2022). Relazione annuale sul 2021. Technical report, Regulatory Authority for Energy, Networks and Environment, Rome.

ARERA (2023). Relazione annuale sul 2022. Technical report, Regulatory Authority for Energy, Networks and Environment, Rome.

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.

Banca d'Italia (2021). Rapporto sulla stabilità finanziaria 2021/2. Accessed: 2026-05-20.

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.

Blonz, J. A. (2022). Making the best of the second-best: Welfare consequences of time-varying electricity prices. *Journal of the Association of Environmental and Resource Economists*, 9(6):1087–1126.

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. (2024). Revisiting event-study designs: robust and efficient estimation. *Review of Economic Studies*, page rdae007.

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).

Byrne, D. P., Nauze, A. L., and Martin, L. A. (2021). An experimental study of monthly electricity demand (in) elasticity. *The Energy Journal*, 42(2):205–222.

Caballero, R. J. and Engel, E. M. R. A. (1999). Explaining investment dynamics in U.S. manufacturing: A generalized (S, s) approach. *Econometrica*, 67(4):783–826.

Cahuc, P., Carcillo, S., and Zylberberg, A. (2014). *Labor economics*. MIT press.

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

Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annu. Rev. Econ.*, 1(1):451–488.

Chetty, R., Looney, A., and Kroft, K. (2009). Salience and taxation: Theory and evidence. *American economic review*, 99(4):1145–1177.

Clarke, D., Pailañir, D., Athey, S., and Imbens, G. (2024). On synthetic difference-in-differences and related estimation methods in stata. *The Stata Journal*, 24(4):557–598.

Cooper, R. W. and Haltiwanger, J. C. (2006). On the nature of capital adjustment costs. *The Review of Economic Studies*, 73(3):611–633.

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.

Costa, F. and Gerard, F. (2021). Hysteresis and the welfare effect of corrective policies: Theory and evidence from an energy-saving program. *Journal of Political Economy*, 129(6):1705–1743.

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

Davis, L. W. (2014). The economic cost of global fuel subsidies. *American Economic Review*, 104(5):581–585.

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.

Davis, S. J., Grim, C., Haltiwanger, J., and Streitwieser, M. (2013). Electricity unit value prices and purchase quantities: Us manufacturing plants, 1963–2000. *Review of Economics and Statistics*, 95(4):1150–1165.

De Chaisemartin, C. and d'Haultfoeuille, 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.

Doms, M. and Dunne, T. (1998). Capital adjustment patterns in manufacturing plants. *Review of Economic Dynamics*, 1(2):409–429.

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

Emiliozzi, S. and Favero, F. (2025). Unveiling natural gas consumption sectoral price elasticities. Technical Report 941, Bank of Italy.

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.

Fontagné, L., Martin, P., and Orefice, G. (2024). The many channels of firms adjustment to energy shocks: evidence from france. *Economic Policy*, 39(117):5–43.

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. (2024). Energy tax exemptions and industrial production. *The Economic Journal*, page ueae048.

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):1658–1688.

Harstad, B. and Holtsmark, K. (2024). The gas trap: Outcompeting coal vs. renewables. Technical report, National Bureau of Economic Research.

Hatamyar, J., Kreif, N., Rocha, R., and Huber, M. (2023). Machine learning for staggered difference-in-differences and dynamic treatment effect heterogeneity. *arXiv preprint arXiv:2310.11962*.

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

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

Jessoe, K. and Rapson, D. (2014). Knowledge is (less) power: Experimental evidence from residential energy use. *American Economic Review*, 104(4):1417–1438.

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*.

Lamorgese, A., Patnaik, M., Linarello, A., and Schivardi, F. (2024). Management practices and resilience to shocks: Evidence from covid-19. *Management Science*.

Leicher, J., Giese, A., and Wieland, C. (2024). Electrification or hydrogen? the challenge of decarbonizing industrial (high-temperature) process heat. *J*, 7(4):439–456.

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

Mallapragada, D. S., Dvorkin, Y., Modestino, M. A., Esposito, D. V., Smith, W. A., Hodge, B.-M., Harold, M. P., Donnelly, V. M., Nuz, A., Bloomquist, C., et al. (2023). Decarbonization of the chemical industry through electrification: Barriers and opportunities. *Joule*, 7(1):23–41.

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.

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.

Polo, M. and Scarpa, C. (2003). The liberalization of energy markets in europe and italy.

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.

Roth, J. (2022). Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights*, 4(3):305–322.

Ruhnau, O., Stiewe, C., Muessel, J., and Hirth, L. (2023). Natural gas savings in germany during the 2022 energy crisis. *Nature Energy*, pages 1–8.

Scarazzato, F. (2025). How (in) elastic is the short-run demand for electricity? *Economics Letters*, 254:112453.

Sgaravatti, G., Tagliapietra, S., and Zachmann, G. (2021). National policies to shield consumers from rising energy prices. *Bruegel Datasets*, 4.

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$_id = 1$, .

Sikl, V., Havránková, Z., Kudela, P., and Kudelova, A. (2025). Price elasticity of electricity revisited: A meta-analysis. IES Working Paper 17/2025, Prague.

Stantcheva, S. (2023). How to run surveys: A guide to creating your own identifying variation and revealing the invisible. *Annual Review of Economics*, 15(1):205–234.

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.

Wehrhöfer, N. (2026). Energy prices and inflation expectations: Evidence from households and firms. *Journal of Political Economy Macroeconomics*, 4(1):135–163.

Wei, M., McMillan, C. A., and de la Rue du Can, S. (2019). Electrification of industry: Potential, challenges and outlook. *Current Sustainable/Renewable Energy Reports*, 6(4):140–148.

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

Zettelmeyer, J., Tagliapietra, S., Zachmann, G., and Heussaff, C. (2022). Beating the european energy crisis.

Appendices

Appendix A Background

A.a The 2021-22 energy crisis in Italy

In 2020, one year before the beginning of the energy crisis, 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. Figure A.1 plots the evolution of the wholesale price of natural gas (solid red line) and electricity (dashed black line) in Italy.[45] The price of gas 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 surged to around 250 in February when Russia invaded Ukraine, but rapidly declined thereafter and stabilized below 100 until summer 2022. At that point, the price climbed again quite rapidly, reaching a historic peak of 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.

A.b The Italian retail energy market

Following a broader trend in the European Union, the Italian retail markets for electricity and natural gas have been gradually liberalized in the late 1990s (Polo and Scarpa, 2003). According to the Regulatory Authority for Energy, Networks and Environment (ARERA), both the market for electricity and for gas are not very concentrated (ARERA, 2022).[46]

The contracts offered in the liberalized market are highly heterogeneous and customizable in terms of conditions and prices, but can be broadly split between fixed-price and variable-

[45]The two time series of gas and electricity prices are indeed highly correlated (98% at the daily frequency). In levels, the price of gas is lower than electricity because gas is an input in the production function of power.

[46]Households and very small businesses – outside the scope of our analysis – can still participate in a "protected" market (*mercato tutelato*), where ARERA periodically sets a controlled price. As of 2024, the transition to a fully liberalized market and the elimination of the "protected" market are almost complete but still ongoing.

Figure A.1: Wholesale gas and electricity price at the daily frequency

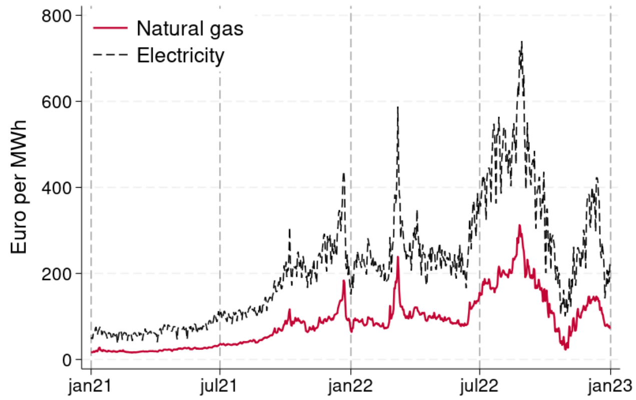

Source: Italian Power Exchange (GME). *Note*: The figure shows the spot price of natural gas and electricity traded on the Italian day-ahead wholesale market.

price contracts. Fixed-price contracts typically have a standard duration of 12, 24, or –less frequently– 36 months. Each supplier can offer as many contracts as it sees fit. However, for households and firms with a low-voltage connection for electricity (a minority in our sample) and an annual gas consumption below 200,000 standard cubic meters (one third of our sample), the supplier is required to also offer two standardized contracts - one with variable and one with fixed price. While the conditions for these standardized contracts are designed by ARERA, suppliers compete on price.[47] All the contracts posted on the market, whether standardized or not, are published daily on a website managed by a government agency.[48] The portal does not include information on large and/or energy-intensive firms, which usually negotiate *ad hoc* contracts directly with their supplier.

A.c Additional information on the retail energy markets

Electricity In 2021, all firms with more than 50 employees, that is those analyzed in this paper, had to purchase electricity in the liberalized market. Most likely a minority of firms in our sample have a low-voltage connection and thus must be offered the standardized contracts designed by the authority. According to ARERA, in 2021 45% of all business customers were on a fixed-price contract, while the rest were on a variable price-contract, linked to the fluctuation of the wholesale price or of other price indexes. On average in 2021 the unit cost for energy for outstanding (both newly signed and signed in previous years) fixed-price contracts was 30% lower than those on a variable price contract. In 2022,

[47] Standardized contracts should allow unsophisticated customers to easily compare prices across sellers.
[48] https://www.ilportaleofferte.it/portaleOfferte/.

the proportion of business customers on fixed-price contracts remained largely unchanged (47%); in that year, they enjoyed a 80% lower unit cost compared to variable-price contracts, which were completely exposed to the increase in wholesale prices (ARERA, 2023). Notice that fixed-price contracts were continuously supplied by some retailers throughout the crisis, at least for low-voltage connection (for which ARERA publishes data). However, fixed-price contracts on sale became rarer, and, of course, more expensive during the crisis. The average unit cost for outstanding (both newly signed and signed in previous years) fixed-price contracts (both standardized and not) increased from an average of 91 euro/MWh in 2021 to 171 the next year.

Natural gas All firms must purchase natural gas on the free market; those with an annual consumption below 200,000 standard cubic meters - approximately 33% in our sample - must be offered the standardized contracts. In 2021, business customers were almost evenly split between those on a fixed-price contract (44%) and those on a variable price contract (56%) (ARERA, 2022). On average the unitary cost of energy was 30% higher in the latter case relative to the former; the difference is driven by the increase in wholesale price that materialized over the second semester. In 2022, the share of business customers on a variable contract increased to 63% (ARERA, 2023). The cost spread between the two types of contracts increased to 50%, due to the further stark increase in the wholesale price. In the case of natural gas, the number of new fixed-price contracts for sale on the free market dropped to less than 5 in the first semester of 2022, and to zero in the second. However, firms with an annual consumption below 200,000 standard cubic meters could still purchase one of the fixed-price standardized contracts. The average unitary cost for outstanding (both newly signed and signed in previous years) fixed-price (both standardized and not) contracts increased from an average of 34 cents/cubic meter in 2021 to 77 the next year, driven by new more expensive contracts.

Appendix B Tariff structure and the fixed-charge share

Italian non-household energy bills comprise four components: (i) the commodity/energy supply cost, set by the retail contract; (ii) network charges for transmission, distribution, and metering, regulated by the Italian Regulatory Authority for Energy, Networks and Environment (ARERA); (iii) system charges (*oneri generali di sistema*), also regulated; and (iv) taxes (excise duties and VAT). For both electricity and gas, retail contracts specify a *quota fissa* (a fixed annual charge, in €/year, covering commercialization costs) and a *quota energia* or *quota variabile* (a per-unit volumetric charge, in €/kWh or €/Smc). For electricity, some contracts also include a *quota potenza* (a capacity charge in €/kW/year). Network charges have a fixed component per delivery point and a dominant volumetric component; for electricity, a regulated capacity charge also applies. Taxes are entirely proportional to consumption or to the total bill.

Data. We quantify the fixed-charge share using open data from ARERA's *Portale Offerte*, which publishes all retail energy offers in a standardized XML format.[49] We retain both fixed-price (TIPO_OFFERTA = 1) and variable-price (TIPO_OFFERTA = 2) contracts. The two contract types differ in what the XML records for the commodity component: for fixed-price contracts, the XML contains the full per-unit energy charge; for variable-price contracts, it contains only the vendor spread above a wholesale index (PUN for electricity, PSV for gas), which we supplement with the wholesale benchmark.

Construction of the fixed-charge share. For each offer we compute $s \equiv F/(F + p \cdot Q)$, where F is total annual fixed charges, p is the total per-unit variable charge, and Q is annual consumption set at a given percentile of the Invind sample (semester observations annualized by multiplying by two).

Fixed charges (F). We include all fixed invoice components:

- Vendor-set fixed charge (€/year) and, where applicable, vendor-set capacity charge (€/kW/year), both extracted from the Portale Offerte XML (MACROAREA = 01).
- ARERA-regulated network fixed charge τ_1 (distribution quota fissa plus metering quota fissa): from delibera 621/2021/R/eel (TIT Tabella 3 and TIME Tabella 1) for electricity, and from delibera 620/2021/R/GAS (Allegato A Tabella 1) for gas. The gas τ_1 varies across Italian distribution zones (*ambiti tariffari*) and meter size classes; we use 437 €/PDR/yr,

[49] Data downloaded from https://www.ilportaleofferte.it/portaleOfferte/opendata. We restrict to non-domestic offers (TIPO_CLIENTE = 02). Pricing components were extracted and classified from XML files containing 361 electricity and 172 gas offers (January 2021).

the median across the six mainland Italian zones for the meter size class appropriate to the industrial consumption levels in our sample (G10–G40, corresponding to annual consumption of 20,000–800,000 Smc/yr); Sardinia is excluded as an outlier due to its anomalously sparse distribution network.

- ARERA-regulated capacity charge τ_2 (€/kW/year, electricity only): from delibera 621/2021/R/eel, TIT Tabella 3. Since contracted power is not recorded in the Portale Offerte XML, we impute it from the Invind survey, which asks firms to report their contracted power in 2022.[50] We regress log contracted power on log annual electricity consumption using firms observed at $t = 1$ ($N = 365$): $\log(\text{kW}) = 6.353 + 0.609 \times \log(Q_{\text{GWh}})$. At the sample median consumption (4,500 MWh/yr), the implied contracted power is 1,452 kW (MTA3 tariff tier, >500 kW).

- *System charges* (*oneri generali di sistema*). For electricity, system charges include two regulated components relevant to fixed charges: ASOS (*Applicazione del Corrispettivo per il finanziamento delle misure ed interventi per la gestione selettiva dei rifiuti urbani e assimilati*) and ARIM (*Agevolazioni alle imprese a forte consumo di energia elettrica*). Both have a quota fissa (a fixed annual charge per delivery point), a quota potenza (a per-kW annual charge), and a quota energia (a per-kWh charge). We include the quota fissa and quota potenza in F and the quota energia in p. For gas, the analogous system charges (GS, RE, RS, UG1) are assessed only on a per-unit basis and are therefore included entirely in p.

  The values we use are taken from delibera 278/2021/R/com (Tabella 1 for ASOS and Tabella 6 for ARIM; Tabella 8 for gas), which sets the rates applicable from 1 July 2021. We apply the Classe 0 rates, corresponding to non-energy-intensive firms (*imprese non energivore*). This is a conservative choice that overstates F: energy-intensive firms (*imprese energivore*), classified under Tabella 2 of the same delibera, face ASOS rates of zero.

  These values represent an upper bound on the system charge contribution to F during our sample period. From 1 January 2022, ARERA set both ASOS and ARIM to zero for all non-household consumers (delibera 35/2022/R/eel), which covers essentially the entire 2022 portion of our estimation window. The true fixed-charge share during 2022 is therefore lower than the figures reported in Table B.1.

Variable charges (p). For fixed-price contracts, p is the full per-unit vendor charge from the XML (which already embeds the commodity price), plus ASOS and ARIM quota energia

[50] The survey question reads (translated from Italian): "What is the contracted power level, even approximately, of your electricity meter? (Answer with reference to the most powerful meter among those available to the firm)." The question elicits the *potenza disponibile*, the capacity reserved on the distribution network under the firm's connection agreement, which is the quantity to which τ_2 applies under the Italian tariff structure.

and excise taxes (0.0125 €/kWh for electricity and 0.0128 €/Smc for gas, under TUA Art. 52 and Art. 26 respectively). For variable-price contracts, p additionally includes the wholesale benchmark (PUN = 0.25 €/kWh for electricity; PSV = 1.20 €/Smc for gas), since the XML records only the vendor spread above the index. IVA (22%) applies proportionally to the entire bill and therefore cancels in the ratio s; it is excluded.

Results. Table B.1 reports the cross-offer distribution of s at the median consumption level of the Invind estimation sample.

For gas, where we find a demand response, the fixed-charge share is minimal. At the sample median consumption (192,000 Smc/yr), the median offer has $s = 0.72\%$, and even the 90th-percentile offer reaches only 1.08%. The average unit cost of gas is therefore an excellent approximation to the per-unit variable price for the firms in our sample.

For electricity, where we find a zero elasticity, the median fixed share at sample median consumption is 8.3%. This is larger than for gas, reflecting the capacity-based structure of Italian electricity tariffs and the inclusion of ASOS and ARIM values, which overstate the actual share during our 2022 estimation window.

Table B.1: Fixed-charge share of the total energy invoice at median sample consumption (upper-bound scenario)

		$s \equiv F/(F + p \cdot Q)$ across offers (%)		
Energy	Consumption [sample P50]	Median	P75	P90
Gas (PSV = 1.20 €/Smc)	192,000 Smc/yr	0.72	0.89	1.08
Electricity (PUN = 0.25 €/kWh)	4,500 MWh/yr	8.29	9.67	10.09

Notes: Distribution across all non-domestic offers (172 gas; 361 electricity; ARERA *Portale Offerte*, January 2021) of the fixed-charge share $s \equiv F/(F + p \cdot Q)$, evaluated at median sample consumption. F and p are constructed as described in the text. Pre-zeroing system charges are included in F, so these figures are upper bounds relative to the actual 2022 invoice.

Appendix C Questionnaires

This section reports survey questions exactly as they appeared to firms. The English translation is carried out by the Bank of Italy.

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

Figure C.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

Appendix D 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 take plausible values. To this end, we implement the following algorithm.

1. First, we exclude from our sample 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.

2. Second, we rely on two benchmarks to cross-check the plausibility of the Invind replies, separately for electricity and natural gas.

   (a) 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.[51] We flag observations for which the unit price – computed by dividing expenditures by quantity – is either below half the minimum or above double the maximum of the reference Eurostat price across consumption classes.[52]

   (b) We examine the ratio between energy costs and turnover. We flag observations above and below the 99th and 1st percentile of the distribution. These correspond to cost-turnover ratios above 50% and below 0.1%, respectively.

3. Combining the criteria in a) and b), we identify 6 error categories for the firm-level replies on electricity and 4 categories for the ones on natural gas. For example, one category is made by those firms that incorrectly report electricity consumption in kWh, instead of MWh; this mistake results in a price (expenditures for electricity over quantity of electricity) which is approximately 1000 times lower than what reported by Eurostat. In Table D.1 we list all the different categories, specifying the reporting mistakes associated with each of them. 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 the same mistake done in both semesters. In these cases, we rescale the values in order to fix the reporting mistake. Carrying on with the previous example, if we believe a firm reports quantities

[51] 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.

[52] In light of the upward trend in prices over time, 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.

in kWh instead of MWh, we divide it by 1000 to convert it back into MWh, and then we calculate the price again. In the final estimation samples, we only include those firms whose answers are consistent with both the a) and b) criteria in all 4 semesters, be it before or after our rescaling.

4. For the subset of energy-intensive firms, we examine differences between input quantities declared in the survey against administrative microdata collected by CSEA. Whenever the gap is larger than 35% for at least one semester, we conservatively drop the firm from the estimation sample. Figure D.1 plots the values from the administrative source against the value from the Invind survey; for the latter, we flag in blue those observations that needed rescaling according to the previous steps. The results are reassuring for both the rescaled observations (in blue) and for those that we left unchanged (in red), as they both lie very close to the 45-degree line. This suggests that our survey data match administrative records and that our correction algorithm works well.

5. As for natural gas quantities, we can perform a further check using firms with plants subject to the EU ETS, for which we observe confidential administrative records from ISPRA. The EU ETS is a plant-level regulation. Thus a multi-plant firm could have a plant subject to the EU ETS, and others not subject to it. As a consequence, we exclude firms for which quantities reported in the survey are smaller than the corresponding EU ETS record; instead, we keep firms reporting larger values than the corresponding EU ETS record. Overall, only 2.8% (4.5%) of the electricity sample and 2.1% (2.5%) of the natural gas sample in 2021 (in 2022) display differences larger than 35% in absolute value between survey and administrative records. We drop these observations.

Table D.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 results from the data validation procedure, 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. Specifically, these two checks allow us to determine the 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 D.1: Consistency with other data sources

(a) Electricity consumption in 2021

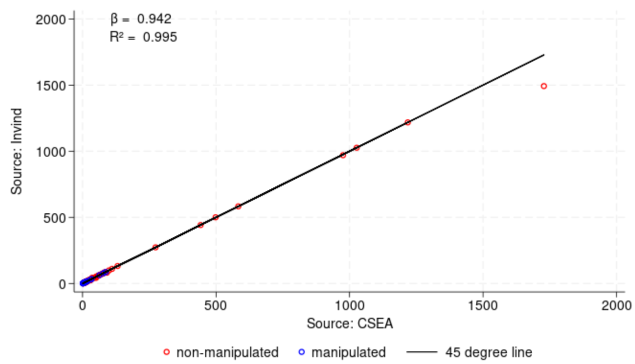

(b) Natural gas consumption in 2021

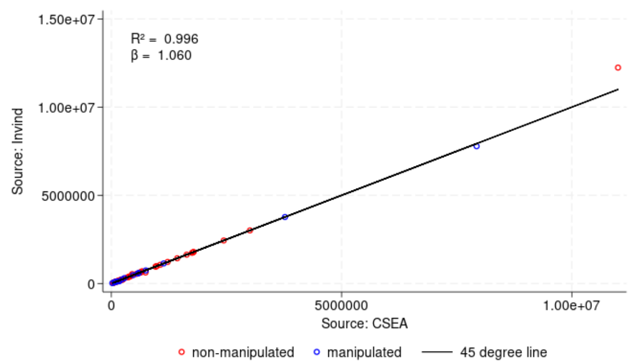

(c) Electricity consumption in 2022

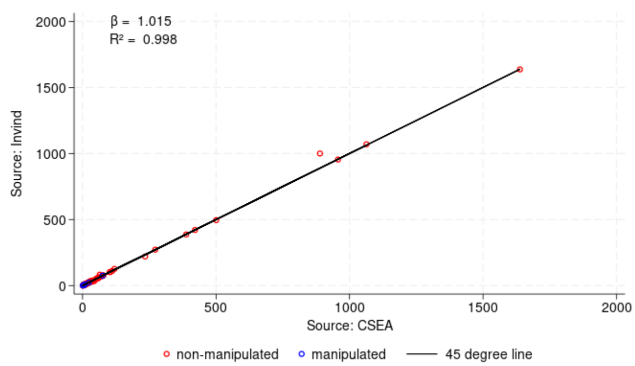

(d) Natural gas consumption in 2022

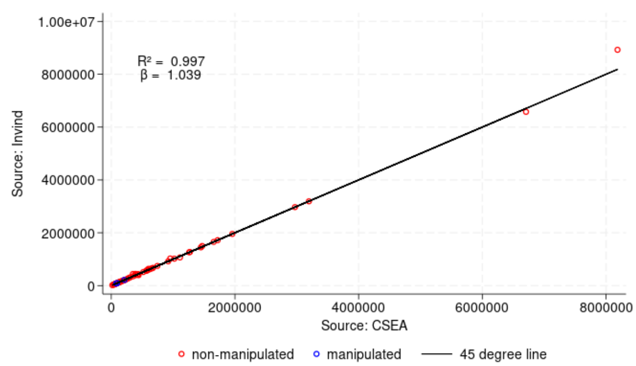

(e) Gas consumption in 2021

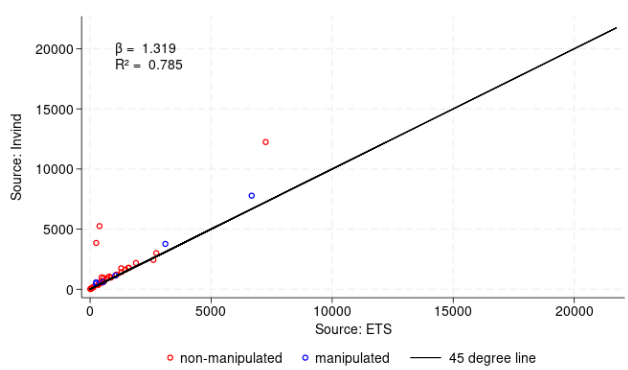

(f) Gas consumption in 2022

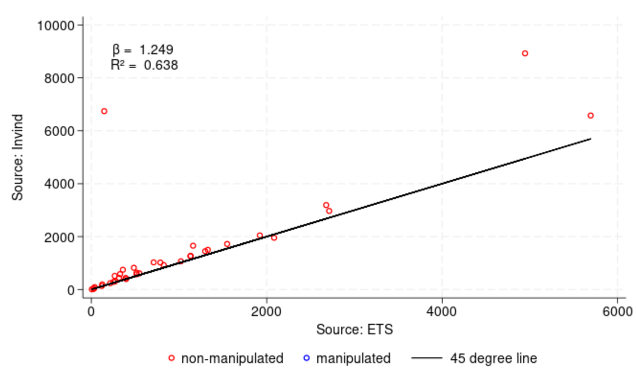

Note: The figures show the consistency between data sourced via the Invind survey and quantities of electricity and natural gas recorded in administrative data for a sub-sample of firms belonging to the *energivore* and *gasivore* lists (panel (a) and (b) for 2021, and panel (c) and (d) for 2022) and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

Appendix E Non-response bias

Each Invind wave related to firms with 50+ employees has approximately 1800 respondents. Response rates to these energy sections is around 50% in both waves, though we drop a small number of observations as detailed in the previous section. Finally, we lose additional observations (approximately 23%) as we only include firms that answered the energy section in both waves.

In this section we examine the robustness of our findings to selective item non-response using "inverse probability weighting" (Wooldridge et al., 2002; Stantcheva, 2023). This method is commonly used to address differential attrition by utilizing the relationships among observed covariates and response rates to re-weight the observed data in a way that approximates the distribution in the full data set (Stantcheva, 2023; 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 13 and 14, where X_i include covariates measured at baseline for all observations and the dependent variable is equal to one for firms belonging to the estimation sample, and zero for other firms surveyed in Invind, including firms dropped due to energy variable misreporting (see previous sections of this Appendix).[53]

$$\mathbf{1}(\text{Electricity sample}_i) = X_i'\beta^e + \varepsilon_i \tag{13}$$

$$\mathbf{1}(\text{Gas sample}_i) = X_i'\beta^g + \xi_i \tag{14}$$

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

In Figure E.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.

[53]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 E.1: Common support of propensity score

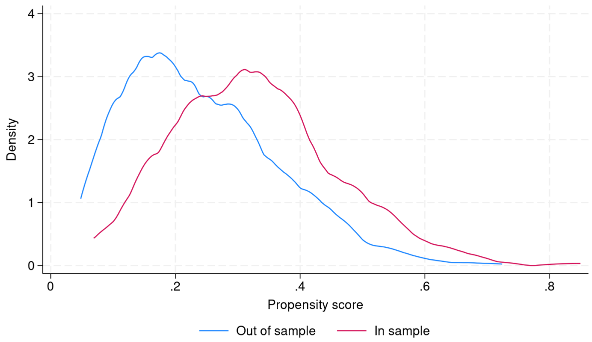

(a) Natural gas sample

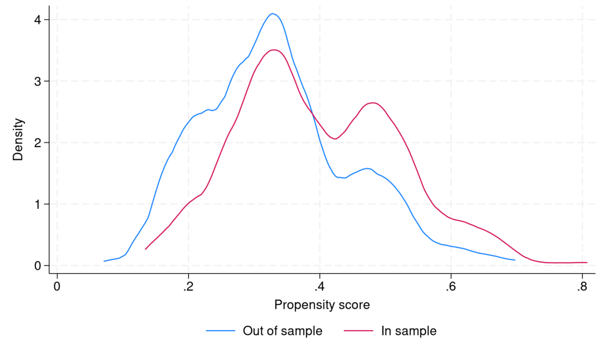

(b) Electricity sample

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

Figure E.2: Inverse-probability-weighted estimates

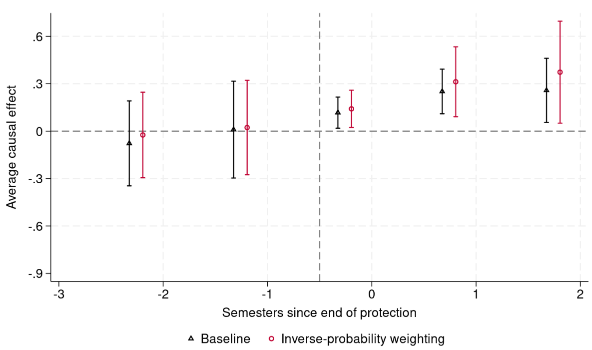

(a) Price of natural gas

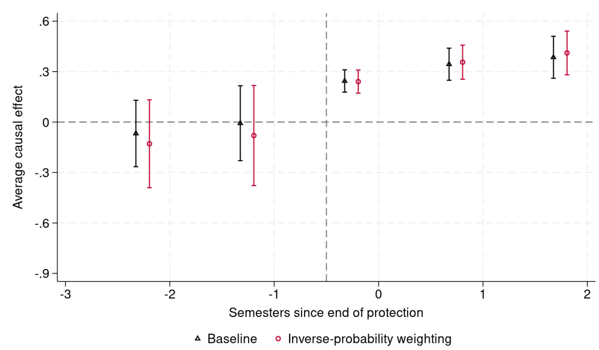

(b) Price of electricity

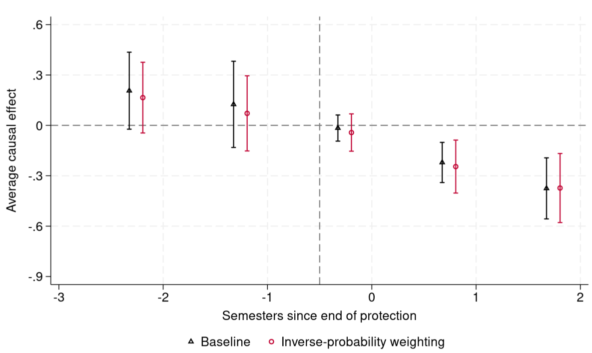

(c) Gas demand

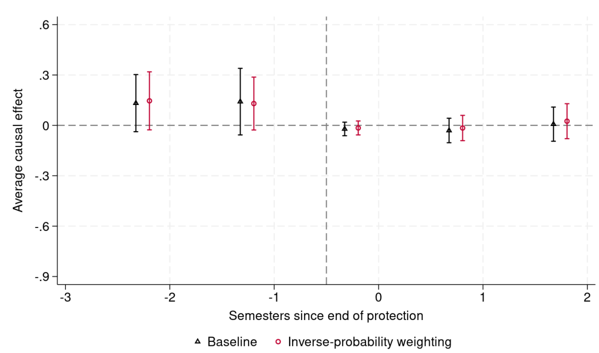

(d) Electricity demand

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, 2023).

Appendix F Firm Exit During the Sample Period

To assess whether conditioning on a balanced panel introduces selection bias, we matched every firm in our *2021* electricity ($N = 860$) and gas ($N = 660$) samples to the Infocamere administrative registry, which records the dates of non-liquidatory insolvency procedures: court-supervised administration, debt restructuring agreements, extraordinary administration, *concordato preventivo*), liquidatory procedures: *fallimento*, judicial liquidation, compulsory administrative liquidation, and related events, non-insolvency exits: voluntary liquidation, and definitive cessations. We flag any event with a date falling in 2021 or 2022.

Table F.1: Exit and distress events among sample firms, 2021–2022

Event type	Electricity sample ($N = 860$)	Gas sample ($N = 660$)
Liquidatory procedure	0	0
Definitive cessation	0	0
Voluntary non-insolvency exit	1	1
Non-liquidatory distress	3	1
Any exit or distress event	4	2

Note: Source: Infocamere administrative registry matched to the 2021 Invind electricity and gas samples. Events are flagged if the relevant date falls within 2021–2022. Non-liquidatory distress procedures include court-supervised administration, debt restructuring agreements, extraordinary administration, and *concordato preventivo*.

Table F.1 confirms that firm exit is negligible in our setting. The absence of exit pressure is particularly notable for energy-intensive firms, which faced the largest cost shocks. Rather than exiting, these firms adjusted via output-price pass-through: a Banca d'Italia analysis documents that they raised selling prices by approximately 23% in 2022 – roughly twice the rate of other industrial firms – while keeping volumes essentially unchanged, and their gross operating margin-to-turnover ratio slightly *increased* (from 9.5 to 9.8%), with government support measures preventing a decline of approximately 1.5 percentage points (Banca d'Italia, *Relazione annuale sul 2024*, Riquadro "L'impatto del costo dell'energia sul settore industriale"). These findings are directly consistent with the near-zero electricity demand elasticity and the output-price pass-through documented in our paper (Section 4.c), and confirm that the balanced-panel restriction does not introduce material survivor bias.

Appendix G Summary statistics

In Table G.1, we present summary statistics for our energy demand estimation samples: an electricity sample and a natural gas sample. Firm characteristics are measured prior to the crisis, specifically in the first half of 2021. The sample comprises industrial firms with at least 50 employees, as determined by the Invind sampling design. All statistics are weighted by survey weights, and the number of observations can be obtained by multiplying the number of firms by four (representing four semesters).

The electricity and natural gas samples share similar characteristics. Over half of the firms belong to the metalworking industry, while the remainder are distributed relatively evenly across other sectors, except for non-metallic minerals and water & waste, which are minor categories. Geographically, more than 80% of firms are located in Northern Italy, with only 57% based in the South or Islands.

On average, firms in the electricity sample pay €0.16 per kWh and consume approximately 6,000 MWh of electricity per semester. For natural gas, the average price is €10 per GJ, with a typical consumption of 63 million standard cubic meters. About one-third of firms in both samples are either *elettrivore* (electricity-intensive) or *gasivore* (gas-intensive) firms, while around 5% operate plants subject to the EU ETS. Additionally, about half of the firms in the natural gas sample consider gas to be an "essential" input for production.

Treatment cohorts are unevenly distributed. Only 14% (electricity sample) and 18% (gas sample) belong to the pure control group. The largest cohorts are the early treated (exposed in 2021h2) and mid treated (exposed in 2022h1), while the late treated (exposed in 2022h2) form a small minority (1% and 3%, respectively).

Finally, the sample firms are relatively large, with average annual revenues of €8090 million and over 200 employees, though these distributions are highly skewed.

Tables G.2 and G.3 report summary statistics and balancing tests for baseline covariates across the four treatment cohorts, separately for the electricity and gas samples. Each table presents cohort-specific means and standard errors for sector dummies and energy-related firm characteristics, along with p-values from pairwise and joint orthogonality tests.

In both samples, there are statistically significant differences in some covariates across cohorts. For instance, in the electricity sample (Table G.2), the proportion of firms in the Food sector is significantly higher among late-treated firms than among pure controls (p-value < 0.01), and joint orthogonality tests reject balance at conventional significance levels for several sector variables. Similar patterns emerge in the gas sample (Table G.3), where treatment groups differ notably in their likelihood of being gas-dependent or classified as energy-intensive.

Despite these statistical differences, most of the covariate imbalances are small in magnitude. Differences in means are typically within a few percentage points across groups. This suggests that, while the null hypothesis of perfect balance is sometimes rejected, the economic significance of these differences is limited. Taken together, these results support the plausibility of the identification strategy and suggest that the estimation of treatment effects is unlikely to be driven by stark differences in pre-treatment characteristics.

To further mitigate concerns about residual imbalance, in Appendix H we probe robustness by allowing for flexible differential trends, by interacting time fixed effects with sector dummies and energy-intensity indicators. These non-parametric specifications absorb sector-specific and energy-dependence-specific dynamics over time. Reassuringly, results are virtually identical to the baseline estimates, with point estimates remaining effectively unchanged.

Table G.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 5. 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.

Table G.2: Summary statistics and p-values comparing treatment groups. Electricity sample

	(1) Early treated	(2) Mid treated	(3) Late treated	(4) Pure control	(5) (1) vs (2) p-val	(6) (1) vs (3) p-val	(7) (1) vs (4) p-val	(8) (2) vs (3) p-val	(9) (2) vs (4) p-val	(10) (3) vs (4) p-val	(11) Joint orthog-onality test p-val
Food and beverages	0.080	0.127	0.182	0.158	0.005	0.019	0.000	0.295	0.198	0.688	0.001
	(0.010)	(0.014)	(0.059)	(0.021)							
Textiles & apparel	0.074	0.107	0.000	0.118	0.038	0.062	0.021	0.022	0.595	0.016	0.010
	(0.010)	(0.013)	(0.000)	(0.019)							
Chem., pharma., rubber	0.136	0.153	0.182	0.171	0.385	0.399	0.154	0.615	0.492	0.860	0.479
	(0.013)	(0.015)	(0.059)	(0.022)							
Non-metallic minerals	0.034	0.073	0.091	0.039	0.001	0.054	0.673	0.669	0.046	0.129	0.005
	(0.007)	(0.011)	(0.044)	(0.011)							
Metalworking industry	0.523	0.400	0.364	0.461	0.000	0.041	0.070	0.635	0.082	0.228	0.000
	(0.019)	(0.020)	(0.073)	(0.029)							
Wood, paper, furniture	0.108	0.113	0.182	0.026	0.758	0.133	0.000	0.175	0.000	0.000	0.000
	(0.012)	(0.013)	(0.059)	(0.009)							
Water & waste	0.045	0.027	0.000	0.026	0.073	0.149	0.153	0.273	0.975	0.278	0.116
	(0.008)	(0.007)	(0.000)	(0.009)							
ETS coverage	0.074	0.067	0.182	0.092	0.613	0.010	0.326	0.005	0.171	0.068	0.033
	(0.010)	(0.010)	(0.059)	(0.017)							
Energy-intensive firm	0.284	0.340	0.455	0.355	0.030	0.016	0.024	0.124	0.649	0.203	0.014
	(0.017)	(0.019)	(0.076)	(0.027)							
Gas essential input	0.503	0.553	0.636	0.603	0.080	0.087	0.005	0.285	0.174	0.675	0.019
	(0.019)	(0.021)	(0.073)	(0.030)							
Labour force (yearly avg)	332.227	304.427	452.091	1034.829	0.346	0.166	0.000	0.081	0.000	0.364	0.000
	(20.352)	(21.236)	(115.903)	(242.808)							
Sales (million euro)	158.528	121.912	166.633	588.517	0.064	0.902	0.000	0.245	0.000	0.279	0.000
	(16.177)	(9.888)	(44.005)	(147.681)							

Notes: Table reports means and standard errors (in parentheses) by treatment cohort. Columns (5)(10) report p-values from pairwise orthogonality tests. Column (11) reports p-values from joint F-tests across treatment arms.

Table G.3: Summary statistics and p-values comparing treatment groups. Gas sample

	(1) Early treated	(2) Mid treated	(3) Late treated	(4) Pure control	(5) (1) vs (2) p-val	(6) (1) vs (3) p-val	(7) (1) vs (4) p-val	(8) (2) vs (3) p-val	(9) (2) vs (4) p-val	(10) (3) vs (4) p-val	(11) Joint orthog-onality test p-val
Food and beverages	0.055	0.105	0.143	0.163	0.003	0.054	0.000	0.527	0.034	0.784	0.000
	(0.010)	(0.014)	(0.067)	(0.026)							
Textiles & apparel	0.102	0.073	0.000	0.102	0.103	0.076	0.985	0.140	0.201	0.077	0.120
	(0.013)	(0.012)	(0.000)	(0.022)							
Chem., pharma., rubber	0.133	0.194	0.286	0.184	0.009	0.023	0.087	0.235	0.766	0.205	0.019
	(0.015)	(0.018)	(0.087)	(0.028)							
Non-metallic minerals	0.031	0.040	0.143	0.082	0.438	0.002	0.004	0.012	0.027	0.290	0.002
	(0.008)	(0.009)	(0.067)	(0.020)							
Metalworking industry	0.516	0.468	0.429	0.367	0.129	0.370	0.000	0.687	0.016	0.533	0.005
	(0.022)	(0.022)	(0.095)	(0.035)							
Wood, paper, furniture	0.141	0.089	0.000	0.082	0.010	0.033	0.033	0.100	0.766	0.118	0.006
	(0.015)	(0.013)	(0.000)	(0.020)							
Water & waste	0.023	0.032	0.000	0.020	0.395	0.414	0.809	0.335	0.403	0.448	0.596
	(0.007)	(0.008)	(0.000)	(0.010)							
ETS coverage	0.109	0.056	0.143	0.204	0.002	0.584	0.001	0.063	0.000	0.448	0.000
	(0.014)	(0.010)	(0.067)	(0.029)							
Energy-intensive firm	0.281	0.282	0.714	0.429	0.972	0.000	0.000	0.000	0.000	0.004	0.000
	(0.020)	(0.020)	(0.087)	(0.035)							
Gas essential input	0.582	0.526	1.000	0.860	0.082	0.000	0.000	0.000	0.000	0.051	0.000
	(0.022)	(0.023)	(0.000)	(0.026)							
Labour force (yearly avg)	337.164	332.331	605.000	1476.000	0.869	0.003	0.000	0.004	0.000	0.383	0.000
	(20.203)	(21.251)	(114.776)	(375.847)							
Sales (million euro)	159.749	133.354	647.554	801.606	0.155	0.000	0.000	0.000	0.000	0.799	0.000
	(14.654)	(11.225)	(226.028)	(225.875)							

Notes: Table reports means and standard errors (in parentheses) by treatment cohort. Columns (5)(10) report p-values from pairwise orthogonality tests. Column (11) reports p-values from joint F-tests across treatment arms.

Appendix H Additional robustness tests

H.a Controlling for non-parametric trends

As a robustness check, we extend the baseline specification by allowing for flexible, non-parametric differential trends across firm types. Specifically, we modify the Borusyak et al. (2024) estimator to include full sets of interactions between time fixed effects and: (i) 11 sector dummies, (ii) an energy-intensive status indicator, (iii) a dummy for EU ETS participation, and (iv) firm size class dummies (five categories based on number of employees). This approach allows each group to follow its own arbitrary evolution over time, mitigating concerns that treatment timing may be correlated with unobserved group-specific dynamics.

Figure H.1 displays the results, together with estimates from our baseline specification. The resulting estimates are virtually unchanged relative to the baseline, both in magnitude and statistical significance. The robustness of the point estimates to the inclusion of these highly flexible controls provides further confidence that our results are not driven by differential underlying trends.

Figure H.1: Robustness to covariate-specific non-parametric trends

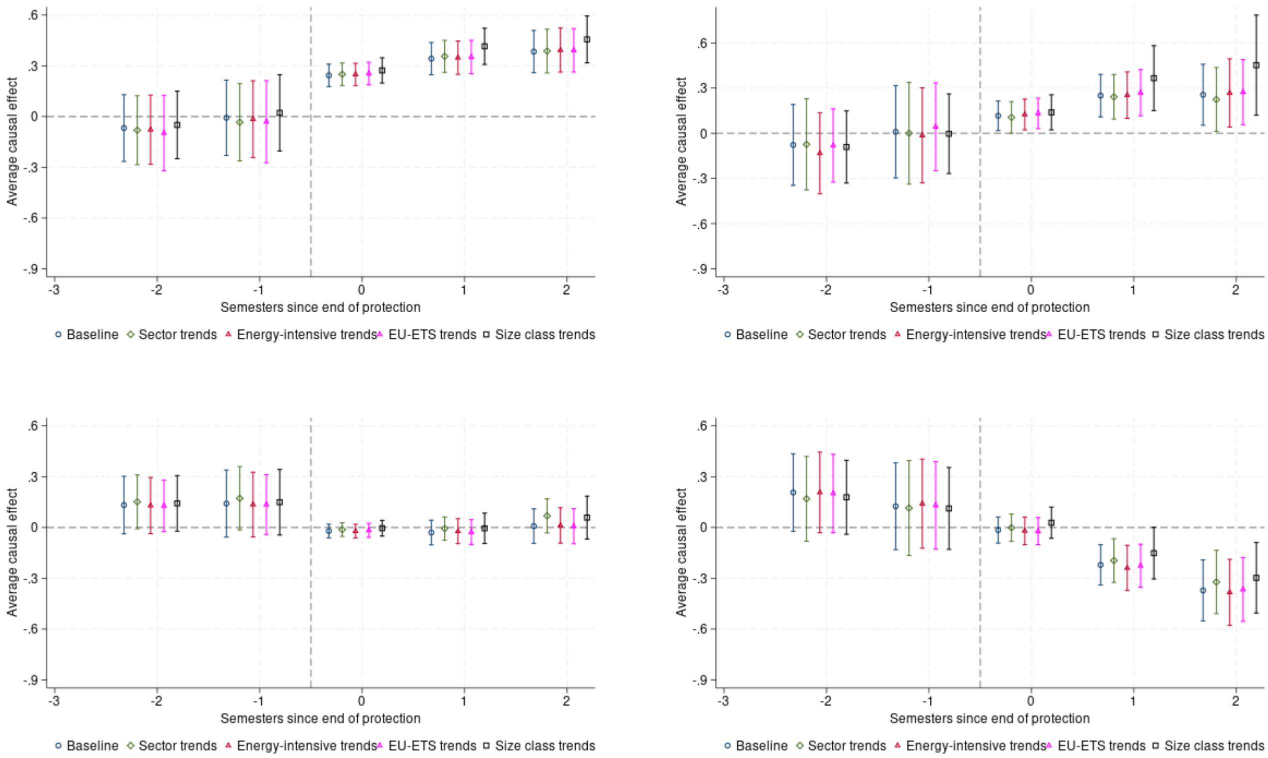

Notes: The figures show average causal effects of the expiration of a fixed-price contract on the average unit cost of electricity and natural gas (panels (a) and (b)) and the corresponding demanded quantities (panels (c) and (d)). Outcome variables are always in logs. Average causal effects before and after the treatment are estimated in two separate regressions, using the "imputation" estimator by Borusyak et al. (2024), as described in Section 4. Blue circles indicate our baseline estimation, identical to the one reported in the main text. Green diamonds include the interaction between calendar time fixed effects and 11 sector dummies. Red triangles include the interaction between calendar time fixed effects and an *energy-intensive* firm dummy. Magenta triangles include the interaction between calendar time fixed effects and a dummy for whether the firm owns a regulated plant under the EU ETS. Black squares include the interaction between calendar time fixed effects and firm size class dummies. Firm size is classified into five categories based on the number of employees: 50-99, 100-199, 200-499, 500-999, and 1,000 or more employees. Confidence intervals are at the 95% level.

Robustness to firm-specific seasonality

To check that our results are not driven by firm-specific seasonal patterns, we add firm × semester fixed effects to the first-stage regression of the Borusyak et al. (2024) estimator. These fixed effects absorb any time-invariant, within-firm difference between the first and second semester of the year, for instance, a firm that persistently consumes more electricity in summer, and are thus more demanding than the calendar-time fixed effects included in the baseline.

A mechanical limitation arises from the short panel. With firm × semester fixed effects, the BJS estimator can impute a unit's H1 (H2) fixed effect only from its untreated H1 (H2) observations. Early-treated firms (E_i = 2021H2) are treated in 2021H2 and 2022H2, every H2 period in the sample, so their firm × H2 fixed effect cannot be imputed. They are therefore excluded at the τ_0 and τ_2 horizons (both H2 periods), but remain in the estimation sample at τ_1 (an H1 period), where 2021H1 serves as their untreated H1 observation. Separately, τ_2 is suppressed for all cohorts because it falls outside the observation window. These two restrictions together account for the sample reduction relative to the baseline (electricity: N = 1,300 vs. 1,652; gas: N = 976 vs. 1,232).[54]

Figure H.2 displays the results alongside the baseline estimates. Across all four outcomes, the τ_0 and τ_1 point estimates are stable in magnitude and the confidence intervals overlap closely with the baseline, confirming that seasonality does not account for our findings.

[54]Mid-treated firms (E_i = 2022H1) and late-treated firms (E_i = 2022H2) retain at least one untreated observation per semester type and are therefore included, though with a single untreated observation per semester per firm in some cases.

Figure H.2: Robustness to firm-specific seasonality

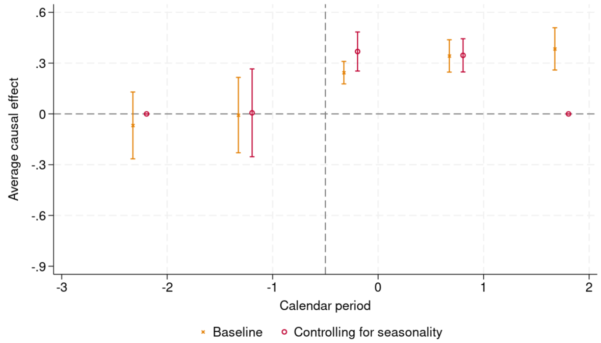

(a) Average cost of electricity

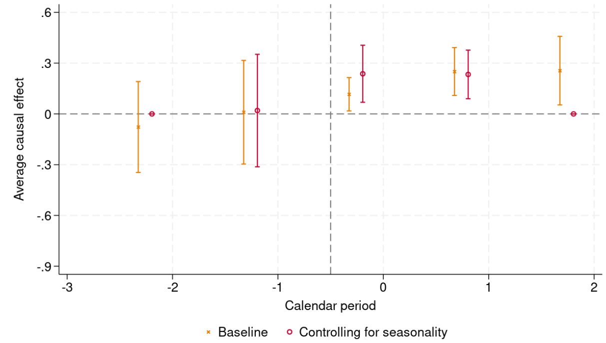

(b) Average cost of natural gas

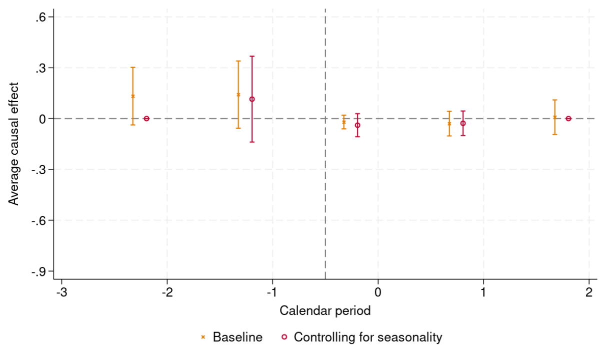

(c) Quantity of electricity

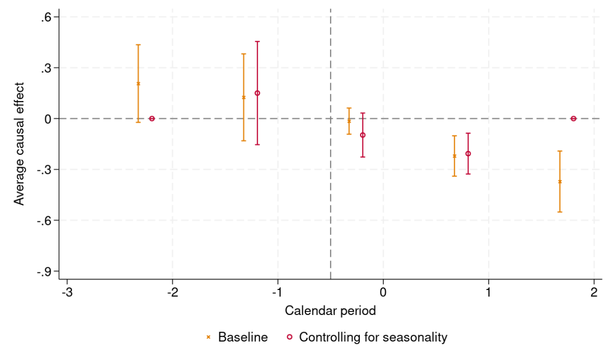

(d) Quantity of natural gas

Notes: The figures replicate the baseline event-study plots (blue circles, identical to Figure 5.1) and overlay estimates from a specification that adds firm × semester fixed effects to the first-stage regression (orange circles). Outcome variables are in logs. Average causal effects are estimated using the imputation estimator of Borusyak et al. (2024). Because early-treated firms (E_i = 2021H2) have no untreated observations in the same semester as their treatment periods, they are excluded from the seasonality-controlled sample; the τ_2 horizon is suppressed for all cohorts as it falls outside the observation window. Sample sizes: electricity N = 1,652 (baseline) vs. N = 1,300 (seasonality); gas N = 1,232 vs. N = 976. Confidence intervals are at the 95% level.

Appendix I Additional results

Figure A.1: The effect of the expiration of a fixed-price contract on average prices and quantities of energy inputs at the firm level, with and without firms that use derivatives or other instruments for protection.

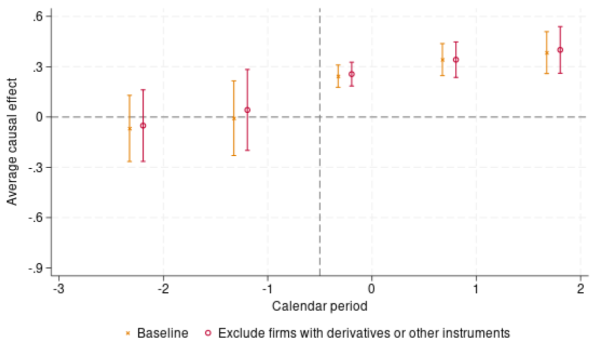

(a) Average cost of electricity

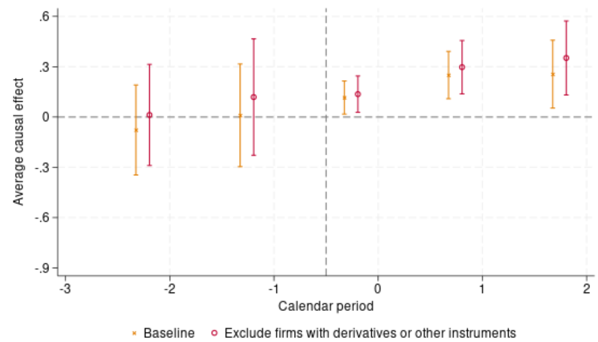

(b) Average cost of natural gas

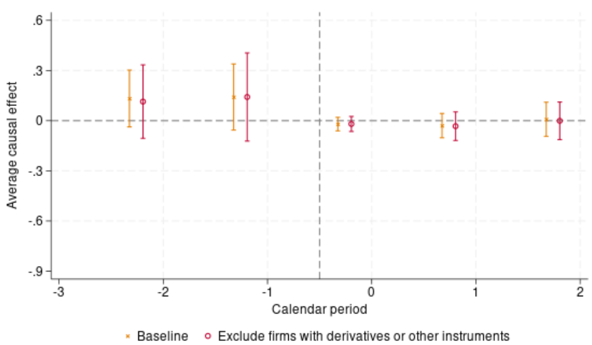

(c) Quantity of electricity

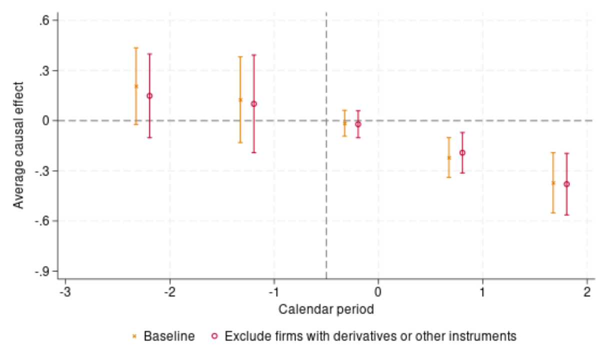

(d) Quantity of natural gas

Note: The figures show average causal effects of the expiration of a fixed-price contract on the average costs of electricity and natural gas (panels (a) and (b)) and the corresponding demanded quantities (panels (c) and (d)). Analysis is carried out separately on the overall sample and on the sample of firms not using derivatives or other financial instruments. All estimates are weighted by survey weights. Standard errors are clustered at the firm level. Confidence intervals are at the 95% level.

Figure A.2: Baseline results with different diff-in-diff estimators

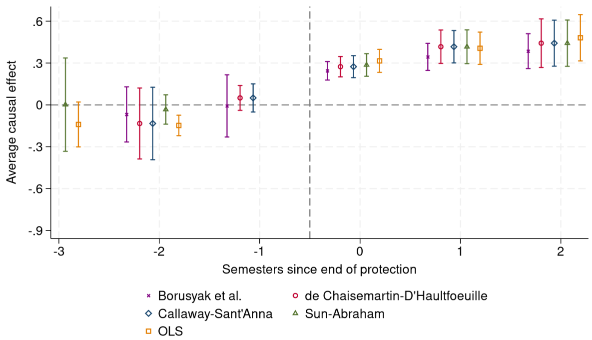

(a) Average cost of electricity

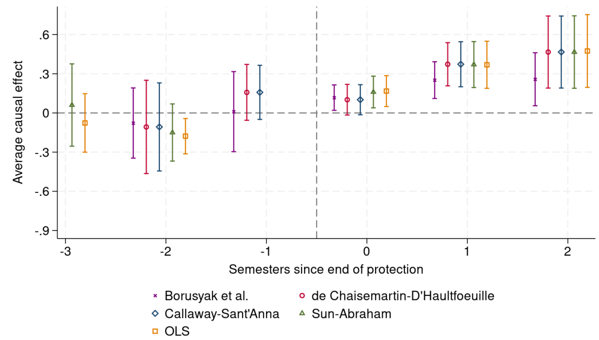

(b) Average cost of natural gas

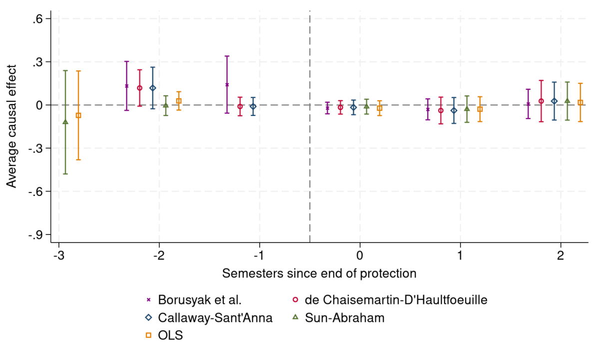

(c) Quantity of electricity

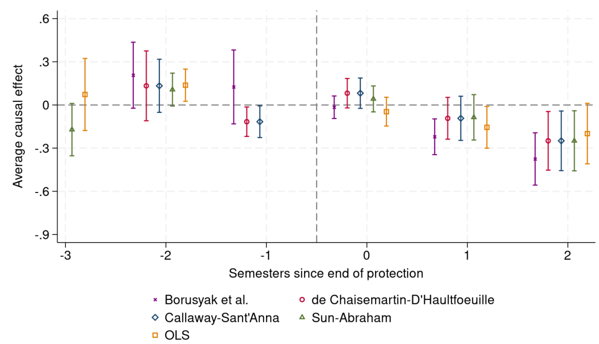

(d) Quantity of natural gas

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.3: Natural gas: heterogeneous effects by cohort

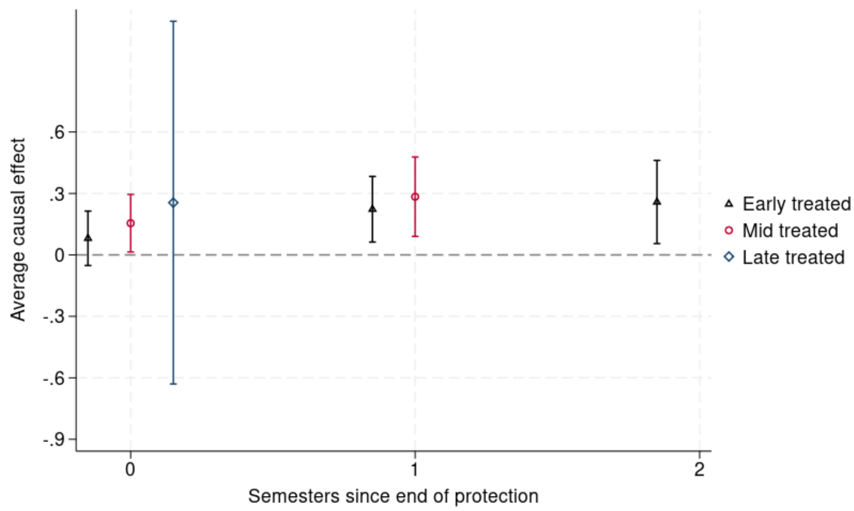

(a) Average cost

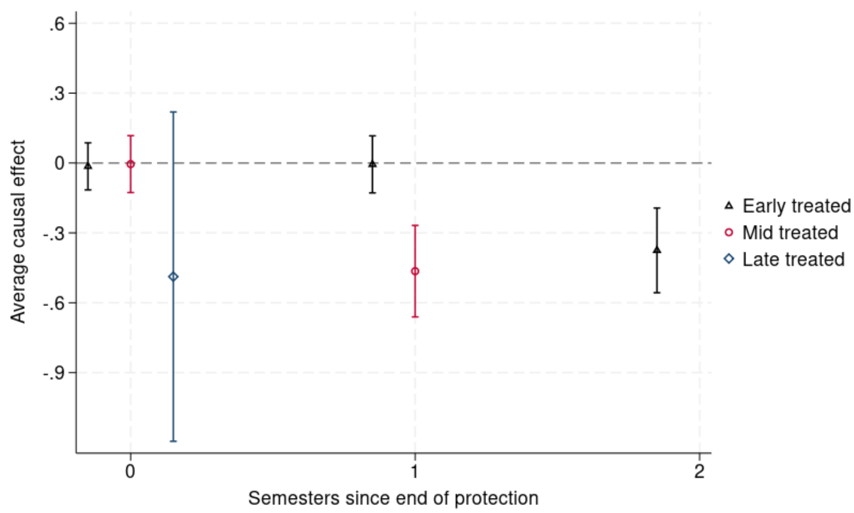

(b) Quantity

Note: the estimates in the upper and lower panel are from the same regression as in panel (b) and (d) of Figure 5.1 respectively; here they are reported by each cohort separately. All estimates are weighted by survey weights. Standard errors are clustered at the firm level. Confidence intervals at the 95% level.

Figure A.4: Output price growth across cohorts of treatment

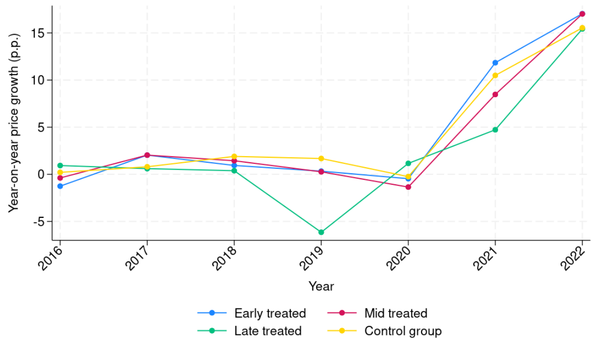

(a) Energy-intensive firms in electricity sample

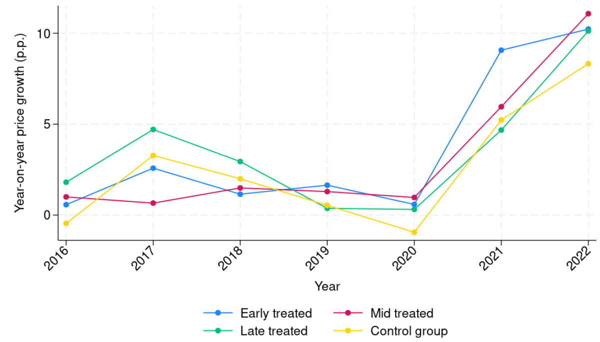

(b) Other firms in electricity sample

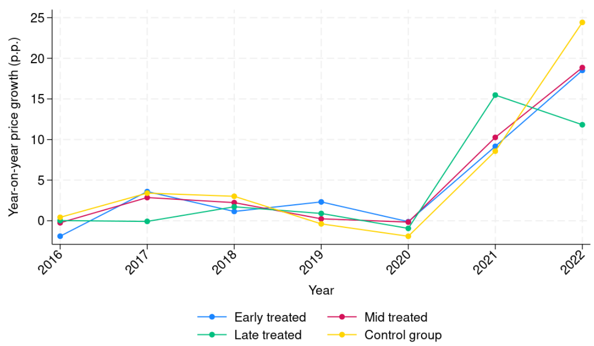

(c) Energy-intensive firms in gas sample

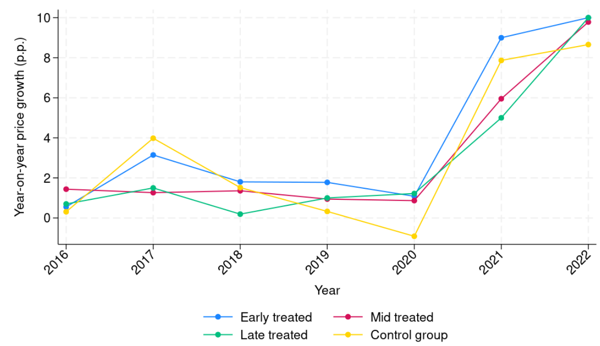

(d) Other firms in gas sample

Note: The figures show averages of yearly growth rate in the prices of goods sold from the Invind survey across the different treatment cohorts, separately for electricity and gas samples, and for energy-intensive and other firms. Averages are weighted by survey weights

Appendix J Background pictures

Figure A.1: Expectations on the economic situation

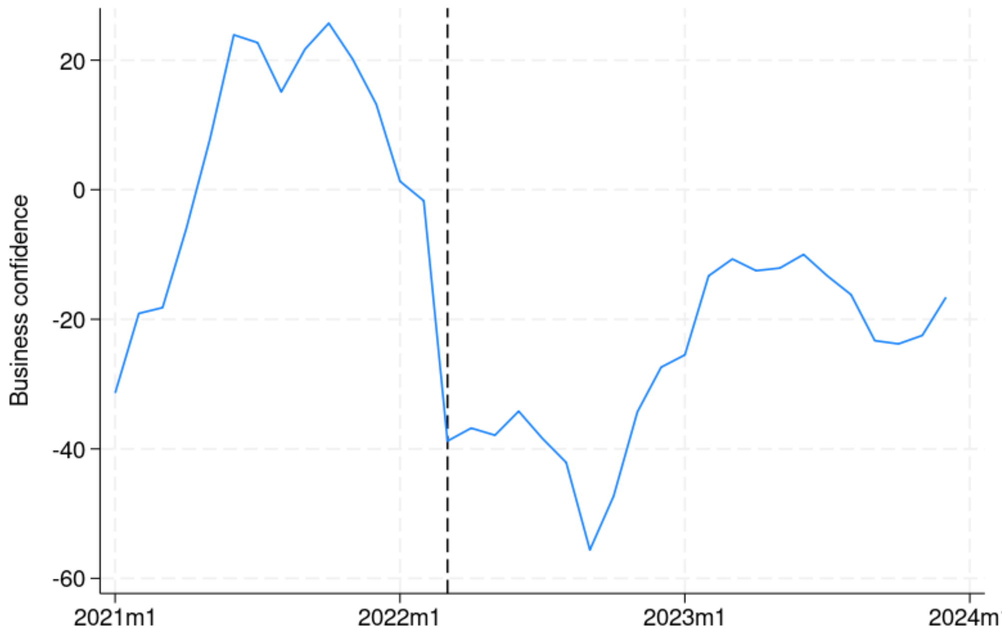

Note:The figure shows the difference between the share of manufacturing firms expecting positive economic situation in the next 3 months and those expecting a negative economic situation. The question refers to the general economic situation of the country. The vertical line refers to March 2022, the first month after Russia invaded Ukraine. Source: Istat.

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

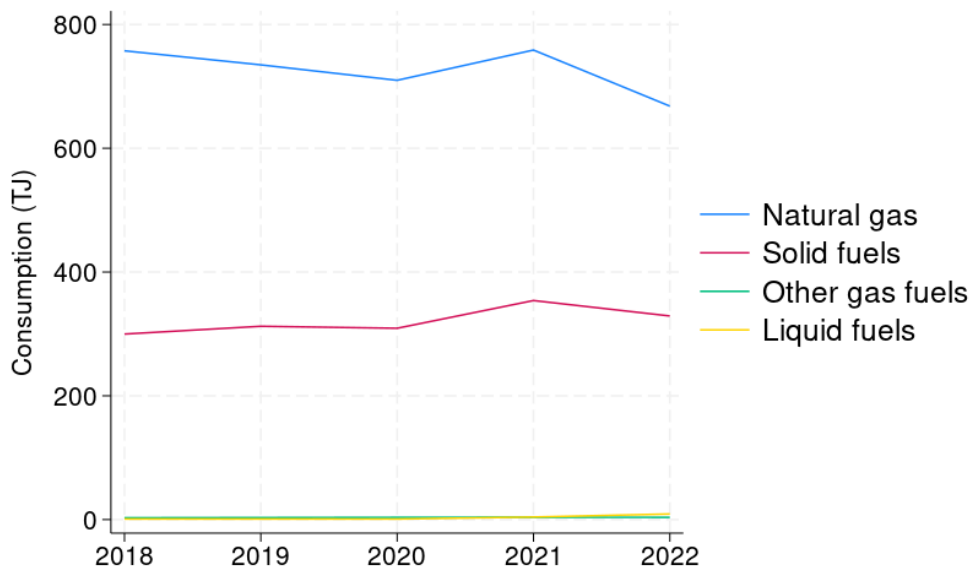

Note: The figure displays the average consumption of different fossil fuels (measured in Terajoules, TJ) for firms belonging to the matched Invind-EU ETS sample.

Figure A.3: Gas substitution in 2022 - Qualitative question

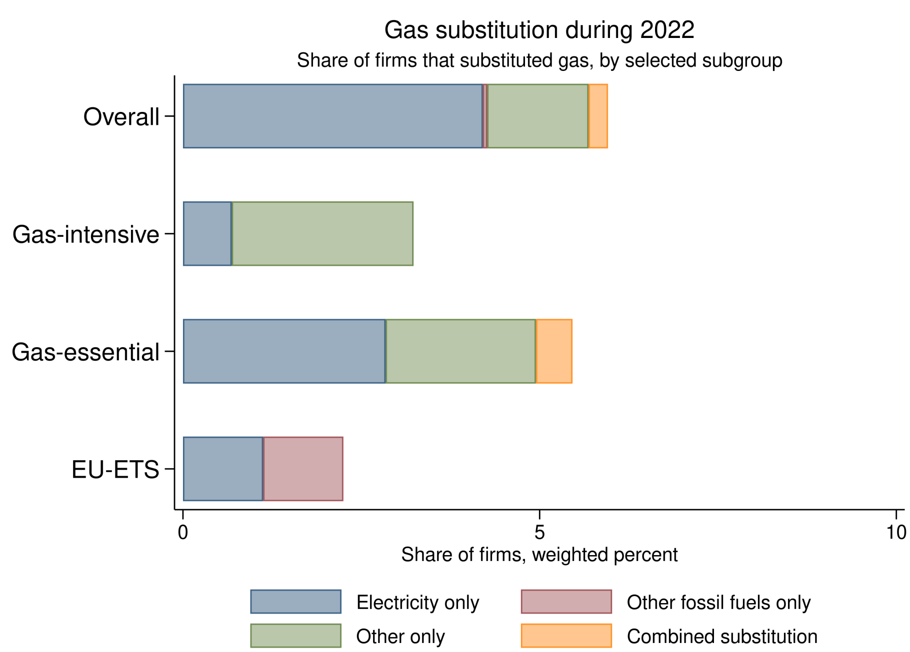

Source: Invind. *Note*: we restrict to the gas sample used for our empirical analysis. The complements to the depicted bars declare that they did not substitute at all.

Appendix K Local vs. arc-elasticities under non-isoelastic demand

Let $Q(P)$ be a smooth, strictly decreasing demand function with local elasticity

$$\epsilon(P) \;\equiv\; \frac{dQ}{dP}\,\frac{P}{Q(P)}. \tag{15}$$

For an isoelastic function, ϵ is constant and the proportional quantity change to any finite price movement is simply $\epsilon \times (\Delta \log P)$. For non-isoelastic demand, the quantity response to a finite price change from P_0 to P_1 is

$$Q(P_1) - Q(P_0) \;=\; \int_{P_0}^{P_1} \epsilon(P)\,\frac{Q(P)}{P}\,dP, \tag{16}$$

a weighted average of local elasticities along $[P_0, P_1]$. When $\epsilon(\cdot)$ varies over this interval, the arc-elasticity implied by a large shock generally differs from the local elasticity at P_0.

Example: demand with a subsistence floor. Consider

$$Q(P) = q_{\min} + A\,P^{-k}, \qquad q_{\min} > 0,\; A > 0,\; k > 0.$$

Demand is smooth, strictly decreasing, and its local elasticity is

$$\eta(P) = -\,\frac{k\,A\,P^{-k}}{q_{\min} + A\,P^{-k}} \;\in\; (-k,\,0).$$

As P rises, the subsistence floor $q_{\min}$ dominates and $|\epsilon(P)|$ shrinks toward zero: demand becomes progressively less responsive. A small-shock estimate obtained around a low initial price P_0 would therefore overstate the quantity reduction achievable at crisis-level prices, where firms are constrained near their operational minimum. This pattern is consistent with heavy-industry settings where furnaces must maintain minimum throughput to avoid irreversible damage, and it motivates our direct estimation of the arc-elasticity over the observed price range.

Appendix L A simple static input choice model

Here we present a simple static model of input choice, following the treatment in Cahuc et al. (2014) to show that whether demand for gas for gas-intensive firms is more or less elastic relative to other firms is theoretically ambiguous. Consider a firm that operates with a homogeneous production function that combines natural gas and a bundle of other inputs. Assume the production function is homogeneous of degree θ and that the firm faces an iso-elastic downward sloping demand. It can be shown that the uncompensated own-price elasticity of demand for natural gas (the parameter estimated in our paper) can be written as:

$$\eta^G_{p_G} = -(1-s)\sigma - \frac{\mu}{\mu-\theta}s, \tag{17}$$

where s is the share of gas expenditures in total costs, σ is the elasticity of substitution between natural gas and other inputs, μ is the markup. The first term is the substitution effect (holding output fixed) and the second term the scale effect, both working in the same direction. For energy-intensive firms, the scale effect is more relevant; for other firms, the substitution effect is more relevant.

According to our surveys, almost all gas-intensive firms declare gas to be an essential input, while the share is lower among remaining firms. We can thus assume that σ is heterogeneous in energy-intensity. In particular, let us assume that $\sigma = 0$ among gas-intensive firms, and $\sigma > 0$ for other firms. Under this assumption, the demand elasticity for the two groups reads:

$$\eta^G_{p_G} = \begin{cases} -\frac{\mu}{\mu-\theta}s^{EI}, & \text{if } \sigma = 0, \text{ energy intensive (EI)}, \\ -(1-s^{NEI})\sigma - \frac{\mu}{\mu-\theta}s^{NEI}, & \text{if } \sigma > 0, \text{ non energy intensive (NEI)}, \end{cases} \tag{18}$$

where $s^{EI} > s^{NEI}$ by definition. The demand for gas among gas intensive firms is less elastic relative to that of other companies if the elasticity of substitution among the latter group is sufficiently high:

$$\sigma > \frac{\mu}{\mu-\theta}\frac{(s^{EI}-s^{NEI})}{1-s^{NEI}}. \tag{19}$$

Appendix M Effect heterogeneity with machine learning

In order to study heterogeneity in treatment effect of fixed-price contract expiration 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 sub-samples: a learn sub-sample (60 per cent of the overall sample) and a test sub-sample (40 per cent). Second, we build a forest of 5,000 trees using the learn sub-sample only. Each tree can pick only a random half of the considered covariates, which are: a (self-reported) dummy for whether gas is an essential input in production[55], which can be thought as a proxy of low elasticity of substitution between gas and other inputs; a gas intensive dummy[56]; sector dummies; an EU ETS dummy; and employment in 2021. 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 A.1 plots the distribution of individual treatment effects as estimated by the Borusyak et al. (2024) 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.[57] 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., 2024), as individual treatment effects contain the error term ϵ_{ijt} in equation 4. 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.[58] 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 to zero. This means that the forest predicts that some treated firms would decrease gas consumption by as much as 60%, 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 A.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;

[55]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.

[56]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.

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

[58]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.

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.[59] The results are confirmed in the test sample when using as an outcome variable the out-of-sample-ML-predicted treatment effects (Table A.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. (2024) 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. (2024) estimates in the test sample were never used to train the forest. Results are presented in Table A.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.

[59]The above-average treatment effect for wood and paper is consistent with greater scope for substitution toward biomass fuels; see Section 5.f for evidence on fuel switching in the EU ETS subsample.

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

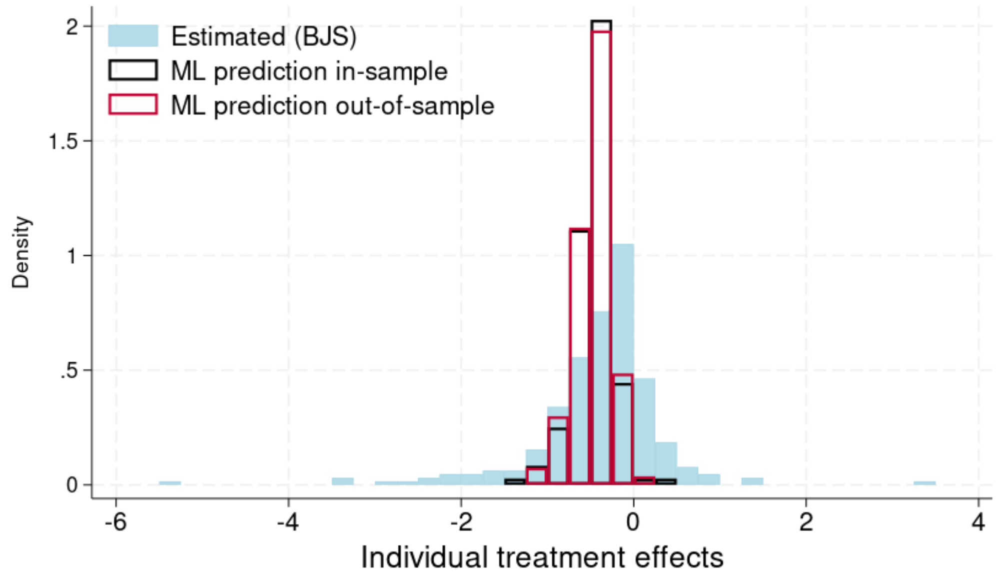

Note: The figure shows the distribution of individual treatment effects in 2022h2. The outcome is natural gas consumption.

Table A.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) are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.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) are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.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. (2024).

Appendix N Price elasticity of energy demand

Throughout the paper, we compute the price elasticity of energy demand by combining our estimates of the effect of fixed-price contract expirations on (log) prices and on (log) quantities obtained from the Borusyak et al. (2024) estimator. In practice, we compute a Wald-type elasticity, where we scale the average treatment effect estimated in the quantity equation by the average treatment effect estimated in the price equation.[60] Under standard IV-LATE assumptions, we can construct an estimate of the elasticity θ as follows:

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

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.

We construct standard errors using the delta method. The 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. (2024).

[60] 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.

The variance of the elasticity is thus

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

Figure A.1 plots the estimated price elasticity of demand by calendar period for the overall Invind sample, and for selected sub-sample.

Figure A.1: Price elasticity of demand by calendar period

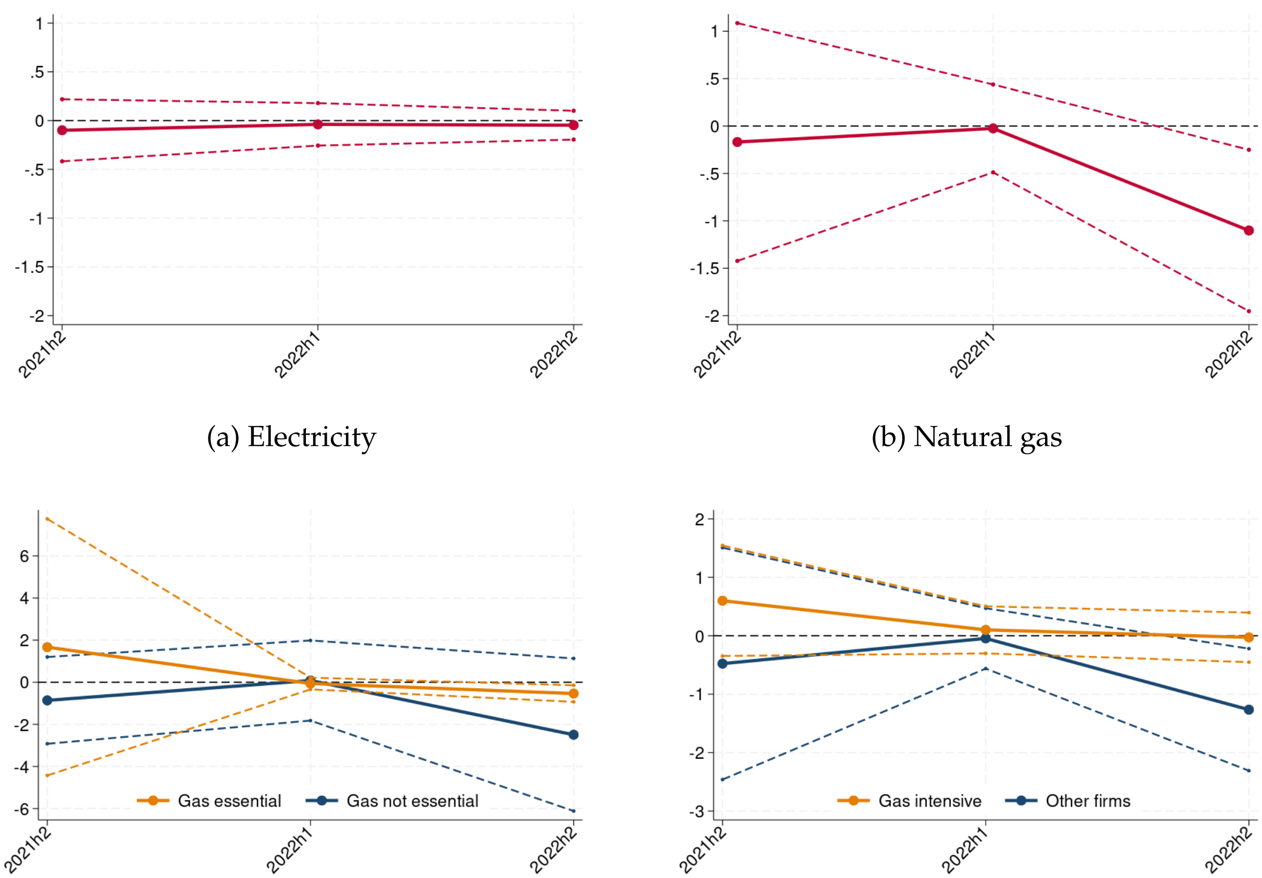

(a) Electricity

(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 5.1; 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.

Appendix O Alternative policy instruments

Per-unit subsidies and pay-for-reductions rewards were not the only policy instrument available to governments during the crisis. We now consider two other alternatives using the same calibration: targeted lump-sum transfers earmarked to baseline consumption and retail price caps.[61]

Targeted lump-sum transfers. A lump-sum transfer provides income support without changing the marginal price of gas. To make its fiscal scale comparable to the per-unit policies, let group g receive a transfer

$$T_g(\tau_g) = \tau_g Q_{g0}, \tag{20}$$

where τ_g is expressed in euros per unit of baseline gas consumption and Q_{g0} is fixed at the pre-policy level. In the calibration below we set $\tau_g = 0.50$ € /Smc, the same rate used for the per-unit subsidy, but the transfer is anchored to baseline rather than current consumption. The market-clearing condition is unchanged, so

$$\frac{dP}{dT_g} = 0, \qquad \frac{dQ}{dT_g} = 0.$$

Lump-sum transfers are therefore the natural benchmark when the objective is compensation while minimizing distortions in gas demand. Their limitation is equally clear: because they do not increase the marginal cost of gas, they cannot actively induce conservation.

Retail price caps. A retail price cap fixes the consumer price paid by a designated group at an administratively chosen level $\bar{p}_a$. If the cap binds, demand by the capped group depends on $\bar{p}_a$ rather than on the producer price P. Let b denote the uncapped group. Market clearing is

$$BP^{\eta} = A_a \bar{p}_a^{-\varepsilon_a} + A_b(\alpha_b + P)^{-\varepsilon_b}. \tag{21}$$

[61] These stylized examples mimic actual policies implemented by EU governments during the crisis: Germany's *Gaspreisbremse* capped the gas price at €0.07/kWh for industrial consumers, but only up to 80 % of each firm's *previous year's consumption*; consumption above this baseline faced prevailing market prices. France's *bouclier tarifaire* froze regulated gas tariffs from October 2021 and capped electricity price increases at 4% in early 2022 – a retail price cap applying to all consumption. See Sgaravatti et al. (2021) for a cross-country survey of national policy designs.

Differentiating with respect to the cap gives

$$\frac{dP}{d\bar{p}_a} = -\frac{\varepsilon_a \dfrac{Q_a}{\bar{p}_a}}{S' + D_b'}. \tag{22}$$

Lowering the cap increases demand from the capped group and raises the producer price. The denominator contains the supply slope (S') and the demand slope of the uncapped group only (D_b'). The capped group no longer helps equilibrate the market through higher wholesale prices, because its retail price is fixed. This makes a cap more distortionary than an equivalent per-unit subsidy: for a given reduction in the price paid by the protected group, the producer-price response is larger because fewer users remain price-sensitive.

If the government compensates suppliers for the difference between the market retail price and the cap, fiscal outlays are

$$C = (\alpha_a + P - \bar{p}_a)Q_a. \tag{23}$$

The fiscal cost is endogenous to the equilibrium price response: tightening the cap both widens the wedge between the market price and the capped price and raises the producer price itself.

Table A.1 summarizes the calibrated comparison across instruments. On top of the considerations made in the main body of the paper, we see that price caps are the least attractive support instrument in this framework because they combine support with strong equilibrium price feedbacks. Under a uniform retail price cap, the price of natural gas increases by 62%, while quantities do so by 12%. The government spends almost €8 billion.

Table A.1: Comparison of policy instruments

Instrument	**Target**	**$\Delta P/P$** (%)	**$\Delta Q/Q$** (%)	**Fiscal cost** (€ bn)
Support instruments				
Per-unit subsidy	Uniform	+26.3	+5.3	3.53
	GI only	+2.2	+0.4	2.71
	NGI only	+24.1	+4.8	0.85
Lump-sum transfer	Uniform	0	0	3.35
	GI only	0	0	2.68
	NGI only	0	0	0.67
Retail price cap	Uniform	+62.1	+12.4	7.90
	GI only	+2.3	+0.5	2.82
	NGI only	+51.0	+10.2	1.42
Conservation incentives				
Pay-for-reductions (PFR)	Uniform	−26.3	−5.3	0.18
	GI only	−2.2	−0.4	0.03
	NGI only	−24.1	−4.8	0.18
Policy mix				
PFR to NGI + subsidy to GI		−21.9	−4.4	2.93

Notes: All instruments are scaled to €0.50/Smc. Lump-sum transfers are based on baseline consumption; pay-for-reductions payments are based on verified reductions; price-cap payments compensate the gap between the capped and market price. The policy mix consists of a per-unit subsidy for gas-intensive firms and a pay-for-reduction payment to non-gas-intensive firms. Fiscal costs are in € billions.