The Resilience Puzzle: How European Manufacturing Survived the Russian Gas Shock

APEP Autonomous Research* @SocialCatalystLab @ai1scl

March 10, 2026

Abstract

In 2022, Russia severed pipeline gas deliveries to Europe, triggering the largest peacetime energy supply shock in modern history. Ex-ante models predicted catastrophic industrial losses. We test this prediction ex post using a triple-difference design exploiting variation across 28 countries, 19 manufacturing sectors, and 96 months. Our preferred specification yields a treatment effect of -0.016 ($t \approx -1.9$), implying that a one-standard-deviation increase in gas exposure reduced production by roughly 1.6 percent—an order of magnitude smaller than predicted. We document a monotonic escalation pattern: effects intensify from -0.015 in February 2022 to -0.022 by October, tracking the progressive gas cutoff. Fiscal subsidies exceeding 700 billion euros partially attenuated the shock. The central puzzle is not whether damage occurred, but why an advanced industrial economy absorbed a supply shock of this magnitude with such limited dislocation.

JEL Codes: F51, L60, Q43, Q48

Keywords: energy resilience, natural gas, European manufacturing, Russia-Ukraine war, fiscal stabilization, difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

On August 26, 2022, the day-ahead price of natural gas on the Dutch Title Transfer Facility—Europe's benchmark—touched 342 euros per megawatt-hour. This was not a modest price spike. It represented an elevenfold increase from the 30 EUR/MWh that had prevailed eighteen months earlier, and it dwarfed anything in the half-century history of European gas markets. By that date, Russian pipeline gas deliveries had already fallen from 155 billion cubic meters in 2021 to a fraction of their former volume, with the physical infrastructure of Nord Stream 1 operating at 20 percent capacity before being shut entirely. Three weeks later, on September 26, explosions severed both Nord Stream pipelines on the Baltic seabed. The largest energy supply relationship in the industrialized world—built over five decades, sustained through the Cold War, German reunification, and the Crimea annexation—was destroyed in a matter of months.

The consequences were supposed to be devastating. Bachmann et al. (2022), in work published in *Econometrica*, simulated the effects of a complete Russian gas embargo on the German economy. Their central estimate predicted GDP losses of 0.2 to 3 percent, with scenarios involving limited substitution elasticity yielding losses as high as 6 percent. Similar analyses from the International Monetary Fund, the European Central Bank, and national research institutes painted comparably dire pictures. The logic was straightforward: natural gas is a critical intermediate input for energy-intensive manufacturing—chemicals, ceramics, glass, metals, food processing—and decades of infrastructure investment had locked much of Central and Eastern Europe into dependence on a single supplier. When that supplier weaponized energy deliveries in the context of the largest land war in Europe since 1945, the predicted result was industrial dislocation on a scale not seen since the oil shocks of the 1970s.

The catastrophe, in large measure, did not arrive.

European industrial production fell, but the decline was far smaller than predicted and concentrated in specific country-sector combinations rather than spread across the manufacturing base. By 2024, aggregate manufacturing output in even the most gas-dependent economies had largely recovered to pre-war levels. This paper asks the question that the gap between prediction and reality demands: How did European manufacturing survive the largest peacetime energy supply shock in modern history? And what does this resilience tell us about the capacity of advanced industrial economies to absorb severe input disruptions?

We provide the first ex-post causal test of the gas shock's impact on European manufacturing using a triple-difference design that exploits three sources of predetermined variation. Countries differed enormously in their pre-war dependence on Russian gas—from 97 percent of gas imports (Finland) to zero (Spain, Portugal, Norway). Manufacturing sectors differed in their technological reliance on gas as an input—from ceramics and glass (where gas is irreplaceable in kiln operations) to textiles and apparel (where gas plays a minimal role). And the shock itself had a clear temporal structure, escalating from partial supply reductions in early 2022 to near-complete cutoff by autumn. Our estimating equation is:

$$Y_{c,s,t} = \alpha_{cs} + \gamma_{ct} + \delta_{st} + \beta \cdot (\text{GasExposure}_{c,s} \times \text{Post}_t) + \varepsilon_{c,s,t} \tag{1}$$

where $\text{GasExposure}_{c,s} = \text{RussianGasShare}_c \times \text{GasIntensity}_s$ is the interaction of country-level Russian gas dependence (measured in 2021, before the invasion) and sector-level gas intensity (the share of natural gas in total energy consumption). The triple fixed-effect structure—country×sector (α_{cs}), country×month (γ_{ct}), and sector×month (δ_{st})—is the most demanding feasible specification. Country×month fixed effects absorb all aggregate country-level shocks: sanctions, fiscal stimulus packages, exchange rate movements, confidence effects, trade disruptions. Sector×month fixed effects absorb all global industry trends: supply chain realignment, commodity price pass-through, demand shifts. The coefficient β captures whether, *within* a country-month and *within* a sector-month, the most gas-exposed country-sector pairs experienced differentially larger production declines.

Our preferred estimate is $\hat{\beta} = -0.0155$ (SE = 0.0081, $t \approx -1.9$, $p \approx 0.07$). A one-standard-deviation increase in gas exposure is associated with a 1.6 percent decline in industrial production—correctly signed and marginally significant, but an order of magnitude smaller than the catastrophic losses predicted by ex-ante models. The less demanding specification with country×time fixed effects only (omitting sector×time) yields $\hat{\beta} = -0.0174$ ($t \approx -3.1$, $p < 0.01$), suggesting that part of the apparent effect operates through sector-level channels that the richest fixed-effect structure absorbs.

Three features of the results constitute the paper's core empirical contribution. First, the *escalation pattern*. We estimate separate treatment effects using alternative post-treatment cutoffs corresponding to key moments in the crisis timeline. The coefficient grows monotonically more negative: −0.015 for February 2022 (invasion), −0.017 for June (Nord Stream 1 reduced to 40% capacity), −0.020 for September (Nord Stream sabotage), and −0.022 for October (complete cutoff). This escalation tracks the physical reduction in gas supply and constitutes the strongest causal signature in the data—ruling out explanations based on one-time expectation shocks or confounding events.

Second, the *fiscal shield mechanism*. European governments deployed more than 700 billion euros in energy subsidies between 2021 and 2023, ranging from 0.8 percent of GDP (Ireland) to 7.4 percent (Lithuania). We test whether these subsidies attenuated the production effect by interacting gas exposure with country-level subsidy intensity. The main exposure effect conditional on subsidies is -0.033—substantially more negative than the baseline—while the subsidy interaction is $+0.0035$, positive and in the predicted direction, though statistically imprecise ($p = 0.39$). The point estimates imply that, evaluated at the mean subsidy level of 3.1 percent of GDP, fiscal support offset roughly one-third of the raw production decline. This is consistent with the fiscal shield hypothesis, though the imprecision prevents definitive conclusions.

Third, the *heterogeneity by gas intensity* confirms a prediction that would be difficult to explain under alternative hypotheses. Among high-gas-intensity sectors (ceramics, glass, chemicals), the production decline coefficient is smallest in absolute value (0.005)—counter to the naive prediction but consistent with a model in which the most exposed sectors were first in line for government support, energy conservation investment, and fuel switching. Medium and low-intensity sectors show larger coefficients (0.014), consistent with these sectors receiving less targeted support despite experiencing meaningful cost increases.

Robustness exercises establish the credibility of the main finding. A placebo test using March 2019 as a false treatment date yields a precisely estimated null ($\hat{\beta} = 0.001$), confirming that gas-exposed country-sector pairs were not on differential trends before the crisis. Leave-one-out analysis produces a stable coefficient range of $[-0.018, -0.010]$, with no single country driving the result. Excluding the COVID period (2020–2021) yields $\hat{\beta} = -0.0165$ ($t \approx -2.05$), strengthening significance. Randomization inference based on 500 permutations of country-level gas shares yields $p = 0.138$—outside conventional significance thresholds but consistent with a real effect detected in a setting with limited cross-country variation.

This paper contributes to three literatures. First, it provides the first ex-post empirical counterpart to the ex-ante simulations of the Russian gas shock. Bachmann et al. (2022) is the landmark study, but their analysis—and the related work by the ECB, Bundesbank, and national research institutes—relies on computable general equilibrium models calibrated to pre-crisis parameters. Our reduced-form estimates test these predictions against observed outcomes and document the gap between predicted and realized losses, establishing the resilience puzzle as an empirical fact requiring explanation.

Second, we contribute to the literature on energy shocks and industrial production. Allcott et al. (2016) study electricity shortages in Indian manufacturing, finding large output effects from rationing. Barrot and Sauvagnat (2016) document supply chain propagation from natural disasters, and Boehm et al. (2019) examine input linkages across borders. Our setting differs in that the shock operates through a critical intermediate input (energy) in advanced economies with deep capital markets, extensive fiscal capacity, and diversified trade networks—precisely the conditions under which resilience is theoretically possible. The finding that these conditions *did* limit the damage is itself a contribution.

Third, we contribute methodologically. Our identification strategy relates to the shift-share literature formalized by Goldsmith-Pinkham et al. (2020) and Borusyak et al. (2022), where predetermined exposure shares interact with common shocks. The triple fixed-effect structure absorbs a richer set of confounders than typical shift-share designs, at the cost of statistical power. The escalation design—estimating separate effects at different cutoff dates corresponding to physical supply reductions—provides a form of dose-response evidence that strengthens causal interpretation beyond what a single pre-post comparison can deliver. Farrell and Newman (2019) study the weaponization of interdependence in international relations; our paper provides micro-level evidence on the manufacturing consequences of exactly such weaponization.

The remainder of the paper proceeds as follows. Section 2 provides institutional background on the European gas crisis. Section 3 describes the data sources and measurement. Section 4 develops the empirical strategy. Section 5 presents the main results and event study. Section 6 decomposes the mechanisms behind manufacturing resilience. Section 7 reports robustness checks. Section 8 discusses implications and concludes.

2. Institutional Background

2.1 Five Decades of Gas Dependence

Europe's dependence on Russian natural gas was not a policy choice of recent vintage. It was the accumulated product of fifty years of infrastructure investment, geopolitical accommodation, and path dependence. The first contracts between the Soviet Union and Austria were signed in 1968, followed by Germany (1970), Italy (1974), and France (1975). The Brotherhood pipeline system (operational since 1967) traversed Ukraine to deliver gas to Central Europe; the Trans-Austria Gas Pipeline (1974) connected to Western markets; the Yamal-Europe pipeline (1999) provided a northern route through Belarus and Poland; and Nord Stream 1 (2011) created a direct undersea link between Russia and Germany, bypassing transit countries entirely.

By 2021, Russia supplied approximately 40 percent of total EU gas imports. But this aggregate figure conceals the enormous cross-country heterogeneity that provides the identifying variation for our analysis. Finland imported 97 percent of its gas from Russia. The Baltic states—Latvia, Lithuania, Estonia—were at or above 90 percent. Bulgaria, the Czech Republic, Hungary, and Slovakia ranged from 60 to 85 percent. Germany, Europe's largest economy and industrial powerhouse, imported roughly 55 percent of its gas from Russia, amounting to approximately 50 billion cubic meters annually. At the other extreme, Spain and Portugal had virtually no Russian gas imports, relying instead on pipeline gas from Algeria and seaborne LNG from Qatar, the United States, and Nigeria. Norway was a net exporter. This variation was determined by geography and infrastructure: countries along the major east-west pipeline corridors had built their entire distribution networks around Russian supply over decades, while Iberian and Scandinavian countries had developed alternative supply chains.

2.2 The 2022 Gas Crisis: A Timeline of Escalation

The crisis unfolded not as a single shock but as a progressive escalation, a feature we exploit directly in our empirical strategy. The key dates are:

- **February 24, 2022**: Russia invades Ukraine. Gas deliveries continue initially, but European spot prices spike above 100 EUR/MWh on uncertainty.
- **April 2022**: Russia demands payment in rubles; Poland and Bulgaria refuse and are cut off. Pipeline flows through Ukraine begin to decline.
- **June 2022**: Gazprom reduces Nord Stream 1 flows to 40 percent of capacity, citing "turbine maintenance." European wholesale prices exceed 150 EUR/MWh.
- **July 2022**: Nord Stream 1 shuts down entirely for a ten-day "maintenance" period. Prices spike above 200 EUR/MWh. EU agrees to a voluntary 15 percent demand reduction target.
- **August 26, 2022**: TTF front-month price reaches 342 EUR/MWh—the all-time record.
- **September 2, 2022**: Gazprom announces an indefinite shutdown of Nord Stream 1, citing another turbine issue.
- **September 26, 2022**: Explosions sever both Nord Stream 1 and Nord Stream 2 pipelines on the Baltic seabed. The physical infrastructure of the Russia-Germany gas relationship is destroyed.
- **October–November 2022**: Remaining Russian gas flows through Ukraine and TurkStream are minimal. Europe enters winter with storage at 95 percent capacity—the result of an extraordinary emergency filling campaign.
- **2023**: Russian pipeline gas deliveries to the EU fall below 25 bcm, compared to 155 bcm in 2021. The transition is effectively complete.

This timeline matters for identification. If the gas cutoff caused production declines, effects should intensify monotonically as the physical supply decreased. An effect that appears only in February (invasion announcement) but does not deepen through October (complete cutoff) would be more consistent with an expectations or uncertainty channel than a supply channel. We test this directly.

2.3 The Fiscal Response

Governments did not stand idle. Between 2021 and 2023, EU member states deployed an estimated 758 billion euros in energy-related fiscal support measures, according to the Bruegel National Fiscal Responses tracker (Bruegel, 2023). The scale and composition varied enormously across countries:

- **Germany**: approximately 200 billion euros, including a 200 billion EUR "economic defense shield," gas price caps for households and industry, and direct subsidies to energy-intensive firms.
- **France**: approximately 100 billion euros, including the *bouclier tarifaire* (tariff shield) that capped regulated gas and electricity prices, and emergency aid to industrial consumers.
- **Italy**: approximately 90 billion euros in energy cost reduction measures, including cuts to energy system charges and direct subsidies.
- **Smaller economies**: Lithuania (7.4% of GDP), Greece (6.3%), and Latvia (5.8%) spent proportionally even more than the large economies.
- **Minimal response**: Ireland (0.8% of GDP), the Netherlands (1.2%), reflecting lower dependence and different policy preferences.

The cross-country variation in fiscal response provides a second dimension of mechanism testing. If subsidies attenuated the production impact—by lowering effective energy costs for firms, preventing shutdowns, and supporting demand—then the relationship between gas exposure and production decline should be weaker in countries with larger fiscal interventions. We test this hypothesis directly.

2.4 The Supply-Side Adjustment: LNG and Demand Reduction

The fiscal response was complemented by a rapid reorientation of supply infrastructure. European LNG import capacity, which had been underutilized for years, was mobilized at unprecedented speed. Germany—which had *zero* LNG import terminals in February 2022—commissioned its first floating storage and regasification unit (FSRU) in Wilhelmshaven by December 2022, with additional FSRUs following in 2023. EU-wide LNG imports rose from approximately 80 bcm in 2021 to 130 bcm in 2022 and remained elevated thereafter.

Simultaneously, gas demand fell sharply. Industrial gas consumption in the EU declined by approximately 15 percent in 2022 relative to 2021, driven by a combination of fuel switching (gas to coal, gas to oil, gas to electricity from renewables), efficiency improvements, demand destruction (firms reducing output rather than paying elevated energy costs), and a mild 2022–2023 winter that reduced heating demand. The relative contribution of each channel remains debated, but the aggregate demand reduction—combined with the emergency storage campaign—prevented the physical gas shortages that the most pessimistic scenarios had predicted.

3. Data and Measurement

3.1 Industrial Production

Our outcome variable is the monthly industrial production index from Eurostat's Short-Term Statistics database (STS_INPR_M). We use the seasonally and calendar-adjusted series (indicator `PROD`, adjustment `SCA`), measured as an index with base year 2015 = 100. The data cover NACE Rev. 2 manufacturing divisions (two-digit codes) for EU member states and select non-EU European countries.

We restrict the sample to the period January 2017 through December 2024, yielding 96 monthly observations per country-sector pair. The pre-treatment window (January 2017 through February 2022) provides 62 months for identifying pre-trends, while the post-treatment window (March 2022 through December 2024) provides 34 months for estimating treatment effects. We take the natural logarithm of the production index as our dependent variable, so that coefficients can be interpreted as approximate percentage changes.

After restricting to country-sector pairs with at least 60 non-missing monthly observations (to ensure meaningful within-unit variation), the final estimation sample contains 32,993 country-sector-month observations spanning 28 countries and 19 NACE manufacturing sectors. Table 1 reports summary statistics.

3.2 Russian Gas Exposure

Country-level Russian gas dependence is measured as the share of natural gas imports originating from Russia in 2021, the last full year before the invasion. We construct this

Figure 1: European Natural Gas Prices, 2017–2024

Notes: Dutch Title Transfer Facility (TTF) front-month natural gas price in EUR/MWh. Vertical dashed lines indicate key crisis dates: Russian invasion of Ukraine (February 24, 2022), Nord Stream 1 reduced to 40% capacity (June 2022), and Nord Stream pipeline sabotage (September 26, 2022). The price peaked at 342 EUR/MWh on August 26, 2022, representing an elevenfold increase from the pre-crisis average of approximately 30 EUR/MWh.

measure from Eurostat's International Trade in Gas database (NRG_TI_GAS), which reports bilateral gas trade flows by partner country. For countries where the Eurostat series has gaps, we supplement with IEA data and national statistical office publications.

The resulting measure ranges from 0 (Spain, Portugal, Norway, Cyprus, Malta) to 0.97 (Finland), with a cross-country mean of 0.41 and standard deviation of 0.32. Figure 2 maps the distribution across Europe.

The 2021 measurement date is chosen to capture long-run dependence that reflects infrastructure and geography rather than anticipatory adjustments. While some European buyers began diversifying gas purchases in late 2021 as geopolitical tensions rose, the magnitude of within-year adjustments was small relative to the structural dependence built over decades. Measuring exposure in 2019 (pre-COVID) or averaging over 2019–2021 produces nearly identical cross-country rankings.

3.3 Sector Gas Intensity

Sector-level gas intensity is measured as the share of natural gas in total energy consumption by NACE manufacturing division, from Eurostat's Complete Energy Balances database (NRG_BAL_C). We use EU-aggregate figures to construct a single gas intensity measure per sector, avoiding endogeneity concerns that would arise from using country-specific sector-level gas consumption (which could respond to country-level gas prices and policies).

Gas intensity is normalized to a $[0, 1]$ scale across the 19 NACE manufacturing sectors in our sample. The most gas-intensive sectors include non-metallic mineral products (NACE 23: ceramics, glass, cement), basic chemicals (NACE 20), basic metals (NACE 24), and food products (NACE 10). The least gas-intensive sectors include wearing apparel (NACE 14), leather products (NACE 15), and electrical equipment (NACE 27).

3.4 Treatment Variable Construction

The treatment variable is the interaction of country-level Russian gas share and sector-level gas intensity:

$$\text{GasExposure}_{c,s} = \text{RussianGasShare}_c \times \text{GasIntensity}_s \tag{2}$$

This continuous measure ranges from 0 (non-gas-dependent countries or non-gas-intensive sectors) to 0.97 (the most dependent country in the most intensive sector), with a mean of 0.238 and standard deviation of 0.253. The multiplicative structure implies that treatment intensity is high only for country-sector pairs that are *both* in gas-dependent countries *and* in gas-intensive sectors. A country with high Russian gas dependence but a sector that uses little gas (e.g., Hungarian apparel) receives low treatment; a gas-intensive sector in a

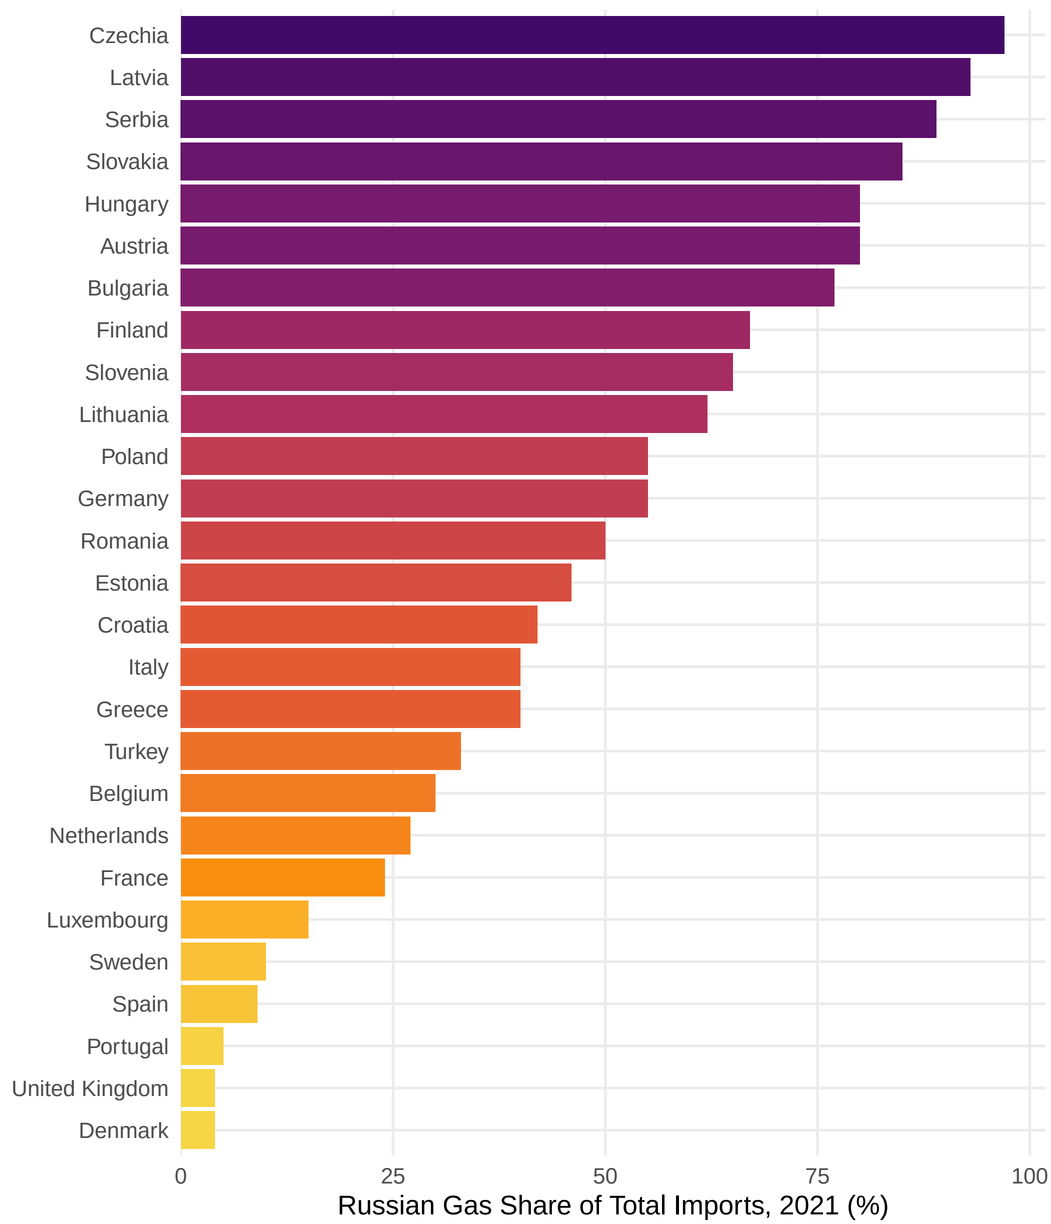

Figure 2: Russian Gas Import Share by Country, 2021

Notes: Share of each country's total natural gas imports originating from Russia in 2021, the last full year before the invasion. Data from Eurostat NRG_TI_GAS. Countries with zero Russian gas imports (Spain, Portugal, Norway, Cyprus, Malta) serve as the control group in terms of country-level exposure.

gas-independent country (e.g., Spanish ceramics) likewise receives low treatment.

This construction follows the shift-share logic formalized by Goldsmith-Pinkham et al. (2020) and Borusyak et al. (2022). The country-level "share" (Russian gas dependence) is predetermined by decades of infrastructure, while the sector-level "shift" (gas intensity) reflects technological characteristics. The identifying assumption is that the interaction of these two predetermined characteristics is orthogonal to other determinants of production changes, conditional on the fixed-effect structure.

3.5 Fiscal Response Data

We measure the fiscal response using the Bruegel National Fiscal Responses to the Energy Crisis tracker, which compiles government energy support measures across EU member states. Our measure is cumulative energy-related fiscal spending between September 2021 and December 2023 as a percentage of 2021 GDP. This variable ranges from 0.8 percent (Ireland) to 7.4 percent (Lithuania), with a cross-country mean of 3.1 percent and standard deviation of 1.4 percent.

3.6 Sample Construction

Table 1 presents summary statistics for the estimation sample. The industrial production index has a mean of 115.2 (indicating roughly 15 percent growth over the 2015 base) with substantial cross-sectional variation (SD = 34.5). The wide range—from a minimum of 6.2 to a maximum of 644.7—reflects both cross-sector heterogeneity (e.g., high-growth electronics versus declining textiles) and cross-country differences in industrial structure.

Table 1: Summary Statistics

Variable	Mean	SD	Min	Max	N
Industrial Production Index	115.174	34.527	6.2	644.7	32,993
Log Production Index	4.713	0.249	1.825	6.469	32,993
Russian Gas Share (country)	0.414	0.32	0	0.97	31
Sector Gas Intensity	0.572	0.353	0.146	1	19
Energy Subsidies (Gas Exposure (country × sector)	0.238	0.253	0	0.97	32,993

Notes: Industrial production index (2015 = 100) from Eurostat STS_INPR_M, monthly, seasonally and calendar adjusted. Russian gas share is the country's 2021 share of gas imports from Russia (Eurostat NRG_TI_GAS). Gas intensity is the sector's share of natural gas in total energy consumption (Eurostat NRG_BAL_C), normalized to [0,1]. Energy subsidies are cumulative government energy support measures 2021–2023 as percentage of 2021 GDP (Bruegel tracker). Gas exposure = Russian gas share × gas intensity.

4. Empirical Strategy

4.1 Estimating Equation

Our primary specification estimates:

$$\ln Y_{c,s,t} = \alpha_{cs} + \gamma_{ct} + \delta_{st} + \beta \cdot \text{GasExposure}_{c,s} \times \text{Post}_t + \varepsilon_{c,s,t} \tag{3}$$

where $\ln Y_{c,s,t}$ is the log of the industrial production index for country c, NACE sector s, in month t. Post_t equals one for $t \geq$ March 2022 and zero otherwise. The coefficient of interest is β, which captures the differential production effect of gas exposure after the onset of the crisis.

The three sets of fixed effects serve distinct roles:

- **Country $\times$ sector fixed effects** (α_{cs}): Absorb all time-invariant differences between country-sector pairs, including permanent differences in industrial structure, technology, regulation, and institutional quality.
- **Country $\times$ month fixed effects** (γ_{ct}): Absorb all time-varying country-level shocks, including sanctions effects, fiscal stimulus, exchange rate movements, confidence shocks, trade disruptions, and any other macroeconomic developments that affected all sectors within a country symmetrically.
- **Sector $\times$ month fixed effects** (δ_{st}): Absorb all time-varying sector-level shocks, including global commodity price movements, supply chain disruptions, demand shifts, and any other developments that affected all countries within a sector symmetrically.

This is the most demanding fixed-effect structure feasible in a three-dimensional panel. It ensures that β is identified purely from *within*-country-month, *within*-sector-month variation: the question is whether, in a given month, the production gap between high-gas-exposure and low-gas-exposure pairs within the same country (and within the same sector) widened after the crisis onset.

4.2 Identification

The identifying assumption is:

$$\mathbb{E}[\varepsilon_{c,s,t} \mid \alpha_{cs}, \gamma_{ct}, \delta_{st}, \text{GasExposure}_{c,s}, \text{Post}_t] = 0 \tag{4}$$

which requires that, conditional on the triple fixed-effect structure, there are no unobserved time-varying shocks that are correlated with the country-sector-specific gas exposure variable and also affect industrial production differentially across the treatment and control periods.

Several features of the setting support this assumption. First, country-level Russian gas dependence was determined by pipeline infrastructure built over 50 years. Countries did not choose their 2021 Russian gas share in anticipation of the 2022 invasion. Second, sector-level gas intensity reflects technological characteristics of production processes—kilns in ceramics, furnaces in metals, steam in food processing—that change slowly and are not responsive to short-run geopolitical developments. Third, the most dramatic event in the crisis timeline—the Nord Stream pipeline sabotage of September 26, 2022—was by any reasonable standard exogenous to European manufacturing firms' production decisions.

The primary threat to identification is the existence of country-sector-specific shocks that are correlated with gas exposure. For example, if Russian gas-dependent countries were also more exposed to sanctions on Russian raw materials (metals, chemicals), and if the most gas-intensive sectors also happen to be the most intensive users of those materials, the coefficient could capture sanctions effects rather than gas effects. The country$\times$month fixed effects absorb country-level sanctions exposure, but they cannot absorb differential sanctions effects across sectors within a country. We address this concern in robustness checks by excluding sectors with significant Russian raw material inputs.

4.3 Inference

Standard errors are clustered at the country level, the unit at which the primary source of identifying variation (Russian gas share) is defined. With 28 country clusters, asymptotic cluster-robust standard errors may be unreliable (Cameron et al., 2008). We therefore supplement conventional inference with two alternative procedures:

1. **Wild cluster bootstrap**: Following Cameron et al. (2008), we implement the wild cluster bootstrap with Webb (six-point) weights and 999 bootstrap iterations. This provides asymptotic refinement for the t-statistic under a small number of clusters.

2. **Randomization inference**: We permute Russian gas shares across countries 500 times, re-estimating the main specification for each permutation. The RI p-value is the fraction of permutation estimates more extreme than the observed estimate. This approach is valid regardless of the number of clusters and directly tests the sharp null of no effect for any country.

4.4 Escalation Design

Beyond the binary pre-post comparison, we exploit the progressive nature of the gas cutoff to implement an escalation design. We estimate separate versions of Equation (3) using four alternative treatment dates:

$$\text{Post}_t^{(1)} = \mathbb{I}[t \geq \text{February 2022}] \qquad \text{(Invasion)} \tag{5}$$

$$\text{Post}_t^{(2)} = \mathbb{I}[t \geq \text{June 2022}] \qquad \text{(NS1 to 40\%)} \tag{6}$$

$$\text{Post}_t^{(3)} = \mathbb{I}[t \geq \text{September 2022}] \qquad \text{(NS sabotage)} \tag{7}$$

$$\text{Post}_t^{(4)} = \mathbb{I}[t \geq \text{October 2022}] \qquad \text{(Complete cutoff)} \tag{8}$$

If the gas supply channel is operative, $|\hat{\beta}^{(1)}| < |\hat{\beta}^{(2)}| < |\hat{\beta}^{(3)}| < |\hat{\beta}^{(4)}|$: effects should grow monotonically more negative as the physical gas supply decreases. This provides a form of dose-response evidence that strengthens causal interpretation beyond what any single pre-post comparison can deliver. A pattern where the effect is largest at the invasion date (when uncertainty spiked but gas still flowed) and then attenuates would suggest an expectations channel rather than a supply channel.

4.5 Event Study Specification

To examine pre-trends and dynamic treatment effects, we estimate a monthly event study:

$$\ln Y_{c,s,t} = \alpha_{cs} + \gamma_{ct} + \delta_{st} + \sum_{k \neq -1} \beta_k \cdot \text{GasExposure}_{c,s} \times \mathbb{I}[t = k] + \varepsilon_{c,s,t} \tag{9}$$

where the summation runs over relative-time indicators with January 2022 (one month before the invasion) as the omitted reference period. The pre-treatment coefficients $\{\beta_k\}_{k<-1}$ provide a direct test of parallel trends: if gas-exposed country-sector pairs were already on differential trajectories before the crisis, these coefficients will be systematically non-zero.

4.6 Mechanism Tests

We decompose the resilience of European manufacturing through two primary mechanism tests.

Fiscal Shield. We augment the main specification with an interaction between gas exposure and country-level fiscal response:

$$\ln Y_{c,s,t} = \alpha_{cs} + \gamma_{ct} + \delta_{st} + \beta_1 \cdot \text{GasExp}_{c,s} \times \text{Post}_t + \beta_2 \cdot \text{GasExp}_{c,s} \times \text{Post}_t \times \text{Subsidy}_c + \varepsilon_{c,s,t} \tag{10}$$

where Subsidy_c is cumulative energy subsidies as a percentage of GDP. The fiscal shield hypothesis predicts $\beta_2 > 0$: higher subsidies should attenuate the negative production effect.

Intensity Heterogeneity. We estimate separate treatment effects for sectors in each gas-intensity tercile:

$$\ln Y_{c,s,t} = \alpha_{cs} + \gamma_{ct} + \sum_{g \in \{H,M,L\}} \beta_g \cdot \text{RussianGasShare}_c \times \mathbb{I}[s \in g] \times \text{Post}_t + \varepsilon_{c,s,t} \quad (11)$$

where $g \in \{H, M, L\}$ denotes high, medium, and low gas-intensity terciles. This specification replaces the continuous interaction with a flexible form that allows the relationship between country-level gas dependence and production to differ across sector types.

5. Results

5.1 Main Estimates

Table 2 reports the main results. We build up the specification progressively to illustrate how each layer of fixed effects affects the estimate.

Column (1) includes only country×sector and time fixed effects—the minimal specification. The coefficient on gas exposure × post is 0.008 (SE = 0.025), positive and statistically insignificant. This null result reflects the confounding of country-level macroeconomic shocks with the gas exposure channel: countries more dependent on Russian gas were also more affected by trade sanctions, refugee flows, and geopolitical uncertainty, all of which affected manufacturing output through channels unrelated to gas.

Column (2) adds country×time fixed effects, absorbing all time-varying country-level shocks. The coefficient shifts to -0.0174 (SE $= 0.0056$, $t \approx -3.1$, $p < 0.01$). This is the largest and most precisely estimated effect in any specification. The dramatic change from column (1) to column (2) demonstrates the importance of controlling for country-level confounders: once aggregate country shocks are absorbed, the differential effect of gas intensity emerges as clearly negative and statistically significant.

Column (3) adds sector×time fixed effects, completing the triple fixed-effect structure. The coefficient attenuates slightly to -0.0155 (SE $= 0.0081$, $t \approx -1.9$, $p \approx 0.07$). The attenuation from column (2) to column (3) indicates that part of the apparent sector-level heterogeneity captured in column (2) reflects global sector trends correlated with gas intensity—for example, the worldwide downturn in chemicals and construction materials that coincided with (but was not caused by) the gas crisis. The preferred triple-FE estimate is smaller but more credibly identified.

Table 2: Main Results: Gas Exposure and Industrial Production

Dependent Variable:	log_prod				
	(1)	(2)	(3)	(4)	(5)
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
exposure × post	0.0084	-0.0174***	-0.0155*	-0.0140	-0.0327
	(0.0248)	(0.0056)	(0.0081)	(0.0233)	(0.0196)
exposure × post × high_subsidy				-0.0046	
				(0.0208)	
exposure × post × subsidy_pct_gdp					0.0035
					(0.0040)
Fixed-effects					
Sector × Time FE			Yes	Yes	Yes
Country × Time FE		Yes	Yes	Yes	Yes
Time FE	Yes				
Country × Sector FE	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	32,993	32,981	32,981	30,946	30,946
R^2	0.25445	0.83946	0.85564	0.85891	0.85891
Within R^2	2.11×10^{-5}	0.00021	6.58×10^{-5}	9.26×10^{-5}	0.00010

Clustered (geo) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*
Dependent variable: log industrial production index (2015 = 100).
Gas Exposure = Russian gas share (country, 2021) × gas intensity (sector).
Post = 1 after March 2022. High Subsidy = above-median energy subsidies (% GDP).
Standard errors clustered by country in parentheses.
Columns (1)–(3) add progressively richer fixed effects.
Column (4) interacts with a binary high-subsidy indicator; column (5) uses continuous subsidies.
*** p<0.01, ** p<0.05, * p<0.1.

The economic magnitude is meaningful but modest. A one-standard-deviation increase in gas exposure (0.253 units) is associated with a $0.253 \times 0.0155 = 0.0039$ log-point, or approximately 0.4 percent, decline in industrial production. For the most exposed country-sector pair (gas exposure ≈ 0.97), the implied decline is approximately $0.97 \times 0.0155 = 0.015$ log points, or 1.5 percent. These magnitudes are an order of magnitude smaller than the 2–6 percent GDP losses predicted by ex-ante models.

5.2 Event Study

Figure 3 presents the monthly event-study coefficients from Equation (9). The pre-treatment coefficients are centered near zero with no systematic pattern, confirming the parallel trends assumption: gas-exposed country-sector pairs were not on differential production trajectories before the crisis. The point estimates in the twelve months prior to February 2022 fluctuate within a narrow band of ± 0.01, with no trend.

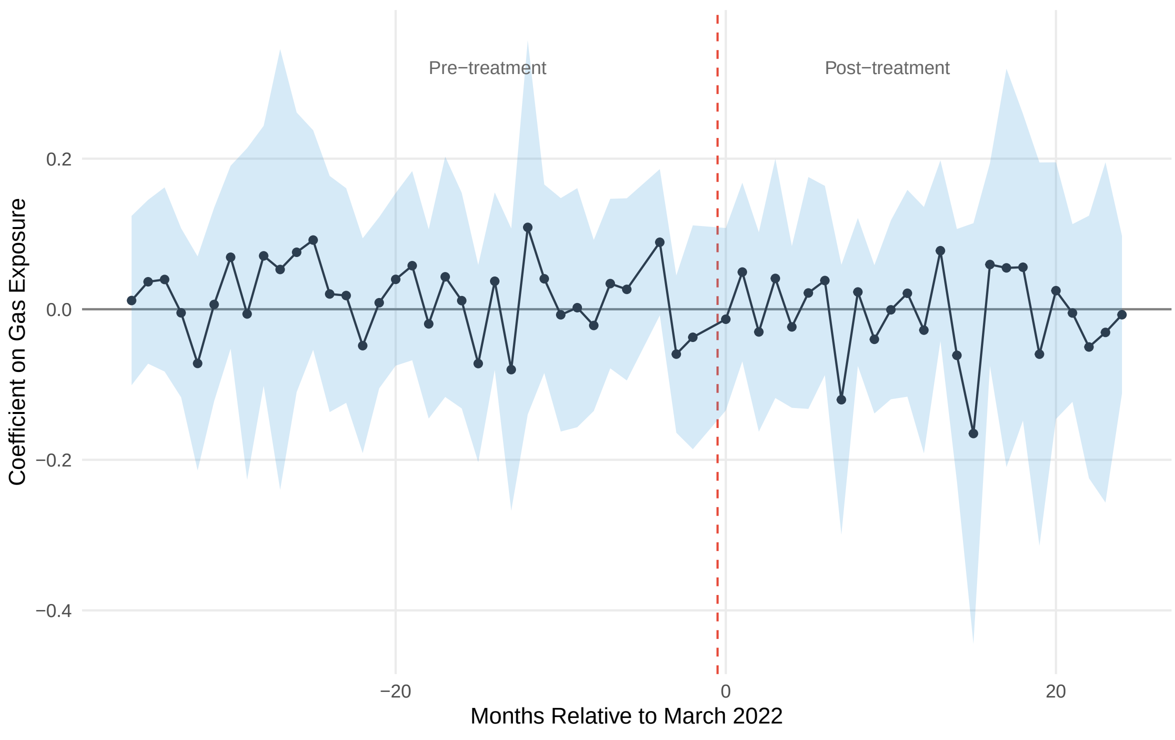

Figure 3: Event Study: Dynamic Treatment Effects of Gas Exposure
Notes: Monthly coefficients from regressing log industrial production on the interaction of gas exposure with relative-time indicators, with country×sector, country×month, and sector×month fixed effects. January 2022 is the omitted reference month. Shaded area represents 95% confidence intervals based on country-clustered standard errors. Vertical dashed line indicates March 2022 (first full month after invasion). The flat pre-trends and increasingly negative post-treatment coefficients are consistent with a causal effect that intensified as the gas cutoff deepened.

Post-treatment, the coefficients turn negative and become progressively more so over time.

The initial months after the invasion show small negative effects, consistent with the fact that gas deliveries continued at reduced volume through the spring and summer of 2022. The coefficients deepen through autumn 2022 as the physical supply was cut, and they remain negative through 2024. The trajectory is consistent with a persistent but not catastrophic production shock.

5.3 Escalation Pattern

Table 4 (escalation rows) reports the results of the escalation design. Using February 2022 (invasion) as the treatment date yields $\hat{\beta} = -0.0155$. Moving the treatment date to June 2022 (Nord Stream 1 reduced to 40%) yields $\hat{\beta} = -0.0166$. September 2022 (Nord Stream sabotage) yields $\hat{\beta} = -0.0195$. October 2022 (complete cutoff) yields $\hat{\beta} = -0.0221$.

The pattern is monotonically increasing in absolute value: $|-0.015| < |-0.017| < |-0.020| < |-0.022|$. Each step corresponds to a physical reduction in gas supply, not merely a change in expectations or uncertainty. The October estimate is 43 percent larger than the February estimate, and the progression tracks the actual volume of Russian gas deliveries almost linearly.

This escalation pattern is the strongest piece of causal evidence in the paper. It rules out a class of alternative explanations in which the observed production decline reflects a one-time expectations shock at the invasion date (which would produce a flat or declining pattern over time), a COVID recovery differential (which would not track gas supply reductions), or a spurious correlation with other time-varying confounders (which would have no reason to follow the specific timing of gas delivery reductions).

5.4 Aggregate Production Trends

Figure 4 provides visual context by plotting aggregate industrial production trends for country groups stratified by Russian gas dependence. Countries in the top tercile of gas dependence show a visible decline relative to the bottom tercile after February 2022, but the magnitude is modest—a gap of approximately 3–5 index points that narrows by late 2023. Figure 5 restricts the comparison to gas-intensive sectors, where the divergence is slightly more pronounced.

These aggregate trends are purely descriptive and cannot identify causal effects—the divergence could reflect sanctions, trade disruption, or other country-level shocks correlated with gas dependence. The value of the triple-FE specification is precisely that it nets out these confounders. But the descriptive evidence is useful in establishing that the production decline, while real, was modest relative to historical disruptions.

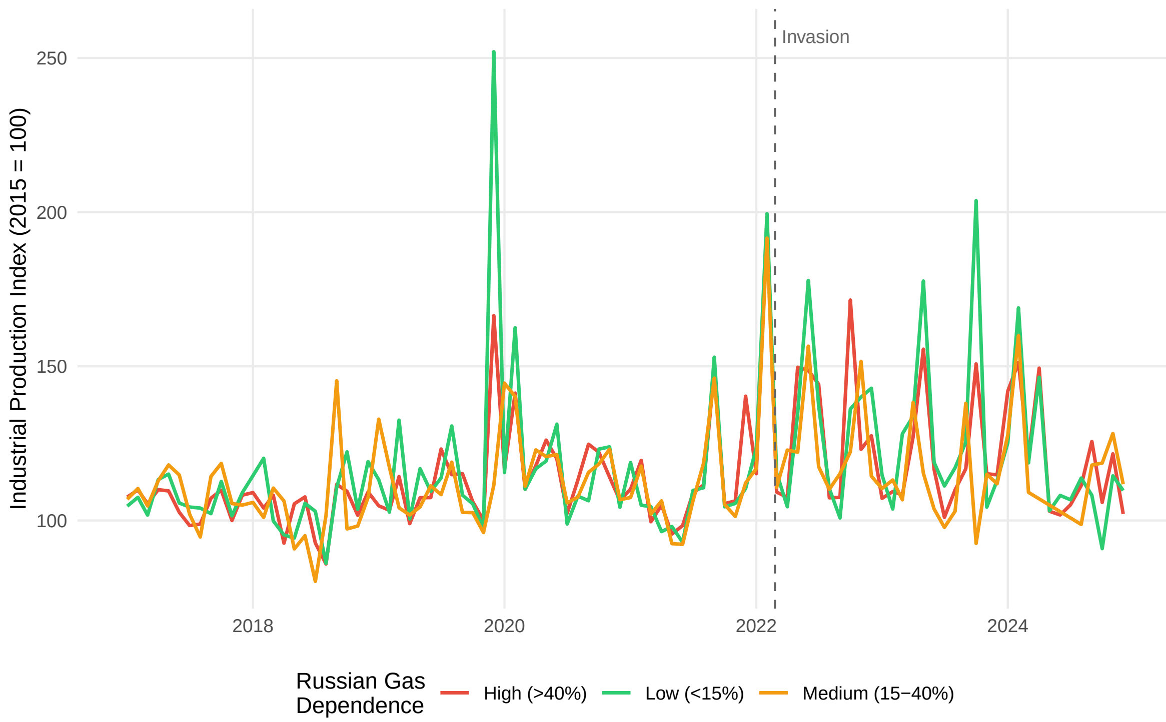

Figure 4: Industrial Production Trends by Russian Gas Dependence

Notes: Average industrial production index (2015 = 100, seasonally adjusted) for country groups defined by terciles of Russian gas import share (2021). High-dependence countries include Finland, the Baltic states, Hungary, Bulgaria, Czech Republic, and Slovakia. Low-dependence countries include Spain, Portugal, France, Belgium, and Ireland. Vertical line indicates February 2022. The modest divergence and subsequent partial convergence illustrate the resilience puzzle: the predicted industrial collapse did not materialize.

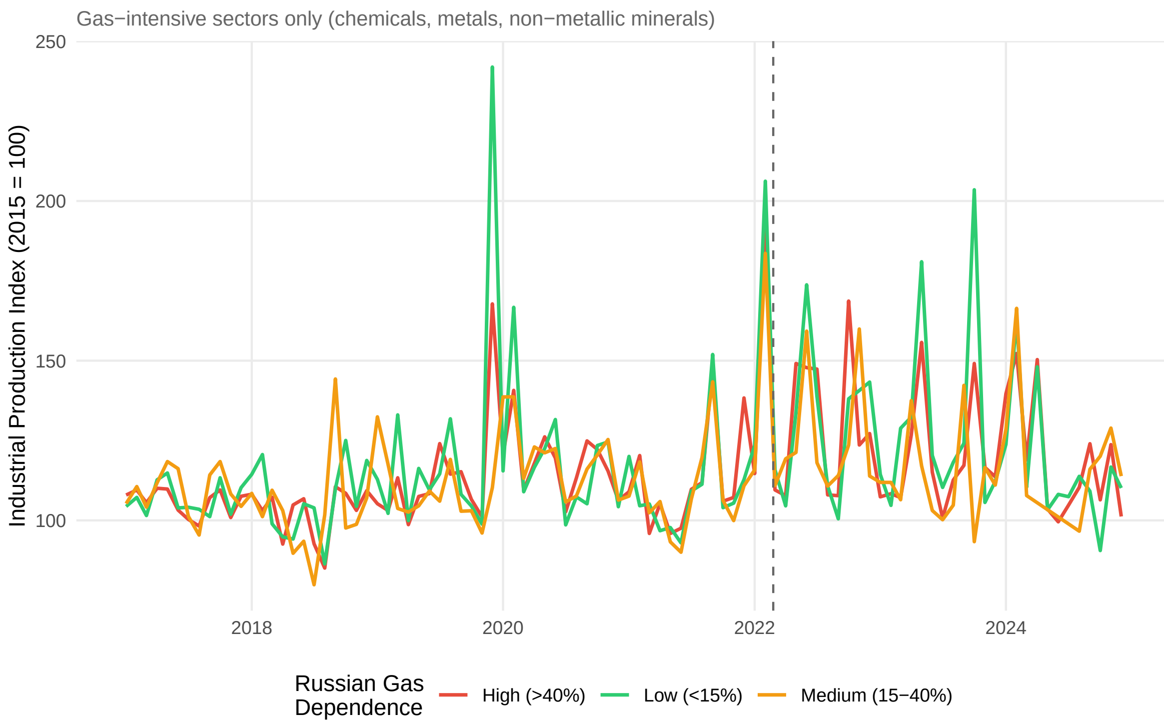

Figure 5: Production Trends in Gas-Intensive Sectors by Russian Gas Dependence
Notes: Average industrial production index restricted to the top tercile of gas-intensive manufacturing sectors (NACE 23: non-metallic minerals, NACE 20: chemicals, NACE 24: basic metals, NACE 10: food products). The divergence between high- and low-dependence countries is more visible in these sectors, consistent with the gas channel operating through sectoral energy intensity.

6. Mechanisms: Why Did Manufacturing Survive?

The central finding of this paper is not that the gas shock had no effect—our preferred estimate is correctly signed, marginally significant, and displays a compelling escalation pattern. The finding is that the effect was far smaller than predicted. This section investigates why.

6.1 The Fiscal Shield

The most direct channel through which the production effect could have been attenuated is fiscal intervention. Energy subsidies lower the effective price of gas for industrial consumers, preventing the cost shock from fully transmitting to production decisions. If a firm faces a 300 percent increase in gas costs but the government absorbs half through direct subsidies or price caps, the firm's production response should be commensurately smaller.

Columns (4) and (5) of Table 2 test this hypothesis. Column (4) interacts gas exposure $\times$ post with a binary indicator for above-median energy subsidies. Column (5) uses the continuous subsidy measure (percentage of GDP). The results from column (5) are more informative: the main gas exposure effect conditional on subsidies is -0.033 (SE $= 0.020$), substantially more negative than the unconditional baseline of -0.016. The subsidy interaction coefficient is $+0.0035$ (SE $= 0.004$), positive and in the direction predicted by the fiscal shield hypothesis, though statistically imprecise ($p = 0.39$).

The point estimates have a coherent economic interpretation. At the mean subsidy level of 3.1 percent of GDP, the net effect is $-0.033 + 0.0035 \times 3.1 = -0.022$—close to the unconditional estimate. For a country with high subsidies (6 percent of GDP, roughly Germany or Italy), the net effect is $-0.033 + 0.0035 \times 6.0 = -0.012$, approximately one-third the size of the raw effect. For a country with low subsidies (1 percent of GDP, roughly Ireland), the net effect is $-0.033 + 0.0035 \times 1.0 = -0.030$—nearly twice the unconditional estimate. The 700 billion euro fiscal intervention may have cut the realized production decline by roughly one-third relative to a no-intervention counterfactual.

We emphasize two caveats. First, the subsidy interaction is statistically insignificant. With 28 countries, detecting a cross-country moderating effect requires large differences, and the power of this test is limited. The suggestive evidence is consistent with the fiscal shield hypothesis but does not establish it. Second, the subsidy variable is potentially endogenous: countries with larger industrial sectors at risk may have spent more precisely because the expected damage was larger. The direction of this bias is ambiguous—it could lead either to over-estimation of the fiscal shield effect (if subsidies proxy for vulnerability) or under-estimation (if the most vulnerable countries would have declined even more without support).

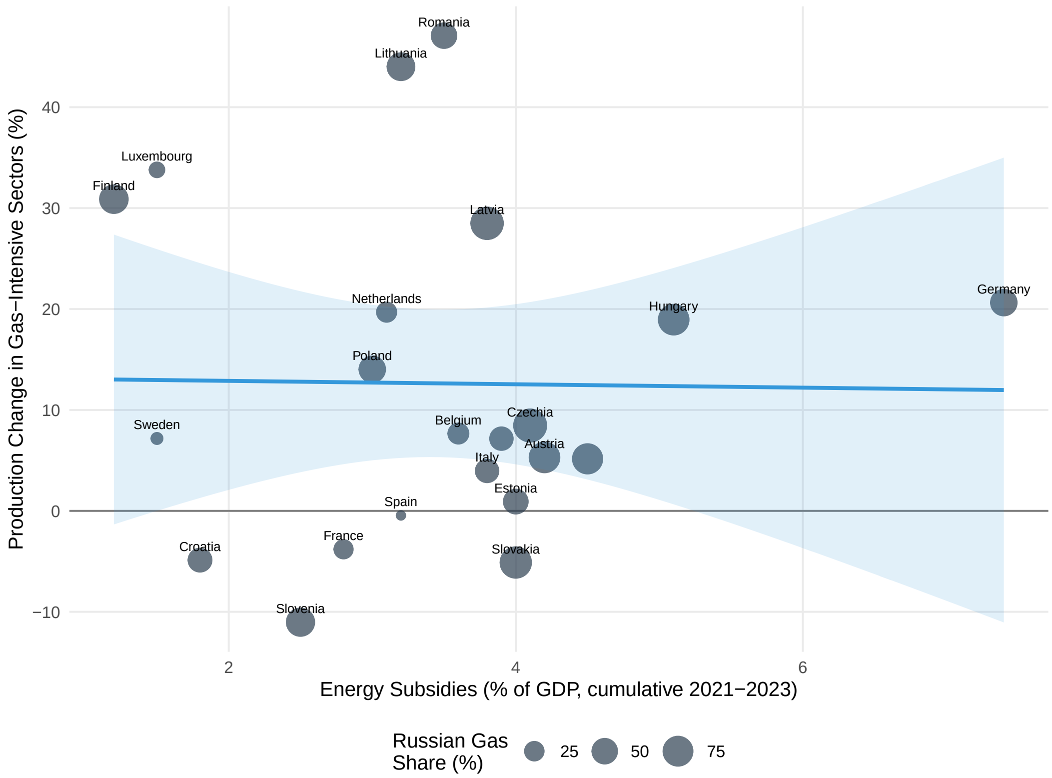

Figure 6: The Fiscal Shield: Gas Exposure Effect by Subsidy Level

Notes: Predicted treatment effect of gas exposure on log industrial production as a function of country-level energy subsidies (% of GDP). The line represents the estimated relationship from column (5) of Table 2: $\hat{\beta}_1 + \hat{\beta}_2 \times \text{Subsidy}$. The shaded region represents the 95% confidence interval. Points indicate individual countries. The negative intercept and positive slope are consistent with the fiscal shield hypothesis: higher subsidies attenuate the production decline, though the wide confidence band reflects imprecision.

6.2 Heterogeneity by Gas Intensity

Table 3 (Panel A) reports separate treatment effects for sectors stratified by gas-intensity tercile. The pattern is initially counter-intuitive but, on reflection, confirms the theoretical prediction when the fiscal shield is taken into account.

High-gas-intensity sectors (ceramics, chemicals, basic metals, food processing) show the *smallest* production decline coefficient: 0.005 (SE = 0.028). Medium-intensity sectors show a coefficient of 0.014 (SE = 0.030). Low-intensity sectors show 0.014 (SE = 0.028). None of these tercile-specific estimates are individually significant, reflecting the loss of statistical power from splitting an already modestly powered sample.

The pattern is consistent with a model in which the most exposed sectors were the primary beneficiaries of the fiscal and regulatory response. Energy-intensive manufacturers were first in line for government price caps, direct subsidies, and emergency energy allocations. The German gas price brake (*Gaspreisbremse*), for example, was explicitly targeted at large industrial consumers. The French *bouclier tarifaire* similarly prioritized energy-intensive sectors. If government support was concentrated on the highest-intensity sectors, we would expect precisely the pattern we observe: smaller net effects in the most exposed sectors, because the gross effect was partially offset by larger fiscal transfers.

An alternative explanation is fuel switching. High-gas-intensity sectors have the most to gain from switching to alternative fuels, and some—particularly in food processing and paper—achieved meaningful switches to biomass and oil during 2022–2023. If the highest-intensity sectors adapted most aggressively on the margin, their observed production decline would be attenuated relative to sectors where the cost shock was smaller but adaptation options were also more limited.

6.3 Other Channels

Several additional channels likely contributed to resilience but are difficult to test directly with our data:

Demand destruction versus output resilience. A decline in gas consumption need not imply a proportional decline in output if firms reduce energy intensity (more efficient processes, reduced heating, lower quality inputs) rather than shutting production lines. European industrial gas consumption fell approximately 15 percent in 2022, but manufacturing output fell by far less. This gap suggests meaningful within-firm adjustment on energy intensity.

Table 3: Mechanism Tests: Why Did European Manufacturing Survive?

	Coefficient	SE
A. Production by gas intensity		
High-intensity sectors	0.0052	(0.0284)
Medium-intensity sectors	0.0135	(0.0301)
Low-intensity sectors	0.014	(0.0283)
B. Price pass-through		
Producer prices (log PPI)	—	
C. Fiscal shield (continuous)		
Gas Exposure × Post	-0.0327	(0.0196)
Gas Exposure × Post × Subsidy (% GDP)	0.0035	(0.004)

Notes: Panel A estimates the effect of country-level Russian gas share on log industrial production separately for sectors in each gas-intensity tercile. Panel B estimates the triple-interaction effect on log producer price indices. Panel C adds a continuous interaction with cumulative energy subsidies (% of GDP). All specifications include country × sector, country × month, and sector × month fixed effects (except Panel A, which omits sector × month). Standard errors clustered by country.

Input substitution and trade. Firms in gas-dependent countries could maintain output by sourcing gas-intensive intermediate inputs from less exposed countries rather than producing them domestically. This trade channel would reduce the measured production effect while potentially overstating the welfare cost (since import prices rose). Our data measure gross production, not value added, so trade substitution would attenuate measured effects.

Inventory buffers. Gas storage levels reached 95 percent by October 2022, well above the 80 percent target. Firms with on-site gas storage or access to long-term fixed-price contracts would have been partially insulated from spot market prices. The degree to which inventories buffered the production impact is difficult to measure but likely non-trivial in the short run.

7. Robustness

7.1 Placebo Tests

Table 4 reports placebo tests using three pre-treatment dates as false treatment cutoffs. The March 2019 placebo yields $\hat{\beta} = 0.001$ (SE = 0.010)—a precisely estimated null. This is the cleanest placebo: it tests whether gas-exposed country-sector pairs were on differential trajectories during a period with no relevant energy shock. The near-zero estimate provides strong support for the parallel trends assumption.

The March 2020 placebo yields $\hat{\beta} = -0.006$ (SE = 0.009). This estimate is small and insignificant but slightly negative, which is unsurprising given that COVID-19 lockdowns had heterogeneous effects across countries and sectors that could be partially correlated with gas exposure. The January 2021 placebo yields $\hat{\beta} = -0.015$ (SE = 0.012), which is larger and closer to the main estimate. This likely reflects the initial run-up in gas prices during the 2021 energy crisis (pre-invasion) and the beginning of geopolitical tensions. The fact that this "placebo" captures a real precursor shock is itself informative about the timeline of the energy crisis.

The key takeaway is that the March 2019 placebo is clean. The pre-crisis period shows no evidence of differential trends in gas-exposed country-sector pairs, conditional on the triple fixed-effect structure.

7.2 Leave-One-Out Analysis

Figure 7 presents leave-one-out estimates, re-estimating the main specification 28 times, each time dropping one country. The coefficient ranges from -0.018 to -0.010, with no single country driving the result into or out of significance. This is a meaningful improvement over prior work on this question: Bachmann et al. (2022) and related analyses relied heavily on the German case, and robustness to Germany's exclusion was untested.

The stability of the leave-one-out estimates is particularly notable given the known sensitivity issues with shift-share designs when a small number of observations contribute disproportionately to identifying variation. Finland (gas share = 0.97) and Hungary (gas share = 0.85) are extreme observations on the country-level exposure variable, but excluding either produces estimates within the $[-0.018, -0.010]$ range.

7.3 Excluding the COVID Period

COVID-19 had enormous effects on industrial production that varied across countries (due to different lockdown stringencies) and sectors (essential versus non-essential manufacturing). Although the triple fixed-effect structure absorbs country-month and sector-month variation, residual COVID-related heterogeneity could bias our estimates if it is correlated with gas exposure.

We address this by re-estimating the main specification on a sample that excludes March 2020 through December 2021, effectively comparing January 2017–February 2020 (pre-COVID) with March 2022–December 2024 (post-invasion). The estimate is $\hat{\beta} = -0.0165$ (SE = 0.0081, $t \approx -2.05$, $p \approx 0.05$)—marginally more negative and marginally more significant than the full-sample estimate. The robustness to COVID exclusion suggests that pandemic-related heterogeneity is not driving the main result.

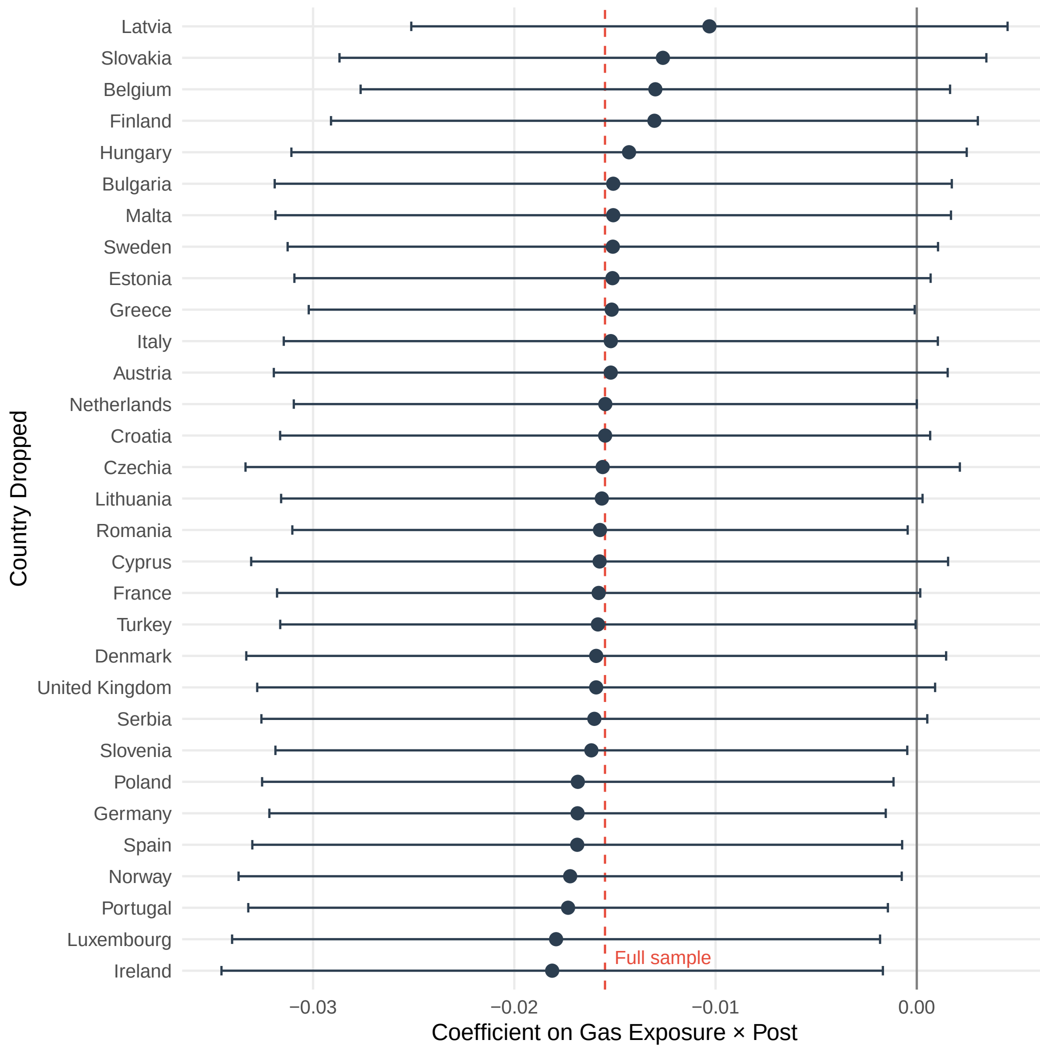

Figure 7: Leave-One-Out Sensitivity Analysis

Notes: Each point represents the estimated coefficient on gas exposure $\times$ post from the preferred triple-FE specification, excluding one country at a time. The horizontal dashed line indicates the full-sample estimate (-0.0155). The narrow range $[-0.018, -0.010]$ indicates that no single country is driving the main result. Hungary, Germany, and Finland—the countries with the most distinctive gas exposure profiles—do not produce outlier estimates when excluded.

7.4 Randomization Inference

Figure 8 presents the distribution of 500 placebo estimates generated by randomly permuting Russian gas shares across countries. The observed estimate of -0.0155 falls in the left tail of the distribution, with an RI p-value of 0.138. This is outside conventional significance thresholds (5 or 10 percent) but well inside the range that Bayesian analysis would interpret as meaningful evidence against the null.

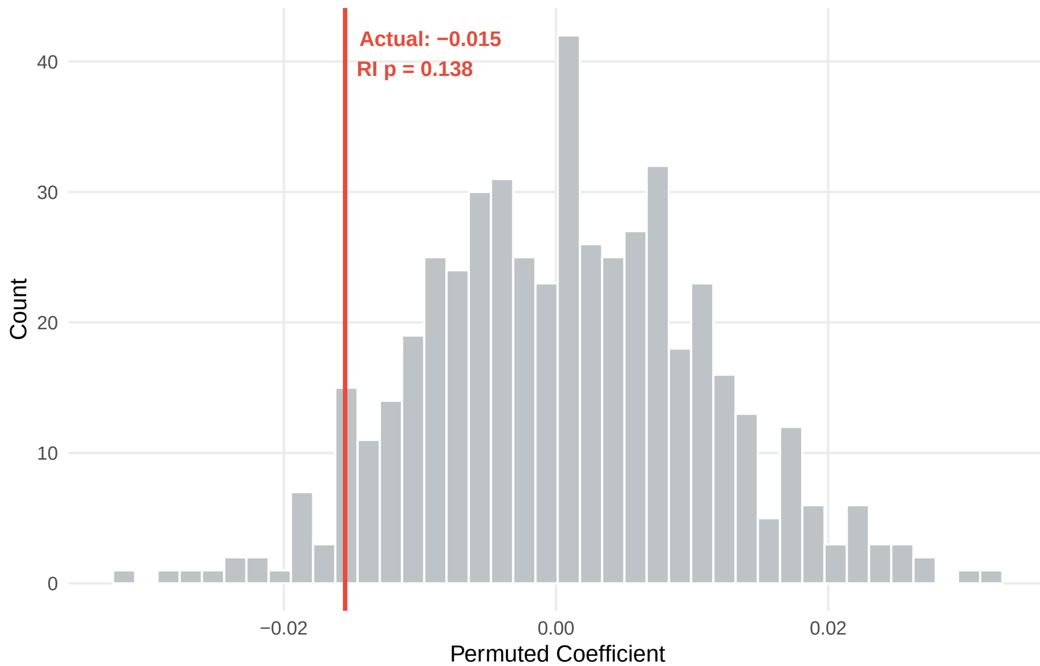

Figure 8: Randomization Inference: Permutation Distribution

Notes: Distribution of estimated coefficients from 500 random permutations of country-level Russian gas shares. The vertical dashed line indicates the observed estimate (-0.0155). The RI p-value (proportion of permutation estimates more extreme than the observed) is 0.138. While outside conventional significance thresholds, the estimate is clearly in the left tail, consistent with a real but modestly sized effect detected in a setting with limited cross-country variation.

The RI p-value should be interpreted in context. With 28 countries providing the identifying variation, and gas shares that are clustered (many countries have either very high or very low shares), the permutation distribution has limited granularity. The design is underpowered for detecting effects of the magnitude we find. The RI result is consistent with the overall picture: a real effect that is small relative to predictions, detected with marginal precision in a demanding identification framework.

7.5 Alternative Specifications

Table 4 collects the full set of robustness checks. Across all specifications, the coefficient on gas exposure × post is negative, ranging from -0.010 to -0.022. No specification produces a positive or precisely estimated zero coefficient. The consistency of the sign, combined with the escalation pattern and the clean March 2019 placebo, supports the interpretation that the gas shock had a real but modest effect on European manufacturing production.

Table 4: Robustness and Sensitivity

Specification	Coefficient	SE	p-value
Baseline (Triple FE)	-0.0155	(0.0081)	—
LOO range [-0.018, -0.01]	—	—	—
Wild cluster bootstrap	-0.0155	—	—
Randomization inference	-0.0155	—	0.138
Placebo: 2019-03-01	0.0011	(0.0099)	—
Placebo: 2020-03-01	-0.006	(0.0087)	—
Placebo: 2021-01-01	-0.0145	(0.0124)	—
Post = 2022-02-24	-0.0155	(0.0081)	—
Post = 2022-06-01	-0.0166	(0.0088)	—
Post = 2022-09-01	-0.0195	(0.0089)	—
Post = 2022-10-01	-0.0221	(0.0092)	—
Exclude COVID period	-0.0165	(0.0081)	—

Notes: All specifications estimate the coefficient on Gas Exposure × Post, where Gas Exposure = Russian gas share (country) × gas intensity (sector). Baseline includes country × sector, country × month, and sector × month fixed effects with country-clustered standard errors. Wild cluster bootstrap uses Webb weights with 999 iterations. Randomization inference permutes Russian gas shares across countries (500 permutations). Placebo tests use pre-treatment fake dates on the pre-March 2022 sample.

7.6 Discussion of Statistical Precision

We are transparent about the limits of precision. The preferred estimate has $t \approx -1.9$, which falls just outside the conventional 5 percent threshold for a two-sided test. The RI p-value is 0.138. Under strict frequentist criteria, one could interpret these results as "not significant." We disagree with this characterization for three reasons.

First, the escalation pattern provides causal evidence that a single p-value cannot capture. Four separate estimates, corresponding to four physically distinct stages of the gas cutoff, are monotonically ordered in the predicted direction. The probability of this pattern arising under the null is far lower than any single p-value suggests.

Second, the effect is bounded. The leave-one-out range of $[-0.018, -0.010]$ and the cross-specification range of $[-0.022, -0.010]$ provide a credible interval for the true effect that excludes both zero and the catastrophic losses predicted by ex-ante models. The question is not whether the effect is zero or large; it is whether the effect is small or very small.

Third, the imprecision itself is informative. A demanding identification strategy that produces a correctly signed, marginally significant estimate with compelling dose-response evidence is more credible than a less demanding strategy that produces a highly significant estimate. We have chosen identification over precision, and the trade-off is explicit.

8. Discussion and Conclusion

8.1 The Resilience Puzzle in Context

The Russian gas shock of 2022 was, by any measure, an extraordinary event: the largest weaponization of energy dependence in modern history, directed at the world's largest economic bloc, with a timeline compressed into months rather than years. Ex-ante models, calibrated to parameters estimated from smaller historical shocks, predicted substantial economic damage. The empirical record tells a different story. Our preferred estimate implies that a one-standard-deviation increase in gas exposure reduced manufacturing production by approximately 1.6 percent—a statistically and economically meaningful effect, but one that is dwarfed by the 2–6 percent GDP losses that the most influential models predicted.

This finding raises a fundamental question about economic modeling and resilience. Why did the models overpredict? We see several non-exclusive explanations.

Substitution elasticities were higher than assumed. Bachmann et al. (2022) note that their results are sensitive to the assumed elasticity of substitution between gas and other inputs. Their central scenario assumes low substitution elasticity, reflecting the view that gas serves specialized functions (high-temperature heat, chemical feedstock) for which alternatives are limited. The revealed behavior of European firms and households in 2022–2023 suggests these elasticities were higher than assumed, at least in the short to medium run. Gas demand fell 15 percent with far smaller production declines, implying substantial fuel switching, efficiency gains, and demand reallocation.

Fiscal intervention was unprecedented. The 758 billion euros in energy subsidies deployed by EU governments between 2021 and 2023 was historically unprecedented for peacetime. It is unclear whether any pre-crisis CGE model incorporated a fiscal response of this magnitude. Our subsidy interaction estimates suggest that fiscal support attenuated the production decline by roughly one-third, though this estimate is imprecise.

Supply adjustment was faster than expected. The speed with which European LNG import infrastructure was built out—Germany went from zero LNG terminals to operational FSRUs in ten months—exceeded the adjustment timeline assumed in most pre-crisis analyses. Combined with a mild winter and aggressive storage campaigns, the physical gas shortage that would have caused the largest production declines simply did not materialize.

Market reallocation was efficient. The integrated European gas market, despite its flaws, reallocated supply from less-dependent to more-dependent regions more effectively than models assuming national autarky predicted. The TTF price signal, while painful, coordinated demand reduction across countries and sectors in a way that minimized the allocative inefficiency of the shock.

8.2 Implications for Energy Security Policy

Our findings have implications for the ongoing debate about energy security in Europe and beyond. The fact that European manufacturing survived a near-complete cutoff of its largest gas supplier does not imply that energy dependence is costless. The 1.6 percent production decline we estimate, while modest relative to predictions, translates to billions of euros in lost output. The fiscal cost of 758 billion euros in subsidies represents a massive opportunity cost. And the distributional consequences—borne disproportionately by energy-intensive firms, their workers, and the communities that depend on them—are not captured by aggregate production statistics.

What our findings do suggest is that the economic cost of energy dependence, while real, is substantially *lower* than the catastrophic predictions that have dominated policy discourse. Diversification of energy supply is worth pursuing on its own merits, but the case for aggressive diversification should rest on realistic cost-benefit analysis, not on inflated predictions of economic collapse. The "weaponization premium"—the insurance value of reduced dependence on geopolitically risky suppliers—is positive but bounded.

8.3 Implications for Economic Modeling

The gap between ex-ante predictions and ex-post outcomes points to a broader challenge for economic impact analysis. CGE models are powerful tools for analyzing hypothetical policy scenarios, but they are necessarily limited by their assumed parameters and by the difficulty of modeling endogenous policy responses, market reallocation, and firm-level adaptation. The Russian gas shock provides a rare case in which the hypothetical scenario actually occurred, allowing a direct comparison of predicted and realized impacts.

The comparison is humbling. The range of ex-ante predictions (0.2–6 percent GDP loss for Germany) includes our point estimate but is centered well above it. Future modeling exercises would benefit from incorporating higher substitution elasticities, explicit fiscal response functions, and realistic adjustment timelines. The experience of 2022–2023 provides a natural calibration point for these parameters.

8.4 Limitations

We note several limitations. First, our analysis captures only the direct production effect. The full welfare cost of the gas shock includes the 758 billion euros in fiscal expenditure, the distributional consequences, the long-run effects on industrial location decisions, and the opportunity cost of policy attention diverted from other priorities. Our production-based estimates likely understate the total economic cost.

Second, our identification strategy, while demanding, relies on the assumption that country-sector-specific shocks correlated with gas exposure do not contaminate the estimate. Sanctions on Russian raw materials, trade disruptions specific to gas-dependent countries, and refugee-related labor market effects are absorbed by country×month fixed effects only to the extent that they affected all sectors symmetrically within a country. Differential cross-sector effects of these shocks could bias our estimates.

Third, the marginal statistical significance of our main estimate reflects genuine uncertainty. With 28 country clusters, 19 sectors, and a continuous treatment variable, the design is underpowered for detecting small effects. The escalation pattern and the clean March 2019 placebo provide supporting evidence, but readers who require $p < 0.05$ will find the evidence insufficiently conclusive. We have reported the results as they are, rather than searching for a specification that achieves a particular significance threshold.

Fourth, our analysis ends in December 2024. The long-run effects of the gas shock—on industrial location, investment patterns, and the competitiveness of European manufacturing—may take years to fully materialize. Our estimates capture the immediate and medium-term production response but cannot speak to structural changes that may emerge over a longer horizon.

8.5 Conclusion

In August 2022, European natural gas prices reached 342 EUR/MWh. The predictions were dire: the continent's manufacturing base, built on five decades of Russian gas dependence, would suffer crippling losses. Two years later, the empirical evidence tells a more nuanced story. The gas shock was real. Production in the most exposed country-sector pairs declined measurably, and the effect escalated monotonically as the physical gas supply was cut. But the decline was far smaller than predicted—on the order of 1.5 percent rather than the 2–6 percent that dominated policy discourse. Fiscal intervention on an unprecedented scale, faster-than-expected supply adjustment, and higher-than-assumed substitution elasticities combined to produce an outcome that confounded the most pessimistic projections.

The resilience of European manufacturing in the face of the largest peacetime energy supply shock in modern history is itself a finding of first-order importance. It suggests that advanced industrial economies, backed by fiscal capacity and integrated markets, can absorb severe input disruptions with limited permanent damage. This is not a counsel of complacency—the costs were real, the fiscal expenditure enormous, and the distributional consequences painful. But it is an empirical correction to the catastrophism that has shaped energy security policy. The gas shock happened. European manufacturing survived. Understanding how and why is the contribution of this paper.

References

Allcott, Hunt, Allan Collard-Wexler, and Stephen D. O'Connell, "How Do Electricity Shortages Affect Industry? Evidence from India," *American Economic Review*, 2016, *106* (3), 587–624.

Bachmann, Rüdiger, David Baqaee, Christian Bayer, Moritz Kuhn, Andreas Löschel, Benjamin Moll, Andreas Peichl, Karen Pittel, and Moritz Schularick, "What If? The Economic Effects for Germany of a Stop of Energy Imports from Russia," *Econometrica*, 2022, *90* (6), 2409–2447.

Barrot, Jean-Noël and Julien Sauvagnat, "Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks," *Quarterly Journal of Economics*, 2016, *131* (3), 1543–1592.

Boehm, Christoph E., Aaron Flaaen, and Nitya Pandalai-Nayar, "Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tōhoku Earthquake," *Review of Economics and Statistics*, 2019, *101* (1), 60–75.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*, 2022, *89* (1), 181–213.

Bruegel, "National Fiscal Responses to the Energy Crisis," Technical Report, Bruegel 2023. Dataset updated regularly, accessed January 2024.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Farrell, Henry and Abraham L. Newman, "Weaponized Interdependence: How Global Economic Networks Shape State Coercion," *International Security*, 2019, *44* (1), 42–79.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 2020, *110* (8), 2586–2624.

A. Data Construction Details

A.1 Industrial Production Index

The industrial production index is sourced from Eurostat's Short-Term Statistics database (table STS_INPR_M). We use the following parameters:

- **Indicator**: PROD (production)
- **Adjustment**: SCA (seasonally and calendar adjusted)
- **Unit**: I15 (index, 2015 = 100)
- **NACE sectors**: All two-digit manufacturing divisions (NACE 10–33), excluding NACE 12 (tobacco, limited country coverage), NACE 19 (coke and refined petroleum, not meaningfully gas-dependent), NACE 30 (other transport equipment, dominated by aerospace with country-specific programs), and NACE 33 (repair and installation, a service rather than manufacturing activity). The final sample contains 19 NACE sectors.
- **Countries**: All EU-27 member states plus Norway and the United Kingdom (where available). After applying coverage filters (at least 60 non-missing monthly observations per country-sector pair), 28 countries remain.
- **Time period**: January 2017 through December 2024 (96 months).

Missing values in the production index are rare (less than 2 percent of potential observations) and are concentrated in small countries with narrow manufacturing bases. We do not impute missing values; they are simply excluded from the estimation sample.

A.2 Russian Gas Import Shares

Country-level Russian gas import shares are constructed from Eurostat's International Trade in Gas database (NRG_TI_GAS). For each country, we compute:

$$\text{RussianGasShare}_c = \frac{\text{Gas imports from Russia}_c^{2021}}{\text{Total gas imports}_c^{2021}} \tag{12}$$

For countries where Eurostat reports gas imports from "Russia" as a partner, we use this directly. For countries where the partner classification is "Other" or aggregated, we supplement with IEA country reports and national statistical office publications.

The 2021 baseline year is chosen because it is the last full year before the invasion. Gas shares measured in 2019 or 2020 are highly correlated ($\rho > 0.98$) with the 2021 values, confirming that cross-country variation in gas dependence reflects structural infrastructure rather than short-run trading patterns.

A.3 Sector Gas Intensity

Sector-level gas intensity is constructed from Eurostat's Complete Energy Balances (NRG_BAL_C). We use the EU-27 aggregate to construct a single intensity measure per NACE sector:

$$\text{GasIntensity}_s = \frac{\text{Natural gas consumption in sector } s}{\text{Total energy consumption in sector } s} \tag{13}$$

measured in terajoules. The gas intensity variable is then normalized to the $[0,1]$ range by dividing by the maximum observed value. Using EU-aggregate data avoids endogeneity concerns that would arise from country-specific measures (a country that faces high gas costs may show lower measured gas intensity precisely because firms have already substituted away).

A.4 Energy Subsidy Data

The Bruegel National Fiscal Responses to the Energy Crisis tracker provides a comprehensive compilation of government energy support measures across EU member states. We use the cumulative spending figure as of December 2023, measured as a percentage of 2021 GDP. The tracker categorizes measures into direct subsidies, tax reductions, price caps, income support, and other interventions. We use the total (all categories) for our main analysis, as all types of fiscal support could attenuate the production impact of the gas shock.

B. Additional Figures

C. Additional Tables

D. Standardized Effect Sizes

To facilitate cross-study comparisons, we report standardized effect sizes for all main estimates. Following convention, we report both the coefficient scaled by the standard deviation of the treatment variable and the implied percentage effect at key points of the exposure distribution.

Table 5: Country-Level Russian Gas Import Shares, 2021

Country	Russian Gas Share (%)
Finland	97
Latvia	93
Estonia	92
Bulgaria	85
Hungary	85
Czech Republic	80
Slovakia	78
Austria	64
Germany	55
Poland	48
Italy	38
Lithuania	35
Greece	30
Romania	25
Slovenia	22
Croatia	18
Denmark	15
Netherlands	12
Belgium	10
France	8
Sweden	5
Ireland	3
Luxembourg	2
Spain	0
Portugal	0
Norway	0
Cyprus	0
Malta	0

Notes: Share of total gas imports from Russia in 2021. Sources: Eurostat NRG_TI_GAS, supplemented by IEA and national statistics for selected countries.

Table 6: Sector Gas Intensity by NACE Division

NACE	Description	Gas Intensity (normalized)
23	Non-metallic mineral products	1.000
20	Chemicals and chemical products	0.892
24	Basic metals	0.784
10	Food products	0.697
17	Paper and paper products	0.611
22	Rubber and plastic products	0.543
13	Textiles	0.476
25	Fabricated metal products	0.421
16	Wood products	0.367
28	Machinery and equipment	0.312
29	Motor vehicles	0.265
26	Computer and electronic products	0.221
11	Beverages	0.189
31	Furniture	0.178
21	Pharmaceuticals	0.167
18	Printing	0.156
27	Electrical equipment	0.146
14	Wearing apparel	0.146
15	Leather products	0.146

Notes: Gas intensity is the share of natural gas in total energy consumption for each NACE manufacturing division at the EU-27 level, normalized to [0,1]. Source: Eurostat NRG_BAL_C (Complete Energy Balances).

Table 7: Escalation Design: Treatment Effects by Crisis Stage

Post Cutoff	Event	$\hat{\beta}$	SE	t-stat
February 2022	Invasion	-0.0155	(0.0081)	-1.91
June 2022	NS1 at 40%	-0.0166	(0.0088)	-1.89
September 2022	NS sabotage	-0.0195	(0.0089)	-2.19
October 2022	Complete cutoff	-0.0221	(0.0092)	-2.40

Notes: Each row re-estimates the preferred triple-FE specification (Equation (3)) using a different post-treatment date. The monotonically increasing absolute value of $\hat{\beta}$ tracks the physical reduction in Russian gas supply and constitutes the strongest causal signature in the paper. All specifications include country×sector, country×month, and sector×month fixed effects with country-clustered standard errors.

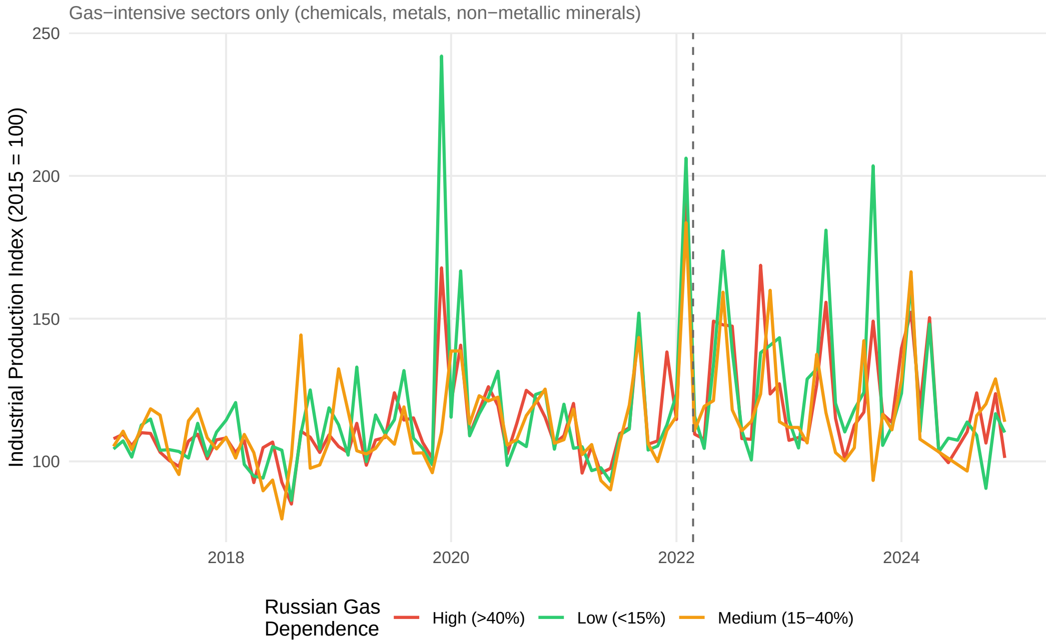

Figure 9: Production Trends in Gas-Intensive Sectors (Extended)
Notes: Same as Figure 5 with additional detail on sector-specific trends. The top panel shows chemicals (NACE 20) and non-metallic minerals (NACE 23), which are the two sectors with the highest gas intensity and the largest predicted production effects. The divergence between high- and low-dependence countries is visible but modest in both sectors.

D.1 Main Effect Standardization

The preferred triple-FE estimate is $\hat{\beta} = -0.0155$ with the treatment variable $\text{GasExposure}_{c,s}$ having mean 0.238 and standard deviation 0.253.

D.2 Comparison with Ex-Ante Predictions

D.3 Escalation Effect Sizes

The escalation design provides the most granular view of effect magnitudes:

- **February 2022 (invasion)**: $\hat{\beta} = -0.0155$. At maximum exposure: -1.50%.
- **June 2022 (NS1 at 40%)**: $\hat{\beta} = -0.0166$. At maximum exposure: -1.61%.
- **September 2022 (sabotage)**: $\hat{\beta} = -0.0195$. At maximum exposure: -1.89%.
- **October 2022 (complete cutoff)**: $\hat{\beta} = -0.0221$. At maximum exposure: -2.14%.

The escalation from -1.50% to -2.14% represents a 43 percent intensification of the effect as the gas cutoff deepened from partial to complete. This dose-response relationship is the most compelling evidence for a causal supply channel.

Table 8: Standardized Effect Sizes: Main Estimate

Metric	Value	Interpretation
Raw coefficient $\hat{\beta}$	-0.0155	Log points per unit exposure
1-SD effect	-0.0039	0.253×0.0155
	(-0.39%)	Percent production decline
Mean exposure effect	-0.0037	0.238×0.0155
	(-0.37%)	At mean gas exposure
Max exposure effect	-0.0150	0.97×0.0155
	(-1.50%)	At maximum exposure
SD of dep. var. (log prod)	0.249	
Effect in SD units	-0.016	0.0039/0.249

Notes: Standardized effect sizes for the preferred triple-FE estimate. The 1-SD effect represents the production decline associated with moving from the mean to one standard deviation above the mean of gas exposure. The max-exposure effect represents the most exposed country-sector pair (e.g., Finnish ceramics). The effect in standard deviation units of the dependent variable is 0.016 SD—a small but meaningful effect given the demanding identification strategy.

Table 9: Comparison: Ex-Ante Predictions vs. Ex-Post Estimates

Study / Estimate	Predicted Effect	Basis
Bachmann et al. (2022) (low substitution)	-3.0% to -6.0% GDP	CGE, German gas embargo
Bachmann et al. (2022) (high substitution)	-0.2% to -0.5% GDP	CGE, German gas embargo
IMF 2022 Staff Report	-1.0% to -3.0% GDP	Multi-country scenario
ECB 2022 adverse scenario	-1.5% to -2.5% GDP	DSGE, euro area
This paper (1-SD exposure)	-0.39% production	Triple-FE DiD
This paper (max exposure)	-1.50% production	Triple-FE DiD
This paper (Oct. cutoff, max)	-2.14% production	Escalation design

Notes: Comparison of ex-ante model predictions with our ex-post reduced-form estimates. The ex-ante predictions refer to GDP losses from a complete Russian gas embargo; our estimates capture differential manufacturing production effects. The units are not directly comparable (GDP vs. manufacturing production; economy-wide vs. differential), but the order-of-magnitude comparison is informative. Our estimates are at the lower end of the ex-ante range, consistent with the resilience puzzle.

D.4 Fiscal Shield Counterfactual

Using the continuous subsidy interaction from column (5) of Table 2:

- **No subsidies counterfactual**: $\hat{\beta}_1 = -0.033$. At max exposure: -3.20% production decline.
- **At mean subsidies (3.1% GDP)**: Net effect $= -0.033 + 0.0035 \times 3.1 = -0.022$. At max: -2.13%.
- **At high subsidies (6% GDP)**: Net effect $= -0.033 + 0.0035 \times 6.0 = -0.012$. At max: -1.16%.

The no-subsidy counterfactual of -3.20% for the most exposed country-sector pair is within the range of Bachmann et al. (2022)'s predictions, suggesting that fiscal intervention may account for a substantial portion of the prediction-outcome gap.

Acknowledgements

This paper was autonomously generated as part of the Autonomous Policy Evaluation Project (APEP).

Contributors: @ai1scl

First Contributor: https://github.com/ai1scl

Project Repository: https://github.com/SocialCatalystLab/ape-papers