Supporting firms in energy crises: lessons from 2021–22 *

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

[a]Bank of Italy

October 2026

Abstract

Government support during energy crises shields firms from price shocks but weakens incentives to conserve scarce energy. When demand responses differ across firms, the equilibrium effects of support depend on the consumption-weighted average price elasticity of demand. We estimate this parameter by exploiting the staggered expiration of fixed-price contracts among Italian industrial firms in the 2021–22 crisis. Electricity demand does not respond to price increases, while gas consumption falls only in the second half of 2022. Gas-intensive firms, which account for 80% of industrial gas use, have almost perfectly inelastic demand; other firms have an elasticity of −1.3. Embedding these estimates in a simple incidence framework, we compare uniform and targeted subsidies, price caps, and voluntary demand-reduction incentives. The average-firm elasticity overstates the price and quantity effects of a uniform subsidy by about 45%. Subsidizing gas-intensive firms while offering demand-reduction incentives to other firms would reduce gas consumption by 4.4%, whereas the average-firm elasticity predicts an increase.

Keywords: Energy crisis, price elasticity of energy demand, incidence, targeted subsidies

*A previous version which analyzed data on 2021 only was circulated under the title "The effects of the 2021 energy crisis on medium 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, Maximilian Auffhammer, Fabio Bertolotti, Marco Bottone, Federico Cingano, Lucas Davis, Guido de Blasio, Antoine Dechezleprêtre, Domenico Depalo, Alberto Felettigh, Sara Formai, Yuriy Gorodnichenko, Rossella Greco, Andrea Lamorgese, Andrea Linarello, Francesca Lotti, Giovanni Marin, Chiara Motta, Matteo Paradisi, Enrico Sette, Joseph Shapiro, Francesco Vona, Nils Wehrhöfer 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. Luca Citino is grateful to UC Berkeley for hosting him during the revision of this paper. 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

When wholesale energy prices rose tenfold in 2021–22, European governments allocated more than half a trillion euros to shield households and firms from rising energy costs, through support measures such as bill subsidies, VAT and excise cuts, tax credits, and retail price caps (Sgaravatti et al., 2021; McWilliams et al., 2024).[1] After the 2026 disruption in the Strait of Hormuz, governments again introduced sizeable support, much of it untargeted (McWilliams et al., 2026). Such support involves a trade-off: it cushions firms against large cost shocks, but it also weakens their incentive to save energy when it is scarce, and by sustaining demand it puts further upward pressure on energy prices (Gros, 2022a; Signorini, 2022; Zettelmeyer et al., 2022). Targeting can soften this trade-off by restricting support to fewer firms, and thus international institutions increasingly call for selective support.[2] How large is this trade-off, and how much can targeting reduce it?

The answer hinges on firms' price elasticity of energy demand, but the estimates available in the literature (Labandeira et al., 2017) are ill suited in this context, for two reasons. First, elasticities estimated in normal times may not carry over to large shocks. Crisis-sized price increases can make demand more elastic, because they are salient and draw managerial attention to neglected adjustments,[3] or less elastic, because once the cheapest margins are exhausted, essential uses are harder to compress. Second, crisis responses may differ across firms, so the average-firm elasticity may be a poor guide to how market demand responds to support. This is particularly important in energy markets where the distribution of gas and electricity consumption is highly skewed, even within narrowly defined industries (Lyubich et al., 2018). Despite their importance, we know very little about the such elasticities in times of crisis and more generally about how they vary across firms.

In this paper we provide new estimates on these elasticities by studying how the electricity and natural gas demand of Italian industrial firms responded to the 2021–22 crisis, the most severe energy shock in Europe since the oil crises of the 1970s.[4] We then use these estimates

[1] Spain and Portugal were a notable exception, intervening *directly* in the wholesale electricity market through the so-called Iberian mechanism. Fabra et al. (2025) study the market and distributional effects of this alternative policy response.

[2] The European Central Bank argues that shrinking fiscal space makes selectivity "not just desirable but necessary" (Lagarde, 2026); the European Commission's 2026 crisis package calls for measures that are "timely, targeted and temporary" (European Commission, 2026); and the IMF argues that support should be "well targeted" and should "not blunt the incentive to reduce energy consumption" (International Monetary Fund, 2026).

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

[4] The European Commission describes the crisis that followed Russia's invasion of Ukraine as "comparable to the oil crises in the 1970s" (European Commission, 2024); the International Energy Agency calls it "the first truly global energy crisis" and "a shock of unprecedented breadth and complexity" (International Energy Agency,

to quantify how large the trade-off between shielding firms and preserving their incentive to save energy is, and how much targeting can reduce it.

To fix ideas, we start from a simple incidence framework that relates different targeted and untargeted support measures to firms to equilibrium energy 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, in the spirit of Deryugina et al. (2020) and Hahn and Metcalfe (2021). The framework delivers two results. First, for untargeted measures, what matters is the consumption-weighted elasticity, which weights the elasticity of each firm by its share of aggregate consumption. Second, for targeted measures *both* the elasticity of eligible firms and the consumption-weighted elasticity matter separately. Intuitively, targeted support raises demand among eligible firms, but also raises the market price, which in turn reduces the equilibrium consumption of *all* firms in the market, even those that do not receive support. Thus, three empirical inputs determine how support is transmitted to equilibrium prices and quantities: group-specific demand elasticities, each group's share of aggregate consumption, and the supply elasticity.

To estimate demand elasticities, we use a new dataset that combines survey responses with detailed administrative data. We designed a dedicated energy module for the Bank of Italy's survey of industrial and service firms (Invind), which collected half-yearly electricity and gas expenditures and physical consumption for 2021 and 2022 from industrial firms with at least 50 employees.[5] These data allow us to measure average unit costs at the firm level, separately for electricity and natural gas. Because the survey also records the tax credits firms received, we can measure the average unit cost each firm actually paid, net of government support. We view this as the policy-relevant price, because crisis support lowered both the fixed and the per-unit components of energy bills, so a measure of marginal tariffs alone would miss part of the relevant effect.[6] Linking the survey to administrative records on energy-intensive firms and on plants covered by the EU Emissions Trading System (ETS) allows us to validate the survey measures and to study substitution between fuels.

Identifying firms' energy demand responses during a crisis is challenging because wholesale prices rise simultaneously for all firms. A large literature exploits time-series variation to study the aggregate effects of energy shocks (Kilian, 2009; Känzig, 2021; Alessandri and

2022).

[5]Our sample includes only firms for which energy is an input to production. We therefore leave out the energy sectors themselves: manufacture of coke and refined petroleum products (NACE 19) and electricity, gas, steam and air conditioning supply (NACE 35).

[6]Evidence from residential electricity consumers also suggests that energy users respond to average rather than marginal prices (Ito, 2014).

Gazzani, 2023; Emiliozzi and Favero, 2025). While these approaches are well suited to studying aggregate responses, they do not directly identify the *distribution* of firm-level demand elasticities needed to assess the effects of targeted support. Firm-level studies instead exploit differential exposure to common energy-price shocks, typically measured by the share of energy in total costs (Ganapati et al., 2020; Marin and Vona, 2021). However, because treatment doses are determined by energy intensity, these designs make it difficult to separately identify how price elasticities vary with firms' energy dependence (Callaway et al., 2024), which is precisely the dimension along which support is often targeted.

In the ideal experiment, some firms would face a large increase in energy costs while others would be shielded from it, for reasons unrelated to their potential demand response during the crisis. We approximate this experiment by exploiting the staggered expiration of fixed-price energy contracts. In our survey, firms report whether they were protected by fixed-price contracts signed well before the crisis began and for how long protection lasted in 2022. Because expiration dates were set when contracts were signed, and expiring contracts could not be renewed on pre-crisis terms, we can compare firms that lose protection during our sample period with otherwise similar firms whose protection expires later or not at all during the observation period. We estimate the effects of contract expiration on log unit costs and log quantities with the event-study estimator of Borusyak et al. (2024), and compute elasticities as Wald-type ratios of the percent change in quantities to the percent change in unit costs. The results are robust to alternative estimators for staggered treatment timing (De Chaisemartin and d'Haultfoeuille, 2020; Callaway and Sant'Anna, 2021; Sun and Abraham, 2021), to covariate-specific non-parametric time trends by industry, energy intensity, firm size, and ETS participation, and to synthetic difference-in-differences (Arkhangelsky et al., 2021), which reweights control firms and pre-treatment periods to match the pre-treatment trajectories of treated firms.

We find that, on average, contract expiration raises firms' unit costs of energy by up to 45 percent for electricity and 30 percent for natural gas. Quantity responses, however, differ sharply between the two fuels. Firms do not adjust their electricity consumption. By contrast, they reduce natural gas consumption by 34 percent, but only in the second half of 2022, irrespective of when protection had ended. The implied *average-firm* elasticity of gas demand is close to zero in the second half of 2021 and the first half of 2022, and reaches −1.1 in the second half of 2022. The timing coincides with increasingly pessimistic market expectations in the summer of 2022, when spot gas prices reached record highs and futures markets projected prices of around €200 per MWh well into 2023, which suggests that firms undertook costly adjustments only once they stopped perceiving the shock as temporary. Energy demand is

thus very inelastic in the first months of the crisis, and gas demand becomes elastic only once the shock appears persistent.

This average, however, masks substantial heterogeneity. Using random forests to detect treatment effect heterogeneity, we show that gas-intensive companies reduced gas consumption significantly less than other firms. Confidential administrative data on monthly gas consumption, which are available only for gas-intensive firms, confirm that their response is small. In the second half of 2022, the elasticity is −0.03 for gas-intensive firms, compared with −1.3 for other firms. This difference appears to reflect a low elasticity of substitution between gas and other inputs among gas-intensive firms. Two pieces of evidence support this interpretation. First, nearly all gas-intensive producers classify natural gas as an *essential* input in the survey. Second, plant-level data for EU ETS facilities, which include the largest industrial gas consumers, reveal little switching from gas to other fossil fuels. As industrial gas consumption is highly concentrated, such heterogeneity in elasticity matters for aggregation. Weighting group elasticities by their share of aggregate consumption gives an elasticity of only about −0.3, roughly a quarter of the average-firm elasticity of −1.1.

In the last part of the paper, we return to the framework and use our estimates to quantify, in partial equilibrium, how alternative support policies affect gas prices, consumption, and fiscal costs in the second half of 2022. We first consider an untargeted subsidy of €0.50 per cubic meter of natural gas, about half the price in the first half of 2022. The predicted effects depend heavily on which elasticity is used. With the consumption-weighted elasticity, the subsidy raises gas consumption by 5.3 percent and the producer price by 26 percent. With the "naive" average-firm elasticity, the same subsidy would raise them by 7.7 and 38 percent, overstating both effects by about 45 percent. Second, we consider a targeted subsidy of the same size, but restricted to gas-intensive firms. In this case, the subsidy raises consumption by only 0.4 percent and the price by about 2 percent. The targeted subsidy still costs €2.7 billion, about three-quarters of the €3.5 billion spent on the uniform subsidy, but almost all of it is a transfer on consumption that would have occurred anyway and thus does not create distortions through additional demand. We then turn to considering other policies. We show that voluntary demand reduction incentives[7] work best when restricted to non-gas-intensive firms, as their demand is elastic. A combination of voluntary demand reduction incentives for non-gas-intensive firms and subsidies to gas-intensive firms would still lower gas consumption by 4.4 percent, at a cost of €2.9 billion. Using the average-firm elasticity, the same combination would instead appear to raise gas consumption, by 4.6 percent, because gas-intensive firms

[7]For example, Germany's *Gasminderungsausschreibungen* and Italy's gas demand reduction auctions. See Sgaravatti et al. (2021) for a survey.

would then wrongly appear responsive to the subsidy. Finally, we consider retail price caps, whereby firms pay a fixed price for gas, and the government pays the difference between the capped price and the resulting new equilibrium price (received by producers). We find these to be more distortionary than subsidies: a uniform cap leads to changes in price and quantity that are twice the size of those resulting from a subsidy of the same generosity. For electricity, where demand does not respond, support works mostly as a transfer in all cases.

The paper makes two contributions. First, we provide causal estimates of firms' energy demand elasticities during an energy crisis and show how these elasticities vary across firms. Household responses are well documented,[8] including during the 2021–22 crisis. Closest to our work, Ahlvik et al. (2025) exploit the expiration of fixed-price contracts among Finnish households to show how low-income households are less responsive than high-income households. No corresponding evidence exists for industry, as previous firm-level estimates rely on the smaller price movements of normal times (von Graevenitz and Rottner, 2022; Fontagné et al., 2024; Gerster and Lamp, 2024; Blonz, 2022), while studies that look at recent energy crises estimate sectoral elasticities with respect to wholesale prices (Emiliozzi and Favero, 2025). Crisis support, however, acted on firms' *retail* prices rather than on wholesale prices, and evaluating targeted support requires knowing how responses vary across firms. Our new data and design allow us to observe the prices each firm actually paid and to estimate the elasticities that these policy evaluations require.

Second, we show how heterogeneity in quantity responses shapes the incidence of energy support. Building on the sufficient-statistics approach in public finance, we adapt standard incidence formulas to heterogeneous consumers and targeted policies. The effects of untargeted support depend on the consumption-weighted elasticity, whereas those of targeted support depend on *both* this elasticity and that of the targeted group. While the relevance of quantity-weighted elasticities for policy incidence has been noted by Chetty (2009) in the context of income taxation, we are the first to show that this matters quantitatively in the context of energy markets, where elasticities vary widely across consumers and energy use is highly concentrated (Lyubich et al., 2018). To the contrary, studies of energy demand, energy subsidies, and crisis support have typically summarized demand with a single elasticity (Deryugina et al., 2020; Davis, 2014; Gros, 2022b).[9]

The rest of the paper proceeds as follows. In Section 2 we present the incidence framework.

[8] See, among others, Reiss and White (2005), Jessoe and Rapson (2014), Auffhammer and Rubin (2018), Hahn and Metcalfe (2021), Costa and Gerard (2021), Byrne et al. (2021), and Deryugina et al. (2020).

[9] Relatedly, Martinsson et al. (2024) document heterogeneity in the carbon price elasticity of emission abatement and how this correlates with emission-intensity, but do not study its implications for policy incidence, as we do.

In Section 3 we describe the institutional setting and the source of variation, and in Section 4 the data, measurement of key variables, and estimation procedures. We report average effects in Section 5 and study heterogeneity, mechanisms, and aggregation in Section 6. In Section 7 we return to the framework to quantify the equilibrium effects of alternative support policies. Section 8 concludes.

2 Incidence with heterogeneous energy consumers

In this section we develop a stylized incidence framework in partial equilibrium that shows which empirical quantities are needed to evaluate energy-support policies when consumers are heterogeneous. For ease of exposition, in the following paragraphs we abstract from pre-existing per-unit wedges between producer and consumer prices, such as network charges and taxes, and from differences in baseline prices across groups, which are small in our empirical application.[10] In this section, we always express demand elasticities as positive magnitudes. In Section 2.a we describe the environment and derive a general incidence formula, which we apply to untargeted and targeted support in Sections 2.b and 2.c.

2.a Setup

Consider a competitive market for natural gas. On the demand side, industrial consumers belong to groups $g = 1, \ldots, G$. Group g consumes $Q_g = D_g(p_g)$, where p_g is the retail price, and has demand elasticity

$$\varepsilon_g \equiv -\frac{dD_g}{dp_g}\frac{p_g}{Q_g} > 0.$$

Aggregate demand is $Q \equiv \sum_g Q_g$, and

$$\omega_g \equiv \frac{Q_g}{Q}$$

denotes group g's share of aggregate consumption.

On the supply side, producers receive price P and supply $Q^S(P)$, with short-run price elasticity

$$\eta \equiv \frac{dQ^S}{dP}\frac{P}{Q^S} > 0.$$

Because industrial firms account for only part of gas demand, η should be read as the elasticity of residual supply facing industrial users, net of the responses of other consumers such as

[10] Allowing consumer prices to differ from P, or across groups, leaves the logic unchanged but replaces consumption weights with weights proportional to Q_g/p_g, where p_g is the baseline price paid by group g.

households and power generators.

Absent policy, every group pays the producer price, $p_g = P$. A per-unit subsidy s_g, which may differ across groups and be zero for some, drives a wedge between the two prices, so that group g pays

$$p_g = P - s_g. \tag{1}$$

The market clears when aggregate demand equals supply:

$$\sum_g D_g(p_g) = Q^S(P). \tag{2}$$

We start from an initial equilibrium without policy. Totally differentiating equation (2) and dividing by Q/P, a small policy change must satisfy

$$\eta \, dP = -\sum_g \omega_g \varepsilon_g \, dp_g, \tag{3}$$

so that the change in supply on the left matches the change in demand on the right.

Under a subsidy, the price a group pays moves one-for-one with the producer price: from equation (1), $dp_g = dP - ds_g$. Substituting into equation (3) separates two demand forces:

$$\underbrace{\eta \, dP}_{\text{supply response}} = \underbrace{\sum_g \omega_g \varepsilon_g \, ds_g}_{\text{direct demand impulse}} - \underbrace{\sum_g \omega_g \varepsilon_g \, dP}_{\text{equilibrium demand response}}. \tag{4}$$

Define the consumption-weighted elasticity, which is the elasticity of aggregate market demand, as

$$\varepsilon_D \equiv \sum_g \omega_g \varepsilon_g. \tag{5}$$

Solving equation (4) for dP gives

$$dP = \frac{\underbrace{\sum_g \omega_g \varepsilon_g \, ds_g}_{\text{direct demand impulse}}}{\underbrace{\eta}_{\text{supply response}} + \underbrace{\varepsilon_D}_{\text{market demand response}}}, \qquad \frac{P}{Q} dQ = \eta \, dP, \tag{6}$$

where the second expression holds because, in equilibrium, the change in consumption equals the change in supply. The numerator is the demand response that would arise if the producer

price did not move. The denominator is the subsequent equilibrium adjustment: producers expand supply as P rises, while all consumers reduce demand in response to the higher market price.

2.b Untargeted subsidies

Consider first a uniform per-unit subsidy s, offered to all groups. Setting $ds_g = ds$ for every g in equation (6) gives

$$\frac{dP}{ds} = \frac{\varepsilon_D}{\eta + \varepsilon_D}, \qquad \frac{P}{Q}\frac{dQ}{ds} = \frac{\eta\varepsilon_D}{\eta + \varepsilon_D}. \tag{7}$$

These are the standard incidence formulas, with the consumption-weighted elasticity as the demand elasticity. For untargeted support, therefore, heterogeneity in demand responses matters only through ε_D.

The consumption-weighted elasticity weights each firm's elasticity by its share of consumption, whereas the elasticity of the average firm weights every firm equally. The two coincide when elasticities are homogeneous, or when they are uncorrelated with consumption.[11] They diverge when consumption is concentrated and the firms that account for most of it respond systematically more, or less, than the others. Whether this is the case is an empirical question, which we take up in Section 6.

2.c Targeted subsidies

Now suppose that the subsidy is offered only to some groups. For any set S of groups, let

$$\varepsilon_S \equiv \sum_{g \in S} \omega_g \varepsilon_g$$

denote their contribution to the consumption-weighted elasticity, so that ε_D is the contribution of all groups. A common subsidy s offered only to a set T of eligible groups corresponds to $ds_g = ds$ for $g \in T$ and $ds_g = 0$ otherwise. Equation (6) becomes

$$\frac{dP}{ds} = \frac{\overbrace{\varepsilon_T}^{\text{response of eligible firms}}}{\underbrace{\eta + \varepsilon_D}_{\text{adjustment of the whole market}}}, \qquad \frac{P}{Q}\frac{dQ}{ds} = \frac{\eta\varepsilon_T}{\eta + \varepsilon_D}. \tag{8}$$

[11] With firms $i = 1, \ldots, N$ consuming q_i, average consumption $\bar{q} = Q/N$, and unweighted mean elasticity $\bar{\varepsilon}$, we have $\varepsilon_D = \sum_i (q_i/Q)\,\varepsilon_i = \bar{\varepsilon} + \mathrm{Cov}(q_i/\bar{q}, \varepsilon_i)$.

Only eligible firms contribute to the direct demand impulse in the numerator. Because ε_T weights each eligible group's elasticity by its consumption share, a subsidy to groups that consume a lot but respond little generates a small impulse: it mainly transfers resources on consumption that would have occurred anyway. The equilibrium adjustment in the denominator, however, depends on the demand response of the entire market, so the effects of a targeted subsidy are not confined to its recipients. Eligible firms keep only part of the subsidy, because the rest goes to producers through the higher producer price, and increase their consumption. Non-eligible firms receive no subsidy and face only the higher producer price, so they reduce theirs. Part of the additional consumption of eligible firms therefore comes at the expense of non-eligible firms.

2.d Additional policies

The same logic extends to the other two instruments used during the crisis. A voluntary demand reduction incentive, a per-unit payment r_g for reducing consumption below a fixed baseline, raises the marginal price: consuming an additional unit entails forgoing the payment, so group g pays $P + r_g$ at the margin. A retail price cap instead fixes the price paid by capped consumers at $\bar{p}$, whatever happens to the producer price; producers continue to receive P, and the government covers the difference $P - \bar{p}$. We set the cap below the initial producer price P_0, $\bar{p} = P_0 - c$, so that $c \geq 0$ is the cut in the price paid by capped consumers.

Let C denote the set of capped groups. Uncapped groups may receive subsidies and voluntary demand reduction incentives, so that $dp_g = dP - ds_g + dr_g$ for $g \notin C$, while capped groups face $dp_g = -dc$ for $g \in C$. Substituting into equation (3) gives the generalized incidence formula

$$dP = \frac{\sum_{g \notin C} \omega_g \varepsilon_g \, (ds_g - dr_g) + \varepsilon_C \, dc}{\eta + \varepsilon_D - \varepsilon_C}, \qquad \frac{P}{Q} dQ = \eta \, dP, \tag{9}$$

which reduces to equation (6) when no group is capped and no voluntary demand reduction incentives are offered. In what follows, and in Section 7, we consider one instrument at a time.

Voluntary demand reduction incentives. A voluntary demand reduction incentive enters equation (9) exactly as a subsidy with the opposite sign, so the results of Sections 2.b and 2.c carry over with signs reversed. In particular, a voluntary demand reduction incentive offered to a set T of groups lowers the producer price by $\varepsilon_T / (\eta + \varepsilon_D)$ per unit of incentive, and saves the most gas when it is offered to groups with a large contribution ε_T to the consumption-weighted elasticity.

Price caps. Caps differ from subsidies in one respect. The cut c enters the numerator of equation (9) exactly as a subsidy of the same size would: capped groups raise their demand in proportion to ε_C. Because their price no longer moves with P, however, capped groups drop out of the equilibrium adjustment in the denominator, which falls on supply and on uncapped consumers alone. For a cap limited to some groups,

$$\frac{dP}{dc} = \frac{\varepsilon_C}{\eta + \varepsilon_D - \varepsilon_C}. \tag{10}$$

When the cap applies to all groups, $\varepsilon_C = \varepsilon_D$ and

$$\frac{dP}{dc} = \frac{\varepsilon_D}{\eta}, \qquad \frac{P}{Q}\frac{dQ}{dc} = \varepsilon_D. \tag{11}$$

Compared with the uniform subsidy in equation (7), ε_D remains in the numerator but disappears from the denominator. Only supply is left to absorb the additional demand, so the producer price rises by more than under a subsidy of the same size, and by more than the cut c itself whenever demand is more elastic than supply, while consumption rises by the full demand response to the price cut. For the same reason, while a cap is in force it amplifies the producer-price effects of any other policy or supply shock, because the denominator of equation (9) is smaller.[12] In Section 7.d we calibrate these effects, for uniform caps and for caps limited to some groups.

2.e Empirical objects

Mapping the framework to the data thus requires the group elasticities ε_g, the consumption shares ω_g, and the supply elasticity η. In Sections 3–5 we estimate firms' energy demand responses to the unit-cost increases induced by the expiration of fixed-price contracts. In Section 6 we estimate how ε_g varies across firms and measure the shares ω_g, and in Section 7 we combine them with an external value of η.

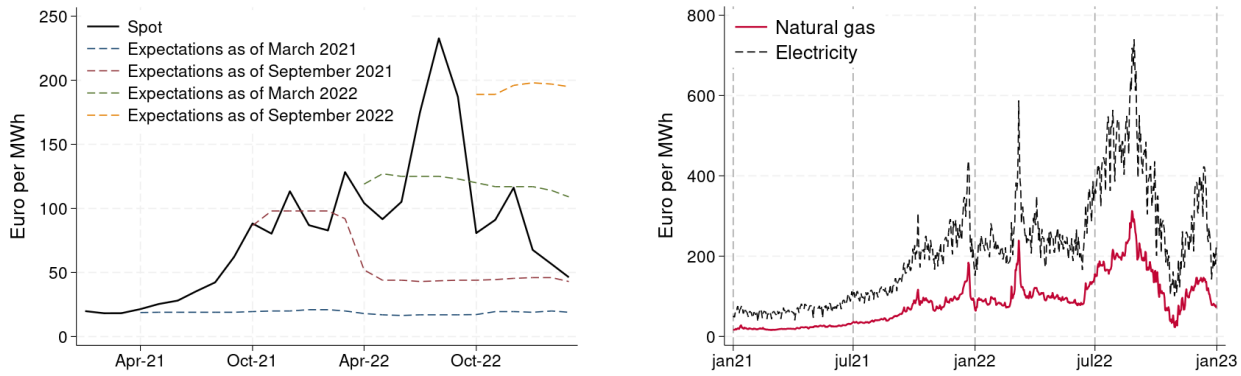
[12] Fixed-price contracts work like a temporary, firm-specific version of a cap. Because energy purchased under retail contracts cannot be resold (Section 3), a fixed-price contract fixes the marginal price a firm pays, so its demand does not respond to market prices until the contract expires and it contributes zero to the consumption-weighted elasticity. While contracts are in force, they shift market adjustment onto unhedged firms, but they leave the logic of the framework unchanged. The staggered expiration of these contracts is the source of variation we exploit in Section 3.c.

3 Institutional setting and sources of variation

In Section 2 we showed that evaluating crisis support requires group demand elasticities ε_g, defined with respect to the price firms pay. Estimating them during a crisis is difficult because wholesale prices rise for all firms at once, so aggregate price movements coincide with other shocks that hit all firms at the same time. Retail contracts break this simultaneity: firms holding fixed-price contracts signed before the crisis were exposed to crisis prices only when those contracts expired. In this section we describe the wholesale and retail developments behind this variation, limiting ourselves to the elements needed to understand the design and interpret the results; further details are in Appendix A.

3.a Wholesale market developments in 2021–22

Figure 3.1: 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 3.1).

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[13]. 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 3.1). These forecasts proved overly optimistic: spot prices increased further following Russia's 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 3.1).

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 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 3.1). 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 3.1).

3.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).[14] By law, energy purchased under these contracts cannot be resold[15].

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

[13]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.

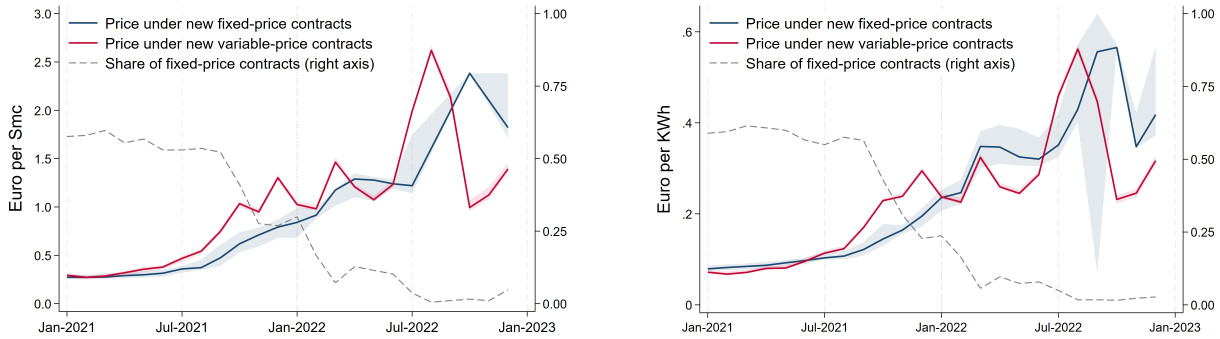
[14]Small firms, which are outside the scope of our analysis, still had access to regulated tariffs; see Appendix A for details.

[15]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 registration 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.

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[16] signed before the onset of the crisis in mid-2021 were fully shielded from the energy price shock until the pre-determined 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 a large increase in energy costs.

Figure 3.2: 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 shows 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 (Figure 3.2). 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.

[16] 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.

3.c Contract expiration as a source of price variation

To identify how firms' energy demand responds to the crisis price shock, we exploit the predetermined, staggered expiration dates of fixed-price contracts signed before the crisis. Three features of the setting make this variation useful. First, contracts in place at the beginning of 2021 were priced before markets anticipated the crisis (panel (a) of Figure 3.1). Second, their expiration dates were fixed at signature, so the timing of each firm's exposure was determined before the shock and cannot reflect firms' responses to it. Third, once the crisis began, an expiring contract could not be replaced on pre-crisis terms: new contracts, whether fixed or variable, embedded the higher wholesale prices (Figure 3.2). Otherwise similar firms therefore lost their protection, and faced the crisis price shock, at different points in time.

4 Data, measurement, and estimation

In Section 3 we explained why the staggered expiration of pre-crisis fixed-price contracts generates variation in the timing of firms' exposure to crisis prices. Exploiting this variation requires three firm-level variables that balance-sheet data do not provide: half-yearly energy quantities, the unit costs firms actually paid for energy inputs net of government support, and the timing of contract expiration. In this section we describe how we collected and measured them, and how we estimate firms' demand responses.

4.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. Invind has been used to study, among other topics, 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).

For our purposes, Invind provides detailed and timely information on firms' energy expenditures, consumption, and hedging strategies in 2021 and 2022. 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.[17] This section gathered half-yearly electricity and gas expenditures (in €) and consumption (in megawatt-hours, MWh, and standard cubic meters, Smc) for the previous year, excluding self-produced energy (Section 4.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 4.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.[18] The response rate to the energy section is around 50% and our results are robust to addressing selective item non-response with inverse-probability weighting (Appendix C).

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).

4.b Administrative data and other sources

Our main estimates, the effects of contract expiration on unit costs and quantities and the demand elasticities derived from them, are based entirely on the Invind survey and use the full survey sample, not only the firms that also appear in administrative records. The sources described below serve three other purposes. First, they classify surveyed firms, for instance as gas-intensive or as subject to the EU ETS, for the heterogeneity analysis and the robustness checks. Second, for the subsets of firms they cover, they provide independent information, such as administrative records of energy use, which we use to validate the survey estimates and to study mechanisms (Sections 6.c and 6.d); in these exercises, treatment status still comes from the survey. Third, they measure how gas consumption is distributed across firms, an input to the framework of Section 2 (Section 6.b).

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.[19] 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

[17] According to Istat's Business Register, in 2021 this group accounted for 49% of sectoral employment and 56% of sectoral turnover. Its share of industrial energy consumption is larger still; we document the distribution of energy consumption across firms in Section 6 and Appendix O.

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

[19] 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.

2022. We also use the CSEA records to measure the share of industrial energy consumption accounted for by energy-intensive firms (Section 6.b).

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.

4.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 half-yearly electricity and gas consumption (in megawatt-hours, 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). In Appendix D we show that the survey data are reliable.

The second outcome variable is the (log of) average unit cost, which we construct from expenditures and quantities, net of all government support. Starting from the second half of 2021, the Italian government cushioned the effects of the energy crisis with measures, such as reductions in system charges and VAT, that lowered firms' energy bills directly; these are already netted out of the expenditures firms report. In 2022, the government also introduced tax credits covering a fraction of energy expenditures. These do not appear on energy bills, but the survey asks about them, so we can subtract them. For firm i, energy input j and half-year t, the average unit cost is therefore

$$p_{ijt} = \frac{\text{Expenditure}_{ijt} - \text{Tax credit}_{ijt}}{\text{Quantity}_{ijt}}, \tag{12}$$

which measures the unit cost firms actually paid, net of government support. We use Eurostat data on average unitary cost by consumption bracket for industrial consumers to validate our survey-based measures (see Appendix D). Again, we find our data to be reliable.

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 E 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 a significant 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%). We discuss the implications for our estimates in Section 5.[20]

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.[21] 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.

4.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?"

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).[22]

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 3). Thus the question does not capture firms' endogenous responses to the crisis. Second, due to space constraints, the question does not distinguish between electricity and gas: a firm protected for *either* carrier answers "yes", so a firm protected for only one of the two could be misclassified for the other. Using aggregate information from ARERA, we estimate that most firms had either fixed-price contracts for

[20] 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.

[21] 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.

[22] 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 Figure F.1.

both energy carriers at the same time, or for none; this suggests that this misclassification is quantitatively small.[23] Consistent with this, we find that this variable strongly predicts changes in average unit costs for both inputs between the first and second half of 2021 (Alpino et al., 2023).

In the 2022 wave of the survey, the question was asked separately for electricity and natural gas:

> "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 $j = \{\text{electricity, natural gas}\}$, 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.

Restricting our attention to firms present in both waves[24], we constructed an input-specific treatment cohort variable E_i^j that identifies the half-year 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:

[23] 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 the 2021 protection question (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 misclassification in the treatment cohorts and suggesting that the treatment variable derived from this question is a reliable proxy for protection against both electricity and gas price shocks. Formally, using the inclusion-exclusion identity:

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

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:

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

[24] Exit is essentially absent from our sample during the crisis period (Appendix H).

$$E_i^j = \begin{cases} 2021h2, & \text{if } I^{2021} = 0 \text{ and } I^{j,2022} = 0 \\ 2022h1, & \text{if } I^{2021} = 1 \text{ and } I^{j,2022} = 0 \\ 2022h2, & \text{if } I^{2021} = 1 \text{ and } I^{j,2022} = 1 \text{ and } m_i^j = 6 \\ 0, & \text{if } I^{2021} = 1 \text{ and } I^{j,2022} = 1 \text{ and } m_i^j = 12 \end{cases} \tag{13}$$

All firms are untreated in $2021h1$ since the wholesale market crisis had not yet started (Section 3.a). Firms with fixed-price contracts signed before the crisis remain untreated in later periods until their contracts expire.[25] 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 both firms on variable-price contracts and those with fixed-price contracts expiring during 2021h2; we discuss the implications of this mixed composition in Section 4.f.

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 (Section 3.b).

Firms with $I^{2021} = 1$ and $m_i^j = 6$ are protected for the first six months of 2022 (late treated). We interpret these as firms holding a fixed-price contract already in place at the beginning of 2021 and expiring in mid-2022, and set $E_i^j = 2022h2$. We validate this timing in Section 5.c: unit costs for this cohort rise sharply in 2022h2.

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 firms in 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.[26] We excluded two

[25] 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 3).

[26] In the case of electricity, going from the early treated to the pure control group, the number of firms in each

subgroups of firms: (a) firms protected for only part of a half-year 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.

Balance tables in Appendix I show some level differences across cohorts. Cohort membership depends on two features of a firm's contract, both fixed before the crisis: when the contract was last signed and how long it lasts. Cohorts can therefore differ because firms renewed their contracts at different times, even with similar contract lengths, or because they chose contracts of different lengths. Firm fixed effects absorb level differences from either source, so the relevant concern is whether cohort membership correlates with differential trends in outcomes, whether pre-existing or ones that would have occurred even absent contract expiration. Our robustness checks address this concern without taking a stand on its source. In Appendix F we show that results are virtually unchanged when we include covariate-specific non-parametric time trends, by industry, energy intensity, firm size, and participation in the EU ETS. In Section 5.c we show that they are also robust to synthetic difference-in-differences (Arkhangelsky et al., 2021), which reweights control units and pre-treatment periods to match the pre-treatment trajectories of treated firms. This considerably loosens the requirement of parallel trends between treated and control units and addresses concerns related to pre-trend testing (Roth, 2022).

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

cohort is 176, 150, 11, 76. As for gas, the same figures are 128, 124, 7, 49.

follows:

$$E_i^j = \begin{cases} 2021q3, & \text{if } I^{2021} = 0 \text{ and } I^{j,2022} = 0 \\ 2022q1, & \text{if } I^{2021} = 1 \text{ and } 0 \leq m_i^j < 3 \\ 2022q2, & \text{if } I^{2021} = 1 \text{ and } 3 \leq m_i^j < 6 \\ 2022q3, & \text{if } I^{2021} = 1 \text{ and } 6 \leq m_i^j < 9 \\ 2022q4, & \text{if } I^{2021} = 1 \text{ and } 9 \leq m_i^j < 12 \\ 0, & \text{if } I^{2021} = 1 \text{ and } m_i^j = 12 \end{cases} \tag{14}$$

For the electricity-intensive case, the balanced matched sample includes 279 firms[27] over the period 2020-2022. For the gas-intensive case, our balanced matched sample includes 126 firms[28] in 2021 and 2022.

4.e Estimation

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

$$\log y_{ijt} = \alpha_i + \gamma_t + \sum_{k=0}^{2} \beta_{ijk} \cdot \mathbf{1}(t - E_i^j = k) + \epsilon_{ijt}, \tag{15}$$

where y_{ijt} is the quantity (q) or average unit cost (p) of energy input j for firm i in half-year t. As we explained in Section 4.c, p is the ratio of expenditures, net of government support, to quantities (equation (12)). 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).[29] 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 $\log(q)$.[30] Let $\hat{\tau}^q$ and $\hat{\tau}^p$ denote the estimated average treatment effects on $\log(q)$ and $\log(p)$ in a given half-year.

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

[27] Going from the earliest treated to the latest treated, cohorts have the following number of firms: 106, 99, 5, 5 and 19. The pure control group includes 45 firms.

[28] 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.

[29] See Section 4.d on how we construct this variable.

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

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 we show 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 event-study estimator of Borusyak et al. (2024) (BJS), 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. Following BJS, we estimate pre-treatment coefficients in a *separate* regression that uses only untreated firm-year observations. This has three advantages: (i) it separates testing the parallel trends assumption from estimation under that assumption; (ii) it improves efficiency, since all untreated observations are used to estimate the counterfactual outcomes of treated firms; and (iii) it removes the correlation between treatment effect estimates and pre-trend coefficients that otherwise introduces bias when researchers condition on passing a pre-trend test (Roth, 2022).

The elasticities we report are Wald-type ratios of the implied percent changes, $\theta = (e^{\hat{\tau}^q} - 1)/(e^{\hat{\tau}^p} - 1)$, where $\hat{\tau}^q$ and $\hat{\tau}^p$ are the estimated ATT for $\log(q)$ and $\log(p)$, respectively. Our baseline confidence sets follow Fieller (1954): they collect the values r for which the hypothesis $\theta = r$, written without division as $(e^{\tau^q} - 1) - r\,(e^{\tau^p} - 1) = 0$, is not rejected, and remain valid however imprecise the effect on unit costs is (Appendix J). Because delta-method approximations can be inaccurate when the statistic of interest is highly nonlinear (Lafontaine and White, 1986), we report delta-method and wild cluster bootstrap intervals (Cameron et al., 2008) only for comparison. Because all estimates use survey weights, these ratios describe the response of the average firm in the population of industrial firms with at least 50 employees. In Section 6.b we contrast them with the consumption-weighted elasticity.

4.f Identifying assumptions

Parallel trends assumption. In Section 3.c we explained why expiration dates are plausibly unrelated to firms' crisis-time demand: they were fixed at contract signature, before markets anticipated the crisis, and could not be extended once prices rose. Our treatment definition preserves this feature by conditioning on protection already in place at the beginning of 2021. Pre-trend tests (Section 4.e) do not reject parallel trends for unit costs and electricity

quantities; for gas quantities, the test rejects. We address this in Section 5.c with synthetic difference-in-differences (Arkhangelsky et al., 2021) and with specifications that include covariate-specific non-parametric time trends.

One concern involves the early-treated cohort ($E_i^j = 2021h2$), which mixes firms whose fixed-price contracts expired in 2021h2 with firms on variable-price contracts (Section 4.d). The latter may have had different characteristics, such as weaker preferences for price protection or different expectations about future demand. In Section 5.c and Appendix G we show that effects are similar across cohorts, which mitigates this concern. Additionally, early-treated firms never serve as controls for other cohorts.

Another concern is that cohorts differ in ways related to their contracts. Cohort membership reflects both when firms last signed their contracts and how long these contracts last, and contract length, in particular, could reflect unobserved preferences for price certainty. Since firm-level fixed effects control for level differences, this only threatens validity if such differences correlate with abrupt changes in outcomes specifically at contract expiration. This is unlikely, and the absence of pre-trends in unit costs and electricity quantities, together with the robustness of gas estimates to synthetic difference-in-differences and to covariate-specific non-parametric time trends (Section 5.c), 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.c and Appendix G 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.

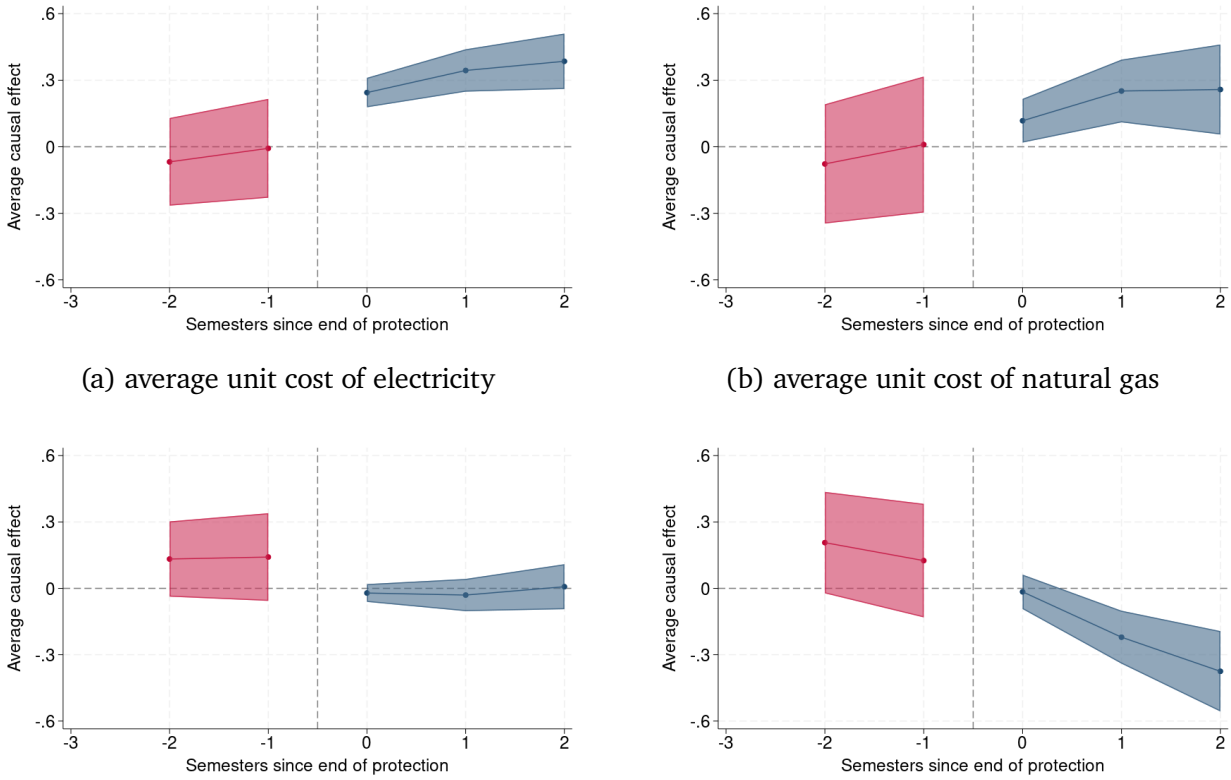
5 Average effects of contract expiration

In this section we report the average effects of losing fixed-price protection on firms' energy unit costs and consumption. We first show that contract expiration raises the prices firms pay, then examine how and when consumption responds, and finally show that these results are robust to alternative estimators and to confounding shocks. We also combine the price and quantity responses into the average-firm elasticity of gas demand.

5.a Unit costs and quantities

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 average 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) average unit cost of electricity

(b) average 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 event-study estimator of Borusyak et al. (2024), as described in Section 4.e. All estimates are weighted by survey weights. Standard errors are clustered at the firm level. Confidence intervals are at the 95% level.

Unit costs. Contract expiration exposes firms to a large price shock. In panel (a) we observe a significant increase in the average unit cost of electricity following the expiration of fixed-price contracts, with no evidence of pre-trending effects before the treatment occurs. One year after expiration, the price increase is 45% relative to the baseline. All three post-expiration coefficients are precisely estimated and strongly significant: $\hat{\tau}_0 = 0.245$ ($p < 0.001$), $\hat{\tau}_1 = 0.344$ ($p < 0.001$), and $\hat{\tau}_2 = 0.371$ ($p < 0.001$), corresponding to price increases of 28%, 41%, and 45% in levels. The pre-trend joint test does not reject parallel trends (F= 1.67, $p = 0.19$). A similar pattern is seen in panel (b) for the average unit cost of natural gas. One year after expiration, the increase in the gas price is 30% relative to the baseline. The post-expiration coefficients are $\hat{\tau}_0 = 0.121$ ($p = 0.018$), $\hat{\tau}_1 = 0.262$ ($p < 0.001$), and $\hat{\tau}_2 = 0.261$ ($p = 0.013$), corresponding to increases of 13%, 30%, and 30% in levels. The pre-trend joint test does not reject parallel trends (F= 1.20, $p = 0.30$).[31]

Electricity consumption. In panel (c) we analyze the (log-)quantity of electricity purchased. Despite the large price increase, this outcome does not respond to the treatment. The post-expiration coefficients are small and statistically insignificant ($\hat{\tau}_0 = -0.021$, $\hat{\tau}_1 = -0.032$, $\hat{\tau}_2 = 0.003$), and the confidence intervals rule out reductions larger than 10.2% at any horizon.[32]

This result differs from most 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.[33] We offer two explanations. First, most micro-level estimates are based on smaller, normal-time price variation (Fontagné et al., 2024; von Graevenitz and Rottner, 2022), whereas our design identifies quantity changes over the crisis-relevant price range. When demand is not isoelastic, local elasticities estimated around normal prices need not extrapolate to very large shocks (Appendix K). For example, if electricity demand has a minimum-throughput floor (machines that cannot be ramped below a threshold without irreversible damage), the local elasticity shrinks toward zero as the price rises and the floor becomes binding. Second, moving away from electricity usually requires

[31] 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).

[32] This result also resolves the measurement question raised in Section 4.c. As documented in Appendix E, 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 and the elasticity with respect to the marginal price are both zero.

[33] 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.

new equipment. Firms can sometimes switch between combustible fuels with minor changes to existing equipment, but replacing electricity with a fuel, or a fuel with electricity, typically means installing different machinery (Wei et al., 2019; Mallapragada et al., 2023; Leicher et al., 2024). Such investments 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; Cooper and Haltiwanger, 2006). Reducing electricity use should then come at the cost of lower output, for which we find little evidence (Section 6.d).

Gas consumption. 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 $\hat{\tau}_2 = -0.367$ ($p < 0.001$), which implies that gas consumption decreases by 31% compared to the counterfactual with no price increase.

The joint pre-trend test for the quantity of gas rejects the null of no pre-trends (F= 4.45, $p = 0.01$). We address this concern in Section 5.c, where we use synthetic difference-in-differences (Arkhangelsky et al., 2021) and obtain very similar estimates. In Appendix F we further show that our results are robust to including covariate-specific non-parametric time trends.

5.b Heterogeneity of gas effects across cohorts and 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 Figure F.3). 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 by calendar time rather than event time, the ATT on gas consumption in 2022h2 is −34%, with 95% confidence interval between −44% and −24%. Combined with the 31% increase in gas unit costs in the same half-year, this implies a gas demand elasticity of −1.1, with 95% confidence interval [−3.9, −0.6] (Table 6.1); in earlier half-years the elasticity is close to zero

(Appendix Figure J.1).[34] Because our estimates use survey weights, this is the elasticity of the *average firm* in the population of industrial firms with at least 50 employees, of which our sample is representative (Section 4.a and Appendix C); we refer to it as the average-firm elasticity.

The increasing gas-demand response, *irrespective* of time since treatment, could reflect changing expectations about the energy crisis. Early on, when financial markets expected the crisis to be short-lived, firms likely followed 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 3.1 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). The existing literature has established that long-run elasticities exceed short-run ones (Deryugina et al., 2020; Labandeira et al., 2017; Sikl et al., 2025); our evidence suggests that firms hit by a shock undertake costly adjustments only once they stop perceiving it as temporary.

One might worry that the 2022h2 differential response reflects macro shocks correlated with treatment timing rather than contract expiration per se. We address this concern in Section 5.c, where we include covariate-specific non-parametric time trends by sector, size, ETS participation, and energy intensity. Results are very robust to the introduction of these controls.

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, and their scale expanded substantially over time, peaking in 2022h2. However, this interpretation is at odds with our price results. As we show in Appendix Figure F.3, panel (a), contract expiration triggered a substantial increase in average unit costs already in 2021h2, precisely because our price measure is net of all support received (Section 4.c). If bailout anticipation were suppressing quantity responses, we would expect this to matter most in 2022h2, when the cost shock, even net of support, was largest and government support most generous. Yet this is precisely the period in which 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.

[34] Using unit costs gross of government support instead changes the 2022h2 gas demand elasticity from −1.1 to −1.0.

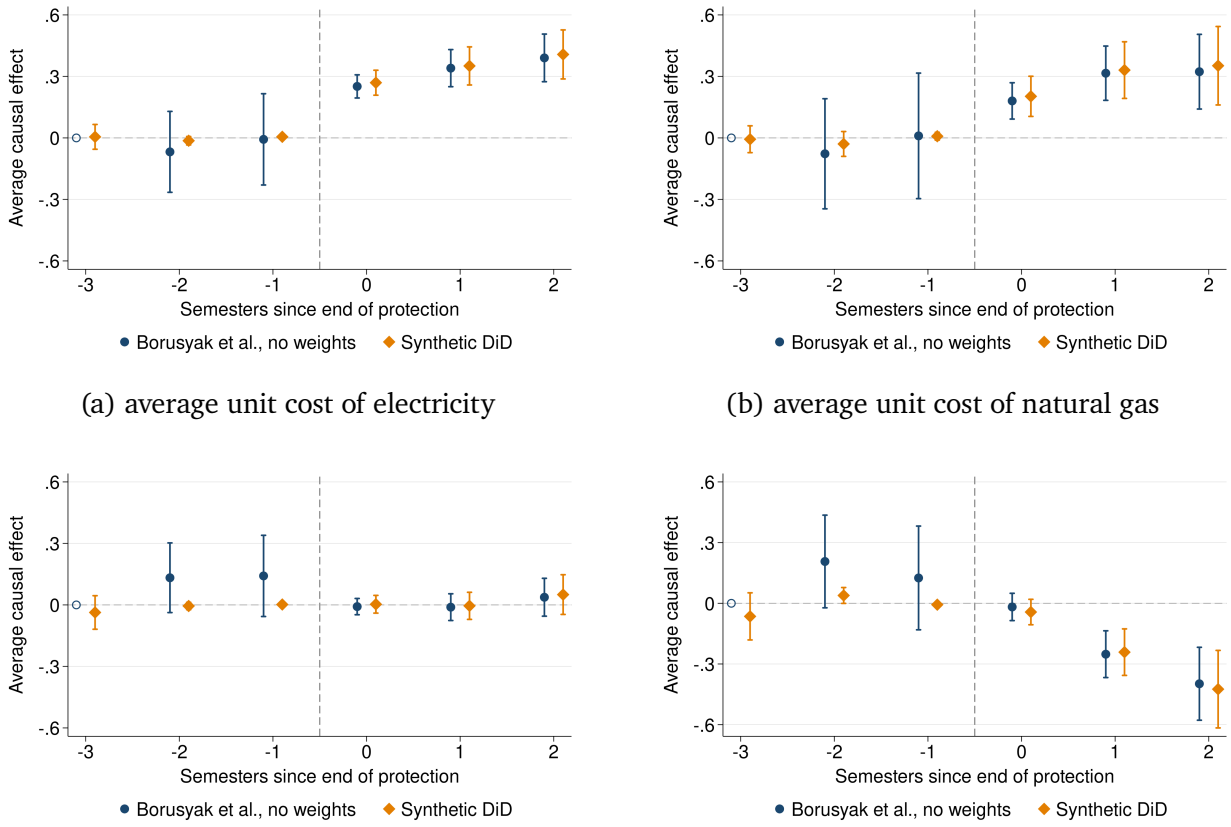
5.c Robustness

Synthetic difference-in-differences. To address the gas pre-trend and, more generally, to relax the parallel trends assumption, we implement a synthetic difference-in-differences (SDID) estimator (Arkhangelsky et al., 2021). This method reweights control units and pre-treatment periods to match the pre-treatment trajectories of treated firms, which considerably loosens the requirement of parallel trends and addresses concerns related to pre-trend testing (Roth, 2022). We use the version for staggered adoption of Clarke et al. (2024), in which each cohort is compared with a synthetic control drawn from the pure control group, and we aggregate the cohort-specific estimates by event time as in Ciccia (2024). Inference is based on a bootstrap that resamples firms. Because the SDID estimator does not support survey weights, we compare it with BJS estimates that also omit them, so that the two estimators differ only in how they construct the counterfactual. Because we observe only four half-years, the reweighting can use more pre-treatment half-years for later cohorts than for earlier ones.

Figure 5.2 reports the two sets of estimates side by side, for unit costs (panels (a) and (b)) and quantities (panels (c) and (d)). For gas, the SDID estimates track the BJS ones: consumption barely changes in the half-year of expiration (−4%, against −2% with BJS) and falls by 21% and 35% one and two half-years later (22% and 33% with BJS). Electricity consumption does not respond under either estimator, and unit costs rise from the half-year of expiration onward under both.

Cohort-level SDID estimates (Appendix G) confirm that the gas response is concentrated in 2022h2, regardless of when contracts expired. In 2022h2, gas consumption falls by 35%, 32%, and 33% in the early, mid, and late cohorts. Pooling the cohorts already treated in each half-year, it falls by 4% in 2021h2, 6% in 2022h1, and 34% in 2022h2, with 95% confidence interval between −42% and −24%, practically the same as the survey-weighted BJS estimate in Section 5.b. For the late cohort, unit costs rise by 58% for gas and 79% for electricity in 2022h2, which supports our coding of its expiration date. The only cohort-level estimate at odds with the pooled results is a fall of about 10% in the electricity consumption of the late cohort in 2022h2. This cohort includes only 11 firms; the early and mid cohorts, also treated in 2022h2, show no electricity response.

Figure 5.2: Synthetic difference-in-differences and BJS estimates of the effect of the expiration of a fixed-price contract on average unit costs and quantities of energy inputs.



(a) average unit cost of electricity

(b) average unit cost of natural gas

(c) Quantity of electricity

(d) Quantity of natural gas

Note: The figure shows average causal effects of the expiration of a fixed-price contract on log average unit costs, net of government support (panels (a) and (b)), and on log purchased quantities (panels (c) and (d)), by half-years since expiration. Circles: Borusyak et al. (2024) estimator, with the same specification as in Figure 5.1; pre-expiration coefficients come from a separate regression on untreated observations, relative to the omitted period $k = -3$ (hollow circle). Diamonds: synthetic difference-in-differences (Arkhangelsky et al., 2021), in the staggered version of Clarke et al. (2024), aggregated across cohorts by event time (Ciccia, 2024). Unlike the baseline estimates in Figure 5.1, which use survey weights, neither estimator here is weighted, because the SDID estimator does not support survey weights. Confidence intervals are at the 95% level, with standard errors clustered at the firm level (BJS) and from 500 bootstrap draws that resample firms (SDID).

Confounders due to macro shocks. The second half of 2022 coincides with major macroeconomic and policy shocks: Russia's 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 F.2 in Appendix F shows that results are robust to the inclusion of covariate-specific non-parametric time *trends*, by sector, size, energy intensity, and ETS

participation.

Confounders due to seasonality. 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 × half-of-year 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 F.4). Because the panel spans only four half-years (2021h1-2022h2), however, early-treated firms (E_i^j = 2021h2) are treated in every h2 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 indicates that seasonal patterns do not drive our results.

Alternative estimators. In Appendix Figure F.5 we confirm robustness using alternative difference-in-differences estimators, including De Chaisemartin and d'Haultfoeuille (2020), Callaway and Sant'Anna (2021), Sun and Abraham (2021), and OLS.

6 Heterogeneity analysis

In the previous section we showed that the average firm cut its gas consumption only in 2022h2, with an elasticity of −1.1. However, theoretical results from Section 2 highlight that the effects of a uniform subsidy rather depend on the consumption-weighted elasticity, which might be different from the simple average if energy demand elasticities are heterogeneous across firms. In this section we therefore ask whether the 2022h2 gas response differs across firms with different observable characteristics.

In doing so, we devote special attention to firm-level gas-intensity. Energy-intensive firms were the earliest and largest recipients of government aid during the 2021-22 crisis in several EU countries (McWilliams et al., 2024). In addition, they represent a large share of total energy industrial consumption (80% of industrial natural gas consumption in our application).[35] Thus, their responsiveness is central to evaluating the effects of government aid in 2022. Whether

[35]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. Gas consumption is far more concentrated across firms than sales or employment, as Lyubich et al. (2018) document for the United States. In our data, the Gini coefficient is 0.90 for gas consumption and 0.70 for sales (Appendix Figure O.1).

their demand should be more or less elastic than that of other firms is theoretically ambiguous. A larger energy cost share strengthens the scale effect of a price increase, which raises the own-price elasticity of energy demand; but energy-intensive processes may offer less scope to substitute away from gas, which lowers it (Marshall, 1920; Hicks, 1932).[36]

In our heterogeneity analysis, we also consider a binary survey-based proxy of the substitutability of natural gas. The dummy variable is derived from the Bank of Italy's Business Outlook Survey of Industrial and Service Firms (Sondtel) conducted in fall 2022 which included the question: "At the beginning of 2022, was gas an essential input for your firm's 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. We therefore expect on average lower energy demand elasticities for those firms that declared gas to be an essential input.

Finally, we also consider other firms' characteristics such as sector and size.

6.a Searching for treatment heterogeneity with machine learning

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 use this estimate as an outcome variable and rely on *random forests* to find its best predictors, in the spirit of Athey and Imbens (2016) and Wager and Athey (2018). We train the model in a random subset of the data, and validate the predictions in the remaining subsample. 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.

We provide more details on our procedure and extensive results in Appendix M. Here, we simply highlight that the random forest selects the following robust predictors of heterogeneity: the self-reported dummy variable indicating whether gas is an essential input in production; the dummy for being a gas-intensive firm according to the Italian legislation; the dummy for

[36] We formalize this ambiguity in Appendix L.

being subject to the EU Emission Trading System (ETS); the dummy for the food industry. Interestingly, the ML algorithm does not detect any additional between-industry treatment heterogeneity. Firm size also does not appear to play a role in this respect.

The four selected characteristics are strongly correlated. Nearly all gas-intensive firms report gas as an essential input, as do most firms in the food industry, and EU ETS plants are among the largest gas consumers. Thus, we interpret them as imperfect proxies for a single dimension of heterogeneity, the importance of gas in production.

6.b Heterogeneous and consumption-weighted elasticities

In what follows, we focus on the gas-intensive classification for two reasons: a) it is the administrative criterion that the government actually used to target support; b) we observe administrative micro data on their energy consumption, which we use to validate the survey-based estimates (Section 6.c) and to measure their share of total industrial gas consumption. In Section 6.d we report the split by whether gas is an essential input, which yields a contrast in the same direction. We compute group effects as subsample averages of the individual-level BJS treatment effects.

The response of gas consumption in 2022h2 is close to zero among gas-intensive firms. The implied elasticity is -0.03 for gas-intensive firms and -1.3 for other firms (Table 6.1; Appendix Figure J.1, panel (d)). The confidence set for other firms, $[-5.7, -0.7]$, extends mostly toward more elastic values; the one for gas-intensive firms, $[-0.38, 0.73]$, rules out elasticities below -0.38.[37] The two elasticities differ by more than an order of magnitude, and the difference in the underlying responses is statistically significant: in 2022h2, gas consumption falls by 2% among gas-intensive firms and by 36% among other firms, and a test of equality of the two effects rejects at the 5% level ($z = 2.38$, $p = 0.017$). We also provide an alternative test, which applies the test behind our Fieller confidence sets to both groups at once (Dufour and Taamouti, 2005): for each candidate value r of a common elasticity, we test jointly that, in both groups, the quantity effect equals r times the price effect. We find that every value of r is rejected ($p = 0.002$; Table 6.1, Panel C): no single elasticity fits the responses of both groups. Administrative data, discussed next, point in the same direction. The contrast does not reflect sector composition: gas-intensive firms are spread throughout all industries except water supply and waste management.

Weighting the group elasticities by their respective consumption shares, as in equation (5),

[37] Using unit costs gross of government support, the elasticity changes from -1.3 to -1.2 for non-gas-intensive firms and from -0.03 to -0.02 for gas-intensive firms.

Table 6.1: Price elasticities of energy demand in 2022h2: estimates and 95% confidence sets

		Baseline	Alternatives	
	Elasticity	Fieller	Delta method	Bootstrap
	(1)	(2)	(3)	(4)
Panel A. Average firm				
Electricity	−0.05	[−0.20, 0.11]	[−0.19, 0.10]	[−0.19, 0.12]
Natural gas	−1.10	[−3.88, −0.61]	[−1.92, −0.29]	[−3.35, −0.59]
Panel B. Natural gas, by group				
Gas-intensive firms	−0.03	[−0.38, 0.73]	[−0.45, 0.40]	[−0.34, 0.65]
Other firms	−1.27	[−5.72, −0.68]	[−2.27, −0.27]	[−4.40, −0.64]
Panel C. Test of equal elasticities (p-values)				
Gas-intensive = other		0.002	0.021	0.008

Notes: Column (1) reports Wald-type elasticities $(e^{\hat{\tau}^q} - 1)/(e^{\hat{\tau}^p} - 1)$ with respect to the average unit cost net of government support, from survey-weighted BJS estimates for 2022h2, aggregated by calendar time. Column (2), our baseline, reports Fieller (1954) confidence sets: the values r for which the test of $(e^{\tau^q} - 1) - r(e^{\tau^p} - 1) = 0$ is not rejected at the 5% level, using the BJS variances and covariance of the two effects. Column (3) reports delta-method intervals. Column (4) reports the 2.5th and 97.5th percentiles of 1,999 draws of a wild cluster bootstrap with Rademacher weights, clustered by firm and applied jointly to the price and quantity effects (Cameron et al., 2008). Panel C reports p-values for the hypothesis that gas-intensive and other firms have the same elasticity; each column uses a different test. Column (2) extends the Fieller approach: for each candidate common elasticity r, we test jointly, in both groups, that $(e^{\tau^q} - 1) - r(e^{\tau^p} - 1) = 0$, with a Wald statistic that is $\chi^2(2)$ under the null (one restriction per group), and report the largest p-value over all r. Equality is thus rejected only if no common elasticity fits both groups; because equality is a single restriction while the test uses two degrees of freedom, the test is conservative. Column (3) reports a z-test of the difference between the two elasticities, with delta-method standard errors. Column (4) reports the p-value implied by the bootstrap percentile intervals: the smallest significance level at which the equal-tailed percentile interval for the difference between the two elasticities excludes zero.

gives a consumption-weighted elasticity of $0.8 \times (-0.03) + 0.2 \times (-1.3) \approx -0.28$. Gas-intensive firms account for four-fifths of consumption but contribute almost nothing to the consumption-weighted elasticity; other firms account for one-fifth of consumption but for almost all of it. The consumption-weighted elasticity is about a quarter of the average-firm elasticity (-1.1; Table 6.1), which is computed from pooled treatment effects rather than as a weighted average of the group elasticities. The gap arises because gas consumption is concentrated among the least responsive firms: weighting firms by their gas use, rather than counting each firm once, moves weight toward gas-intensive firms.

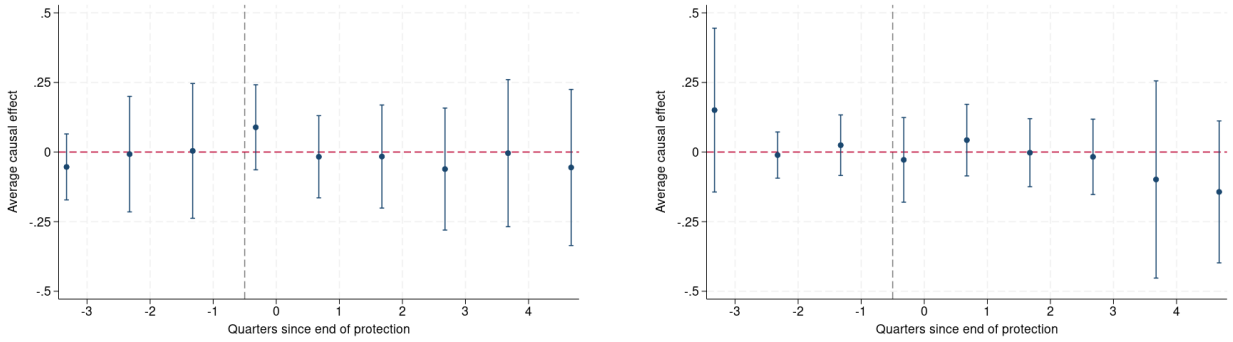
6.c Robustness with administrative data

For energy-intensive firms, we can complement the survey evidence with confidential administrative data on monthly energy consumption, which are available only for these firms (Section 4). The administrative data have two advantages: they do not rely on self-reported consumption, and their higher frequency allows us to define cohorts at the quarterly level. Figure 6.1 presents event-study graphs for the evolution of gas and electricity consumption around the expiration of fixed-price contracts.

For electricity-intensive firms, as for the average industrial firm in Section 5, losing fixed-price protection has a negligible effect on the (log-)quantity of electricity consumed in all treated periods.

For gas-intensive firms, the effects on gas consumption are close to zero in the first four quarters and turn negative, though 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, consistent with the baseline results. Aggregating by half-year, the results indicate an ATT of -6% in the second half of 2022, not statistically different from zero at conventional levels. This is far smaller than the -34% ATT for the average industrial firm in the baseline analysis, and confirms with an independent data source that gas-intensive firms respond little.

Figure 6.1: 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 event-study estimator of Borusyak et al. (2024). Standard errors are clustered at the firm level. Confidence intervals are at the 95% level.

6.d Mechanisms

Gas-intensive firms respond little, even though their larger cost share might have made them more responsive. Three pieces of evidence suggest that an important reason is the limited short-run substitutability of gas. The first two build on characteristics flagged by the forest in Section 6.a.

The first piece of evidence comes from firms' own assessment of their production technology. Splitting the sample by whether firms report gas as an essential input, the ATT on gas consumption in the second half of 2022 is −41% among firms for which gas is *not* essential, and −27% among those for which it *is* essential. Both estimates are significant at the 1% level and their difference is significant at the 10% level. The associated price elasticities of demand are −2.5 and −0.5, respectively (Appendix Figure J.1, panel (c)), consistent with the interpretation of the gas-essential indicator as a proxy for low elasticity of substitution suggested by the formulation of the survey question. Nearly all gas-intensive firms classify gas as an essential input. In the input-demand model of Appendix L, a zero elasticity of substitution leaves only the scale effect, which is small when demand for firms' output is inelastic enough for higher costs to be passed on to prices without large output losses.

Second, firms replaced only part of the gas they saved with other fuels. For plants subject to the EU ETS, administrative records report the use of all fossil fuels. Among these plants, contract expiration reduced gas consumption in 2022 by 89 terajoules, or 14% of the 2018-

2020 average, but total fossil energy consumption by only 56 terajoules: the use of other fuels rose slightly, offsetting only part of the reduction in gas (Appendix N).

Survey evidence points in the same direction. The Invind survey included a 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 N.3).

How firms absorbed the shock. Firms that neither cut energy use much nor switched fuels still had to cope with much higher energy costs. 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 environment, many firms could absorb the shock through nominal and financial margins, by raising output prices, compressing profit margins, or both.

Our evidence is consistent with this interpretation. Firms in our sample increased output prices significantly in 2021 (8%) and even more in 2022 (12%) (Appendix Figure N.4). All cohorts raised prices markedly and almost synchronously, in line with the general inflationary episode and with evidence that firms' price updating depends more on rivals' and market-wide cost shocks than on idiosyncratic ones (Duprez and Magerman, 2018; Amiti et al., 2019; Muehlegger and Sweeney, 2022).[38] Energy-intensive firms increased their prices substantially more than other industrial firms. This suggests that price setting was an important margin of adjustment for them, and helps explain how they remained financially viable despite a near-zero demand elasticity.

Qualitative questions in the 2022 *Business Outlook Survey of Industrial and Service Firms* (Sondtel) give a similar picture (Appendix Figure N.5). 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. Only a relatively small share reported reducing plant operating time or curtailing production, even partially.[39] Investment-based responses, such as installing energy-saving equipment, expanding self-

[38]The near-synchronous price movement across cohorts also explains why a causal event-study design analogous to equation (15) would be uninformative for estimating the pass-through.

[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. 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 3). It also has no implications for the validity of our empirical design, because we classify firms as protected only if they hold a fixed-price contract signed at the *beginning* of 2021, before the crisis (see Section 4.d).

generation, or adapting plants to alternative energy sources, were also uncommon at that stage, although they became notably more frequent when firms were asked about their plans for the coming months. This pattern suggests that technological adjustment was feasible but slow, consistent with the delayed gas response documented in Section 5.b.

7 Equilibrium effects of crisis-support policies

We now return to the framework of Section 2 and use the estimates of Sections 5 and 6 to compare crisis-support policies. We first describe the calibration and how we compute fiscal costs. We then show how much the choice between the average-firm and the consumption-weighted elasticity matters for a uniform subsidy, compare uniform and targeted subsidies, and finally consider voluntary demand reduction incentives, retail price caps, and a combination of instruments.

7.a Calibration and fiscal costs

Table 7.1 lists the inputs. The demand elasticities are our estimates for gas-intensive (GI) firms, $\varepsilon_{\text{GI}} = 0.03$, and for other, non-gas-intensive (NGI) firms, $\varepsilon_{\text{NGI}} = 1.3$ (Section 6.b). The consumption shares, 0.8 for GI firms and 0.2 for NGI firms, come from administrative data (Section 6.b). Prices and industrial gas consumption refer to the first half of 2022, the half-year before the simulated policy: the producer price is from the Italian Wholesale Market Operator (GME), consumer prices are from Eurostat, and consumption is from SNAM. We set the short-run supply elasticity to $\eta = 0.2$, following Abuin (2024) as in Harstad and Holtsmark (2024).

Consumer prices include per-unit charges and taxes from which we abstract in Section 2, and are about 4.5 percent higher than the producer price. We hold these wedges fixed, so that consumer prices move one-for-one with P. As noted in Section 2, the formulas then apply with weights proportional to Q_g/p_g rather than Q_g, but the difference is negligible here.[40]

All instruments are set at €0.50 per standard cubic meter (Smc), close to half of the 2022h1 producer price, and we evaluate their effects at the 2022h1 equilibrium with the first-order formulas of Section 2. Because the policies are large, the results should be read as comparisons across policy designs rather than as precise forecasts.

[40] With these weights the consumption-weighted elasticity is 0.27 rather than 0.28, and the results below change only marginally.

Fiscal costs. A per-unit subsidy to group g costs $s_g Q_g$ at the new equilibrium. We split this into a mechanical transfer on baseline consumption, $s_g Q_{0,g}$, and a behavioral cost on the consumption that the subsidy induces, $s_g \Delta Q_g$. A voluntary demand reduction incentive pays r_g per unit saved relative to baseline consumption, so its cost, $r_g(Q_{0,g} - Q_g)$, is entirely behavioral. A cap costs the government the gap between the market price and the cap, $c + \Delta P$, on all consumption of capped groups, $Q_{0,C} + \Delta Q_C$. Because the cap fixes the consumer price, the government, rather than consumers, pays for the increase in the producer price that the cap itself causes, on all baseline consumption of capped groups. We call this the price-feedback term. We evaluate each cost at the first-order changes in prices and quantities.

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: Prices and consumption refer to 2022h1, the half-year before the simulated policies. Demand elasticities are in absolute value.

7.b A uniform subsidy: average-firm versus consumption-weighted elasticity

We start with a uniform per-unit subsidy. For untargeted policies, the framework shows that demand matters only through the consumption-weighted elasticity (Section 2.b). Table 7.2 shows how much the predicted effects change if we use instead the average-firm elasticity (-1.1; Table 6.1).

The average-firm elasticity overstates the distortions created by the subsidy. It predicts that the producer price rises by 38 percent rather than 26 percent, and that gas consumption rises by 7.7 percent rather than 5.3 percent, an overstatement of about 45 percent. Fiscal costs, by contrast, are similar in the two rows, because they are dominated by the mechanical transfer on baseline consumption.

Table 7.2: A uniform subsidy: average-firm versus consumption-weighted elasticity

Demand elasticity	$\Delta P/P$	$\Delta Q/Q$	**Fiscal cost** (€ bn)		
	(%)	(%)	**Mechanical**	**Behavioral**	**Total**
(1) Average firm ($\varepsilon = 1.1$)	+38.4	+7.7	3.35	0.26	3.61
(2) Consumption-weighted ($\varepsilon = 0.28$)	+26.3	+5.3	3.35	0.18	3.53

Notes: Effects of a uniform subsidy of €0.50/Smc on the producer price, aggregate industrial gas consumption, and fiscal costs, computed from equation (7) with the calibration in Table 7.1. Row (1) uses the average-firm elasticity, row (2) the consumption-weighted elasticity. Fiscal costs are defined in Section 7.a. Components may not add up to totals because of rounding.

7.c Targeted subsidies

Panel (a) of Table 7.3 compares the uniform subsidy with the same subsidy offered only to GI firms or only to NGI firms. From now on, all calculations use the two groups' own elasticities.

A uniform subsidy raises aggregate gas consumption by 5.3 percent and the producer price by 26.3 percent: short-run supply is inelastic, so most of the subsidy goes into a higher price. The same subsidy offered only to GI firms raises the producer price by 2.2 percent and consumption by 0.4 percent. GI firms consume most of the gas, but their demand barely responds to the price they pay, so subsidizing them distorts the market least.

GI-targeted support is nevertheless expensive. Because GI firms consume most of the gas, the mechanical transfer to them is large (€2.7 billion), while the behavioral cost is negligible: the subsidy is almost entirely a transfer on consumption that would have occurred anyway. Its total cost is about a quarter lower than that of the uniform subsidy, but its effect on the market price per euro spent is much smaller. For each billion euros transferred on baseline consumption, the producer price rises by less than 1 percentage point with the GI-targeted subsidy, by about 8 points with the uniform subsidy, and by 36 points with the NGI-targeted subsidy.

Targeting NGI firms gives the opposite pattern. Mechanical outlays are small (€0.67 billion), because NGI firms consume little gas, but the producer price rises by 24.1 percent and consumption by 4.8 percent, almost as much as with the uniform subsidy: in both cases, NGI firms account for nearly all of the demand response. Subsidizing the elastic part of demand raises consumption and the producer price; subsidizing the inelastic part mostly transfers resources to it.

Table 7.3: Equilibrium and fiscal effects of alternative policy designs

Instrument	**Target**	$\Delta P/P$	$\Delta Q/Q$	**Fiscal cost** (€ bn)			
		(%)	(%)	**Mechanical**	**Price feedback**	**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)							
Voluntary demand reduction incentive	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)							
Retail price cap	Uniform	+62.1	+12.4	3.35	4.55	0.98	8.88
	GI only	+2.3	+0.5	2.68	0.14	0.04	2.85
	NGI only	+51.0	+10.2	0.67	0.75	0.81	2.22
Panel (d)							
Mix: Subsidy to GI + voluntary demand reduction incentive to NGI		−21.9	−4.4	2.68	–	0.25	2.93

Notes: Effects of instruments of €0.50/Smc on the producer price, aggregate industrial gas consumption, and fiscal costs, computed from equation (9) with the two groups' own elasticities and the calibration in Table 7.1. In panels (a) to (c), each instrument applies to all firms (Uniform), to gas-intensive firms only (GI only), or to other firms only (NGI only). Panel (d) combines a subsidy to GI firms with a voluntary demand reduction incentive for NGI firms; its fiscal cost is computed from the quantity changes that the two instruments cause together. Fiscal costs are defined in Section 7.a. Components may not add up to totals because of rounding.

7.d Voluntary demand reduction incentives, price caps, and a policy mix

Voluntary demand reduction incentives. Panel (b) of Table 7.3 shows that a voluntary demand reduction incentive has effects on prices and quantities of the same size as the subsidy, with the opposite sign. Its fiscal cost is much lower, because the government pays for gas saved rather than for gas consumed, so outlays go where demand responds. A voluntary demand reduction incentive offered to GI firms is cheap but saves very little gas. The same incentive offered to NGI firms achieves almost all the savings of a uniform scheme, at the same cost. These costs assume that the baseline equals what firms would have consumed without the policy. With a higher baseline, the scheme would also pay for reductions that would have occurred anyway, and would cost more.

Retail price caps. Panel (c) reports the effects of a retail price cap that lowers the price paid by capped firms by €0.50/Smc. For any set of eligible firms, a cap distorts the market more than a subsidy of the same size. Because capped firms do not cut their consumption when the producer price rises (Section 2.d), a uniform cap raises the producer price by 62.1 percent and consumption by 12.4 percent, more than twice as much as a uniform subsidy of the same size. It costs €8.88 billion: on top of the mechanical transfer of €3.35 billion, the government pays €4.55 billion for the price increase that the cap itself causes. A cap limited to GI firms has small effects, like the corresponding subsidy. A cap limited to NGI firms raises the producer price by 51.0 percent and costs €2.22 billion, more than two and a half times the corresponding subsidy.

A policy mix. The instruments can be combined to give relief to inelastic, high-consumption firms and voluntary demand reduction incentives to elastic ones. Panel (d) reports a €0.50/Smc subsidy to GI firms combined with a €0.50/Smc voluntary demand reduction incentive for NGI firms. This mix lowers gas consumption by 4.4 percent and the producer price by 21.9 percent, at a fiscal cost of €2.93 billion, mostly the transfer to GI firms. With the average-firm elasticity for both groups, the same mix would instead raise gas consumption by 4.6 percent: GI firms would then appear responsive, and the extra consumption induced by the subsidy to them would exceed the savings of NGI firms. It obtains gas savings from the elastic part of demand while directing support to the firms least able to adjust. We do not claim that this mix is optimal, which would require a welfare objective (Section 7.e). The calculation shows only that, once heterogeneity is taken into account, the conflict between relief and conservation is smaller than the average-firm elasticity would suggest.

7.e Scope

These calculations quantify the effects of support on the producer price, aggregate consumption, the consumption of each group, and public spending. They do not rank policies by welfare. A welfare analysis would also require assumptions or estimates on firm exit, employment, output losses, competitiveness and strategic considerations, distributional weights, the marginal cost of public funds, and externalities. Lower support to elastic firms, for instance, may come with output or employment losses, or additional fiscal externalities that our calculation does not value. Any such analysis, however, would need the quantities estimated here: how strongly different firms respond to energy prices, how much of the market each group represents, and how support to one group affects the others through the market price.

8 Conclusions

This paper asks how industrial firms respond to large, sudden energy price increases, and what their responses imply for the design of support measures. The 2021–22 energy crisis pushed natural gas prices to unprecedented levels and led governments across the EU to intervene at substantial fiscal cost. With geopolitical tensions over energy supply persisting and European industry still heavily dependent on imported gas, these questions are likely to arise again.

We provide causal firm-level evidence on industrial energy demand during the crisis, exploiting the staggered expiration of fixed-price energy contracts that Italian industrial firms held before it began. Contract expiration sharply raises the unit costs firms pay. Electricity consumption does not respond. Gas consumption falls substantially, but only in the second half of 2022, when markets came to expect a prolonged shortage.

The average gas response hides large differences across firms. Gas-intensive firms, whose demand is almost perfectly inelastic, account for most industrial gas consumption; other firms respond substantially but consume much less. As a result, the average-firm elasticity, about -1.1, is a poor guide to policy incidence: the consumption-weighted elasticity, which standard incidence formulas require for an untargeted policy, is about -0.3.

This gap changes the predicted effects of crisis support. Using the average-firm elasticity would substantially overstate the price and quantity distortions of a uniform subsidy. Subsidies restricted to gas-intensive firms create far smaller equilibrium distortions than untargeted subsidies, although they remain fiscally costly because they transfer resources on large infra-marginal consumption. Voluntary demand reduction incentives, which pay firms for verified reductions, achieve the largest savings where demand is elastic. These conclusions concern

equilibrium prices, quantities, and fiscal costs, conditional on the crisis setting and our calibration; they do not imply that supporting energy-intensive firms is socially optimal, which would depend on margins, such as employment, exit, and the cost of public funds, that we do not model.

The main lesson is not specific to Italy. In markets with heterogeneous consumers, the elasticity of the average consumer can be the wrong statistic for policy incidence. During the European energy crisis, the distinction is quantitatively large because industrial gas consumption is concentrated among firms whose demand is unusually inelastic. Combining causal micro estimates with the appropriate aggregation changes the predicted effects of both untargeted and targeted support. A natural next step is to study the longer-run response, as firms invest in more energy-efficient capital, switch processes, or substitute toward alternative carriers. Whether the 2021–22 crisis led to lasting changes in industrial energy use, and whether current geopolitical risks are speeding up the energy transition by making self-generation and fuel diversification more valuable, are questions we leave for future research.

References

Constanza Abuin. Power decarbonization in a global energy market: The climate effect of US LNG exports. Working paper, 2024.

Lassi Ahlvik, Tuomas Kaariaho, Matti Liski, and Iivo Vehviläinen. Household-level responses to the European energy crisis. CEEPR Working Paper 2025-08, MIT Center for Energy and Environmental Policy Research, 2025.

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

Hunt Allcott, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow. The welfare effects of social media. *American Economic Review*, 110(3):629–676, 2020.

Matteo Alpino, Karen Evelyn Hauge, Andreas Kotsadam, and Simen Markussen. Effects of dialogue meetings on sickness absence—evidence from a large field experiment. *Journal of Health Economics*, 83:102615, 2022.

Matteo Alpino, Luca Citino, and Annalisa Frigo. The effects of the 2021 energy crisis on medium and large industrial firms: evidence from Italy. Questioni di Economia e Finanza (Occasional Papers) 776, Bank of Italy, 2023.

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

Theodore W Anderson and Herman Rubin. Estimation of the parameters of a single equation in a complete system of stochastic equations. *The Annals of Mathematical Statistics*, 20(1): 46–63, 1949.

Monica Andini, Guido De Blasio, Gilles Duranton, and William C Strange. Marshallian labour market pooling: Evidence from Italy. *Regional Science and Urban Economics*, 43(6): 1008–1022, 2013.

Isaiah Andrews, James H Stock, and Liyang Sun. Weak instruments in instrumental variables regression: Theory and practice. *Annual Review of Economics*, 11(1), 2019.

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

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

Dmitry Arkhangelsky, Susan Athey, David A Hirshberg, Guido W Imbens, and Stefan Wager. Synthetic difference-in-differences. *American Economic Review*, 111(12):4088–4118, 2021.

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

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

Maximilian Auffhammer and Edward Rubin. 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, 2018.

Banca d'Italia. Rapporto sulla stabilità finanziaria 2021/2, 2021. URL https://www.bancaditalia.it/pubblicazioni/rapporto-stabilita/2021-2/index.html. Accessed: 2026-05-20.

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

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

Joshua A Blonz. 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, 2022.

Stephen R Bond, Giacomo Rodano, and Nicolas Andre Benigno Serrano-Velarde. Investment dynamics in Italy: Financing constraints, demand and uncertainty. Questioni di Economia e Finanza (Occasional Papers) 283, Bank of Italy, 2015.

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

Kirill Borusyak, Xavier Jaravel, and Jann Spiess. Revisiting event-study designs: robust and efficient estimation. *Review of Economic Studies*, page rdae007, 2024.

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

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

Pierre Cahuc, Stéphane Carcillo, and André Zylberberg. *Labor economics*. MIT press, 2014.

Brantly Callaway and Pedro HC Sant'Anna. Difference-in-differences with multiple time periods. *Journal of econometrics*, 225(2):200–230, 2021.

Brantly Callaway, Andrew Goodman-Bacon, and Pedro HC Sant'Anna. Difference-in-differences with a continuous treatment. Technical report, National Bureau of Economic Research, 2024.

A. Colin Cameron, Jonah B. Gelbach, and Douglas L. Miller. Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics*, 90(3):414–427, 2008.

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

Raj Chetty, Adam Looney, and Kory Kroft. Salience and taxation: Theory and evidence. *American economic review*, 99(4):1145–1177, 2009.

Diego Ciccia. A short note on event-study synthetic difference-in-differences estimators. *arXiv preprint arXiv:2407.09565*, 2024.

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

Russell W. Cooper and John C. Haltiwanger. On the nature of capital adjustment costs. *The Review of Economic Studies*, 73(3):611–633, 2006. doi: 10.1111/j.1467-937X.2006.00389.x.

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

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

Clément De Chaisemartin and Xavier d'Haultfoeuille. Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–2996, 2020.

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

Mark Doms and Timothy Dunne. Capital adjustment patterns in manufacturing plants. *Review of Economic Dynamics*, 1(2):409–429, 1998. doi: 10.1006/redy.1998.0011.

Jean-Marie Dufour. Some impossibility theorems in econometrics with applications to structural and dynamic models. *Econometrica*, 65(6):1365–1387, 1997.

Jean-Marie Dufour and Mohamed Taamouti. Projection-based statistical inference in linear structural models with possibly weak instruments. *Econometrica*, 73(4):1351–1365, 2005.

Cédric Duprez and Glenn Magerman. Price updating in production networks. Technical report, NBB Working Paper, 2018.

Simone Emiliozzi and Filippo Favero. Unveiling natural gas consumption sectoral price elasticities. Questioni di Economia e Finanza (Occasional Papers) 941, Bank of Italy, June 2025. URL https://www.bancaditalia.it/pubblicazioni/qef/2025-0941/QEF_941_25.pdf.

European Commission. Energy prices and costs in Europe. Report from the Commission, COM(2024) 136 final, Brussels, 2024. URL https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52024DC0136.

European Commission. AccelerateEU – Energy Union: Affordable and secure energy through accelerated action. Communication from the Commission, COM(2026) 370 final, Brussels, 22 April 2026, 2026. URL https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0370.

Natalia Fabra, Clement Leblanc, and Mateus Souza. Unpacking the distributional implications of the energy crisis: Lessons from the Iberian electricity market. CEPR Discussion Paper 20593, CEPR Press, 2025.

Edgar C Fieller. Some problems in interval estimation. *Journal of the Royal Statistical Society: Series B (Methodological)*, 16(2):175–185, 1954.

Lionel Fontagné, Philippe Martin, and Gianluca Orefice. The many channels of firm's adjustment to energy shocks: evidence from France. *Economic Policy*, 39(117):5–43, 2024.

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

Andreas Gerster and Stefan Lamp. Energy tax exemptions and industrial production. *The Economic Journal*, page ueae048, 2024.

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

Daniel Gros. *Why gas price caps and consumer subsidies are both extremely costly and ultimately futile*. CEPS, 2022a.

Daniel Gros. The simple economics of consumer subsidies for natural gas (and why it matters for Putin's war). VoxEU.org, CEPR, 2022b. URL https://cepr.org/voxeu/columns/simple-economics-consumer-subsidies-natural-gas-and-why-it-matters-putins-war.

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

Ingar Haaland and Christopher Roth. Labor market concerns and support for immigration. *Journal of Public Economics*, 191:104256, 2020.

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

Daniel S. Hamermesh. *Labor Demand*. Princeton University Press, Princeton, NJ, 1993.

Bård Harstad and Katinka Holtsmark. The gas trap: Outcompeting coal vs. renewables. Technical report, National Bureau of Economic Research, 2024.

Julia Hatamyar, Noemi Kreif, Rudi Rocha, and Martin Huber. Machine learning for staggered difference-in-differences and dynamic treatment effect heterogeneity. *arXiv preprint arXiv:2310.11962*, 2023.

John R. Hicks. *The Theory of Wages*. Macmillan, London, 1932.

International Energy Agency. World Energy Outlook 2022. Technical report, IEA, Paris, 2022. URL https://www.iea.org/reports/world-energy-outlook-2022.

International Monetary Fund. The 2026 energy shock: How to deliver targeted and temporary support while encouraging energy conservation. Learning from the 2022 shock and best practice in 2026. Background note to remarks by O. Celasun at the Eurogroup meeting, Brussels, 4 May 2026, 2026. URL https://www.imf.org/-/media/files/countries/europe/imf-note-on-how-to-design-support-in-the-face-of-the-2026-energy-shock-europe.pdf.

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

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

Diego R Känzig. The macroeconomic effects of oil supply news: Evidence from OPEC announcements. *American Economic Review*, 111(4):1092–1125, 2021.

Lutz Kilian. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American economic review*, 99(3):1053–1069, 2009.

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

Francine Lafontaine and Kenneth J. White. Obtaining any Wald statistic you want. *Economics Letters*, 21(1):35–40, 1986.

Raphaël Lafrogne-Joussier, Julien Martin, and Isabelle Mejean. Cost pass-through and the rise of inflation. *Unpublished Manuscript*, 2023.

Christine Lagarde. The energy shock: Where we stand and what we need to know. Keynote speech at the 75th anniversary reception of the Association of German Banks, Berlin, 20 April 2026, 2026. URL https://www.ecb.europa.eu/press/key/date/2026/html/ecb.sp260420~cdf674023e.en.html.

Andrea Lamorgese, Megha Patnaik, Andrea Linarello, and Fabiano Schivardi. Management practices and resilience to shocks: Evidence from COVID-19. *Management Science*, 2024.

Jörg Leicher, Anne Giese, and Christoph Wieland. Electrification or hydrogen? the challenge of decarbonizing industrial (high-temperature) process heat. *J*, 7(4):439–456, 2024. doi: 10.3390/j7040026. URL https://doi.org/10.3390/j7040026.

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

Eva Lyubich, Joseph S. Shapiro, and Reed Walker. Regulating mismeasured pollution: Implications of firm heterogeneity for environmental policy. *AEA Papers and Proceedings*, 108: 136–142, 2018. doi: 10.1257/pandp.20181089.

Dharik S Mallapragada, Yury Dvorkin, Miguel A Modestino, Daniel V Esposito, Wilson A Smith, Bri-Mathias Hodge, Michael P Harold, Vincent M Donnelly, Alice Nuz, Casey Bloomquist, et al. Decarbonization of the chemical industry through electrification: Barriers and opportunities. *Joule*, 7(1):23–41, 2023.

Giovanni Marin and Francesco Vona. The impact of energy prices on socioeconomic and

environmental performance: Evidence from French manufacturing establishments, 1997–2015. *European Economic Review*, 135:103739, 2021.

Alfred Marshall. *Principles of Economics*. Macmillan, London, 8th edition, 1920.

Gustav Martinsson, László Sajtos, Per Strömberg, and Christian Thomann. The effect of carbon pricing on firm emissions: Evidence from the Swedish CO_2 tax. *The Review of Financial Studies*, 37(6):1848–1886, 2024.

Ben McWilliams, Giovanni Sgaravatti, Simone Tagliapietra, and Georg Zachmann. Europe's under-the-radar industrial policy: intervention in electricity pricing. Technical report, Bruegel, 2024.

Ben McWilliams, Thomas Mramor, Alexander Roth, and Simone Tagliapietra. 2026 European energy crisis fiscal response tracker. Bruegel Datasets, 2026. URL https://www.bruegel.org/dataset/2026-european-energy-crisis-fiscal-response-tracker. First published 28 April 2026.

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

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

Michele Polo and Carlo Scarpa. The liberalization of energy markets in Europe and Italy. IGIER Working Paper 230, IGIER, Bocconi University, 2003.

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

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

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

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

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

Giovanni Sgaravatti, Simone Tagliapietra, and Georg Zachmann. National policies to shield consumers from rising energy prices. *Bruegel Datasets*, 4, 2021.

Luigi Federico Signorini. Sustainable investment choices: emergencies and transition, June 2022. URL https://www.bancaditalia.it/pubblicazioni/interventi-direttorio/int-dir-2022/en-SIGNORINI-11-giugno-2022.pdf?language_id=1. Remarks by Senior Deputy Governor Luigi Federico Signorini at the Centesimus Annus Pro Pontifice Foundation.

Vojtech Sikl, Zuzana Havránková, Peter Kudela, and Anna Kudelova. Price elasticity of electricity revisited: A meta-analysis. IES Working Paper 17/2025, Charles University, Institute of Economic Studies, Prague, 2025. URL https://hdl.handle.net/10419/331536.

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

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

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

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

Max Wei, Colin A. McMillan, and Stephane de la Rue du Can. Electrification of industry: Potential, challenges and outlook. *Current Sustainable/Renewable Energy Reports*, 6(4): 140–148, 2019. doi: 10.1007/s40518-019-00136-1. URL https://doi.org/10.1007/s40518-019-00136-1.

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

Jeromin Zettelmeyer, Simone Tagliapietra, Georg Zachmann, and Conall Heussaff. Beating the European energy crisis. *Finance & Development*, December 2022.

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, combined with marginal pricing in the day-ahead power market, implies that shocks to the wholesale price of natural gas pass almost completely through to the wholesale price of electricity. Panel (b) of Figure 3.1 plots the daily wholesale prices of natural gas and electricity in Italy.[41] The price of gas was stable at around 30 euros per MWh until mid-2021, and then started to rise gradually above historical levels. The first major upswing came in the fall of 2021, when the price exceeded 100 euros per MWh; the second came in December 2021, when it almost reached 200 euros. After a temporary drop, the price surged to around 250 euros when Russia invaded Ukraine in February 2022, but declined rapidly thereafter and stayed below 100 euros until the summer. It then climbed again, reaching a historic peak of over 300 euros in late August 2022. Before the end of the year, the price fell sharply, rose again, and finally declined to about 70 euros.

Figure A.1 shows how manufacturing firms' expectations about the general economic situation evolved during the crisis; they reached their lowest point in the summer of 2022 (Section 5.b).

A.b The Italian retail energy market

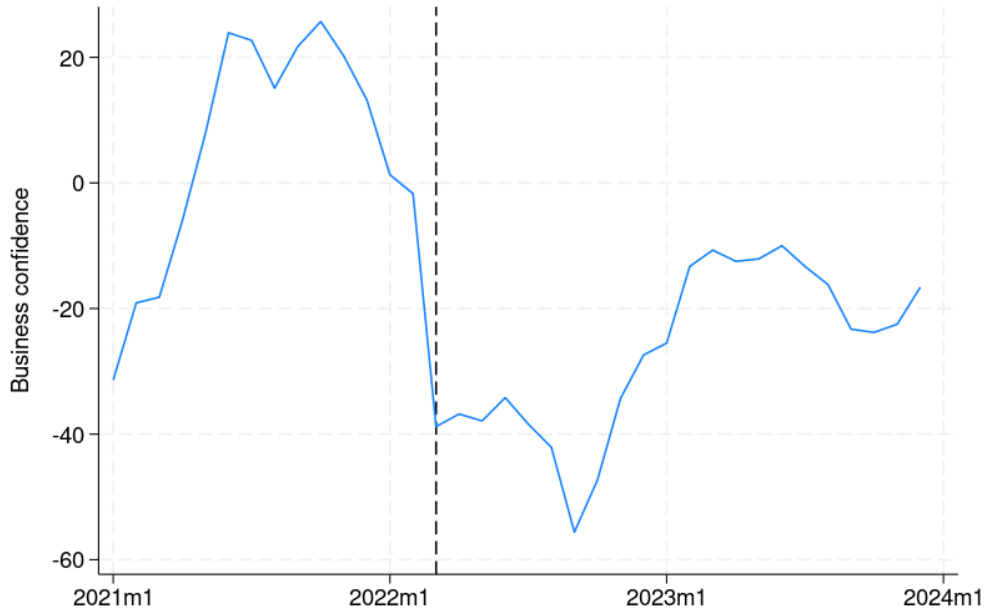
Following a broader trend in the European Union, the Italian retail markets for electricity and natural gas were gradually liberalized starting in the late 1990s (Polo and Scarpa, 2003). According to the Regulatory Authority for Energy, Networks and Environment (ARERA), neither market is highly concentrated (ARERA, 2022).[42]

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-price

[41] The two series are highly correlated (98% at the daily frequency). The price of gas is lower than that of electricity because gas is an input in power generation.

[42] Households and very small businesses, which are outside the scope of our analysis, could buy energy in a "protected" market (*mercato tutelato*), at prices periodically set by ARERA. As of 2024, the transition to a fully liberalized market, which phases out the protected market, was almost complete.

Figure A.1: Manufacturing firms' expectations about the economic situation



Source: Istat. *Note*: The figure shows the difference between the share of manufacturing firms expecting a positive general economic situation in the country over the next three months and the share expecting a negative one. The vertical line marks March 2022, the first full month after Russia invaded Ukraine.

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, to electricity customers with a low-voltage connection (a minority in our sample) and to gas customers consuming less than 200,000 standard cubic meters a year (one third of our sample), suppliers must also offer two standardized contracts, one with a variable and one with a fixed price. ARERA sets the conditions of these standardized contracts, while suppliers compete on price.[43] All contracts on offer, standardized or not, are published daily on a website managed by a government agency.[44] The portal does not cover contracts for large or energy-intensive firms, which usually negotiate *ad hoc* contracts directly with their suppliers.

A.c Fixed- and variable-price contracts in 2021–22

Electricity. In 2021, all firms with at least 50 employees, the firms we study, had to buy electricity in the liberalized market. Probably only a minority of firms in our sample have a low-voltage connection, for which suppliers must offer the standardized contracts designed by

[43] Standardized contracts are meant to let less sophisticated customers compare prices across sellers easily.
[44] https://www.ilportaleofferte.it/portaleOfferte/.

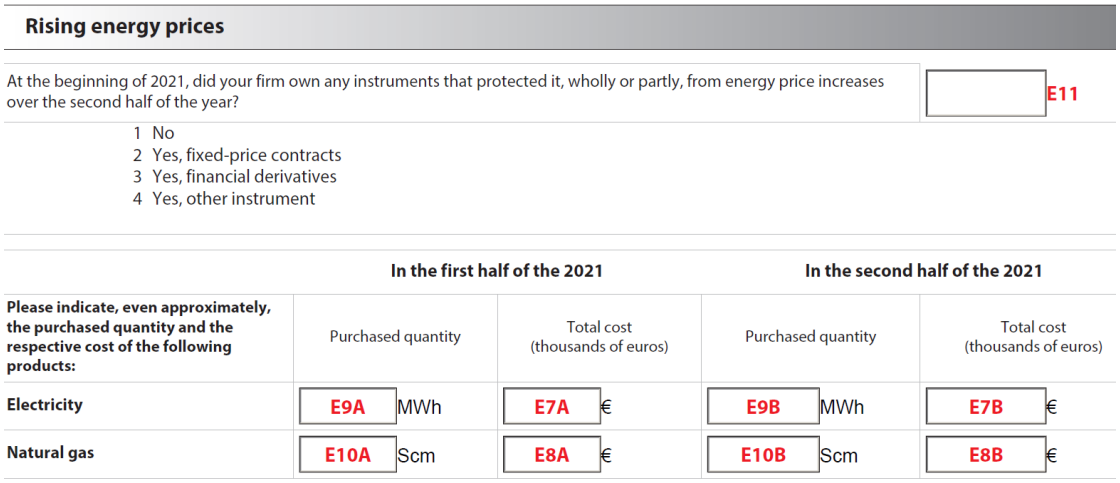
ARERA. According to ARERA, in 2021, 45% of business customers had a fixed-price contract, and the rest had a variable-price contract indexed to the wholesale price or to other price indices. In 2021, the average unit cost of energy under outstanding fixed-price contracts (both newly signed and signed in previous years) was 30% lower than under variable-price contracts. In 2022, the share of business customers with fixed-price contracts remained largely unchanged (47%), and their unit cost was 80% lower than under variable-price contracts, which were fully exposed to the increase in wholesale prices (ARERA, 2023). Some retailers continued to offer fixed-price contracts throughout the crisis, at least for low-voltage connections, for which ARERA publishes data. These offers, however, became rarer and more expensive. The average unit cost of outstanding fixed-price contracts (both newly signed and signed in previous years, standardized or not) rose from 91 euros/MWh in 2021 to 171 euros/MWh in 2022.

Natural gas. All firms must buy natural gas in the liberalized market; those consuming less than 200,000 standard cubic meters a year, about a third of our sample, must be offered the standardized contracts. In 2021, business customers were almost evenly split between fixed-price (44%) and variable-price contracts (56%) (ARERA, 2022). On average, the unit cost of energy was 30% higher under variable-price contracts, because of the increase in wholesale prices in the second half of the year. In 2022, the share of business customers with a variable-price contract rose to 63% (ARERA, 2023), and the cost gap between the two types of contracts widened to 50%, as wholesale prices rose further. The number of new fixed-price gas contracts on offer in the liberalized market fell to fewer than five in the first half of 2022, and to zero in the second. Firms consuming less than 200,000 standard cubic meters a year, however, could still buy the fixed-price standardized contract. The average unit cost of outstanding fixed-price contracts (both newly signed and signed in previous years, standardized or not) rose from 34 cents per cubic meter in 2021 to 77 cents in 2022, driven by new, more expensive contracts.

Appendix B Questionnaires

In this appendix we report the questions of the energy section of the Invind survey, in the English translation by the Bank of Italy (Figure B.1).

Figure B.1: Survey questions for the energy section



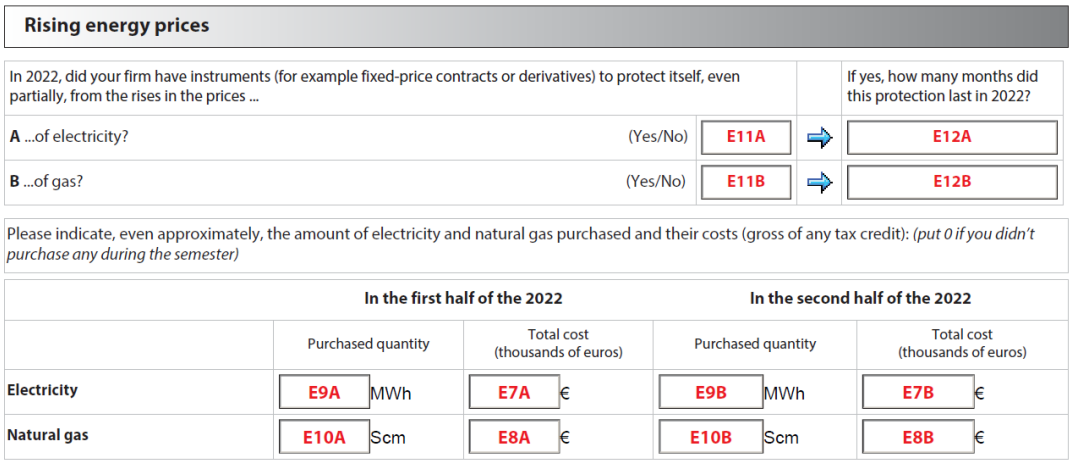
Rising energy prices

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

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

(a) 2021 wave



Rising energy prices

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

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

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

(b) 2022 wave

Note: The figure displays the questions of the energy section of the Invind survey in the 2021 and 2022 waves.

Appendix C Non-response bias

Each Invind wave includes about 1,800 respondents with at least 50 employees. The response rate to the energy section is about 50% in both waves, and we drop a small number of observations that fail the validation checks of Appendix D. We lose further observations (about 23%) because we include only firms that answered the energy section in both waves.

In this appendix we examine whether our findings are robust to selective item non-response, using inverse probability weighting (Wooldridge et al., 2002; Stantcheva, 2023). This method, commonly used to address differential attrition, uses the relationship between observed covariates and response to reweight the observed data so that they approximate the distribution in the full dataset (Stantcheva, 2023; Glynn and Quinn, 2010). In practice, we rerun our baseline specification weighting each observation by the inverse of the probability of being in the respective estimation sample. We estimate this probability with the logit models (16) and (17), where Λ is the logistic function, X_i includes covariates measured at baseline, and the outcome equals one for firms in the estimation sample and zero for the other firms surveyed in Invind, including those dropped for misreporting energy variables (Appendix D).[45]

$$\Pr(\text{Electricity sample}_i = 1 \mid X_i) = \Lambda(X_i'\beta^e) \tag{16}$$

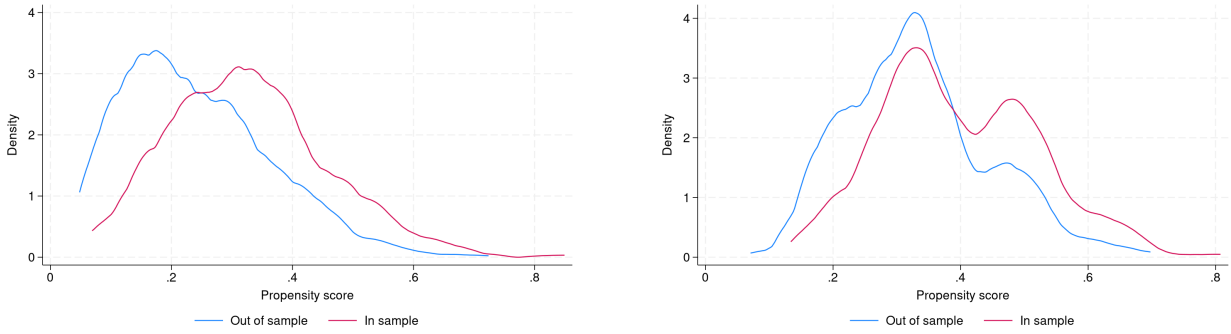
$$\Pr(\text{Gas sample}_i = 1 \mid X_i) = \Lambda(X_i'\beta^g) \tag{17}$$

Figure C.1 shows that the propensity scores of in-sample and out-of-sample firms have common support, for both electricity and gas. Following Becker and Ichino (2002), we verify that covariates are balanced within blocks of the propensity score.

Figure C.2 compares our baseline estimates with those obtained with inverse probability weighting. The two sets of estimates are very similar, which mitigates concerns that item non-response biases our results.

[45] We include total sales, total investment, and dummies for size class, sector, macro-region, EU ETS participation, and electricity- and gas-intensive status.

Figure C.1: Common support of propensity score

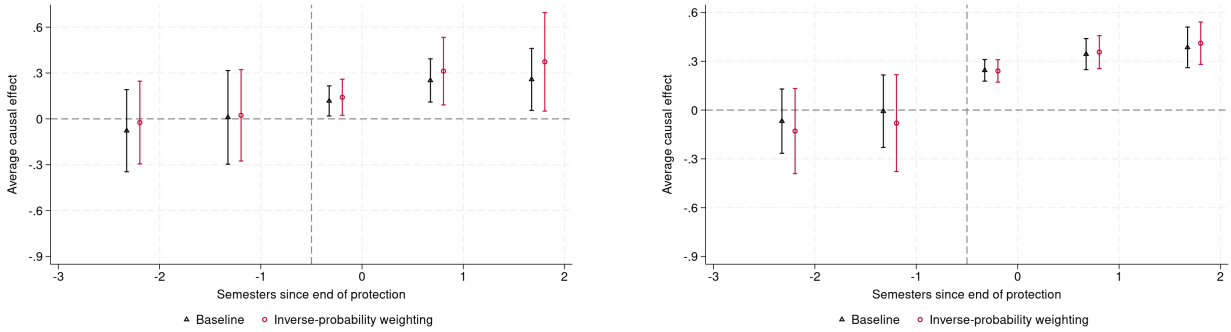


(a) Natural gas sample

(b) Electricity sample

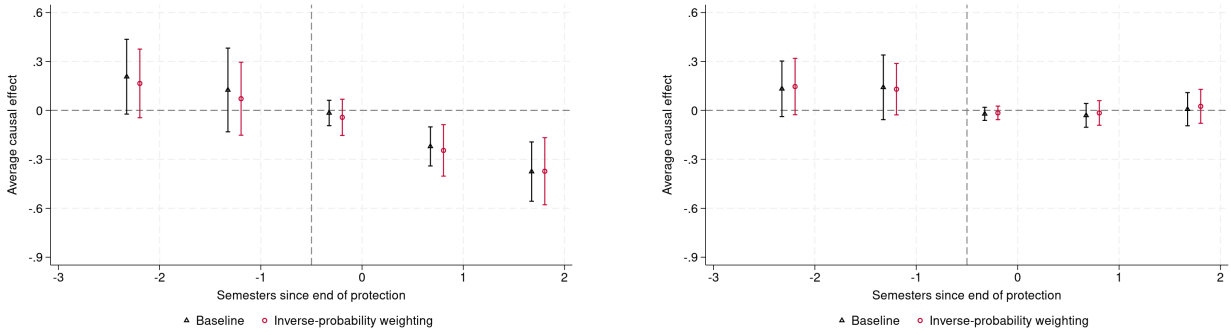
Note: The figure shows the distribution of the propensity score for in-sample and out-of-sample firms.

Figure C.2: Inverse-probability-weighted estimates



(a) Average unit cost of natural gas

(b) Average unit cost of electricity



(c) Quantity of natural gas

(d) Quantity of electricity

Note: The figures show average causal effects of the expiration of a fixed-price contract on the average unit costs of natural gas and electricity (panels (a) and (b)) and the corresponding 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 D Validation of survey answers

In this appendix we describe how we check the quality of the Invind energy data. Because respondents may not be familiar with physical units of measurement, we verify whether quantities and costs of gas and electricity take plausible values, using the following procedure.

1. First, we exclude firms that did not answer *all* the energy-related questions. Quantities must be strictly positive for the unit cost to be defined.

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

   (a) We compute the unit price paid by each firm in each half-year and compare it with the corresponding average price recorded by Eurostat for Italy.[46] We flag observations for which the unit price, computed by dividing expenditures by quantity, is either below half the minimum or above twice the maximum of the reference Eurostat price across consumption classes.[47]

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

3. Combining criteria (a) and (b), we identify 6 error categories for electricity and 4 for natural gas. For example, one category consists of firms that report electricity consumption in kWh instead of MWh; this mistake yields a unit price about 1,000 times lower than the Eurostat price. Table D.1 lists all categories and the reporting mistakes associated with each. We perform this exercise separately for each half-year. In 18.7% and 2.8% of electricity answers and in 21% and 6.4% of gas answers in the 2021 and 2022 waves, respectively, we observe the same mistake in both halves of the year. In these cases, we rescale the values to correct the mistake: in the previous example, we divide the reported quantity by 1,000 to convert it into MWh, and then recompute the unit price. The final estimation samples include only firms whose answers satisfy both criteria in all four half-years, either as reported or after rescaling.

4. For energy-intensive firms, we compare the quantities reported in the survey with administrative microdata collected by CSEA. Whenever the gap exceeds 35% in at least one half-year, we conservatively drop the firm from the estimation sample. Figure D.1

[46] The Invind and Eurostat prices are constructed similarly, as both include levies and taxes. However, Eurostat covers all non-household consumers, whereas our sample includes only industrial firms with at least 50 employees.

[47] Because prices rose over time, the upper bound is specific to each half-year, while the lower bound is constant over time and equal to half the minimum price reported by Eurostat in the first half of 2021.

plots administrative values against survey values, showing in blue the observations rescaled in the previous steps and in red those left unchanged. Both lie very close to the 45-degree line, which suggests that the survey data match administrative records and that our correction procedure works well.

5. For natural gas, we can perform a further check on firms with plants subject to the EU ETS, for which we observe confidential administrative records from ISPRA. Because the EU ETS applies at the plant level, a multi-plant firm may have some plants covered by it and others not. We therefore exclude firms whose survey quantities are smaller than the corresponding EU ETS record, but keep firms whose survey quantities are larger. Overall, only 2.8% (4.5%) of the electricity sample and 2.1% (2.5%) of the natural gas sample in 2021 (2022) show differences larger than 35% in absolute value between survey and administrative records. We drop these observations.

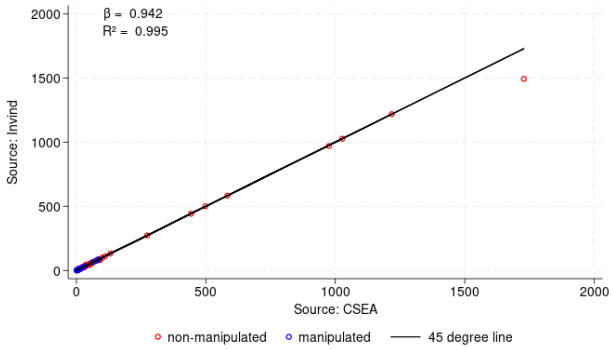
Table D.1: Validation of data quality: units of measurement in quantities and expenditure

(1) **Cost-share criterion**	(2) **Price-range criterion**	(3) **Expenditure**	(4) **Quantity**	(5) **Prevalence** 2021	(6) 2022
Panel A: Natural gas					
✓	✓	'000 €	Smc	70%	90%
✗ - upper tail	✗ - higher price (1,000-fold)	€	Smc	3%	0%
✓	✗ - higher price (1,000-fold)	'000 €	'000 Smc	18%	4.9%
✓	✗ - higher price (million-fold)	'000 €	million Smc	0%	0.7%
✗ - lower tail	✗ - lower price	million €	'000 Smc	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%

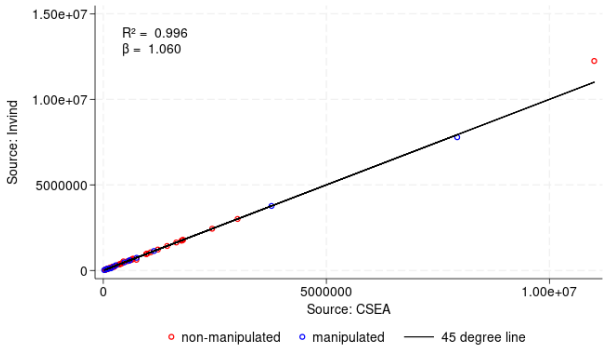
Note: The table presents results from the data validation procedure, separately for gas (Panel A) and electricity (Panel B). Depending on whether unit prices satisfy the two criteria (columns 1 and 2), observations are sorted into mutually exclusive categories of reporting mistakes. The two checks identify the units of measurement that the respondent most likely used (columns 3 and 4). We perform this exercise for each half-year. When we observe the same mistake in both halves, we rescale the values so that all observations are expressed in thousands of euros for expenditure, and in MWh and Smc for quantities of electricity and natural gas, respectively. We apply this correction to 18.7% and 2.8% of electricity answers and to 21% and 6.4% of gas answers in the 2021 and 2022 waves, respectively (columns 5 and 6). Shares are unweighted.

Figure D.1: Consistency with other data sources

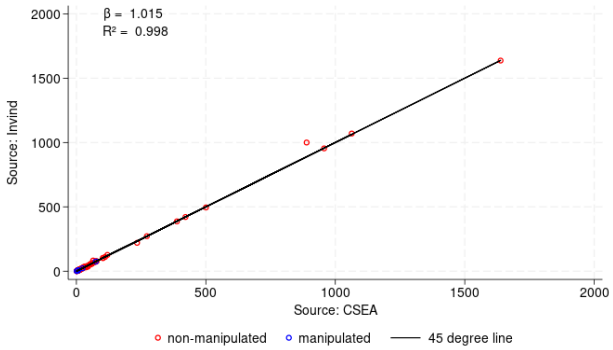
(a) Electricity consumption in 2021 (CSEA)



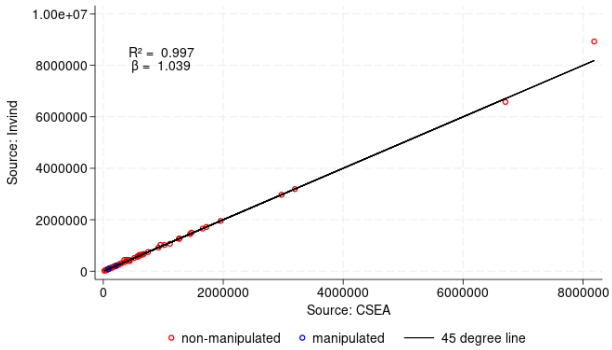
(b) Natural gas consumption in 2021 (CSEA)



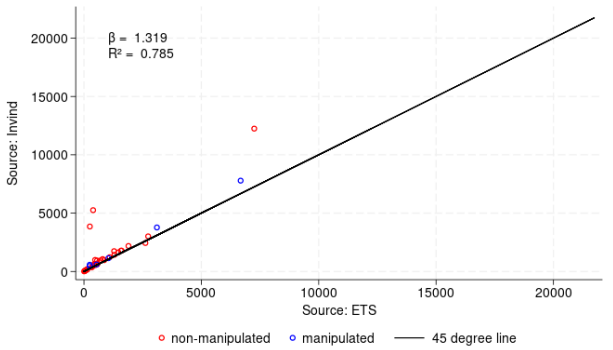
(c) Electricity consumption in 2022 (CSEA)



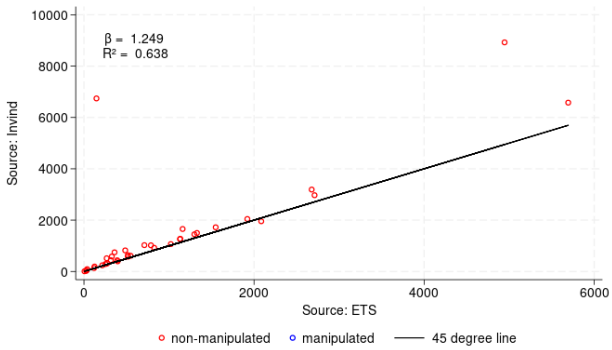
(d) Natural gas consumption in 2022 (CSEA)



(e) Natural gas consumption in 2021 (EU ETS)



(f) Natural gas consumption in 2022 (EU ETS)



Note: The figure compares quantities of electricity and natural gas reported in the Invind survey with those recorded in administrative data, for firms on the CSEA lists of electricity- and gas-intensive firms (*energivore* and *gasivore*; panels (a) and (b) for 2021, (c) and (d) for 2022) and for firms with plants in the EU ETS (panels (e) and (f)).

Appendix E 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.[48] 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 (half-yearly 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,

[48] Data downloaded from https://www.ilportaleofferte.it/portaleOfferte/opendata. We restrict the sample to non-domestic offers (TIPO_CLIENTE = 02) and extract and classify pricing components from XML files containing 1,479 electricity and 813 gas offers.

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.[49] 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, which finances support for renewable energy and cogeneration, and ARIM, which covers the remaining system charges. 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 E.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 and excise taxes (0.0125 €/kWh for electricity and 0.0128 €/Smc for gas, under Articles 52 and 26, respectively, of the consolidated excise law, *Testo unico delle accise*). For variable-price

[49] 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.

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. VAT (22%) applies proportionally to the entire bill and therefore cancels out in the ratio s, so we exclude it.

Results. Table E.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 $s = 0.67\%$, and even the 90th-percentile offer reaches only 1.02%. The average unit cost of gas is therefore a very close 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 E.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	Median	P75	P90
Gas (192,000 Smc/yr = sample P50; PSV = 1.20 €/Smc)	0.67	0.84	1.02
Electricity (4,500 MWh/yr = sample P50; PUN = 0.25 €/kWh)	8.29	9.67	10.09

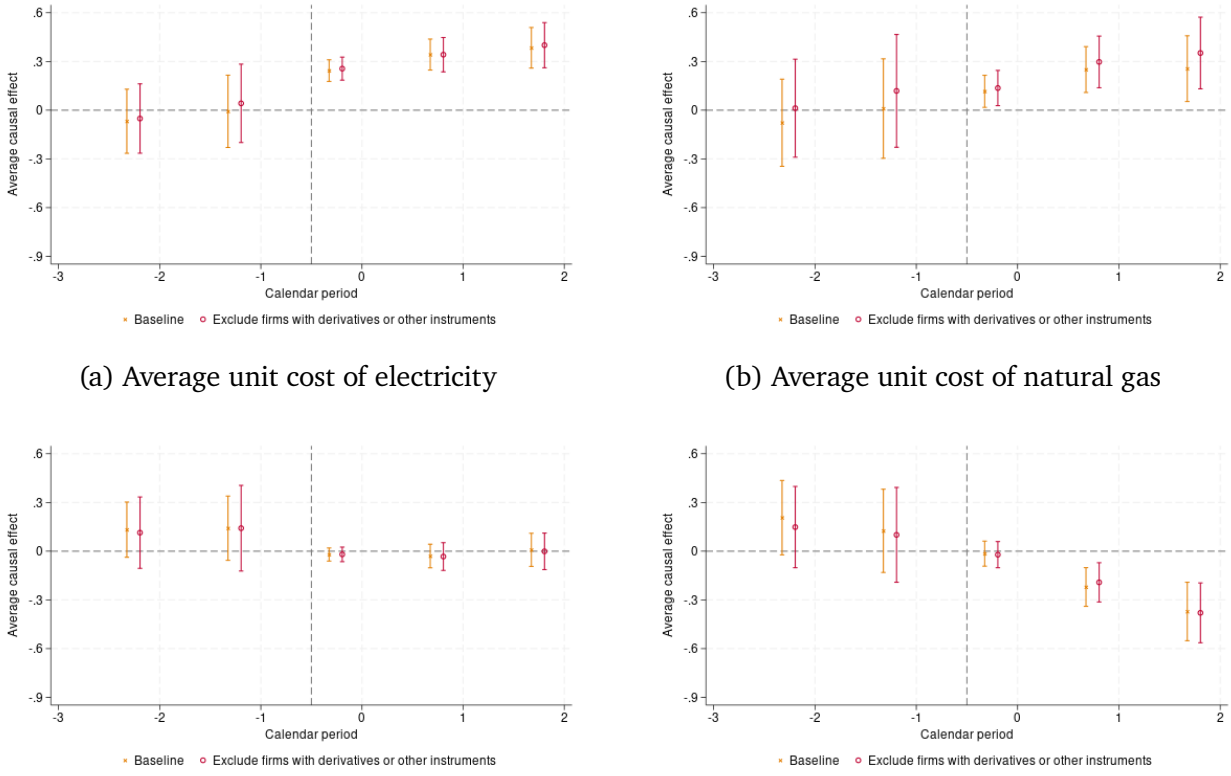
Notes: Distribution across all non-domestic offers (813 gas; 1,479 electricity; ARERA *Portale Offerte*, March 2026) of the fixed-charge share $s \equiv F/(F + p \cdot Q)$, evaluated at median sample consumption. F and p are constructed as described in the text. System charges are included in F at their rates before ARERA set them to zero in 2022, so these figures are upper bounds for actual 2022 invoices.

Appendix F Additional robustness tests

F.a Excluding firms protected by derivatives or other instruments

Some firms reported that they were protected at the beginning of 2021 by financial derivatives or other instruments rather than by fixed-price contracts; they account for 10.2% of the electricity sample and 10.7% of the gas sample (Section 4.d). Figure F.1 shows that excluding them leaves the estimates virtually unchanged.

Figure F.1: Effects of contract expiration, with and without firms protected by derivatives or other instruments



(a) Average unit cost of electricity

(b) Average unit cost of natural gas

(c) Quantity of electricity

(d) Quantity of natural gas

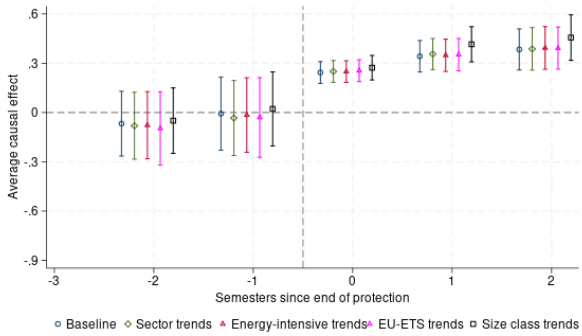
Note: The figure shows average causal effects of the expiration of a fixed-price contract on the average unit costs of electricity and natural gas (panels (a) and (b)) and on the corresponding quantities (panels (c) and (d)), for the full sample and for the sample that excludes firms protected by derivatives or other instruments. All estimates are weighted by survey weights. Standard errors are clustered at the firm level. Confidence intervals are at the 95% level.

F.b Controlling for covariate-specific non-parametric time trends

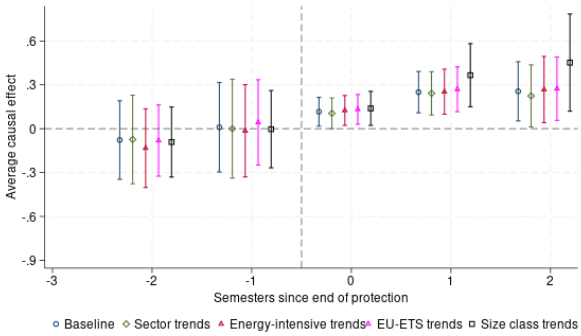
We extend the baseline specification with covariate-specific non-parametric time trends: we add to the Borusyak et al. (2024) estimator full sets of interactions between time fixed effects and (i) 11 sector dummies, (ii) an energy-intensive status indicator, (iii) a dummy for EU ETS participation, and (iv) firm size class dummies (five categories based on number of employees). This lets each group follow its own path over time, which addresses concerns that treatment timing is correlated with group-specific dynamics.

Figure F.2 displays the results, together with estimates from our baseline specification. The estimates are virtually unchanged relative to the baseline, both in magnitude and statistical significance, so differential trends across these groups do not appear to drive our results.

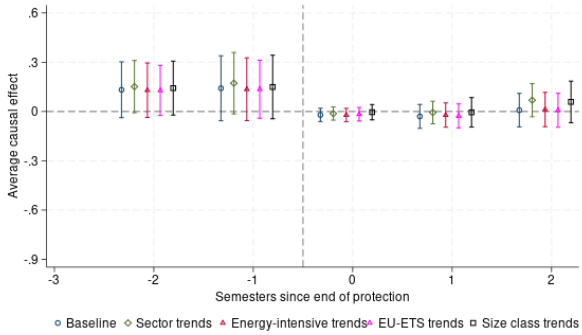
Figure F.2: Robustness to covariate-specific non-parametric time trends



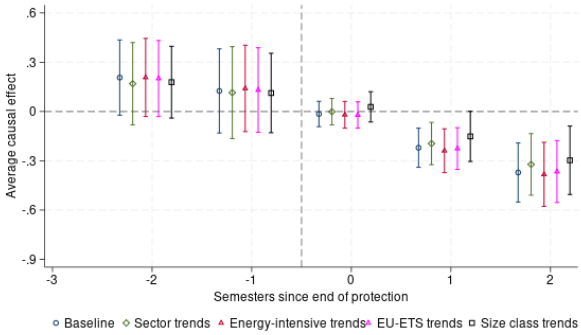
(a) Average unit cost of electricity



(b) Average unit cost of natural gas



(c) Quantity of electricity



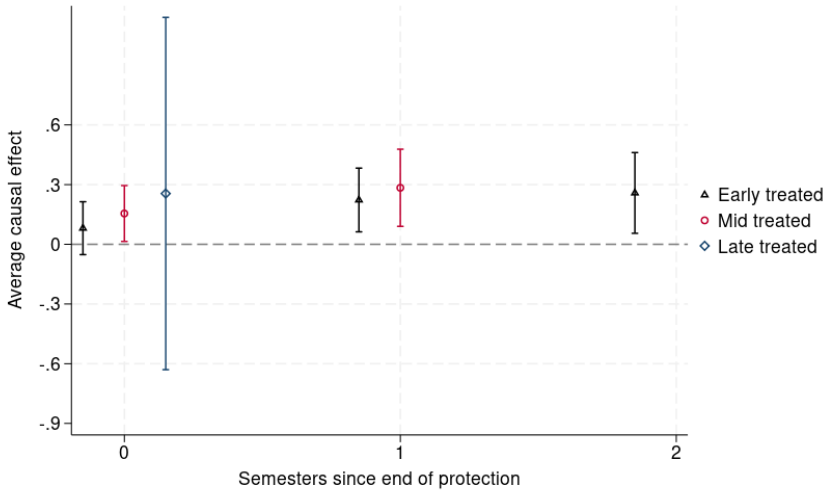
(d) Quantity of natural gas

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 event-study estimator of Borusyak et al. (2024), as described in Section 4.e. 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 plant covered by 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.

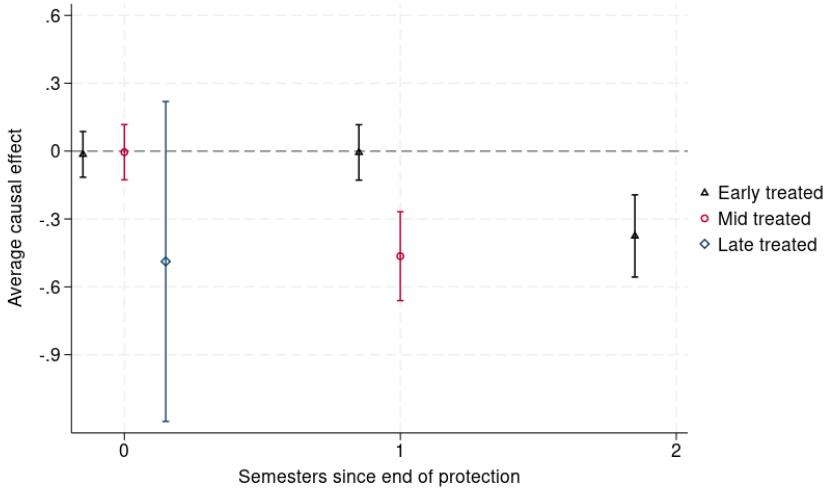
F.c Effects by cohort

Figure F.3 reports the effects of contract expiration on the unit cost and consumption of natural gas separately for each cohort. Gas consumption falls only in 2022h2, in every cohort, regardless of when its contracts expired (Section 5.b). Appendix G reports the corresponding synthetic difference-in-differences estimates for both energy carriers.

Figure F.3: Natural gas: heterogeneous effects by cohort



(a) Average unit cost



(b) Quantity

Note: The estimates in the upper and lower panels come from the same regressions as panels (b) and (d) of Figure 5.1, respectively, and are reported separately by cohort. All estimates are weighted by survey weights. Standard errors are clustered at the firm level. Confidence intervals are at the 95% level.

F.d Robustness to firm-specific seasonality

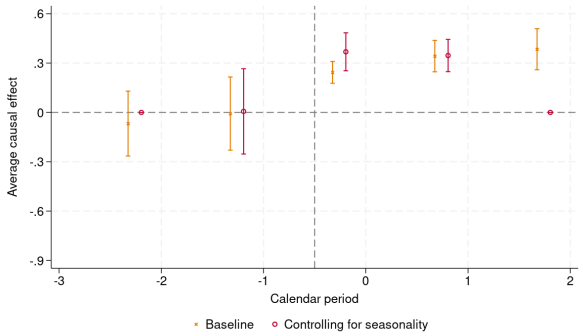
To check that our results are not driven by firm-specific seasonal patterns, we add firm × half-of-year fixed effects to the imputation step (the first stage) of the Borusyak et al. (2024) estimator. These fixed effects absorb any time-invariant difference between a firm's first and second half of the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

A mechanical limitation arises from the short panel. With firm × half-of-year 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).[50]

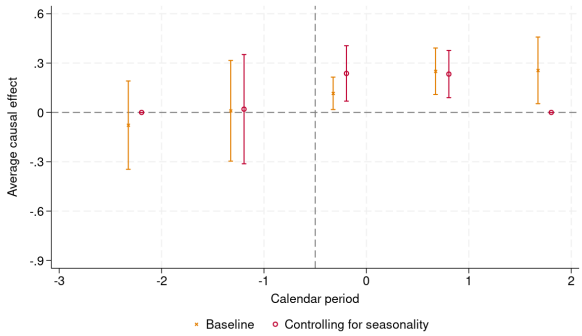
Figure F.4 displays the results alongside the baseline estimates. Across all four outcomes, the τ_0 and τ_1 point estimates are stable in magnitude and the confidence intervals overlap closely with the baseline, confirming that seasonality does not account for our findings.

[50] Mid-treated firms (E_i = 2022h1) and late-treated firms (E_i = 2022h2) retain at least one untreated observation in each half of the year and are therefore included, though with a single untreated observation per half of the year per firm in some cases.

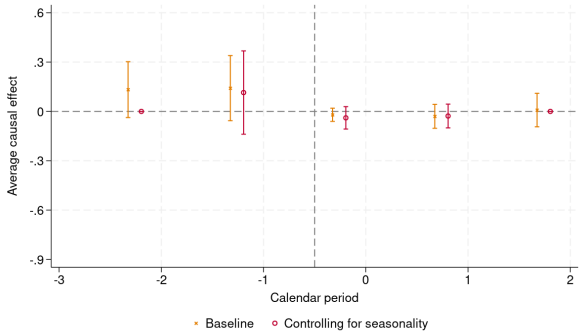
Figure F.4: Robustness to firm-specific seasonality



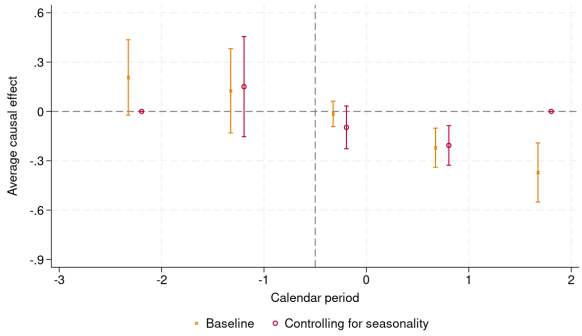
(a) Average unit cost of electricity



(b) Average unit cost of natural gas



(c) Quantity of electricity



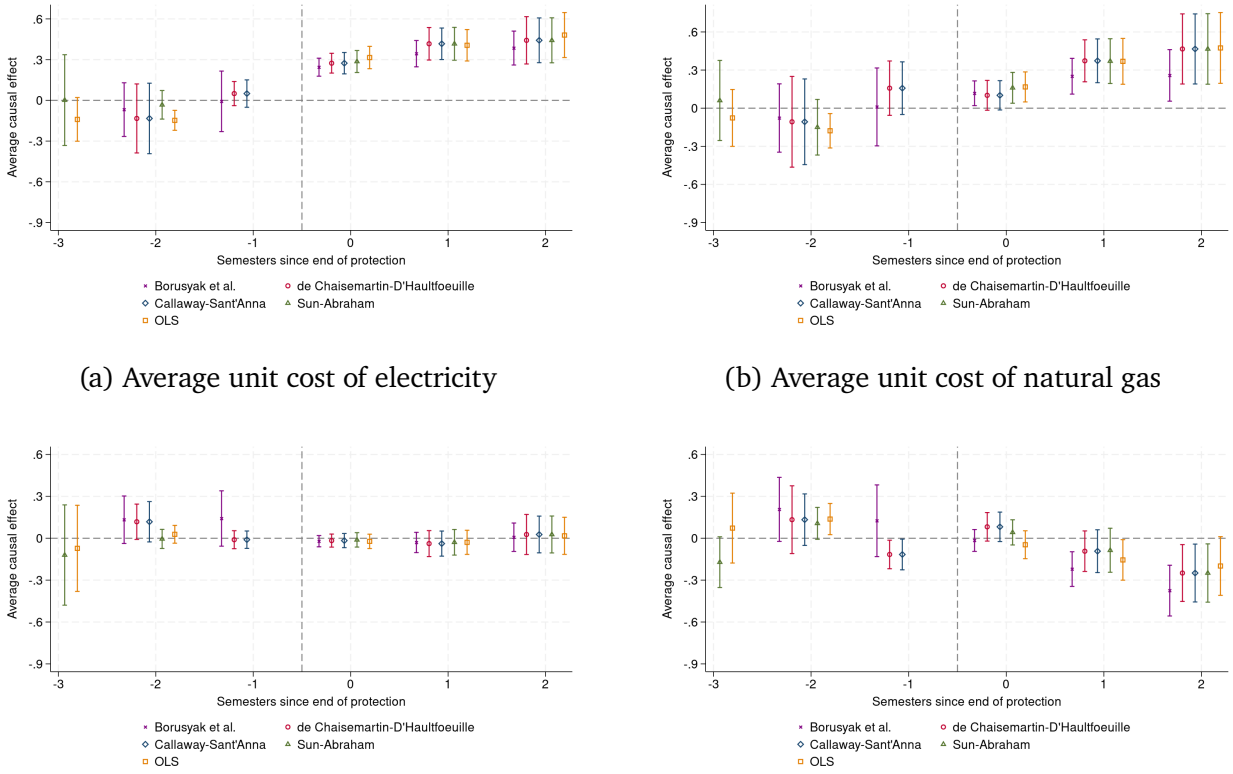
(d) Quantity of natural gas

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 × half-of-year fixed effects to the imputation step (orange circles). Outcome variables are in logs. Average causal effects are estimated using the event-study estimator of Borusyak et al. (2024). Because early-treated firms (E_i = 2021h2) have no untreated observation in the second half of the year, they are excluded at the τ_0 and τ_2 horizons; 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.

F.e Alternative estimators

Figure F.5 compares the baseline estimates with those obtained with the estimators of De Chaisemartin and d'Haultfoeuille (2020), Callaway and Sant'Anna (2021), and Sun and Abraham (2021), and with OLS.

Figure F.5: Baseline results with alternative difference-in-differences estimators



(a) Average unit cost of electricity

(b) Average unit cost of natural gas

(c) Quantity of electricity

(d) Quantity of natural gas

Note: The figure shows average causal effects of the end of price protection on the average unit costs of electricity and natural gas (panels (a) and (b)) and on the corresponding 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.

Appendix G Synthetic difference-in-differences by cohort

This appendix reports the cohort-level synthetic difference-in-differences (SDID) estimates discussed in Section 5.c. Each cohort is compared with a synthetic control drawn from the pure control group (Clarke et al., 2024), and Figure 5.2 averages the same cohort-level estimates by event time. Here we report them by calendar half-year, separately for the early (2021h2), mid (2022h1), and late (2022h2) cohorts, for unit costs (Figure G.1) and quantities (Figure G.2).

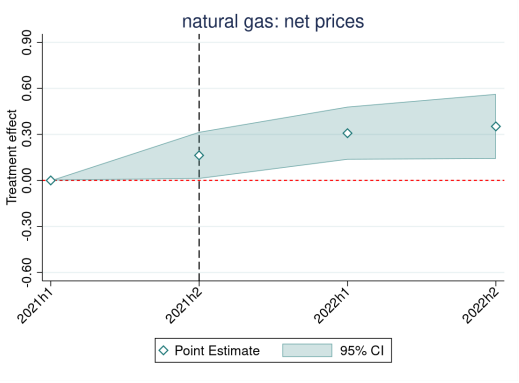
In every cohort, the unit costs of both electricity and natural gas rise in the half-year in which contracts expired, which supports our coding of the expiration dates. Gas consumption, by contrast, falls only in 2022h2, by 35%, 32%, and 33% in the early, mid, and late cohorts. Before 2022h2, the gas estimates are not statistically significant, even for the early cohort, whose contracts had expired a year earlier. Electricity consumption does not respond in the early and mid cohorts. In the late cohort it falls by about 10% in 2022h2; this cohort includes only 11 firms in the electricity sample, and 7 in the gas sample.

Figure G.1: Synthetic difference-in-differences estimates by cohort: average unit costs



(a)

Electricity, early cohort (2021h2)



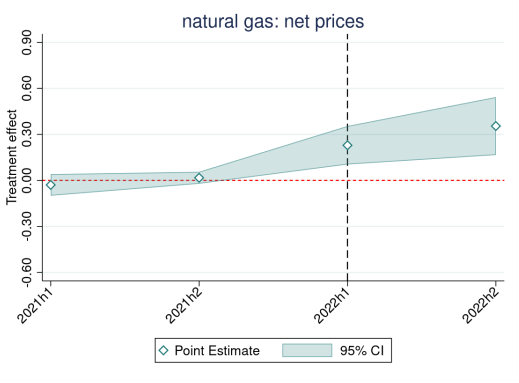
(b)

Natural gas, early cohort (2021h2)



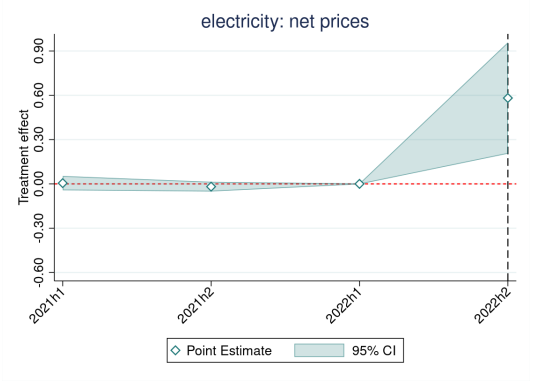
(c)

Electricity, mid cohort (2022h1)



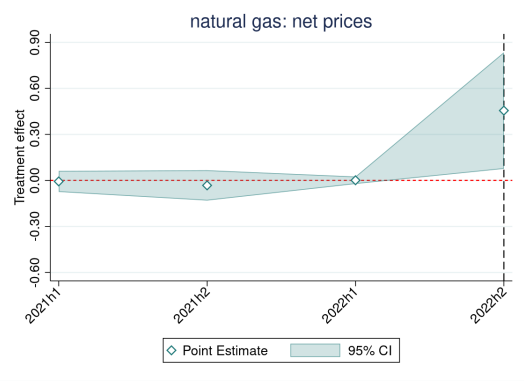
(d)

Natural gas, mid cohort (2022h1)



(e)

Electricity, late cohort (2022h2)

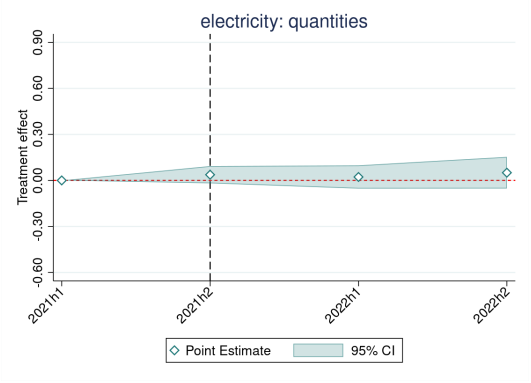


(f)

Natural gas, late cohort (2022h2)

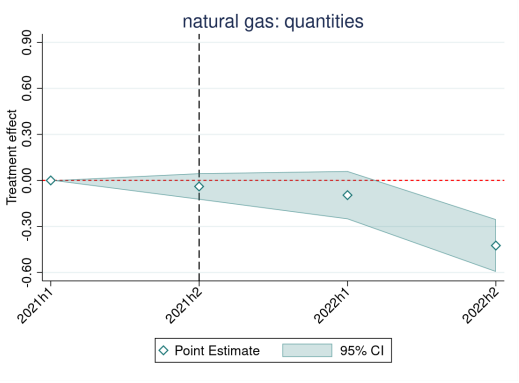
Note: The figure shows the effects of the expiration of a fixed-price contract on the log average unit cost, net of government support, of electricity (left column) and natural gas (right column), separately for each cohort and by calendar half-year. The dashed vertical line marks the half-year in which the cohort's contracts expired. Estimates use the synthetic difference-in-differences estimator (Arkhangelsky et al., 2021) in the staggered version of Clarke et al. (2024), in which each cohort is compared with a synthetic control drawn from the pure control group; Figure 5.2 averages the same estimates by event time. Because the SDID weights are fitted on the pre-expiration half-years, the estimates before expiration are close to zero by construction. The estimates are not weighted, because the SDID estimator does not support survey weights. Confidence intervals are at the 95% level, from a bootstrap that resamples firms.

Figure G.2: Synthetic difference-in-differences estimates by cohort: quantities



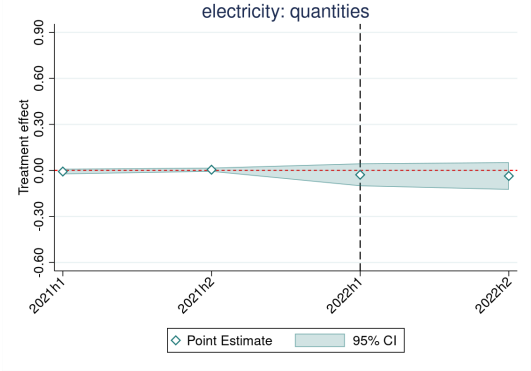
(a)

Electricity, early cohort (2021h2)



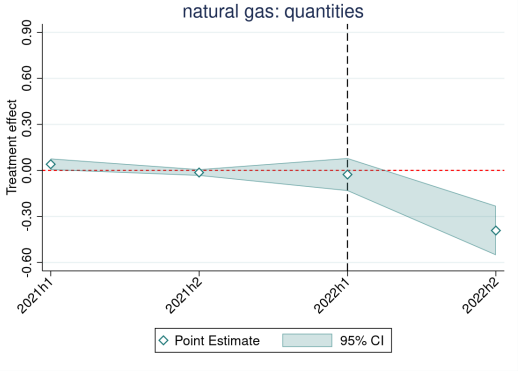
(b)

Natural gas, early cohort (2021h2)



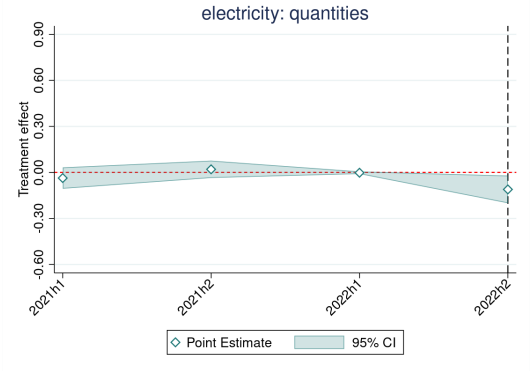
(c)

Electricity, mid cohort (2022h1)



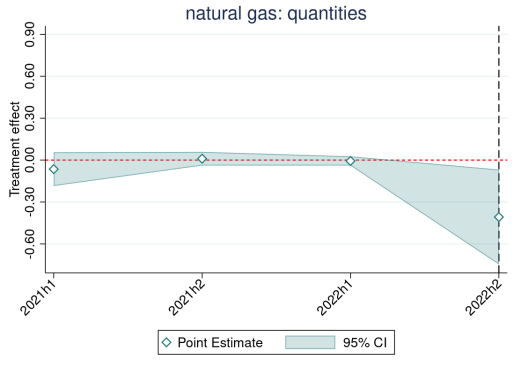
(d)

Natural gas, mid cohort (2022h1)



(e)

Electricity, late cohort (2022h2)



(f)

Natural gas, late cohort (2022h2)

Note: The figure shows the effects of the expiration of a fixed-price contract on the log purchased quantity of electricity (left column) and natural gas (right column), separately for each cohort and by calendar half-year. The dashed vertical line marks the half-year in which the cohort's contracts expired. Estimation and inference are as in Figure G.1.

Appendix H Firm exit during the sample period

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

Table H.1: Exit and distress events among sample firms, 2021–22

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 2021–22. Non-liquidatory distress procedures include court-supervised administration, debt restructuring agreements, extraordinary administration, and *concordato preventivo*.

Table H.1 shows that firm exit is negligible in our setting. This is true even for energy-intensive firms, which faced the largest cost shocks. Rather than exiting, these firms passed costs on to output prices: 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 consistent with the near-zero electricity demand elasticity and the output-price pass-through documented in Sections 5.a and 6.d, and indicate that the balanced-panel restriction does not introduce material survivor bias.

Appendix I Summary statistics

Table I.1 reports summary statistics for our two estimation samples, for electricity and for natural gas. Firm characteristics are measured before the crisis, in the first half of 2021. The samples comprise industrial firms with at least 50 employees, as determined by the Invind sampling design. All statistics are weighted by survey weights. The number of observations in the estimation samples is four times the number of firms, one for each half-year.

The electricity and natural gas samples have similar characteristics. About half of the firms belong to the metalworking industry, while the rest are distributed fairly evenly across other sectors, except for non-metallic minerals and water & waste, which are small. About 80% of firms are located in Northern Italy, and only 5–7% in the South and Islands.

On average, firms in the electricity sample pay €0.16 per kWh and consume approximately 6,000 MWh of electricity per half-year. For natural gas, the average price is €10 per GJ, and average consumption is about 1.6 million standard cubic meters per half-year. About 30% of firms in both samples are energy-intensive (*energivore* or *gasivore*), and 3–5% operate plants subject to the EU ETS. About half of the firms in the gas sample report that gas is an essential input.

Treatment cohorts are unevenly sized. Only 18% of firms in the electricity sample and 13% in the gas sample belong to the pure control group. The largest cohorts are the early treated (exposed in 2021h2) and the mid treated (exposed in 2022h1), while the late treated (exposed in 2022h2) are a small minority (3% and 1%, respectively).

Finally, sample firms are relatively large, with average sales of €85–97 million and more than 200 employees, although both distributions are highly skewed.

In Tables I.2 and I.3 we report balancing tables across different treatment cohorts. We highlight some significant level differences across groups.

Table I.1: Summary statistics for the electricity and gas samples

	Electricity sample (1)	Gas sample (2)
Sectoral composition		
Food and beverages	0.080	0.063
Textiles & apparel	0.120	0.107
Chem., pharma., rubber	0.117	0.146
Non-metallic minerals	0.035	0.036
Metalworking industry	0.515	0.513
Wood, paper, furniture	0.100	0.110
Water & waste	0.033	0.025
Macroarea		
North-West	0.396	0.422
North-East	0.395	0.396
Center	0.141	0.130
South and Islands	0.068	0.052
Energy-related variables		
Price of electricity (€/kWh)	0.160	
Price of natural gas (€/GJ)		10.41
Quantity of electricity (GWh)	6.05	
Quantity of natural gas (million smc)		1.57
Energy-intensive firm (0/1)	0.299	0.286
Subject to the EU ETS (0/1)	0.034	0.051
Gas is an indispensable input* (0/1)		0.535
Treatment cohorts		
Pure control	0.183	0.134
Early treated (exposed 2021h2)	0.437	0.454
Mid treated (exposed 2022h1)	0.350	0.401
Late treated (exposed 2022h2)	0.030	0.011
Other firm characteristics		
Sales (million €)	85.3	96.9
Employees	202	223
Number of firms	413	308

Note: Source: Invind. The table reports means for the estimation samples used in Section 5, weighted by survey weights. Characteristics are measured at baseline, in the first half of 2021, so the number of observations equals the number of firms. *Whether gas is an indispensable (essential) input comes from the Bank of Italy's Business Outlook Survey of Industrial and Service Firms (Sondtel) and refers to the beginning of 2022; it is available for 287 of the 308 firms in the gas sample, and the share is computed over these firms.

Table I.2: Balance across treatment cohorts: 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 orthogonality test p-val
Food and beverages	0.080	0.127	0.182	0.158	0.161	0.242	0.061	0.603	0.521	0.842	0.236
	(0.020)	(0.027)	(0.122)	(0.042)							
Textiles & apparel	0.074	0.107	0.000	0.118	0.301	0.353	0.252	0.256	0.791	0.233	0.421
	(0.020)	(0.025)	(0.000)	(0.037)							
Chem., pharma., rubber	0.136	0.153	0.182	0.171	0.665	0.674	0.477	0.803	0.732	0.931	0.893
	(0.026)	(0.030)	(0.122)	(0.043)							
Non-metallic minerals	0.034	0.073	0.091	0.039	0.113	0.338	0.833	0.832	0.321	0.452	0.369
	(0.014)	(0.021)	(0.091)	(0.022)							
Metalworking industry	0.523	0.400	0.364	0.461	0.027	0.308	0.367	0.813	0.386	0.551	0.148
	(0.038)	(0.040)	(0.152)	(0.058)							
Wood, paper, furniture	0.108	0.113	0.182	0.026	0.878	0.454	0.031	0.500	0.026	0.021	0.115
	(0.023)	(0.026)	(0.122)	(0.018)							
Water & waste	0.045	0.027	0.000	0.026	0.371	0.473	0.477	0.586	0.988	0.591	0.690
	(0.016)	(0.013)	(0.000)	(0.018)							
ETS coverage	0.074	0.067	0.182	0.092	0.801	0.203	0.625	0.162	0.496	0.367	0.538
	(0.020)	(0.020)	(0.122)	(0.033)							
Energy-intensive firm	0.284	0.340	0.455	0.355	0.278	0.231	0.262	0.444	0.821	0.529	0.451
	(0.034)	(0.039)	(0.157)	(0.055)							
Gas essential input	0.503	0.553	0.636	0.603	0.383	0.395	0.166	0.595	0.499	0.836	0.479
	(0.039)	(0.042)	(0.152)	(0.060)							
Employees (yearly average)	332.227	304.427	452.091	1034.829	0.639	0.490	0.032	0.385	0.039	0.653	0.032
	(40.791)	(42.578)	(240.343)	(488.038)							
Sales (million euros)	158.528	121.912	166.633	588.517	0.356	0.951	0.033	0.563	0.029	0.592	0.026
	(32.423)	(19.827)	(91.251)	(296.836)							

Notes: Source: Invind; whether gas is an essential input comes from Sondtel and is averaged over the firms that answered. The table reports means and standard errors of the mean (in parentheses) by treatment cohort, with one observation per firm; characteristics are measured at baseline, in the first half of 2021. Means are not weighted by survey weights, so they describe the estimation sample (Table I.1 reports weighted means). Columns (5)–(10) report p-values from pairwise t-tests of equal means; column (11) reports p-values from F-tests that means are equal across the four cohorts.

Table I.3: Balance across treatment cohorts: 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 orthogonality test p-val
Food and beverages	0.055	0.105	0.143	0.163	0.142	0.340	0.020	0.754	0.292	0.893	0.143
	(0.020)	(0.028)	(0.143)	(0.053)							
Textiles & apparel	0.102	0.073	0.000	0.102	0.417	0.379	0.993	0.464	0.525	0.385	0.694
	(0.027)	(0.023)	(0.000)	(0.044)							
Chem., pharma., rubber	0.133	0.194	0.286	0.184	0.193	0.261	0.395	0.556	0.882	0.534	0.482
	(0.030)	(0.036)	(0.184)	(0.056)							
Non-metallic minerals	0.031	0.040	0.143	0.082	0.699	0.130	0.151	0.210	0.273	0.603	0.301
	(0.015)	(0.018)	(0.143)	(0.040)							
Metalworking industry	0.516	0.468	0.429	0.367	0.449	0.657	0.078	0.841	0.233	0.760	0.366
	(0.044)	(0.045)	(0.202)	(0.070)							
Wood, paper, furniture	0.141	0.089	0.000	0.082	0.198	0.290	0.290	0.414	0.882	0.442	0.377
	(0.031)	(0.026)	(0.000)	(0.040)							
Water & waste	0.023	0.032	0.000	0.020	0.672	0.685	0.904	0.633	0.677	0.709	0.926
	(0.013)	(0.016)	(0.000)	(0.020)							
ETS coverage	0.109	0.056	0.143	0.204	0.130	0.786	0.101	0.357	0.003	0.709	0.037
	(0.028)	(0.021)	(0.143)	(0.058)							
Energy-intensive firm	0.281	0.282	0.714	0.429	0.986	0.015	0.061	0.016	0.064	0.162	0.024
	(0.040)	(0.041)	(0.184)	(0.071)							
Gas essential input	0.582	0.526	1.000	0.860	0.386	0.041	0.001	0.023	0.000	0.339	0.000
	(0.045)	(0.047)	(0.000)	(0.053)							
Employees (yearly average)	337.164	332.331	605.000	1476.000	0.935	0.143	0.017	0.149	0.018	0.668	0.010
	(40.525)	(42.632)	(243.476)	(757.543)							
Sales (million euros)	159.749	133.354	647.554	801.606	0.479	0.003	0.025	0.000	0.021	0.900	0.013
	(29.394)	(22.519)	(479.478)	(455.266)							

Notes: Source: Invind; whether gas is an essential input comes from Sondtel and is averaged over the firms that answered. The table reports means and standard errors of the mean (in parentheses) by treatment cohort, with one observation per firm; characteristics are measured at baseline, in the first half of 2021. Means are not weighted by survey weights, so they describe the estimation sample (Table I.1 reports weighted means). Columns (5)–(10) report p-values from pairwise t-tests of equal means; column (11) reports p-values from F-tests that means are equal across the four cohorts.

Appendix J 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) unit costs and on (log) quantities obtained from the Borusyak et al. (2024) estimator. In practice, we compute a Wald-type elasticity: we convert the average treatment effects estimated in the quantity and price equations, which are in log points, into percent changes, and divide the first by the second.[51] Under the identifying assumptions of Section 4.f, we can construct an estimate of the elasticity θ as follows:

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

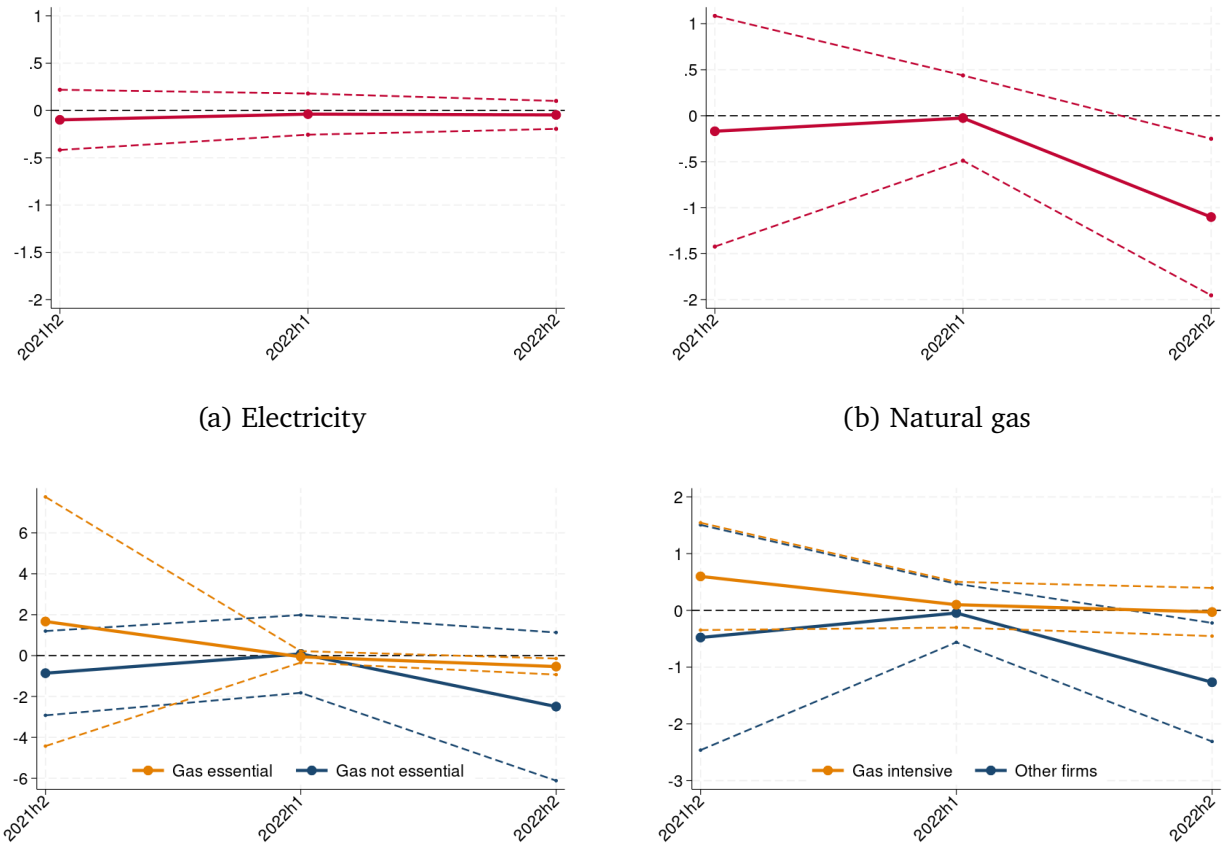
where $\hat{\tau}^q$ and $\hat{\tau}^p$ are the estimates of, respectively, the average treatment effect on the log quantity and on the log price of energy.

Inference. We compute delta-method standard errors from the BJS variance-covariance matrix of the two effects, which includes their covariance (Borusyak et al., 2024, Theorem 3). Because the delta method linearizes the ratio, it can be misleading when the denominator is estimated imprecisely relative to its size (Lafontaine and White, 1986), as it is here: the t-statistic of the 2022h2 effect on gas unit costs is 3.1. Our baseline confidence sets are therefore Fieller (1954) sets, which collect the values r for which the test of $(e^{\tau^q} - 1) - r\,(e^{\tau^p} - 1) = 0$ does not reject. This is the counterpart, for a ratio of two reduced-form effects, of the Anderson and Rubin (1949) test in instrumental-variable estimation (Andrews et al., 2019). These sets are bounded only if the effect on unit costs is significant; otherwise they are unbounded, as any valid confidence set for a ratio must be with positive probability (Dufour, 1997). The test of equal elasticities in two groups applies the same test to both groups jointly, and is conservative (Dufour and Taamouti, 2005). Finally, we report percentile intervals from a wild cluster bootstrap with Rademacher weights, clustered by firm and applied jointly to the price and quantity effects (Cameron et al., 2008); a pairs bootstrap that resamples firms gives similar intervals.

Figure J.1 plots the estimated price elasticity of demand by calendar period for the full sample and for selected subsamples.

[51] An alternative approach would be to relate firm-level treatment effects on quantities to those on unit costs, as in Deryugina et al. (2020). This is infeasible in our application because we construct unit costs as the ratio of expenditures to physical quantities (Section 4), so any measurement error in quantities translates into a measurement error of opposite sign in unit costs, which produces a negative correlation by construction. This is the "division bias" discussed by Borjas (1980) for regressions of hours on wages computed as the ratio of earnings to hours.

Figure J.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 panels (b), (c), and (d) combine effects from panels (b) and (d). The upper panels plot average-firm elasticities; the lower panels plot elasticities for selected subsamples. Confidence intervals are 95% delta-method intervals, which understate uncertainty when the effect on unit costs is imprecise; Table 6.1 reports our baseline Fieller confidence sets for 2022h2.

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

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

$$\epsilon(P) \equiv \frac{dQ}{dP}\frac{P}{Q(P)}, \tag{18}$$

which, unlike in Section 2, we write in signed form. For an isoelastic function, ϵ is constant, and the change in log quantity caused by any finite price change is simply $\epsilon\,\Delta \log P$. For non-isoelastic demand, the change in log quantity caused by a price change from P_0 to P_1 is

$$\log Q(P_1) - \log Q(P_0) = \int_{\log P_0}^{\log P_1} \epsilon(P)\, d\log P, \tag{19}$$

so the arc elasticity $\Delta \log Q / \Delta \log P$ is the average of local elasticities over $[\log P_0, \log P_1]$. 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

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

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

In this appendix we apply the Hicks–Marshall laws of derived demand (Marshall, 1920; Hicks, 1932) to show that it is theoretically ambiguous whether gas-intensive firms' demand for gas is more or less elastic than that of other firms. We use the version of the formula in Cahuc et al. (2014), which allows for non-constant returns to scale and markups. Consider a firm whose production function combines natural gas and a bundle of other inputs and is homogeneous of degree ν, and which faces an isoelastic, downward-sloping demand for its output. The uncompensated own-price elasticity of demand for natural gas, the parameter we estimate, can be written as

$$\varepsilon^{G}_{p_G} = -(1-s)\sigma - \frac{\mu}{\mu-\nu}s, \tag{20}$$

where s is the share of gas in total costs, σ is the elasticity of substitution between natural gas and other inputs, and μ is the markup. The first term is the substitution effect, holding output fixed, and the second is the scale effect; both work in the same direction. With constant returns to scale ($\nu = 1$), $\mu/(\mu-1)$ equals the price elasticity of demand for the firm's output, ϕ, and equation (20) reduces to the textbook formula $\varepsilon^{G}_{p_G} = -(1-s)\sigma - s\phi$ (Hamermesh, 1993). The larger the gas cost share s, the more the scale effect weighs relative to the substitution effect.

In our surveys, nearly all gas-intensive firms report gas as an essential input, while the share is lower among other firms. We therefore let σ differ between the two groups, and assume $\sigma = 0$ for gas-intensive (GI) firms and $\sigma > 0$ for other (NGI) firms. The demand elasticities of the two groups are then

$$\varepsilon^{G}_{p_G} = \begin{cases} -\frac{\mu}{\mu-\nu}s^{GI}, & \text{gas-intensive firms } (\sigma = 0), \\ -(1-s^{NGI})\sigma - \frac{\mu}{\mu-\nu}s^{NGI}, & \text{other firms } (\sigma > 0), \end{cases} \tag{21}$$

where $s^{GI} > s^{NGI}$ by definition. Gas demand is less elastic among gas-intensive firms than among other firms if the elasticity of substitution of the latter is high enough:

$$\sigma > \frac{\mu}{\mu-\nu}\frac{s^{GI}-s^{NGI}}{1-s^{NGI}}. \tag{22}$$

Appendix M Effect heterogeneity with machine learning

To study heterogeneity in the effect of contract expiration on gas consumption in the second half 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 into two random subsamples: a training subsample (60% of the sample) and a test subsample (40%). Second, we grow a forest of 5,000 trees on the training subsample 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[52], which can be thought of as a proxy for a low elasticity of substitution between gas and other inputs; a gas-intensive dummy[53]; 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 subsample. Finally, we test whether the machine learning (ML) predictions carry over to the test subsample. Sample splitting ensures that overfitting does not drive our results.

Figure M.1 plots the distribution of individual treatment effects estimated with the event-study estimator of Borusyak et al. (2024) (in blue), and as predicted by ML in sample (black) and out of sample (red). The average is close to −0.45 in all three distributions.[54] The blue distribution is very dispersed, but this variation may reflect either heterogeneity in treatment effects (along observables or unobservables) or noise, because individual treatment effects contain the error term ϵ_{ijt} of equation (15) (Borusyak et al., 2024). The distribution of in-sample predicted treatment effects (in black) is much less dispersed, but still displays economically relevant heterogeneity, with values around zero for some observations.[55] The distribution of out-of-sample predicted treatment effects (in red) is similar to the black one, but even more compressed, plausibly because in-sample predictions overfit. Still, the red distribution ranges from −1 to zero: the forest predicts that some treated firms cut gas consumption by as much as 60%, while others do not change it at all.

To understand which observables predict heterogeneity, and in which direction, we first regress the in-sample predicted treatment effects on each covariate separately (Table M.1). The forest predicts smaller-than-average reductions in gas consumption for firms in the food sector and in

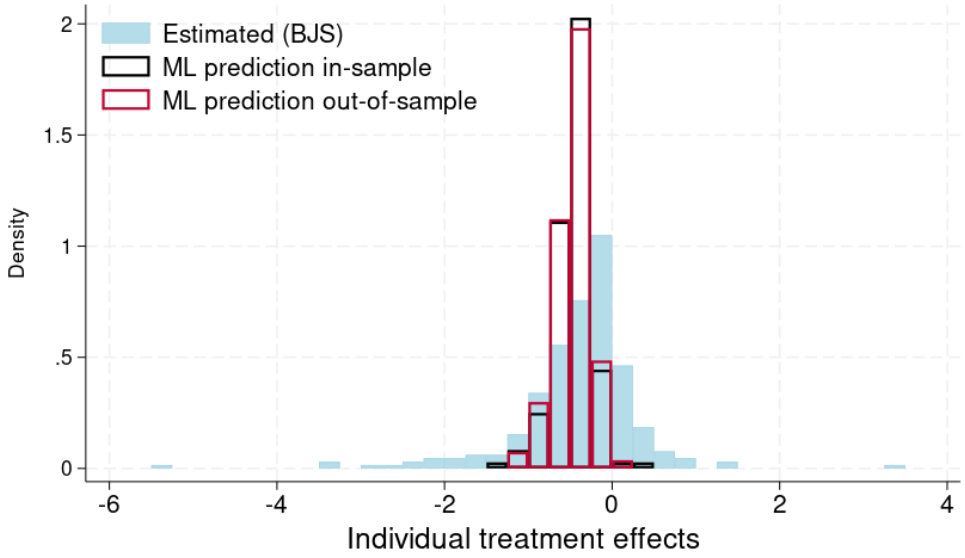
[52]The yes/no question comes from a different Bank of Italy survey (*Business Outlook Survey of Industrial and Service Firms*, Sondtel) run in fall 2022 and reads: "At the beginning of 2022, was gas an essential input for your firm's manufacturing process?" The survey defines inputs as essential when,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

[53]The dummy follows the definition in Italian law: gas-intensive firms must consume more than 1 GWh a year, belong to certain industrial sectors, and have a gas intensity (gas expenditure relative to sales or value added) above certain thresholds.

[54]These are effects on log consumption, not the percent changes $e^{\tau} - 1$ that we report in most of the paper.

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

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



Note: The figure shows the distribution of individual treatment effects on log natural gas consumption in 2022h2: Borusyak et al. (2024) estimates (blue), in-sample ML predictions (black), and out-of-sample ML predictions (red).

the chemicals-pharmaceuticals-rubber sector, firms that report gas as an essential input, firms subject to the EU ETS, and gas-intensive firms as defined by Italian law. It predicts larger-than-average reductions for firms in the wood and paper industry.[56] The same patterns hold in the test subsample when the outcome is the out-of-sample predicted treatment effect (Table M.2). To test whether these patterns reflect true heterogeneity rather than overfitting, we estimate the same regressions in the test subsample with the Borusyak et al. (2024) estimates of the treatment effects as the outcome. If the heterogeneity is real, we would expect the same signs and similar coefficients, because the Borusyak et al. (2024) estimates in the test subsample were never used to train the forest. Table M.3 reports the results. For some covariates, signs differ or coefficients are much smaller, but several predictions are confirmed out of sample, in both sign and size. In particular, four (not mutually exclusive) groups of firms reduce gas consumption less than average: firms in the food industry, firms that report gas as an essential input, firms in the EU ETS, and gas-intensive firms. We take this as evidence of treatment effect heterogeneity along these covariates.

[56]The larger reduction for wood and paper may reflect greater scope for substitution toward biomass fuels; see Appendix N for evidence on fuel switching in the EU ETS subsample.

Table M.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 training subsample. 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 M.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 subsample. 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 M.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 subsample. The outcome variable is the treatment effect (on gas consumption in 2022h2) as estimated with the event-study estimator of Borusyak et al. (2024).

Appendix N Fuel switching and other adjustment margins

In this appendix we report the evidence on fuel switching and on the other adjustment margins discussed in Section 6.d.

N.a Fuel switching among EU ETS plants

For firms subject to the EU ETS, ISPRA records yearly plant-level fuel use (Section 4), which we match to the Invind survey. Because these data are yearly, we aggregate treatment cohorts into a yearly variable for gas only, E_i^{gas}, using the notation of Section 4.d:

$$E_i^{gas} = \begin{cases} 2021, & \text{if } I^{2021} = 0 \text{ and } I^{gas,2022} = 0 \\ 2022, & \text{if } I^{2021} = 1 \text{ and } I^{gas,2022} = 0 \\ 0, & \text{if } I^{2021} = 1 \text{ and } I^{gas,2022} = 1 \text{ and } m_i^{gas} = 12 \end{cases} \tag{23}$$

The matched EU ETS-Invind sample includes 107 plants (66 firms).[57]

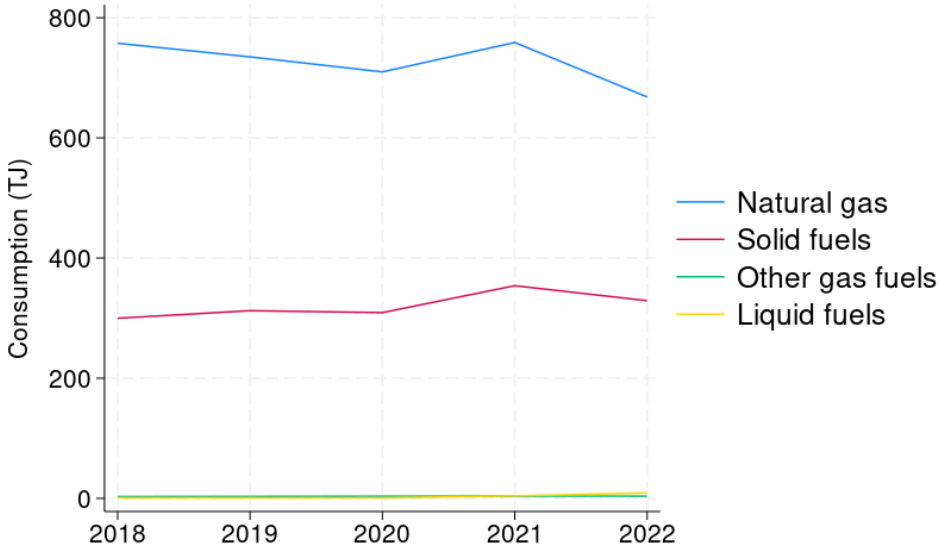
We explore whether the decline in gas consumption was offset by increased reliance on other fossil fuels. The analysis is restricted to firms subject to the EU ETS for which we observe, from administrative sources, energy consumption across all fossil fuels.[58] Natural gas is the main fossil fuel used by EU ETS plants, and an essential input for 90% of these firms, but these plants also use substantial amounts of solid fuels (e.g., coal). Liquid fuels (e.g., kerosene) and other gaseous fuels (e.g., LPG) are much less common (Figure N.1).

Panel (a) of Figure N.2 shows a growing negative effect on gas consumption, with no evidence of a pre-trend ($p = 0.47$). The ATT is −26 terajoules (TJ) in 2021 and −89 TJ in 2022, a decline of 4% and 14%, respectively, relative to the 2018–20 average. We next examine total fossil energy consumption, defined as the combined use of gas and other fuels, including biofuels. If firms had fully replaced gas with other fuels, the ATT would be zero; with no substitution, it would equal the effect on gas. Panel (b) shows that the point estimates fall between these two extremes, with an ATT 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

[57] The 2021 (2022) cohort consists of 53 (20) plants. The pure control group has 34 plants.

[58] 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, most firms in Italy were protected against both electricity and gas prices, or against neither (Section 4.d). Moreover, the tight link between Italian electricity and gas wholesale prices rules out any cost advantage of electrification, unless firms generate their own renewable electricity.

Figure N.1: Fossil fuel mix of firms in the matched EU ETS–Invind sample



Note: The figure displays the average consumption of different fossil fuels (in terajoules, TJ) for firms in the matched EU ETS–Invind sample.

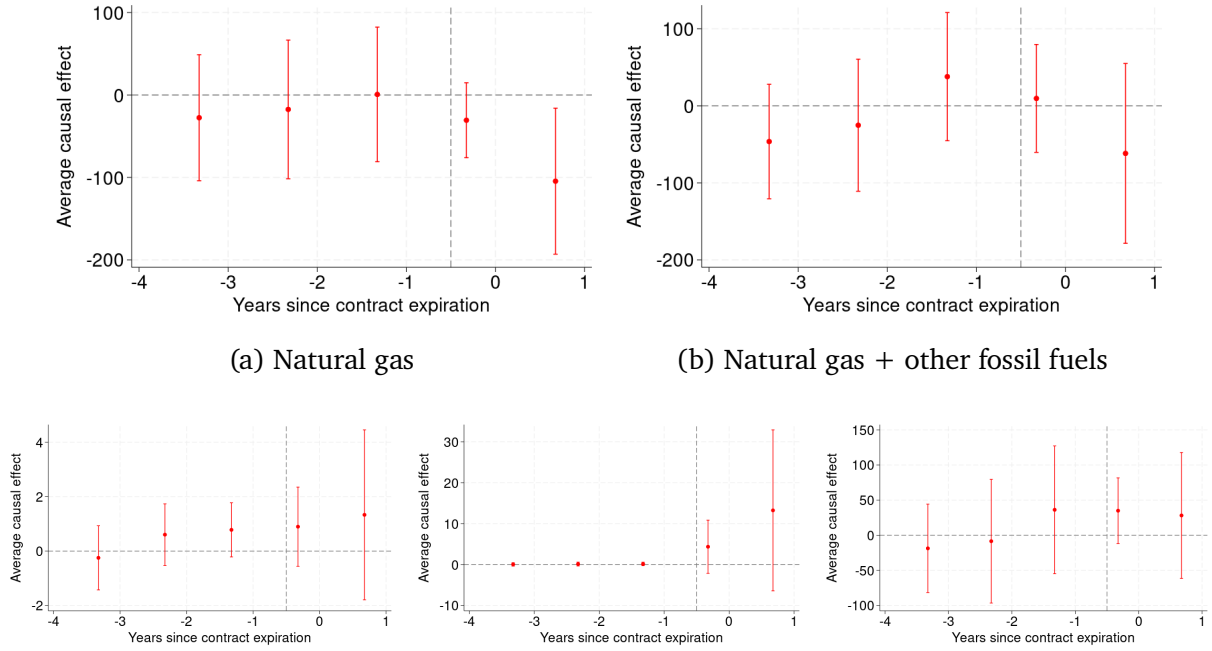
firms.

Panels (c)–(e) of Figure N.2 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 just short of significance at the 10% level. Solid fuels show a larger increase, but it arises earlier and is more pronounced in 2021 than in 2022, which raises doubts about the identifying assumptions for this outcome (pre-trend test $p = 0.02$). Some substitution from gas to other fossil fuels appears to have occurred, but it does not fully offset the reduction in gas consumption among EU ETS plants.

N.b Survey evidence on gas substitution

The 2022 wave of the Invind survey asked firms whether they had replaced natural gas with other inputs. Figure N.3 reports the answers for the gas estimation sample (Section 6.d).

Figure N.2: Fuel switching among EU ETS plants



(a) Natural gas

(b) Natural gas + other fossil fuels

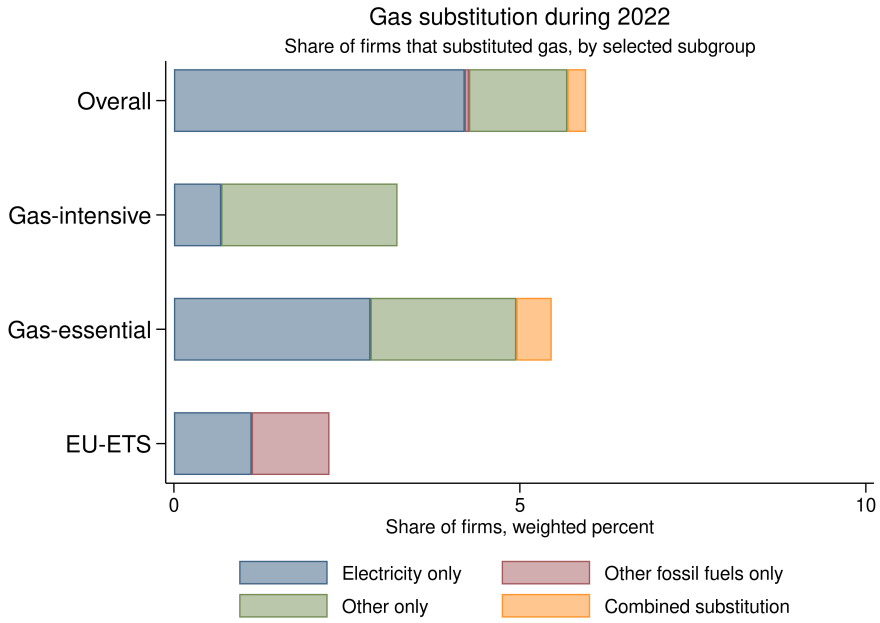
(c) Other fossil gas fuels

(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 levels (terajoules, TJ). The outcome is reported below each panel. 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 event-study estimator of 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.

Figure N.3: Substitution away from gas in 2022 (Invind survey)

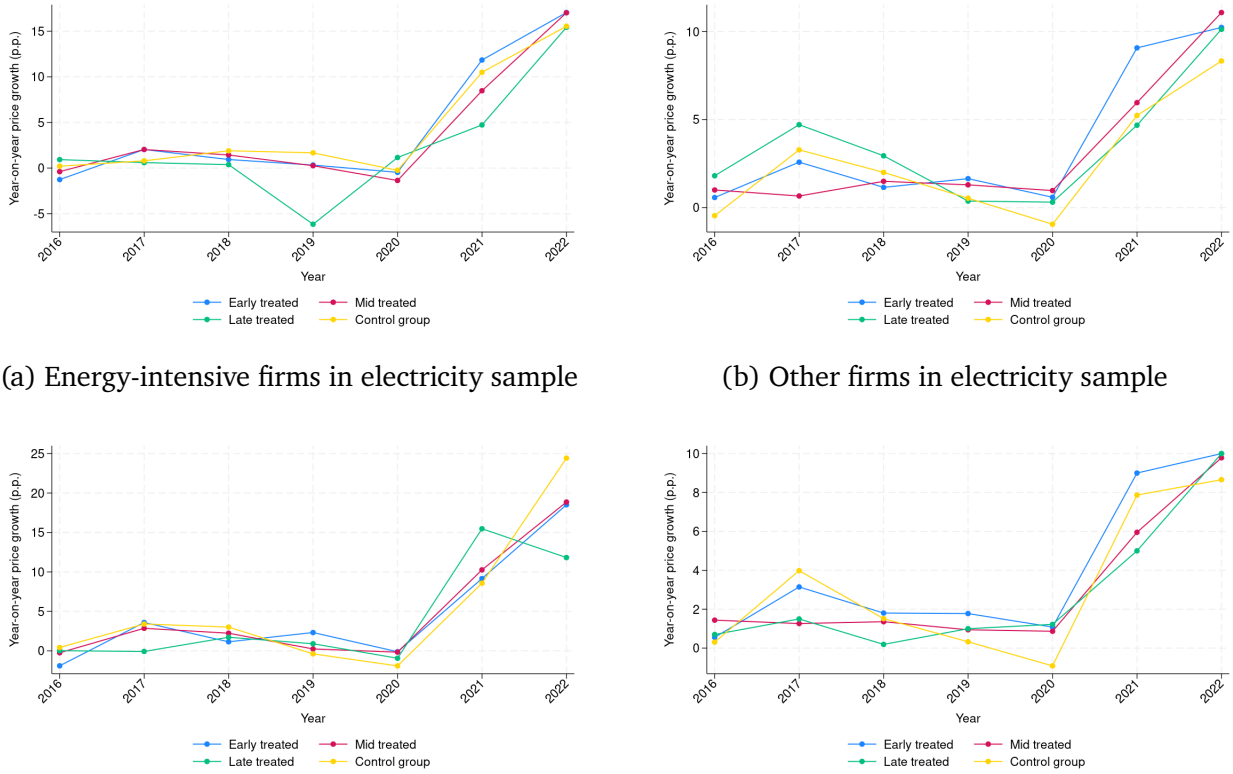


Source: Invind. *Note*: The sample is the gas estimation sample. The remaining firms report no substitution.

N.c Output prices and firms' strategies

Figure N.4 reports the growth of output prices by treatment cohort, and Figure N.5 the strategies that manufacturing firms in the Business Outlook Survey of Industrial and Service Firms (Sondtel) adopted, or planned, in response to higher energy costs in 2022 (Section 6.d).

Figure N.4: Output price growth by treatment cohort



(a) Energy-intensive firms in electricity sample

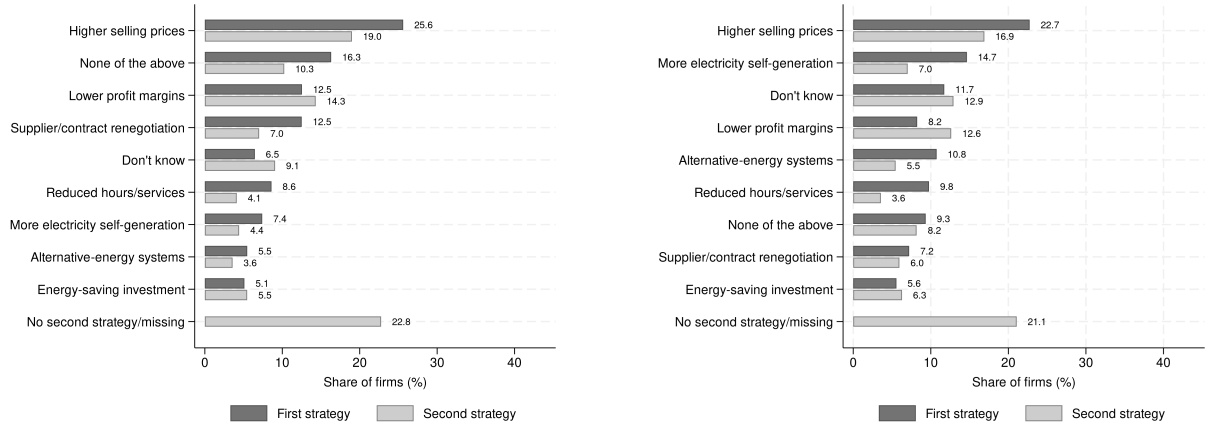
(b) Other firms in electricity sample

(c) Energy-intensive firms in gas sample

(d) Other firms in gas sample

Note: The figure shows the average yearly growth rate of the prices of goods sold, from the Invind survey, by treatment cohort, separately for the electricity and gas samples and for energy-intensive and other firms. Averages are weighted by survey weights.

Figure N.5: Firm strategies in response to energy price increases in 2022



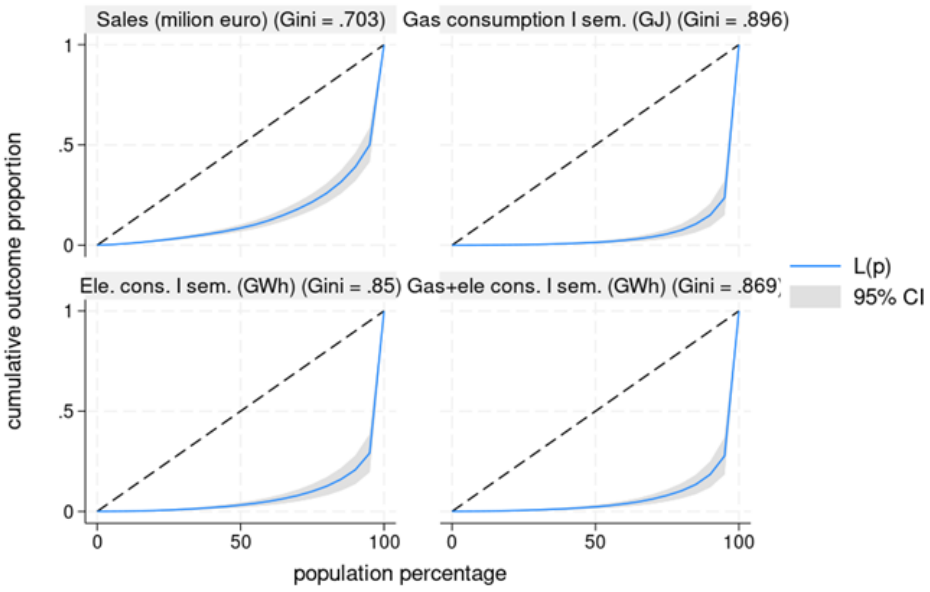
(a) Adopted in the first nine months of 2022 **(b)** Planned for the following six months

Source: Bank of Italy, Business Outlook Survey of Industrial and Service Firms (Sondtel), 2022. *Notes:* Bars report percentages of firms, weighted by 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.

Appendix O Concentration of energy use

Figure O.1 shows Lorenz curves for sales and for energy consumption across firms (Section 6).

Figure O.1: Concentration of sales and energy consumption across firms, with Gini coefficients



Note: The figure shows Lorenz curves, with 95% confidence bands, for sales, gas consumption, electricity consumption, and total gas and electricity consumption across firms. Gini coefficients are reported in the panel titles. Energy consumption refers to the first half of the year.