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

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

[a]Bank of Italy

April 2025

Abstract

We quantify how large, unexpected energy-price hikes affect Italian industrial firms during the 2021-22 energy crisis. For identification we exploit the staggered expiration of fixed-price energy contracts. Contract expiration raises firms' electricity prices by 47 percent and gas prices by 29 percent. Electricity demand does not respond, while gas consumption falls only in the second half of 2022, yielding a short-run price-elasticity of 1.1 overall and close to zero for gas-intensive producers. Heterogeneity and contractual arrangements thus interact to shape the real-economy impact of energy crises. Our estimates imply that broad-based gas bill subsidies could triple industrial demand relative to subsidies targeted at gas-intensive firms.

(*JEL classification*: Q41) *Keywords*: Energy crisis, price-elasticity of energy demand

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

†Contact: matteo.alpino@bancaditalia.it

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

§Contact: annalisa.frigo@bancaditalia.it

1 Introduction

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

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

The relevance of these concerns hinges on a key parameter: the price elasticity of energy demand (Gros, 2022). Despite its importance, there is limited micro-level evidence on this parameter during periods of large price volatility. Existing papers using microdata have typically focused on smaller, less disruptive energy price shocks from earlier periods (von Graevenitz and Rottner, 2022; Fontagné et al., 2024; Gerster and Lamp, 2024), limiting their applicability to understanding how firms navigate severe crises.

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

To estimate effects, we rely on the imputation estimator of Borusyak et al. (2024). However, our results are robust to alternative staggered difference-in-difference estimators (De Chaisemartin and dHaultfoeuille, 2020; Callaway and SantAnna, 2021; Sun and Abraham, 2021) and to the synthetic difference-in-differences method (Arkhangelsky et al., 2021), which allows to considerably loosen the requirement of parallel trends between treated and control units and addresses concerns related to pre-trend testing (Roth, 2022). Results are also robust when controlling for non-parametric time trends that vary at the level of industry affiliation and energy-intensity.

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

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

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

Using random forests to detect treatment effect heterogeneity, we show that gas-intensive companies reduced gas consumption significantly *less* than other firms, a result confirmed also when using confidential administrative data on monthly gas usage among these firms. Gas-intensive firms exhibit a very low price elasticity of -0.03, compared to -1.3 for other firms. This difference appears driven by a low elasticity of substitution between gas and other inputs among gas-intensive firms. Two pieces of evidence support this. First, survey responses indicate that gas-intensive producers almost uniformly classify natural gas as an *essential* input. Second, plant-level data for EU-ETS facilities reveal little switching from gas to alternative fossil fuels. By contrast, these same energy-intensive firms raised their output prices markedly more than other firms, passing a substantial share of the cost shock on to customers, regardless of their cohort of treatment. Taken together, the pricing response and the lack of technological substitution provide a coherent explanation for the very low demand elasticity observed among energy-intensive producers.

Our results carry significant implications for the design of support measures. We perform a simple back-of-envelope calculation calibrated to Italy, where we simulate the effect of a subsidy covering 40% of the retail gas price for firms. For example, consider a subsidy that is targeted to gas-intensive companies, a group accounting for 80% of industrial gas consumption. This would increase demand by 80 million SMC. Conversely, when targeted to non-gas-intensive firms, the same subsidy can increase demand up to 167 million SMC.

We contribute to the literature on energy demand elasticities for firms.[1] From a substantive perspective, we make a unique contribution by estimating elasticities during a period of severe crisis. This is relevant because such parameters may exhibit nonlinearities depending on the size of the shock, rendering local estimates less informative in crisis times (Geerolf, 2022). For instance, studies on electricity demand that examine small price shocks tend to find larger elasticities (0.4-0.5) than those exploiting larger variations (0.1-0.2) (Marin and Vona, 2021; von Graevenitz and Rottner, 2022; Fontagné et al., 2024; Blonz, 2022; Gerster and Lamp, 2024). Using even larger shocks, we find that the elasticity approaches zero. For natural gas, we provide two novel findings. First, elasticities increase as the crisis unfolds, consistently with the idea that long-run elasticities are higher (Deryugina et al., 2020). Second, gas-intensive firms are less elastic compared to the average industrial firm, as they are more likely to declare this input is essential. From a methodological perspective, our study introduces a novel identification strategy based on the staggered expiration of fixed-price contracts. This approach appears more credible than earlier studies, which exploit sectoral or time-series variation (Burke and Abayasekara, 2018; Csereklyei, 2020; Graf and Wozabal, 2013; Davis and Muehlegger, 2010; Hausman and Kellogg, 2015; Faiella et al., 2022).

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

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

2 Data

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

A key advantage of Invind is its detailed, timely data on firms' energy expenditures, consumption, and hedging strategies in 2021 and 2022 – information that is not available in balance-sheet datasets.[2] To collect this data, we introduced an *ad hoc* energy section to the 2021 wave, administered in spring 2022, targeting industrial firms with 50 or more employees. This section gathered semester-level electricity and gas expenditures (in €) and consumption (in MWh and cubic meters) for the previous year, excluding self-produced energy. 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 2.a). The next year wave collected analogous data for 2022, adding information on tax credits under the energy cost relief program implemented by the government.[3]

We ensure self-reported data on energy consumption are plausible and aligned with administrative records, where available, applying a cleaning algorithm to correct systematic errors (e.g., KWh instead of MWh). Validation results, detailed in Appendix C, confirm the data reliability. The response rate to the energy section is high (around 50%) and our results are robust to addressing selective item non-response with inverse-probability weighting (Appendix D). 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).

For the subset of energy-intensive firms, we use confidential administrative records from the Fund for Energy and Environmental Services (CSEA), which we match to our survey through unique tax identifiers. This group includes nearly 3,800 firms, which collectively account for 80% of industrial natural gas consumption, 45% of industrial electricity consumption, and

[2]For background information on the retail energy market in Italy see Appendix A.
[3]We report all the questions in Appendix B.

17% of the sector's value added. For electricity-intensive firms CSEA data include monthly electricity consumption since 2018 and up to 2022. As for gas-intensive firms, we observe monthly gas consumption for the years 2019, 2021, and 2022.[4]

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%.

We use data from the Italian National Institute of Statistics (Istat) on energy intensity by 2-digit NACE industries and Eurostat data on average retail prices and consumption brackets for industrial consumers to validate our survey measures (see Appendix C). Additionally, we use the 2022 Bank of Italy's Business Outlook Survey of Industrial and Service Firms (Sondtel) to identify firms which declare natural gas to be an "essential" input; we consider this a proxy for low elasticity of substitution between gas and other inputs.

2.a 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).

Two things are worth noticing. First, contracts had to be in place *at the beginning of 2021*, prior to any anticipation of the crisis.[5] Thus the question does not capture endogenous firms' responses to the crisis. Second – due to space constraints – the question conflates protection for electricity *and* gas, introducing potential measurement error. Reassuringly, this variable

[4] 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.

[5] As late as March 2021, future markets predicted October TTF natural gas prices to match the preiovus five-year average (around 15 euro), with expectations rising to 36 euro by July, still below the realized price of 76 euro.

strongly predicts changes in unit costs for both inputs between the first and second semester of 2021.

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

> "In 2022, did your firm have instruments (for example fixed-price contracts or derivatives) to protect itself, even partially, from rises in the prices of electricity (natural gas)?"

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

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

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

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

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

All firms are untreated in 2021h1 since the wholesale market crisis had not yet started. Firms with fixed-price contracts signed before the crisis remain untreated in later periods until their contracts expire.[6] We classify firms into four treatment cohorts based on E_i^j.

Firms with $I^{2021} = 0$ and $I^{j,2022} = 0$ are immediately exposed in 2021h2 ($E_i^j = 2021h2$; **early treated**). This cohort includes firms on variable-price contracts and those with fixed-price contracts expiring during 2021h2. In Section 4.a we confirm the absence of differential treatment effects for this group, mitigating concerns that its composition may affect the results.

Firms with $I^{2021} = 1$ and $I^{j,2022} = 0$ are exposed in 2022h1 ($E_i^j = 2022h1$) due to contract

[6] 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 but at higher costs than pre-crisis contracts.

expiration at the end of 2021 (**mid treated**). By 2022, these firms face higher energy costs regardless of whether they purchase a new fixed-price or variable-price contract, as both became more expensive following the start of the crisis in late 2021.

Firms with $I^{2021} = 1$ and $m^j = 6$ are protected until mid-2022, when their fixed-price contracts expire (**late treated**). We assume these are the first 6 months of the year. In particular, we assume that at the beginning of 2021 the firm was holding a fixed-price contract signed at the beginning of 2021 that eventually expired in mid-2022. Thus, we set $E_i^j = 2022h2$ for these firms. We test the validity of this assumption in our event-study analysis by checking that price of input j increases in line with our assumed timing for this specific cohort (see Section 4.a).

Firms with $E_i^j = 0$ are protected throughout 2021 and 2022 (**pure control group**). Note that the formulation of our question in 2022 does not ensure that the pure control group are protected under the exact same contract signed at the beginning of 2021, for which they declared to be protected in the first survey. In principle, they might have purchased a new fixed-price contract in January 2022, but in the results section we show that this is not consistent with our data. In particular we show that 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.[7] We excluded two subgroups of firms: (a) firms protected for only part of a semester in 2022 ($1 \leq m_i \leq 5$ or $7 \leq m_i \leq 11$), which ensures binary treatment status during protection periods and helps interpretability. This subgroup consists of 31 and 20 firms in the case of electricity and gas, respectively; (b) Firms unprotected in 2021 ($I^{2021} = 0$) but protected in 2022 ($I^{2022} = 1$), as this protection may reflect endogenous responses to the shock. These are 14 and 32 firms in the electricity and gas sample, respectively.

In Appendix E we report balancing tables for sectoral and energy-intensity related covariates across different treatment cohorts. One concern could be that firms with higher energy intensity or in particular sectors would select into longer contracts – and this type of selection is related to potential *trends* in outcomes. While we highlight some significant level differences across groups, these do not appear economically relevant. To avoid any concerns, in Appendix F we show that results are virtually unchanged when controlling for sector-specific and energy-intensity specific non-parametric time trends.

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

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 follows:

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

For the electricity-intensive case, the balanced matched sample includes 279 firms[8] over the period 2020-2022. For the gas-intensive case, our balanced matched sample includes 126 firms[9] in 2021 and 2022.

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

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

The matched EU ETS-Invind sample includes 107 plants (66 firms).[10]

3 Empirical strategy

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

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

[9] 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.

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

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

where y_{ijt} is the quantity (q) or retail price (p) of energy input j for firm i in semester t. 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).[11] We are interested in potentially heterogeneous treatment effects (β_{ijk}) of contract expiration on p and q.

We estimate each model on a balanced panel of firms over 2021-2022 using the Borusyak et al. (2024) "imputation estimator", which identifies the ATT under standard parallel trends and no-anticipation assumptions. Appendix G confirms robustness using alternative diff-in-diff estimators, including De Chaisemartin and dHaultfoeuille (2020), Callaway and SantAnna (2021), Sun and Abraham (2021), and OLS. We cluster standard errors at the firm level to avoid known serial correlation issues (Bertrand et al., 2004). Pre-treatment coefficients are estimated *separately* for improved efficiency, reducing bias from correlation between treatment effects and pre-trend estimates (Borusyak et al., 2024).

Parallel trends are plausible because contract expiration dates are predetermined, unrelated to firms' crisis responses, and tied to pre-2021 fixed-price contracts based on stable market forecasts. Futures markets indicated no significant price surge before May 2021 (see Appendix H). Pre-trend tests consistently support this assumption, with no evidence of violations. As for no-anticipation, pre-trend tests and the consistency of effects across cohorts suggest that firms did not anticipate treatment, lending support to this assumption.

One concern involves the "early-treated" cohort ($E_i = 2022h1$), which includes both firms with expiring fixed-price contracts and those with variable-price contracts. The latter set of firms may have had different characteristics, such as weaker preferences for price protection. However, cohort-specific estimates show that effects are similar across cohorts (see 4.a). Additionally, early-treated firms never serve as controls for other cohorts.

Another concern is selection into longer contracts based on unobserved preferences for price certainty. Since firm-level fixed effects control for level differences, this only threatens validity if these preferences correlate with abrupt changes in outcomes specifically at contract expiration. This is unlikely, and the absence of pre-trends (see 4) supports this assumption.

Finally, given that we observe few periods across two years, one may wonder whether results

[11] See Section 2.a on how we construct this variable.

are driven by seasonality. We find that both magnitude and precision of our results are virtually unaffected when including firm $\times$ semester fixed effects, which control for *firm*-specific seasonality.

To further validate our findings, we implement a synthetic diff-in-diff (SDID) exercise (Arkhangelsky et al., 2021), following the procedure in Clarke et al. (2024) for staggered treatment and using bootstrap for inference. This method explicitly matches on pre-trends and addresses concerns related to pre-trend testing (Roth, 2022). This exercise confirms robustness, ensuring our results are not driven by underlying pre-trends.

We estimate separate models to avoid multicollinearity stemming from the high correlation in expiration dates across inputs. This could be problematic in case gas and electricity effects reinforce each other, leading to omitted variable bias. This concern arises in two cases: substitution between the inputs and scale effects. However, in this setting, neither issue is problematic. Substitution was minimal during the crisis, as both gas and electricity prices rose sharply. Regarding scale effects, as shown in Section 4, 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.

4 Results

Figure 4.1 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 of retail price 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. Since many of the τ_k coefficients are relatively large, in the text we describe the magnitude of the effects by commenting on $e^{\tau_k} - 1$.

In panel (a) we observe a significant increase in the retail price of electricity following the expiration of fixed-price contracts, with no evidence of pre-trending effects before the treatment occurs (the p-value of the pre-trend test is 0.13). A year and a half after expiration, the price increase is 47% relative to the baseline, quite precisely estimated. A similar pattern is seen in panel (b) for the retail price of natural gas. Again, there is no evidence of a pre-trend (the p-value of the pre-trend test is 0.39) and estimates are precise. One and a half years after expiration, the increase in gas price is 29% relative to the baseline.[12]

[12] 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

Figure 4.1: The effect of the expiration of a fixed-price contract on average prices and quantities of energy inputs at the firm level.

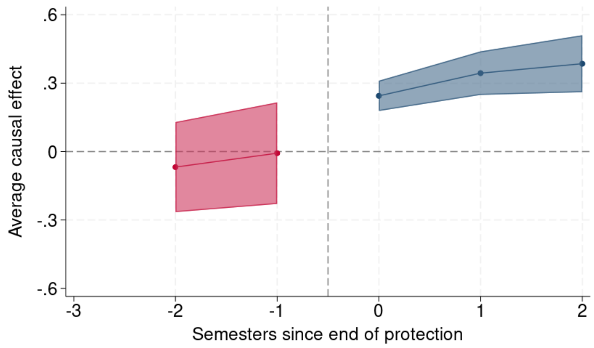

(a) Average cost of electricity

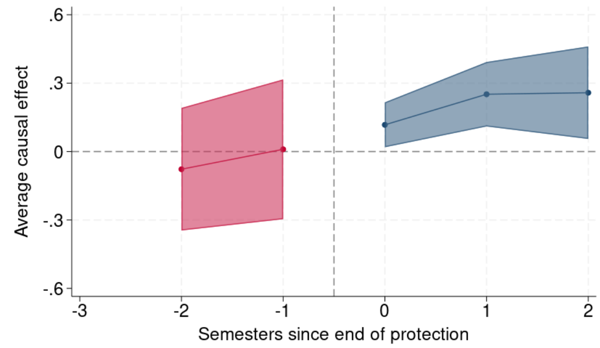

(b) Average cost of natural gas

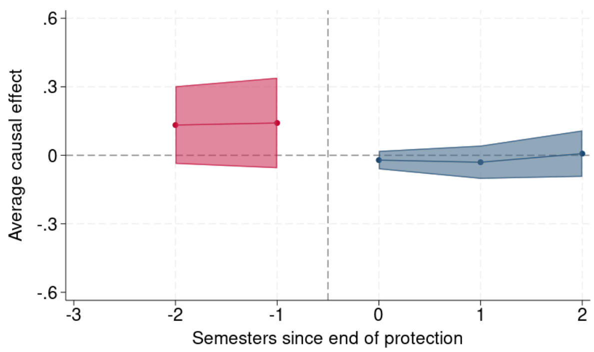

(c) Quantity of electricity

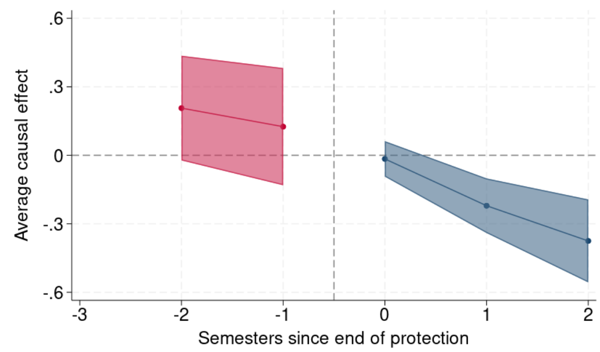

(d) Quantity of natural gas

Note: The figures show average causal effects of the expiration of a fixed-price contract on the average retail prices of electricity and natural gas (panels (a) and (b)) and the corresponding demanded quantities (panels (c) and (d)). Outcome variables are always in logs. Average causal effects before and after the treatment are estimated in two separate regressions, using the "imputation" estimator by Borusyak et al. (2024), as described in Section 3. Confidence intervals are at the 95% level.

In panel (c) we analyze the (log-)quantity of electricity purchased. Despite the large price increase, this outcome does not respond to the treatment. Coefficients are positive, close to zero and confidence intervals are tight. We interpret these as precisely estimated zeros. 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 implies that gas consumption decreases by 32% compared to the counterfactual with no price increase. Results are virtually identical when controlling for sector-specific and energy-intensity-specific non-parametric time trends (Appendix Figure F.1), with point estimates

variation in energy intensity, may suffer from measurement error, as noted by Lafrogne-Joussier et al. (2023).

remaining effectively unchangeda. This addresses concerns that differential composition of different treatment arms could bias our estimates.

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

The increasing gas-demand elasticity, *irrespectively* of time-since-treatment, could reflect changing expectations about the energy crisis. Early on, when the crisis was expected to be short-lived, firms likely employed a "wait-and-see" strategy. In summer 2022, market pessimism peaked: spot prices hit record highs, futures markets projected a prolonged crisis and business confidence reached its trough (Appendix Figures A.1 and A.2). 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).

4.a Robustness: synthetic diff-in-diff

Figure 4.2 reports similar event-study graphs for electricity and natural gas using the synthetic diff-in-diff methodology. We present the results for the three treatment cohorts separately, with the "donor pool" always drawn from the pure control group. Naturally, the pre-trend matching can use more pre-treatment periods for the "late treated" than for the "mid treated" and "early treated", as we only have four periods at our disposal.

Figure 4.2: Synthetic diff-in-diff estimates of the effect of the expiration of a fixed-price contract on electricity and gas prices (P) and quantities (Q) for different treated cohorts.

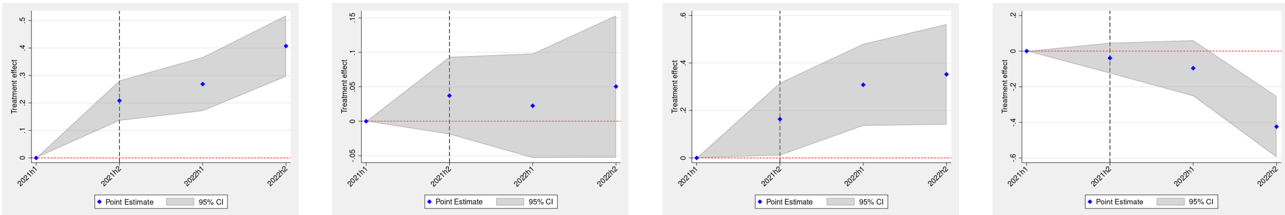

(a) P of electricity; early treated

(b) Q of electricity; early treated

(c) P of natural gas; early treated

(d) Q of natural gas; early treated

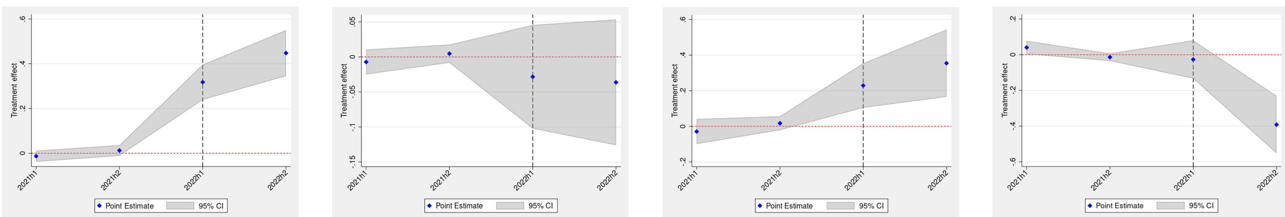

(e) P of electricity for mid treated

(f) Q of electricity for mid treated

(g) P of natural gas; mid treated

(h) Q of natural gas; mid treated

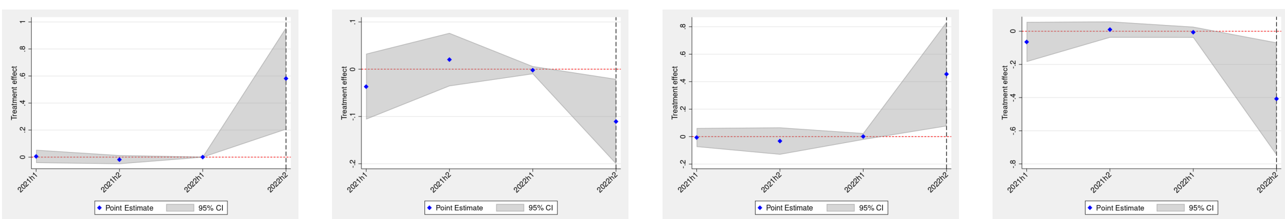

(i) P of electricity; late treated

(j) Q of electricity; late treated

(k) P of natural gas; late treated

(l) Q of natural gas; late treated

Note: The figure shows average causal effects of the expiration of a fixed-price contract on the average retail prices of electricity and natural gas, estimated with the synthetic difference-in-difference method in the staggered case (Arkhangelsky et al., 2021; Clarke et al., 2024). The outcome variable is always in logs. Bootstrapped confidence intervals are at the 90% level.

Across all cohorts, we observe that the retail price increases *exactly* at the time of expiration of the respective fixed-price contract. This reassures us that our coding of the cohort-of-treatment variable is correct. Moreover, the ATT estimates are qualitatively and quantitatively similar to our baseline results, which shows robustness to a relaxation of the parallel trend assumption. As in the staggered diff-in-diff, all three cohorts display very similar treatment effects, confirming results are not driven by a specific cohort.

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

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

4.b Heterogeneity in gas demand response across firms in 2022h2

The ATT in 2022h2 for gas consumption might hide substantial heterogeneity, due to firm characteristics, sectoral specificities, and particularly the large variation in energy intensity across firms. Energy-intensive firms were the earliest and largest government aid recipients during the 2021-22 crisis (McWilliams et al., 2024), but whether their demand is more or less elastic than other companies is theoretically ambiguous, as it depends on several parameters, including the elasticity of substitution between inputs (Cahuc et al., 2014).[13]

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. (2021) is particularly well suited for this, as it provides a treatment effect estimate for each treated observation. We can thus use it as an outcome variable and use *random forests* to find its best predictors, in the spirit of Athey and Imbens (2016) and Wager and Athey (2018).[14] The advantages of this approach relative to a more traditional heterogeneity analysis are twofold: first, the quest for heterogeneity is more efficient and also explores non-linear combinations of the available covariates; second the procedure is less prone to bias arising from multiple hypotheses testing.

The random forest algorithm highlights some key predictors of heterogeneity: a self-reported dummy variable indicating whether gas is an essential input in production;[15] a dummy for being a gas-intensive firm according to the Italian legislation; and the EU ETS dummy. Interestingly, the ML algorithm does not detect a large amount of heterogeneity between industries, with only the food sector dummy being selected as a predictor. Firms in this industry overwhelmingly declare gas to be an essential input.

[13]We provide formulae in Appendix I.

[14]We provide more details and extensive results in Appendix J.

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

When computing heterogeneous treatment effects with the Borusyak et al. (2024) estimator in the sub-samples selected by the ML algorithm and obtain the following results. In the second half of 2022 the ATT is equal to -41% among firms for which gas is *not* considered essential, and to -28% among those for which it *is* essential. Both estimates are significant at the 99% level and their difference is significant at the 10% level. The associated price elasticities of demand are -2.5 and -0.5, respectively, which suggests that the gas-essential dummy can be considered as proxy for low elasticity of substitution, as the formulation of the survey question suggests (see Appendix K for details on how we compute elasticities).

Interestingly, we find the ATT - and consequently the demand elasticity - to be near zero among gas-intensive firms. Since all gas-intensive firms classify gas as an essential input, their lower demand elasticity appears directly linked to the reduced substitutability of gas with other inputs. Note that the heterogeneity unveiled in this section cannot be attributed to sectoral dynamics, or else the ML algorithm would have selected several sector dummies as key predictors. Furthermore, gas-intensive firms are spread throughout all industries except water and waste management.

4.c Energy intensive firms

Previous analyses indicate that energy-intensive firms exhibit inelastic demand for both gas and electricity. Given the policy significance of this group, we conduct a complementary analysis using confidential administrative data on their energy consumption, which is exclusively available for these firms. Figure 4.3 presents event-study graphs for the evolution of gas and electricity consumption around the expiration of fixed-price contracts, at a quarterly frequency. For electricity-intensive firms, the impact of losing fixed-price protection on the (log-)quantity of electricity consumed remains negligible across all treated periods. This finding closely mirrors the patterns observed for the average industrial firm, as detailed in the first part of Section 4.

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

Figure 4.3: Event-study estimates for the (log-)quantity of energy among energy intensive firms

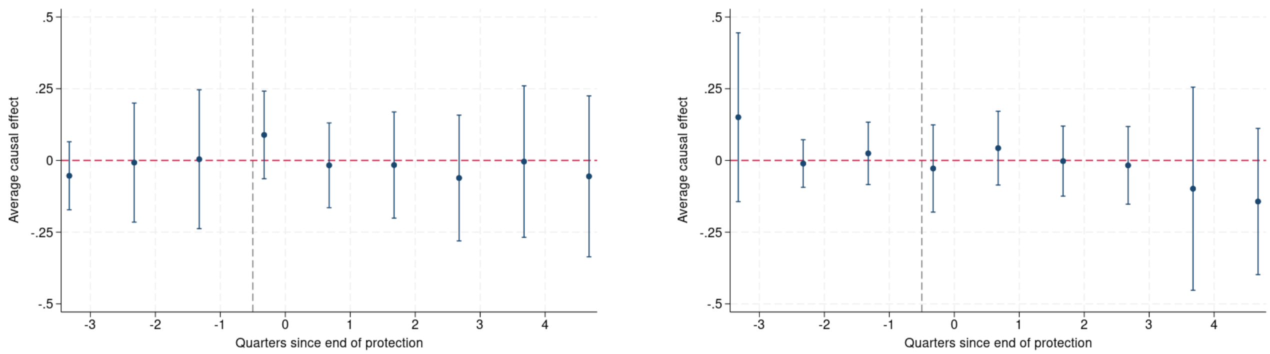

(a) Electricity (b) Natural gas

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). Average causal effects before and after the treatment are estimated in two separate regressions, using the "imputation" estimator by Borusyak et al. (2024). Confidence intervals are at the 95% level.

Pass-through to final good prices The limited quantity response just documented for energy-intensive firms does not imply that they bore the entire cost shock internally. A plausible alternative adjustment channel is output-price pass-through. Using complementary information from the Invind survey, we report cohort averages of year-on-year changes in selling prices in Appendix Figure A.3. Two things are noticeable. First, all cohorts raise prices markedly and almost synchronously during 2022, in line with the general inflationary episode. This is consistent with recent evidence that firms price updating decisions do not depend much on idiosyncratic cost shocks, but rather on rivals shocks (via strategic complementarities) and on market-wide cost shocks (Duprez and Magerman, 2018; Amiti et al., 2019; Muehlegger and Sweeney, 2022). Second, energy-intensive firms increased their prices substantially more than other industrial firms, indicating that price setting constituted an important margin of adjustment for them. By increasing output prices at rates higher than other firms, they were able to accommodate higher energy costs, thereby reconciling their near-zero demand elasticity with continued financial viability during the crisis.

4.d Input substitution among plants subject to EU ETS

Here we explore whether the decline in gas consumption during 2022h2 was partly offset by increased reliance on other fossil fuels. While natural gas is typically the primary fossil fuel consumed by EU ETS plants – serving as an essential input for 90% of these firms – these plants also utilize solid fuels (e.g., coal) to a significant extent. The use of liquid (e.g., kerosene) or gas fuels (e.g., LPG) is much less common.

Figure 4.4: Input substitution test among EU ETS plants

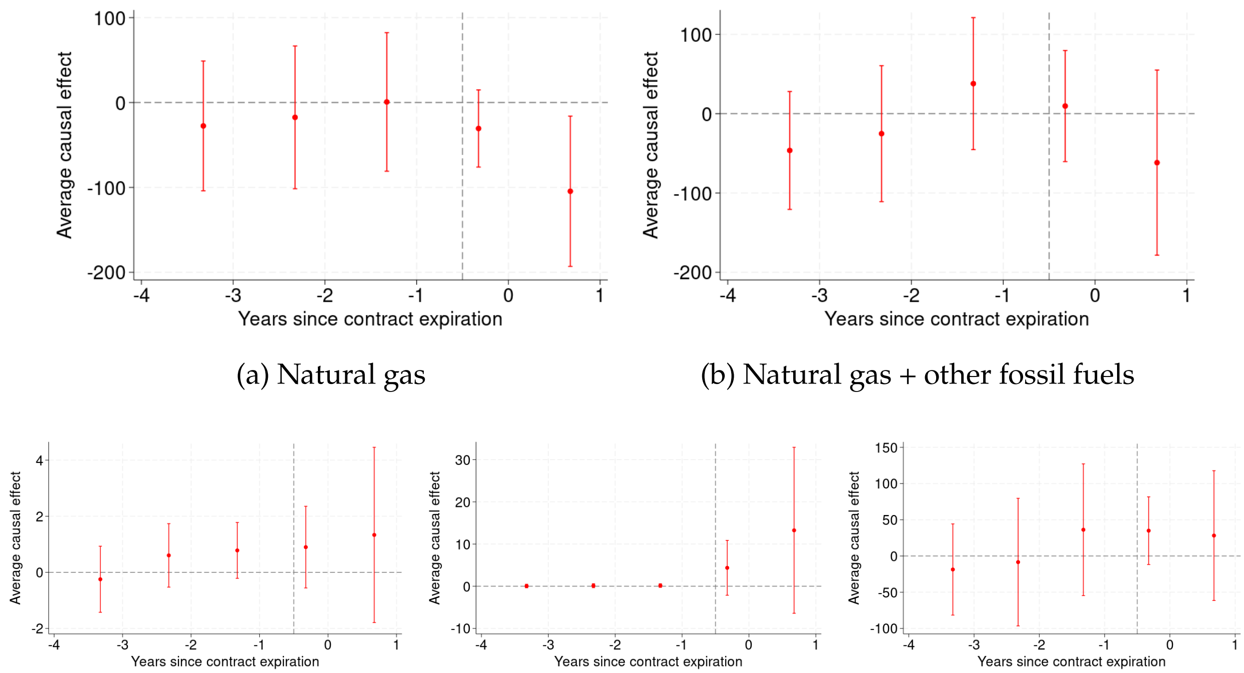

(a) Natural gas

(b) Natural gas + other fossil fuels

(c) Other fossil gas fuels

(d) Other liquid fossil fuels

(e) Other solid fossil fuels

Note: The figures show average causal effects of the the expiration of a fixed-price contract on different outcomes in levels. The outcome is reported underneath each event-study. Average causal effects are estimated on the set of firms belonging to both the Invind and the EU ETS sample. Average causal effects before and after the treatment are estimated in two separate regressions, using the "imputation" estimator by Borusyak et al. (2024). Confidence intervals are at the 95% level.

Panel (a) of Figure 4.4 illustrates a negative effect on gas consumption over time, with no evidence of a pre-trend (p-value: 0.47). The ATT estimates are -26 terajoules (Tj) in 2021 and -89 Tj in 2022, corresponding to declines of -4% and -14%, respectively, relative to the 2018-2020 average. We next examine total fossil energy consumption, defined as the combined use of gas and other fuels. If firms completely substituted gas with other fuels, the ATT would equal zero, while no substitution would produce effects equivalent to those for gas alone. Panel (b) shows that point estimates fall between these two extremes, with ATT estimates of +27 Tj in 2021 and -56 Tj in 2022. The event-study profile is similar but attenuated, with wider confidence intervals. These findings indicate incomplete substitution, with treated firms reducing total fossil energy consumption more than control firms.

Panels (c)-(e) of Figure 4.4 display changes in the consumption of single fuel types. Results show slight increases in the use of all other fuels among treated firms, but substitution remains incomplete. For gaseous fuels, the increase is negligible. For liquid fuels, the effects are modest (6 Tj in 2021 and 10 Tj in 2022) and marginally insignificant at the 10% level. Solid fuels show a larger increase, but the effect arises earlier and is more pronounced in 2021 than in 2022, raising doubts about the validity of identifying assumptions for this particular outcome (pre-trend test p-value: 0.02). While some substitution from gas to other fossil fuels appears to have occurred, it is incomplete and does not fully explain the observed reduction in gas consumption among EU ETS plants. This is consistent with gas-intensive firms declaring that gas is an essential input.

5 Policy implications

We illustrate the implications of our results for the design of support measures to firms during energy crises. EU governments widely implemented such policies in 2021-22 (Sgaravatti et al., 2021; McWilliams et al., 2024), although some economists have raised concerns about their potential impact on incentives to conserve gas (Gros, 2022; Signorini, 2022). Following Deryugina et al. (2020), we provide calculations calibrated to Italy in 2022h2. Assuming a perfectly competitive market for industrial gas, the introduction of a per-unit subsidy s increases equilibrium quantity according to:

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

where ϵ_D and ϵ_S are the absolute demand and supply elasticities and Q and P are the initial quantity and price.

Gas supply is highly inelastic in the short term, but precise estimates are scarce. We consider two scenarios: a very low supply elasticity ($\epsilon_S = 0.06$; Krichene (2002); Albrizio et al. (2022)), and a moderately higher value ($\epsilon_S = 0.2$). Figure 5.1 illustrates the percentage increase in gas consumption for both scenarios, as a function of demand elasticity.

Energy-intensive firms are the most exposed to the crisis, and policymakers may support them for several reasons, such as mitigating strategic dependencies in a challenging geopolitical environment (Draghi, 2024). Our estimates indicate that a 50-cent subsidy (equivalent to 40% of the observed price) for gas-intensive firms ($\epsilon_D = 0.03$) would increase their consumption by approximately 1.5% in 2022h2, with little variation across supply elasticity scenarios. Since this group accounts for 80% of total industrial gas consumption, this translates to an increase of around 80 million Standard Cubic Meters (SMC).

If subsidies are targeted non-gas-intensive firms instead ($\epsilon_D = 1.3$), the same 50-cent subsidy would raise consumption by 4 to 12.5%, depending on the supply elasticity. Given this group represents 20% of industrial gas consumption, this corresponds to an absolute increase between 54 and 167 million SMC.

Figure 5.1: Effect of a subsidy on equilibrium gas consumption

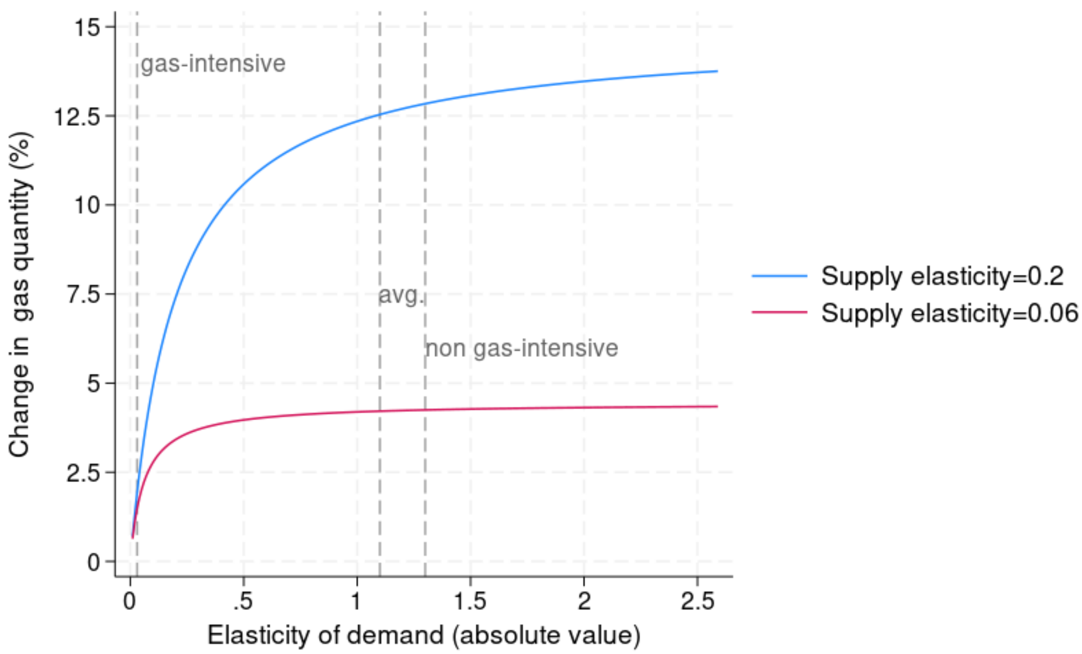

Note: The figure shows the percentage increase in the equilibrium quantity of gas purchases induced by a 50 cent subsidy for different values of demand and supply elasticities. It is calculated according to the formula in (5) and scaled by the baseline quantity; Q is set equal to 6,700 million SCM (average in second semester 2015-19; source: SNAM - network operator) and P to 1.35 euro per SCM (retail price in business sector in 2002h2; source: Eurostat). Vertical dotted lines are drawn in correspondence of some values of the elasticity of demand estimated in the paper: 0.03 (gaas-intensive firms), 1.1 (average) and 1.3 (non gas-intensive firms).

In this context, our results allow a quantification of the unintended consequences of these policies in terms of additional demand. Crucially, extending subsidies to all firms, even those least exposed to the crisis, could triple gas consumption due to their higher demand elasticity.

For electricity, the demand elasticity is near zero, implying that subsidies for electricity would not exacerbate the energy crisis further.

6 Conclusions

This study examines how Italian industrial firms responded to the 2021-2022 energy crisis, using the staggered expiration of fixed-price energy contracts to identify the impact of rising energy costs. We find that electricity prices increased by up to 45% and gas prices by 30%, with limited adjustments in energy demand. Gas consumption decreased only in late 2022, while electricity usage remained unchanged.

Responses varied significantly across firms. Gas-intensive firms exhibited minimal elasticity, reflecting their reliance on natural gas as an essential input. In contrast, non-gas-intensive firms showed greater reductions in gas usage and higher price sensitivity.

These findings suggest that public support to firms during the crisis, while providing financial relief, have reduced incentives for conservation, but less so among gas-intensive firms with inelastic demand.

References

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Borusyak, K., Jaravel, X., and Spiess, J. (2024). Revisiting event-study designs: robust and efficient estimation. *Review of Economic Studies*, page rdae007.

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

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

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

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

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

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

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

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

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

Draghi, M. (2024). The future of european competitiveness. *Report Part B. In-depth analysis and recommendations. September*.

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

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

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

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

Geerolf, F. (2022). The baqaee-farhi approach and a russian gas embargo. *Revue de l'OFCE*, 179(4):143–165.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Appendices

Appendix A Background

A.a The 2021-22 energy crisis in Italy

In 2020, one year before the beginning of the energy crisis, Italy was a net importer of natural gas: imports accounted for 93% of gross inland consumption. Of these imports, Russia accounted for 43%, making it a key supplier. Italy was also a net importer of electricity, and natural gas accounted for approximately half of domestic power generation. The high reliance on natural gas in electricity production coupled with the marginal price system at work in the day-ahead power market implies that shocks to the wholesale price of natural gas almost completely pass-through to the wholesale price of electricity. Figure A.1 plots the evolution of the wholesale price of natural gas (solid red line) and electricity (dashed black line) in Italy.[16] The price of gas was rather stable at low levels until mid-2021 (around 30 euro per MWh). After that, it slowly started to rise above historical levels. The first major upswing occurred in the fall of 2021, when the price went above 100 euros per MWh; the second took place in December of the same year, when it almost reached 200 euros. After a temporary drop, the price surged to around 250 in February when Russia invaded Ukraine, but rapidly declined thereafter and stabilized below 100 until summer 2022. At that point, the price climbed again quite rapidly, reaching a historic peak of over 300 euros in late August 2022. Before the end of the year, the price dropped to much lower levels and then up again, before a final descent to 70.

A.b The Italian retail energy market

Following a broader trend in the European Union, the Italian retail markets for electricity and natural gas have been gradually liberalized in the late 1990s (Polo and Scarpa, 2003). According to the Regulatory Authority for Energy, Networks and Environment (ARERA), both the market for electricity and for gas are not very concentrated (ARERA, 2022).[17]

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-

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

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

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

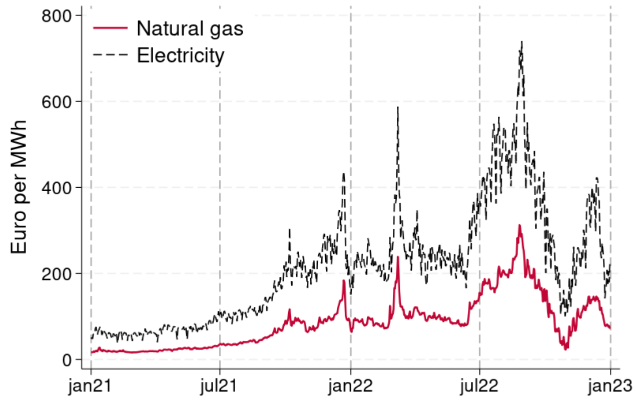

Source: Italian Power Exchange (GME). *Note*: The figure shows the spot price of natural gas and electricity traded on the Italian day-ahead wholesale market.

price contracts. Fixed-price contracts typically have a standard duration of 12, 24, or –less frequently– 36 months. Each supplier can offer as many contracts as it sees fit. However, for households and firms with a low-voltage connection for electricity (a minority in our sample) and an annual gas consumption below 200,000 standard cubic meters (one third of our sample), the supplier is required to also offer two standardized contracts - one with variable and one with fixed price. While the conditions for these standardized contracts are designed by ARERA, suppliers compete on price.[18] All the contracts posted on the market, whether standardized or not, are published daily on a website managed by a government agency.[19] The portal does not include information on large and/or energy-intensive firms, which usually negotiate *ad hoc* contracts directly with their supplier.

A.c Additional information on the retail energy markets

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

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

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

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

Appendix B Questionnaires

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

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

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

(a) 2021 wave

Rising energy prices

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

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

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

(b) 2022 wave

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

Appendix C Validation of survey answers

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

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

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

   (a) We compute the average unitary price paid by firms for each semester and compare it with the corresponding average price recorded by Eurostat for the Italian market.[20] We flag observations for which the unit price – computed by dividing expenditures by quantity – is either below half the minimum or above double the maximum of the reference Eurostat price across consumption classes.[21]

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

3. Combining the criteria in a) and b), we identify 6 error categories for the firm-level replies on electricity and 4 categories for the ones on natural gas. For example, one category is made by those firms that incorrectly report electricity consumption in KWh, instead of MWh; this mistake results in a price (expenditures for electricity over quantity of electricity) which is approximately 1000 times lower than what reported by Eurostat. In Table C.1 we list all the different categories, specifying the reporting mistakes associated with each of them. This exercise is performed for both semesters separately. In 18.7 and 2.8% of the electricity-related replies and 21 and 6.4% of the gas-related replies in the 2021 and 2022 wave respectively, we observe the same mistake done in both semesters. In these cases, we rescale the values in order to fix the reporting mistake. Carrying on with the previous example, if we believe a firm reports quantities

[20]The Invind and the Eurostat prices are constructed similarly, as they both include levies and taxes. However, Eurostat includes all non-household consumers, while we only have industrial firms with at least 50 employees.

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

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

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

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

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

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

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

Figure C.1: Consistency with other data sources

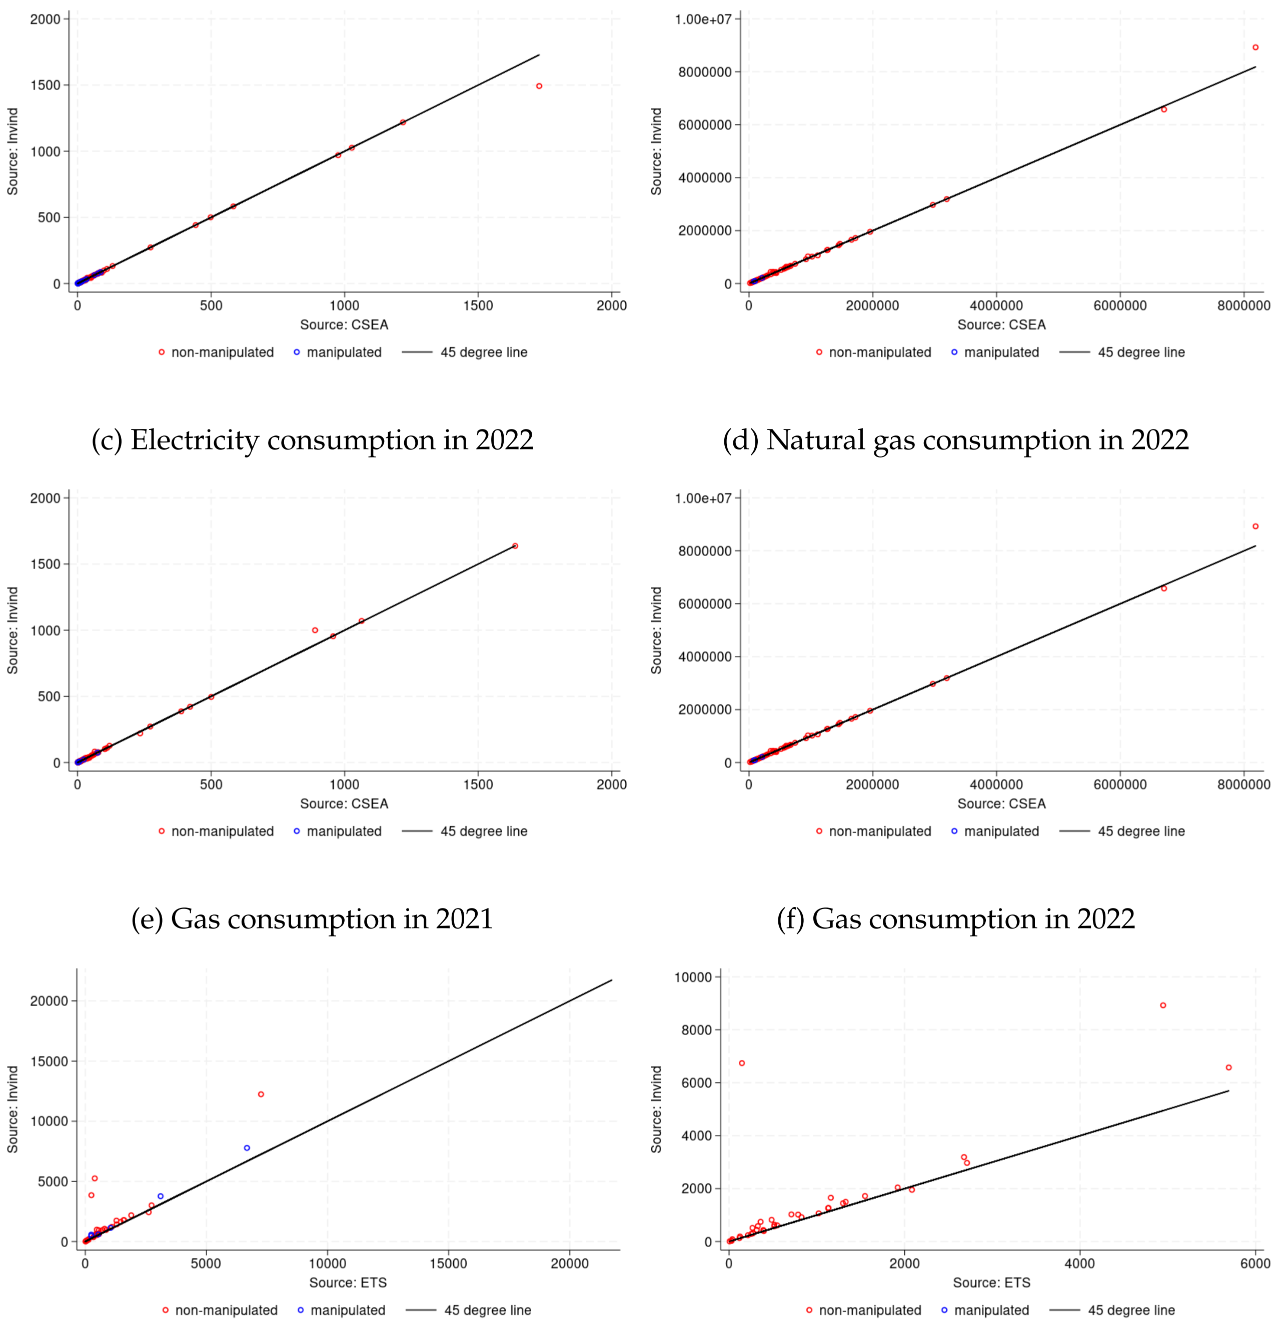

(a) Electricity consumption in 2021

(b) Natural gas consumption in 2021

(c) Electricity consumption in 2022

(d) Natural gas consumption in 2022

(e) Gas consumption in 2021

(f) Gas consumption in 2022

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

Appendix D Non-response bias

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

In this section we examine the robustness of our findings to selective item non-response using "inverse probability weighting" (Wooldridge et al., 2002; Stantcheva, 2022). This method is commonly used to address differential attrition by utilizing the relationships among observed covariates and response rates to re-weight the observed data in a way that approximates the distribution in the full data set (Stantcheva, 2022; Glynn and Quinn, 2010). In practice, we run our baseline specification weighting observations by the inverse of the probability of being part of the respective estimation sample. The latter probability is obtained as the propensity score from estimating by logit equations 6 and 7, where X_i include covariates measured at baseline for all observations and the dependent variable is equal to one for firms belonging to the estimation sample, and zero for other firms surveyed in Invind, including firms dropped due to energy variable misreporting (see previous sections of this Appendix).[22]

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

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

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

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

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

Figure D.1: Common support of propensity score

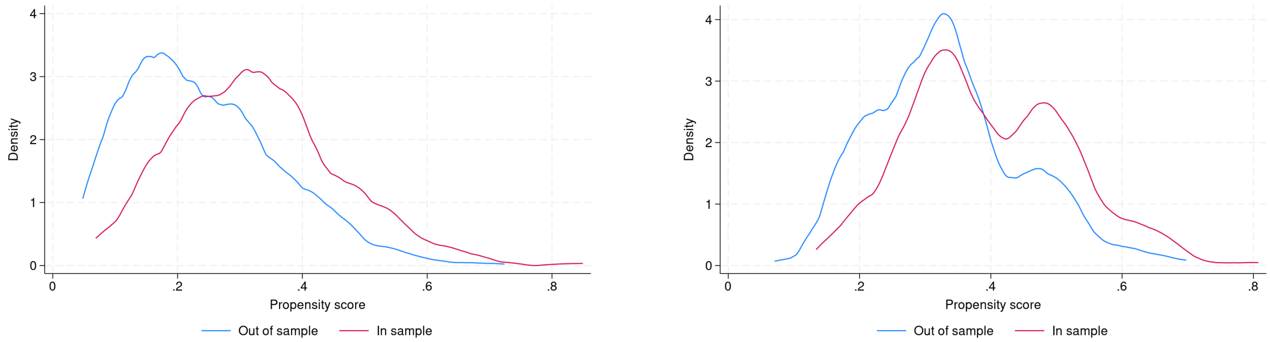

(a) Natural gas sample

(b) Electricity sample

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

Figure D.2: Inverse-probability-weighted estimates

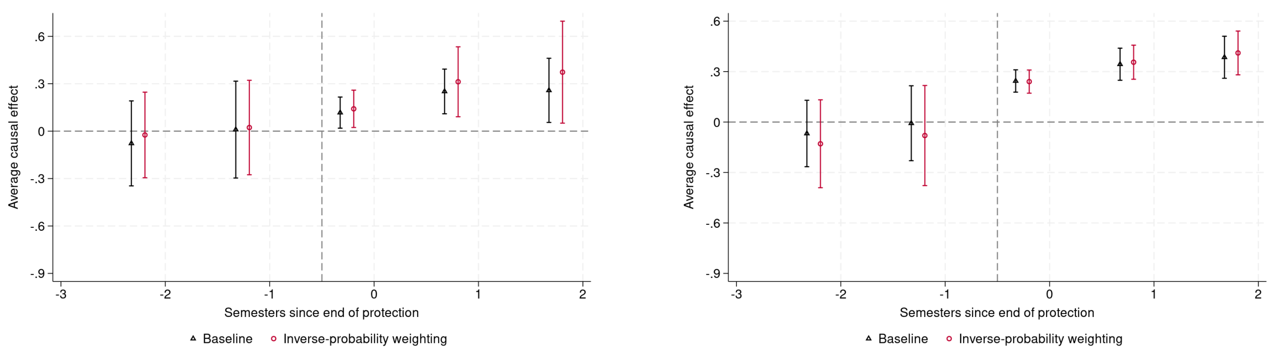

(a) Price of natural gas

(b) Price of electricity

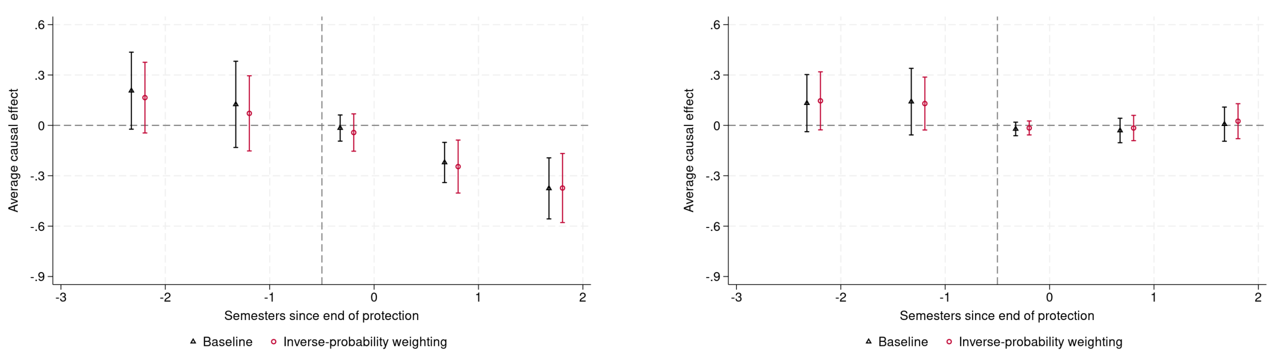

(c) Gas demand

(d) Electricity demand

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

Appendix E Summary statistics

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

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

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

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

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

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

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

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

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

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

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

Note: Invind data. The table reports summary statistics for the energy demand analyses used in Section 4. Characteristics are measured in the first semester of 2021, at baseline (thus the number of observations corresponds to the number of firms). *The variable "Gas is an essential input" is taken from the Business Outlook survey of the Bank of Italy and it refers to the beginning of 2022.

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

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

Notes: Table reports means and standard errors (in parentheses) by treatment cohort. Columns (5)(11) report p-values from pairwise orthogonality tests. Column (12) reports p-values from joint F-tests across treatment arms.

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

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

Notes: Table reports means and standard errors (in parentheses) by treatment cohort. Columns (5)(11) report p-values from pairwise orthogonality tests. Column (12) reports p-values from joint F-tests across treatment arms.

Appendix F Additional robustness tests

As a robustness check, we extend the baseline specification by allowing for flexible, non-parametric differential trends across firm types. Specifically, we modify the first-stage regression of the Borusyak et al. (2024) estimator to include a full set of interactions between time fixed effects and sector dummies, as well as between time fixed effects and indicators for energy-intensive status. This approach allows each sector and energy-intensive firms to follow their own arbitrary evolution over time, mitigating concerns that treatment timing may be correlated with unobserved group-specific dynamics.

Figure F.1 displays the results, together with estimates from our baseline specification. The resulting estimates are virtually unchanged relative to the baseline, both in magnitude and statistical significance. The robustness of the point estimates to the inclusion of these highly flexible controls provides further confidence that our results are not driven by differential sectoral or energy-intensity-specific trajectories unrelated to treatment.

Figure F.1: Robustness to sector-specific and energy-intensive-specific trends

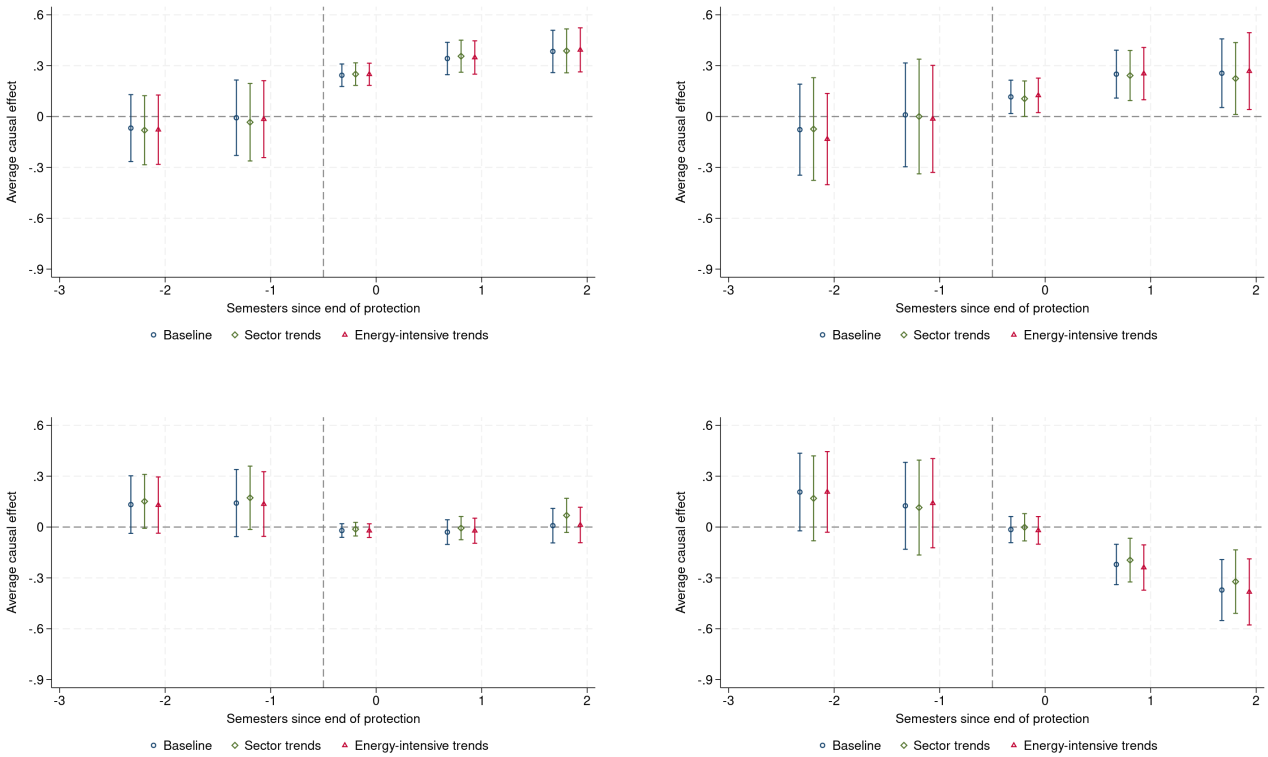

Notes: The figures show average causal effects of the expiration of a fixed-price contract on the average retail prices of electricity and natural gas (panels (a) and (b)) and the corresponding demanded quantities (panels (c) and (d)). Outcome variables are always in logs. Average causal effects before and after the treatment are estimated in two separate regressions, using the "imputation" estimator by Borusyak et al. (2024), as described in Section 3. Blue dots 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. Confidence intervals are at the 95% level.

Appendix G Additional results

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

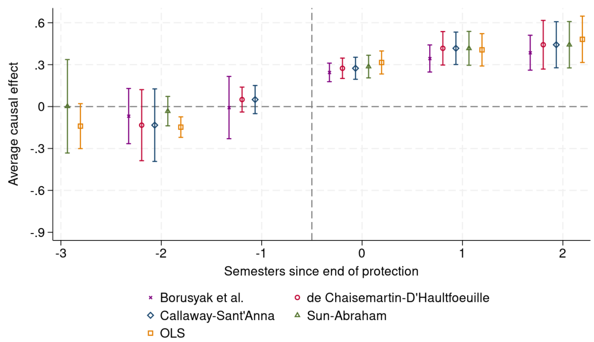

(a) Average cost of electricity

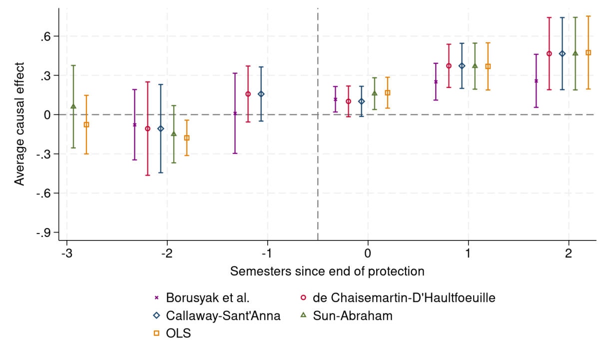

(b) Average cost of natural gas

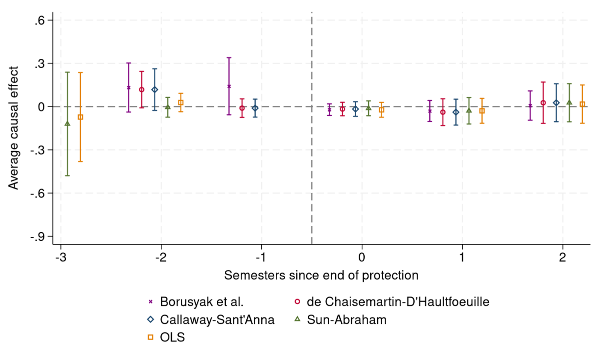

(c) Quantity of electricity

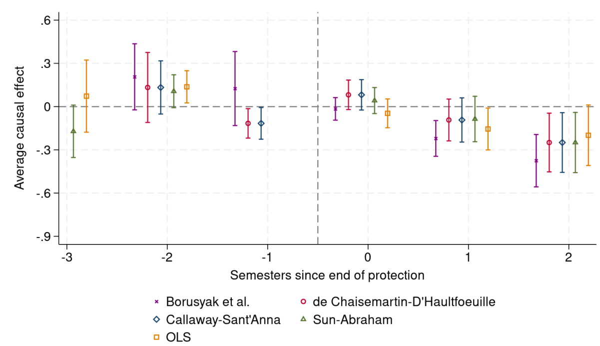

(d) Quantity of natural gas

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

Figure A.2: Natural gas: heterogeneous effects by cohort

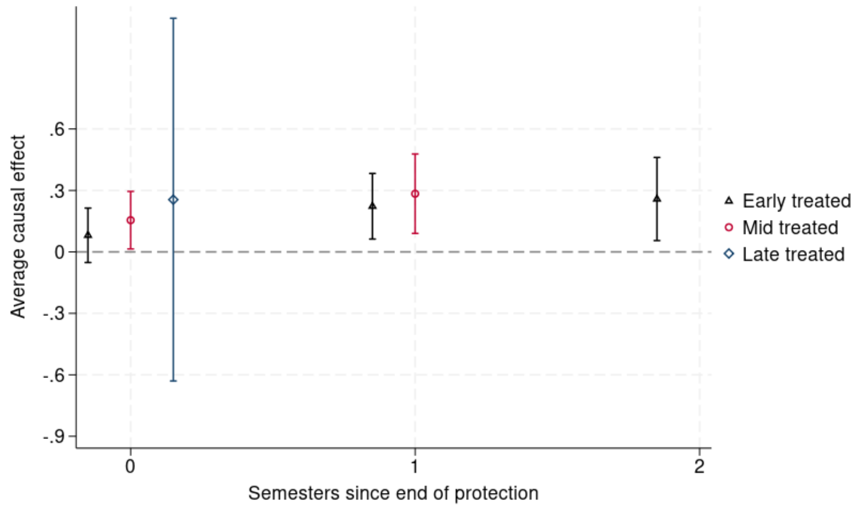

(a) Average cost

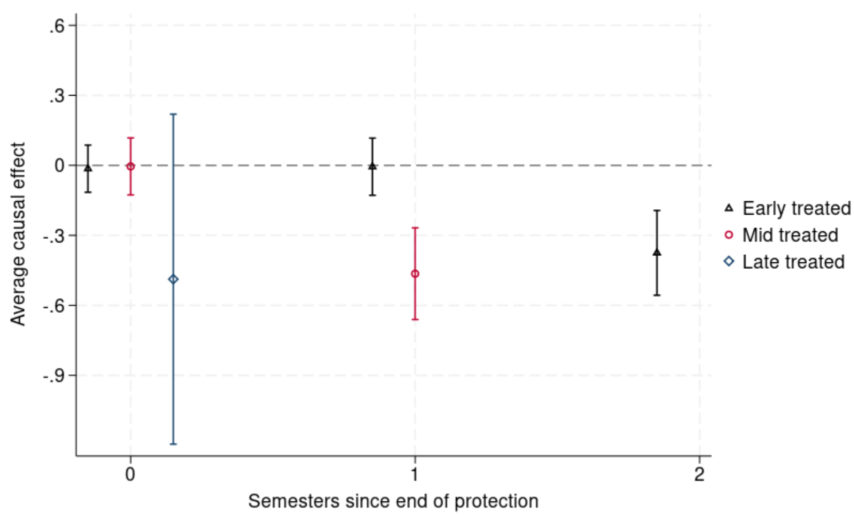

(b) Quantity

Note: the estimates in the upper and lower panel are from the same regression as in panel (b) and (d) of Figure 4.1 respectively; here they are reported by each cohort separately. Confidence intervals at the 95% level.

Figure A.3: Output price growth across cohorts of treatment

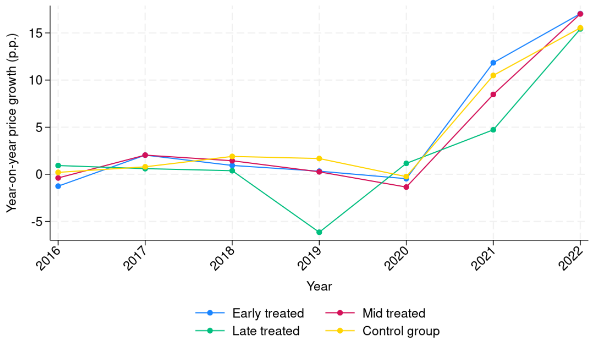

(a) Energy-intensive firms in electricity sample

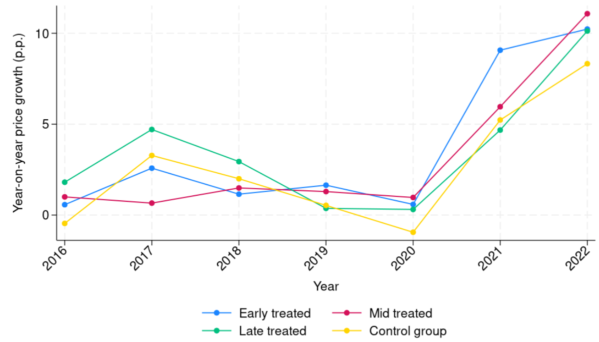

(b) Other firms in electricity sample

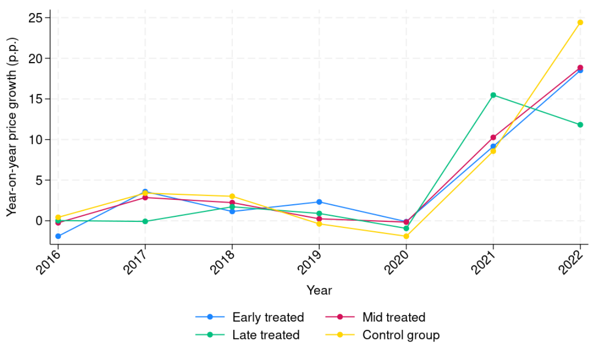

(c) Energy-intensive firms in gas sample

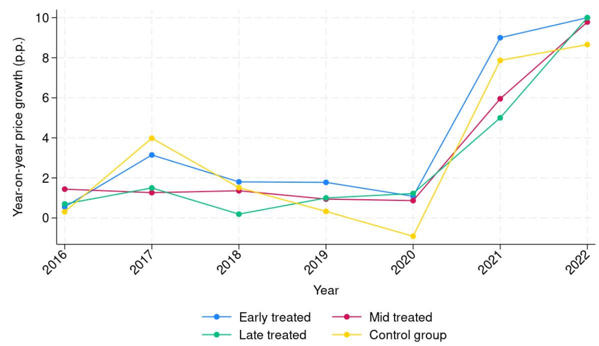

(d) Other firms in gas sample

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

Appendix H Background pictures

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

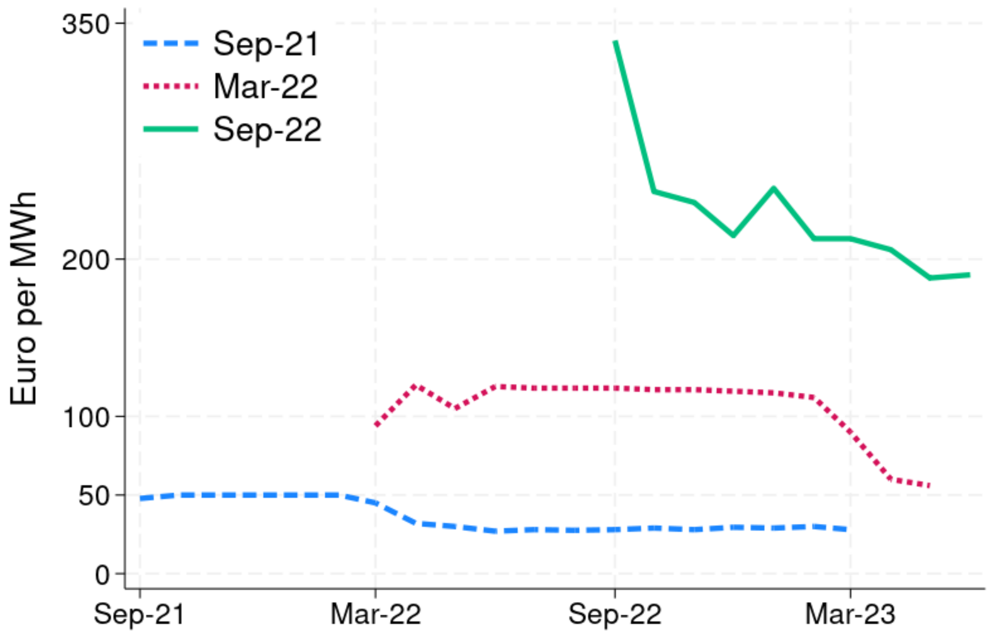

Note:The figure shows futures curve for the Title Transfer Facility (TTF) price at three different points in time: September 2021 (in blue), March 2022 (in red) and September 2022 (in green). The first point of each line is the spot price at that date.

Figure A.2: Expectations on the economic situation

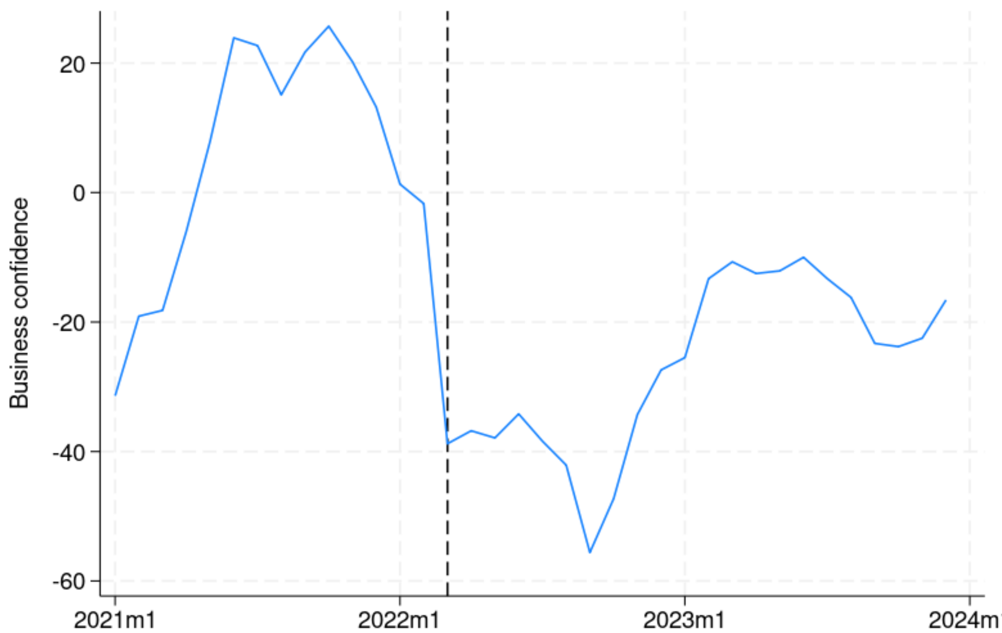

Note:The figure shows the difference between the share of manufacturing firms expecting positive economic situation in the next 3 months and those expecting a negative economic situation. The question refers to the general economic situation of the country. The vertical line refers to March 2022, the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

Figure A.3: Fossil fuel energy mix for firms in the matched EU ETS - Invind sample

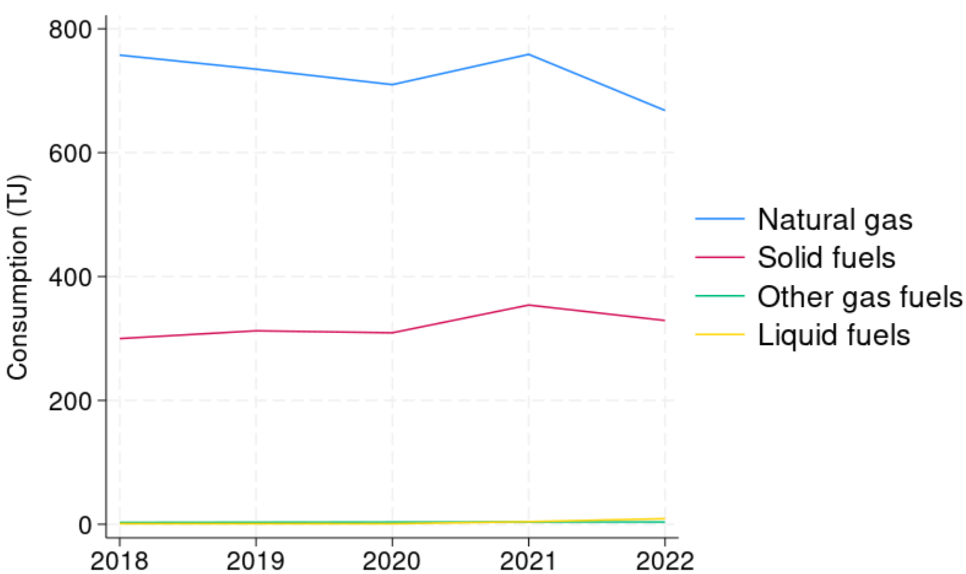

Note: The figure displays the average consumption of different fossil fuels (measured in Terajoules, TJ) for firms belonging to the matched Invind-EU ETS sample.

Appendix I A simple static input choice model

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

$$\eta^G_{p_G} = -(1-s)\sigma - \frac{\mu}{\mu-\theta}s, \tag{8}$$

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

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

$$\eta^G_{p_G} = \begin{cases} -\frac{\mu}{\mu-\theta}s^{EI}, & \text{if } \sigma = 0, \text{ energy intensive (EI)}, \\ -(1-s^{NEI})\sigma - \frac{\mu}{\mu-\theta}s^{NEI}, & \text{if } \sigma > 0, \text{ non energy intensive (NEI)}, \end{cases} \tag{9}$$

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

$$\sigma > \frac{\mu}{\mu-\theta}\frac{(s^{EI}-s^{NEI})}{1-s^{NEI}}. \tag{10}$$

Appendix J Effect heterogeneity with machine learning

In order to study heterogeneity in treatment effect of fixed-price contract expiration on gas consumption in the second semester of 2022, we use random forests in the spirit of Athey and Imbens (2016) and Wager and Athey (2018). First, we split the treated observations in two random sub-samples: a learn sub-sample (60 per cent of the overall sample) and a test sub-sample (40 per cent). Second, we build a forest of 5,000 trees using the learn sub-sample only. Each tree can pick only a random half of the considered covariates, which are: a (self-reported) dummy for whether gas is an essential input in production[23], which can be thought as a proxy of low elasticity of substitution between gas and other inputs; a gas intensive dummy[24]; sector dummies; an EU ETS dummy; and employment in 2021. Third, we use the forest to predict treatment effects out-of-sample in the test sample. Finally, we test whether machine learning (ML) predictions carry over to the test sample. The sample splitting approach ensures that overfitting does not drive our results.

Figure A.1 plots the distribution of individual treatment effects as estimated by the Borusyak et al. (2024) imputation method (in blue), as predicted by ML in sample (black) and out of sample (red). The point estimate is always very close to -0.45 across the three distributions.[25] The blue distribution is very dispersed. However note that this variation could be due to treatment effect heterogeneity (along observables or unobservables) or due to noise (Borusyak et al., 2024), as individual treatment effects contain the error term ϵ_{ijt} in equation 4. The distribution of in-sample-ML-predicted treatment effects (in black) is much less dispersed, but still displaying economically relevant heterogeneity and including values around zero for some observations.[26] The distribution of out-of-sample-ML-predicted treatment effects (in red) is similar to the black one, but even more compressed, with the difference plausibly due to overfitting in the in-sample predictions. Still, the red distribution has a support going from -1 to zero. This means that the forest predicts that some treated firms would decrease gas consumption by as much as 60%, while others would not change it at all.

In order to understand which observables predict heterogeneity in treatment effects, and in which direction, we first regress the in-sample-ML-predicted treatment effect on each covariate separately (see Table A.1). The forest predicts that firms having lower-than-average treatment effects (i.e. a small gas demand reduction) are concentrated: in the food sector; in the chemicals-pharmaceutical-rubber sector; firms declaring natural gas to be an essential

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

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

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

[26]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 A.1: Distribution of treatment effects on gas demand in 2022h2

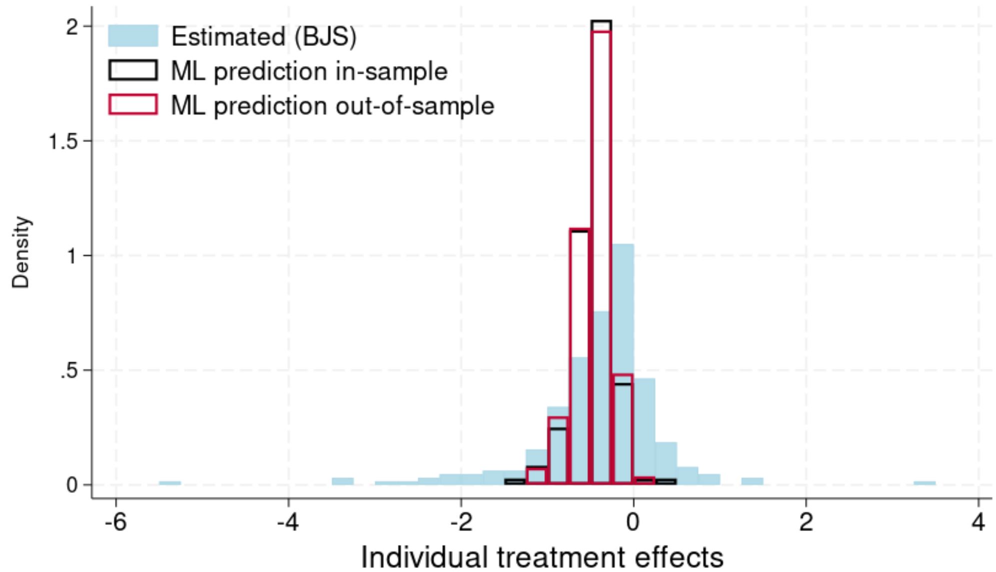

Note: The figure shows the distribution of individual treatment effects in 2022h2. The outcome is natural gas consumption.

input; those subject to the EU ETS; the gas intensive ones according to the Italian state aid regulation. Firms having higher-than-average treatment effects (i.e. a large gas demand reductions) are those in the wood and paper industry. The results are confirmed in the test sample when using as an outcome variable the out-of-sample-ML-predicted treatment effects (Table A.2). In order to test whether these results represent true heterogeneity and not a statistical fluke due to overfitting, we estimate the same regressions in the test sub-sample using as an outcome variable the estimates of the treatment effect obtained with the Borusyak et al. (2024) method. If treatment effect heterogeneity is real, we would expect to see the same signs and similar coefficients in these data, because the Borusyak et al. (2024) estimates in the test sample were never used to train the forest. Results are presented in Table A.2. For some covariates, signs are different and/or coefficients are greatly attenuated, but some of the predictions are confirmed out-of-sample, both in terms of sign and size of the coefficients. In particular, four (non mutually exclusive) groups of observations display lower gas adjustment than the average: firms in the food industry, firms declaring that gas is an essential input, firms in the EU ETS, and gas intensive firms. We take this as evidence that treatment effect heterogeneity exists in this context along these covariates.

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

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

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

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

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

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

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

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

Standard errors in parentheses

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

Note: OLS regressions in the test sub-sample. The outcome variable is the treatment effect (on gas consumption in 2022h2) as estimated by the imputation methods by Borusyak et al. (2024).

Appendix K Price elasticity of energy demand

Throughout the paper, we compute the price elasticity of energy demand by combining our estimates of the effect of fixed-price contract expirations on (log) prices and on (log) quantities obtained from the Borusyak et al. (2021) estimator. In practice, we follow an IV-LATE approach, where we scale the average treatment effect estimated in the quantity equation by the average treatment effect estimated in the price equation.[27] Under standard IV-LATE assumptions, we can construct an estimate of the elasticity θ as follows:

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

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

We construct standard errors using the delta method. The vector of partial derivatives is:

$$\nabla_\theta = \begin{bmatrix} \frac{e^{\hat{\tau}_t^q}}{e^{\hat{\tau}_t^p} - 1} \\ -\frac{(e^{\hat{\tau}_t^q} - 1)e^{\hat{\tau}_t^p}}{(e^{\hat{\tau}_t^p} - 1)^2}. \end{bmatrix}$$

The variance of θ is given by:

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

where

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

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

[27] An alternative approach could be to regress the individual treatment effects from the price equation on the individual treatment effect from the quantity equation following Deryugina et al. (2020). However, this strategy is unfeasible in our application because we construct prices as the ratio between expenditures (in monetary terms) and physical quantity (see section 2). Thus any measurement error in quantity will translate in a measurement error of opposite sign in prices, thus yielding a negative correlation by construction. This issue is reminiscent of the "division bias" discussed by Borjas (1980) in the context of using the ratio of earnings and hours as a proxy for wage in the regression of hours on wages.

The variance of the elasticity is thus

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