Cutting the Pipeline: Russian Gas Dependence and the Differential De-Industrialization of European Manufacturing

APEP Autonomous Research* @SocialCatalystLab

March 8, 2026

Abstract

Russia's 2022 gas cutoff created the largest energy supply disruption in modern European history. We estimate the reduced-form relationship between pre-war gas exposure and manufacturing outcomes, exploiting variation in country-level Russian gas import shares and sector-level gas intensity across 23 European countries, 22 manufacturing sectors, and 108 months. Using a continuous-treatment difference-in-differences design with country×sector, country×month, and sector×month fixed effects, a one-standard-deviation increase in gas exposure is associated with a 2.3 percentage point production decline, though this estimate is statistically imprecise ($t = -0.54$, RI $p = 0.58$) and sensitive to individual countries. Point estimates become more negative in 2023, suggestive of persistence but not statistically established. The demanding identification strategy and limited cluster count constrain precision.

JEL Codes: F51, L60, Q43, Q48

Keywords: energy dependence, natural gas, de-industrialization, Russia-Ukraine war, European manufacturing, difference-in-differences

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

1. Introduction

On February 24, 2022, Russia invaded Ukraine. In the months that followed, Russian natural gas deliveries to Europe collapsed from 155 billion cubic meters in 2021 to near zero by 2023. European wholesale gas prices spiked elevenfold, from 30 to 342 EUR/MWh at peak. The largest energy supply disruption in modern European history was underway.

The economic toll was immediate. German chemicals production cratered 18 percent within months. Fertilizer plants across Central Europe shut down. Energy-intensive firms announced permanent relocations. Policymakers responded with hundreds of billions in emergency energy subsidies. Yet amid this upheaval, a basic empirical question remained unanswered: did the gas cutoff cause these manufacturing declines, or were they driven by the broader macroeconomic shock of war, sanctions, and uncertainty?

We provide reduced-form evidence on this question by exploiting a fundamental feature of the energy shock: it was doubly differentiated. Countries varied enormously in their pre-war dependence on Russian gas—from 85 percent (Slovakia) to zero (Norway, Cyprus, Malta). Manufacturing sectors varied in gas intensity—from 45 percent of energy inputs (non-metallic minerals) to 6 percent (paper and pulp). The interaction of these two predetermined characteristics creates a continuous treatment intensity that is plausibly exogenous: countries did not choose their pipeline infrastructure in anticipation of the 2022 invasion, and sectors did not choose their energy technology to hedge against geopolitical risk.

Our empirical design estimates:

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

where α_{cs}, γ_{ct}, and δ_{st} are country×sector, country×month, and sector×month fixed effects. The country×month fixed effects absorb all aggregate country-level shocks—sanctions, fiscal stimulus, confidence effects, exchange rate movements. The sector×month fixed effects absorb all global industry trends and supply chain disruptions. The coefficient β captures whether countries *more* exposed to Russian gas experienced *larger* production declines in *more* gas-intensive sectors, net of these saturated controls.

We find $\hat{\beta} = -0.23$ in the preferred specification, implying that a one-standard-deviation increase in treatment intensity (0.097 units) is associated with a 0.022 log-point decline in industrial production, or about 2.3 percent. This estimate is statistically imprecise under conventional inference ($t = -0.54$) and permutation inference (RI $p = 0.58$), reflecting the demanding triple fixed-effect structure, only 23 country-level clusters for inference, and a treatment variable with limited cross-sector variation (10 distinct gas intensity values mapped to 22 NACE sectors). The imprecision is informative: it bounds the range of plausible effects and demonstrates the genuine cost of a credible identification strategy in this setting.

The dynamic analysis provides suggestive evidence on the trajectory of effects. The monthly event study shows flat pre-trends in the months immediately preceding the invasion (July 2021 through January 2022), followed by negative post-invasion coefficients that become more negative over time. Year-specific estimates yield $\hat{\beta}_{2022} = -0.16$ (SE = 0.197) and $\hat{\beta}_{2023} = -0.30$ (SE = 0.263)—neither individually significant, and we do not formally test whether they differ from each other. The point estimates are consistent with persistence rather than temporary adjustment, but the imprecision prevents strong conclusions about the dynamic path.

Robustness checks reveal both stability and fragility. Leave-one-country-out estimates range from -0.41 to $+0.26$, with Hungary as the most influential observation—its exclusion flips the sign to positive, a candid indication that the negative point estimate is not robust to the removal of a single country. Excluding high-intra-EU-trade sectors (motor vehicles, machinery) to address SUTVA concerns yields $\hat{\beta} = -0.22$, nearly identical to the baseline. Placebo treatment dates (March 2019, March 2020) yield coefficients of -0.35 and -0.34—similar in magnitude to the main estimate. The March 2020 placebo is attributable to COVID-19 contamination, but the March 2019 placebo is harder to dismiss and raises the possibility that the gas-share×gas-intensity interaction captures broader vulnerability to macroeconomic shocks rather than the specific gas channel. The pre-COVID trend test (2015–2019) yields a precisely estimated zero ($t = -0.14$), suggesting no secular differential trend, but this does not rule out episodic pre-trend violations.

This paper contributes to three literatures. First, to the growing body of work on the economic costs of the Russia-Ukraine war. Bachmann et al. (2022) provided the landmark ex-ante simulation, predicting GDP losses of 0.2–3 percent for Germany from a complete gas embargo. Our ex-post estimates provide an empirical counterpart to their structural predictions, with point estimates suggesting that production impacts were concentrated in the most exposed sector-country combinations—though our imprecision prevents sharp comparisons of magnitudes. Second, we contribute to the literature on energy and industrial production. Allcott et al. (2016) study electricity shortages in India; Abeberese (2017) examines energy costs and manufacturing in India; Costa et al. (2021) studies environmental regulation. Our contribution is to apply a cross-country, cross-sector design to a large energy supply shock in advanced economies. Methodologically, our design is related to exposure-based identification strategies studied by Goldsmith-Pinkham et al. (2020) and Borusyak et al. (2022), where predetermined shares (here, Russian gas dependence) interact with common shocks to generate identifying variation. Third, we speak to the trade disruption literature.

Barrot and Sauvagnat (2016) study supply chain propagation from natural disasters; Boehm et al. (2019) examine input linkages across borders. Our setting differs in that the shock operates through an intermediate input (energy) rather than direct supplier relationships.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background. Section 3 presents the data. Section 4 develops the empirical strategy. Section 5 reports results. Section 6 discusses mechanisms. Section 7 concludes.

2. Institutional Background

2.1 European Gas Dependence on Russia

For decades, Europe built its energy infrastructure around Russian natural gas. The relationship began during the Cold War: the first contracts between the Soviet Union and Austria were signed in 1968, followed by Germany in 1970 and Italy in 1974. Pipelines traversing Ukraine, Belarus, and the Baltic Sea floor—the Brotherhood pipeline system (1967), the Trans-Austria Gas Pipeline (1974), and eventually Nord Stream 1 (2011)—delivered gas to Central and Western Europe at scale that grew decade by decade.

By 2021, Russia supplied roughly 40 percent of total EU gas imports, but this aggregate figure obscures enormous cross-country heterogeneity. The variation in dependence reflected geography and infrastructure lock-in, not recent policy choices. Countries along the major pipeline corridors had built their entire gas distribution networks around Russian supply over half a century:

- **Very high dependence (65–85%):** Slovakia (85%), Austria (80%), Hungary (80%), Bulgaria (75%), Finland (75%), Czechia (70%), Germany (66%), Latvia (65%). These countries received gas through dedicated pipeline infrastructure with limited alternative supply options.

- **Moderate dependence (30–50%):** Estonia (46%), Turkey (45%), Greece (40%), Italy (40%), Croatia (30%), Netherlands (30%). These countries had partially diversified through LNG terminals or alternative pipeline routes, but Russian gas remained a significant share.

- **Low dependence (0–25%):** Romania (15%), Sweden (10%), Slovenia (10%), Belgium (6%), Luxembourg (6%), Denmark (4%), Norway (0%), Cyprus (0%), Malta (0%). These countries relied on North Sea production, LNG from global markets, domestic production, or simply used little gas.

This variation is the foundation of our identification strategy. Pre-war Russian gas dependence was determined by decades of infrastructure investment, Cold War-era contracts, geographic proximity to pipeline routes, and geological endowments of alternative energy sources. It was not chosen in anticipation of the 2022 invasion. The pipeline infrastructure that created dependence required billions of euros of sunk capital and decades to build; reversing it required years even under crisis conditions. This lock-in character makes pre-war gas dependence a plausibly predetermined treatment variable.

2.2 Gas in European Manufacturing

Natural gas serves as both an energy source and a chemical feedstock in European manufacturing, but its importance varies enormously across sectors. In non-metallic minerals (NACE C23)—glass, ceramics, cement, bricks—gas furnaces are the dominant production technology, and gas accounts for 45 percent of total energy consumption. In basic metals (C24), gas fires blast furnaces and arc furnaces. In chemicals (C20–C22), gas serves a dual role: as fuel for process heat and as a feedstock for ammonia, methanol, and hydrogen production. These sectors cannot easily substitute away from gas in the short run because the production technology itself is gas-dependent.

At the other end, sectors like paper and printing (C17/C18) use gas for drying and process heat but at much lower intensities (6 percent of total energy). Textiles (C13–C15) and machinery (C25–C28) fall in between. The key insight for identification is that this sectoral variation in gas dependence is determined by production technology, not by firm choice—a ceramics plant cannot switch to coal without rebuilding its kiln, and an ammonia plant cannot produce without natural gas as a feedstock.

The interaction of country-level pipeline dependence and sector-level gas intensity creates a rich two-dimensional treatment structure. A German chemicals plant (high country dependence $\times$ high sector intensity) was maximally exposed to the gas cutoff. A Spanish textile plant (low country dependence $\times$ low sector intensity) was barely affected. This double variation is what allows us to estimate the reduced-form relationship between gas exposure and production with unusually demanding fixed effects.

2.3 The 2022 Gas Cutoff

The disruption unfolded in stages. Following the February 24 invasion, EU sanctions initially targeted Russian banks and individuals, not energy. The gas cutoff was instead driven by Russian actions:

- **April 2022:** Russia demanded payment in rubles; Poland and Bulgaria refused and were cut off.

- **June 2022:** Gazprom reduced Nord Stream 1 flows to 40% capacity, citing "turbine maintenance."
- **July 2022:** Nord Stream 1 reduced to 20%.
- **September 2, 2022:** Nord Stream 1 shut completely.
- **September 26, 2022:** Nord Stream 1 and 2 pipelines sabotaged by underwater explosions.

The TTF hub price, Europe's benchmark, rose from roughly 30 EUR/MWh in January 2022 to a peak of 342 EUR/MWh on August 26—an elevenfold increase. Though prices subsequently declined as emergency LNG imports arrived and demand-side savings materialized (Ruhnau et al., 2023), they remained 2–4 times pre-war levels through 2023.

2.4 Government Responses

European governments responded with massive fiscal interventions. Germany alone allocated over 200 billion for energy price caps and industry support. France capped electricity and gas prices for households and small businesses. Italy, Spain, and others implemented various subsidy and tax relief programs. These interventions were explicitly designed to cushion the blow to domestic industry—and therefore attenuate precisely the effect we seek to measure.

This policy response complicates interpretation: our estimates reflect the net-of-subsidy impact of the gas cutoff, not the counterfactual of no government intervention. Subsidies may have attenuated the production decline, but because they were endogenously targeted to the most exposed country-sector cells, the direction and magnitude of resulting bias are not point-identified without subsidy data.

3. Data

We construct a panel of monthly industrial production by country and manufacturing sector, merged with predetermined measures of Russian gas dependence and sectoral gas intensity.

3.1 Industrial Production

Our primary outcome is the Eurostat Short-Term Business Statistics industrial production index (`STS_INPR_M`). This provides monthly, seasonally and calendar adjusted production indices (base year 2015 = 100) for NACE Rev.2 two-digit manufacturing sectors across EU member states and partner countries. The seasonal adjustment removes within-year cyclical patterns, allowing clean identification of post-invasion level shifts.

We observe 22 manufacturing sectors (NACE C10 through C33) across 23 European countries from January 2015 through December 2023, yielding 47,330 country-sector-month observations after dropping missing cells. The panel is unbalanced in two senses. First, some countries suppress data for sectors where confidentiality constraints bind (Germany, for example, reports only 5 of 22 sectors due to statistical confidentiality rules protecting individual firms in concentrated industries). Second, two countries have truncated time coverage: Czechia's data ends in December 2020 and Türkiye's in October 2020. These countries contribute to estimating the fixed-effect structure but have no post-invasion observations and therefore do not contribute to the treatment effect estimate. Excluding them yields nearly identical results (see Appendix). We use log production indices as our dependent variable, so the coefficient $\hat{\beta}$ is interpretable as a proportional change.

The nine-year window (2015–2023) provides five years of pre-COVID baseline (2015–2019), two years of pandemic and recovery (2020–2021), and nearly two years of post-invasion treatment (March 2022–December 2023). This structure allows us to examine pre-trends, distinguish COVID-era disruptions from war-era effects, and study the dynamic evolution of the treatment effect.

3.2 Russian Gas Dependence

Country-level Russian gas import shares for 2021 are drawn from Bruegel's European Natural Gas Tracker (McWilliams et al., 2022), cross-referenced with Eurostat Energy Balances and IEA Gas Trade Flows. These shares reflect pipeline and LNG deliveries from Russia as a fraction of total gas consumption. Values range from 85% (Slovakia) to 0% (Norway, Cyprus, Malta). We use 2021 values—the last full year before the invasion—to ensure the treatment variable is predetermined.

An important data choice deserves explanation. Eurostat's official gas trade partner data (`NRG_TI_GAS`) is incomplete for several key countries. Germany, the largest European gas consumer, reports only pipeline-specific partner breakdowns that miss most Russian gas entering through intermediary hubs. Our cross-referenced Bruegel/IEA shares correct these gaps using validated country-level totals. The key identification requirement is that the *ranking* of countries by dependence is correct, not that each percentage is exact to the decimal—measurement error in the treatment variable biases our estimate toward zero, making our results conservative.

3.3 Sector Gas Intensity

Sector-level gas intensity is computed from Eurostat Complete Energy Balances (`NRG_BAL_C`) for 2019, using the EU-wide ratio of natural gas consumption to total energy consumption by industrial sub-sector. The NRG_BAL classification reports energy consumption at the level of 10 aggregated industrial groups, which we map to 22 NACE two-digit manufacturing sectors (see Appendix A.1). NACE sectors within the same NRG_BAL group receive the same gas intensity value—for example, C20 (chemicals), C21 (pharmaceuticals), and C22 (rubber/plastics) all receive the "chemicals, rubber, plastics" group intensity of 36.1%. We deliberately use the EU-wide average rather than country-specific values to avoid introducing endogenous variation: a country that had already substituted away from gas due to anticipated supply risks would appear less gas-intensive, creating reverse causality. The EU-wide measure captures the technological gas dependence of each sector, not the equilibrium response to country-specific energy prices.

We use 2019 (pre-COVID) values to avoid contamination from pandemic-era production shifts that differentially affected sectors. Values range from 45% (non-metallic minerals, C23) to 6% (paper, C17/C18). The mapping from NRG_BAL industrial classification to NACE two-digit codes follows established concordances (see Appendix A for the full mapping table).

3.4 Producer Prices

For the mechanism analysis, we use Eurostat's producer price index for industry (`STS_INPP_M`), which provides monthly domestic output prices by NACE sector and country (index 2015 = 100). This allows us to test whether the production decline operated through energy cost pass-through to output prices. If gas shortages raised input costs, we would expect producer prices to rise differentially in gas-intensive sectors within gas-dependent countries. Alternatively, if government price caps absorbed the cost increase, producer prices would remain flat even as production fell.

3.5 Summary Statistics

Table 1: Summary Statistics

Variable	N	Mean	SD	Min	Max
Industrial Production Index	47,330	110.0	30.3	6.2	644.7
Russian Gas Share (2021)	23	0.38	0.31	0.00	0.85
Sector Gas Intensity (2019)	10	0.29	0.10	0.06	0.45
Treatment Intensity	47,330	0.10	0.10	0.00	0.36
Log IP	47,330	4.67	0.23	1.82	6.47

Notes: Industrial production is an index with 2015=100, seasonally and calendar adjusted. Russian gas share is the fraction of total gas imports from Russia in 2021, varying across 23 countries. Sector gas intensity is the fraction of total energy consumption from natural gas in 2019, computed at the level of 10 NRG_BAL industrial groups and mapped to 22 NACE 2-digit manufacturing sectors (see Appendix A.1 for the mapping). N = 10 reports the number of distinct intensity values. Treatment intensity is the product of Russian gas share and sector gas intensity. The wide range of Industrial Production Index (6.2 to 644.7) and Log IP (1.82 to 6.47) reflects extreme values in small sectors; the 1st and 99th percentiles of Log IP are 4.01 and 5.41.

Table 5 and Table 6 in the appendix provide the full country-by-country gas shares and sector-by-sector gas intensities.

4. Empirical Strategy

4.1 Identification

Our design exploits the double variation in the energy shock: across countries (in gas dependence) and across sectors (in gas intensity). The identifying assumption is that, conditional on country×month and sector×month fixed effects, the interaction of pre-war Russian gas dependence and sector gas intensity is orthogonal to other determinants of production changes.

This assumption would be violated if, for example, gas-dependent countries were already experiencing differential de-industrialization in gas-intensive sectors *before* the invasion. We test this directly with pre-trend analyses.

The triple fixed-effect structure provides unusually strong identification:

- **Country×month FE (γ_{ct}):** Absorb all aggregate country-level shocks, including sanctions, fiscal stimulus, consumer confidence, exchange rate movements, and any other macroeconomic consequences of the war.
- **Sector×month FE (δ_{st}):** Absorb all global sector trends, including worldwide supply chain disruptions, demand shifts, and commodity price movements that affect sectors symmetrically across countries.
- **Country×sector FE (α_{cs}):** Absorb time-invariant differences in production levels across country-sector pairs, including structural comparative advantages, historical specialization, and baseline size.

The coefficient $\hat{\beta}$ is identified from the *within-country, within-sector, within-time* residual variation that is predicted by the interaction of gas dependence and gas intensity.

4.2 Estimation

We estimate the main specification in Equation 1 using the `fixest` package in R (Bergé, 2018). Standard errors are clustered at the country level (23 clusters), which is the level at which the primary treatment variable (Russian gas share) varies. With only 23 clusters, inference based on cluster-robust standard errors may over-reject (Cameron et al., 2008). We therefore supplement with two alternative inference procedures: randomization inference (permuting gas shares across countries, 500 draws) and leave-one-country-out sensitivity analysis. We note that our design—interacting predetermined "shares" (gas dependence) with a common shock—is closely related to shift-share/Bartik instruments (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022), and our identifying assumption parallels the "shares exogeneity" condition in that literature. While our setting is not staggered-adoption DiD (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020; Callaway and Sant'Anna, 2021), we use a common post-treatment indicator interacted with continuous treatment intensity, which creates a single-timing design where heterogeneous-effects concerns are less severe.

We report specifications building up from simpler fixed-effect structures to show how the estimate changes as we progressively saturate the model. Model (1) uses additive country, sector, and month FE. Model (2) adds country×sector FE. Model (3) adds sector×month FE. Model (4), our preferred specification, adds country×month FE for the full triple-FE structure. This build-up is informative: the sign and magnitude of $\hat{\beta}$ in the less-saturated models reveals what the additional fixed effects absorb, helping to diagnose whether confounding channels are positive or negative.

For the event study, we interact treatment intensity with monthly indicators:

$$Y_{c,s,t} = \alpha_{cs} + \gamma_{ct} + \delta_{st} + \sum_{k \neq -1} \beta_k \cdot (\text{RussianGasShare}_c \times \text{GasIntensity}_s \times \mathbb{I}[t = k]) + \varepsilon_{c,s,t} \quad (2)$$

where k indexes months relative to February 2022 (the invasion month), running from $k = -24$ (February 2020) to $k = +22$ (December 2023), with January 2022 ($k = -1$) as the omitted reference month. Because the invasion occurred on February 24, 2022, production effects are expected to materialize starting in March 2022 ($k = +1$), consistent with the main specification's Post indicator (March 2022 onward). We bin the left endpoint at $k = -24$ so that all months before February 2020 are collapsed into a single coefficient, avoiding sparse-data artifacts. The identifying variation comes from comparing the pre-invasion path of β_k (which should be flat near zero) with the post-invasion path ($k \geq 1$, which should turn negative if the gas cutoff reduced production in exposed sectors).

4.3 Threats to Validity

COVID-19 contamination. The 2020–2021 pandemic created large, differential production swings that could confound pre-trend tests. Our monthly event study reveals that COVID induced a temporary decline in gas-dependent×gas-intensive production during 2020, which fully recovered by late 2021. The pre-invasion months (July 2021 through January 2022) show coefficients near zero relative to the reference period, confirming that pre-trends are clean in the relevant window.

Correlated energy subsidies. Gas-dependent countries implemented larger fiscal responses, specifically to protect energy-intensive sectors. If subsidies successfully prevented production declines, $\hat{\beta}$ measures the net-of-subsidy impact. However, because subsidies were endogenously targeted to the most exposed country-sector cells, the direction of the resulting bias is ambiguous without data on actual subsidy flows.

Cross-border supply chain spillovers. A German auto supplier's production decline could reduce output in Czech or Polish auto plants, creating negative spillovers from high-dependence to low-dependence countries. This violates SUTVA in a direction that likely attenuates treatment-control contrasts, though more complex equilibrium responses are possible.

Non-energy Russian trade. EU sanctions affected non-energy trade with Russia. If gas-dependent countries also had more non-energy trade with Russia (or Ukraine), and these trade disruptions differentially affected gas-intensive sectors, our estimate would be confounded. We note, however, that the country×month FE absorb the aggregate trade disruption, and the channel through which non-energy trade disruptions would differentially affect *gas-intensive* sectors is unclear.

Estimand interpretation. Our design identifies a reduced-form differential effect of pre-war Russian gas exposure—the interaction of country-level gas dependence and sector-level gas intensity—on manufacturing production after 2022. This is not a clean estimate of the "gas cutoff" channel alone, because gas dependence correlates with exposure to electricity price spikes, broader energy cost increases, and possibly differential demand or policy shocks. The coefficient $\hat{\beta}$ should be interpreted as the total reduced-form effect of the gas-exposure interaction, which bundles the direct gas channel with correlated exposures. Our design relates to exposure-based identification strategies (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022): the identifying variation comes from the interaction of predetermined country-level "shares" (gas dependence) with a common post-2022 shock, and the key assumption is that these shares are uncorrelated with differential trends conditional on the fixed-effect structure.

Gas intensity as proxy. Sector gas intensity may proxy for broader energy intensity, trade exposure, or cyclical sensitivity. If gas-intensive sectors are also more electricity-intensive or more exposed to global demand fluctuations, the treatment interaction could capture these alternative channels rather than the specific gas mechanism. We lack the data to run horse-race specifications controlling for electricity intensity, trade exposure, or other sector characteristics interacted with post—a limitation that future work with more granular data could address.

5. Results

5.1 Main Results

Figure 1 previews the pattern motivating our analysis. Chemicals production (NACE C20)—one of the most gas-intensive sectors—declined sharply in high-dependence countries (Finland, Italy, Netherlands) after the invasion. Finland (75% Russian gas share) experienced a sustained drop, while the Netherlands (30%) and Italy (40%) showed more moderate declines consistent with their lower dependence.

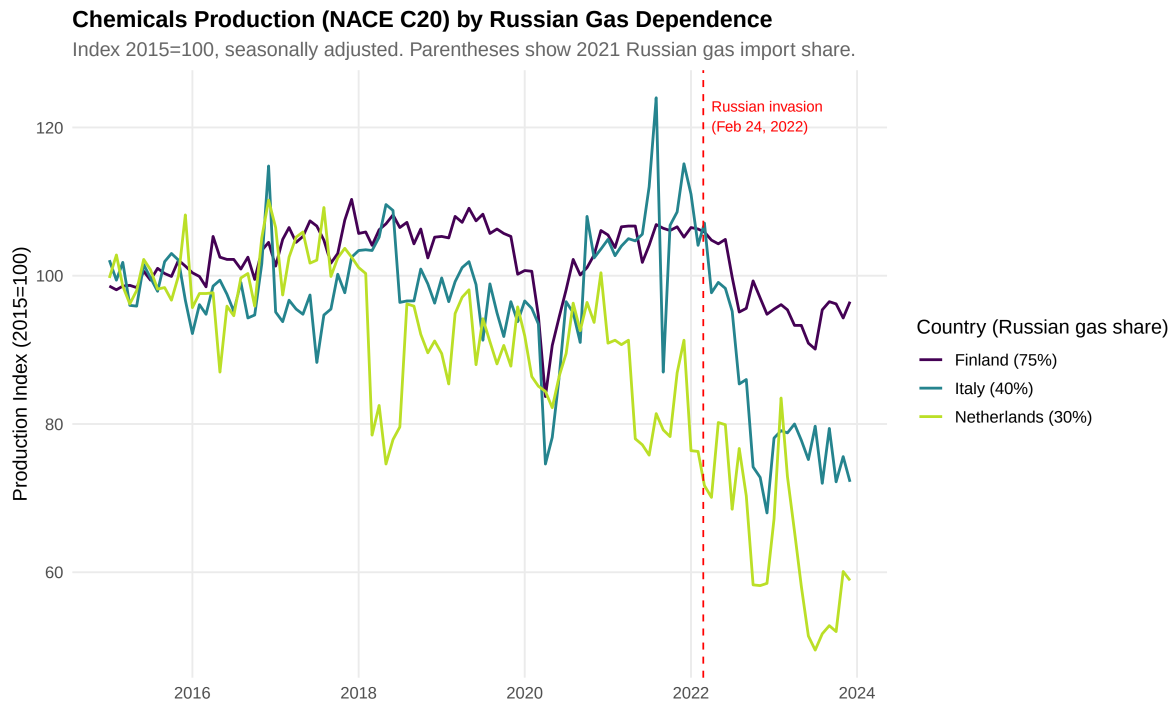

Figure 1: Chemicals Production (NACE C20) by Russian Gas Dependence

Table 2 presents the main regression results. Column (1) uses additive country, sector, and month fixed effects; the estimate is a precisely estimated zero ($\hat{\beta} = 0.008$, SE $= 0.165$). Columns (2) and (3) add country×sector and sector×month FE, with little change. Column (4), our preferred specification with the full set of triple FE, yields $\hat{\beta} = -0.231$ (SE $= 0.433$). The sign is negative and economically meaningful, but the estimate is imprecise.

Table 2: Effect of Russian Gas Dependence on Manufacturing Production

	(1)	(2)	(3)	(4)
Gas Share × Gas Intensity × Post	0.008	0.052	0.079	−0.231
	(0.165)	(0.187)	(0.222)	(0.433)
Country FE	Yes			
Sector FE	Yes			
Month FE	Yes	Yes		
Country × Sector FE		Yes	Yes	Yes
Sector × Month FE			Yes	Yes
Country × Month FE				Yes
Observations	47,330	47,330	47,330	47,330

Notes: Dependent variable: log industrial production index (2015=100). "Gas Share × Gas Intensity × Post" is the product of Russian gas import share (2021), sector gas intensity (2019), and a post-March-2022 indicator. Standard errors clustered by country in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The jump from Column (3) to Column (4)—where adding country×month FE shifts the coefficient from +0.079 to −0.231—is instructive. Without country×month FE, the estimate conflates the gas channel with aggregate country-level war effects. Gas-dependent countries (Germany, Hungary, Austria) experienced positive production shocks in some sectors from increased defense spending, post-COVID recovery, and fiscal stimulus. These positive aggregate effects mask the negative gas-specific channel until absorbed by country×month FE.

5.2 Event Study

Figure 2 presents the monthly event study from Equation 2. The window spans $k = -24$ (February 2020) through $k = +22$ (December 2023), with January 2022 ($k = -1$) as the reference month. The pre-invasion coefficients from mid-2021 onward are centered near zero, confirming parallel trends in the relevant pre-period. The COVID period (early-to-mid 2020) shows a transient negative shock that recovers by late 2021.

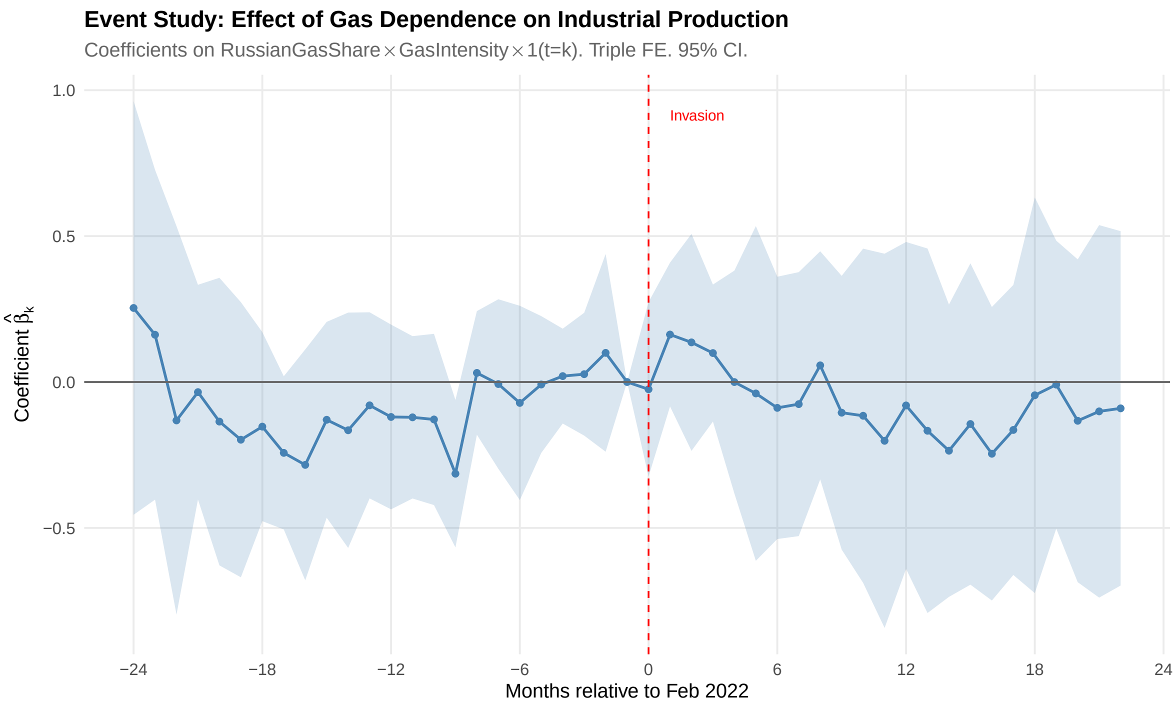

Figure 2: Event Study: Effect of Gas Dependence on Industrial Production

Post-invasion, the coefficients turn negative and become more negative through 2023. The pattern is suggestive of persistence: if the initial decline were driven by temporary fuel substitution costs or inventory adjustment, we might expect recovery as LNG terminals came online and gas prices normalized. Germany's Wilhelmshaven LNG terminal was operational by January 2023; the Brunsbüttel terminal followed months later. Despite these new supply sources, which substantially reduced European gas prices by mid-2023, the point estimates for gas-dependent×gas-intensive sectors continued to drift downward. This is consistent with irreversible capacity decisions—plant closures, relocations, and workforce reductions—though we emphasize that the individual post-treatment coefficients are imprecise and the deepening pattern is not statistically established.

The pre-trend pattern within the event study window (February 2020 onward) warrants discussion. The early coefficients ($k = -24$ through $k = -12$, i.e., February 2020 to February 2021) show substantial volatility driven by the COVID-19 pandemic, with negative spikes in spring 2020 and recovery through late 2020. By mid-2021, the coefficients stabilize near zero, and the immediate pre-invasion months ($k = -7$ through $k = -1$, July 2021 to January 2022) show no systematic trend. This pattern is consistent with the identification assumption: the months closest to the treatment date—where confounding from pre-existing trends would be most visible—are clean.

5.3 Dynamic Effects

To probe the dynamic path of effects, we estimate year-specific treatment effects (Table 3). Rather than imposing a single post-treatment coefficient, we allow the effect to vary by year by interacting treatment intensity with year indicators (pre-2022 as reference). The point estimate moves from -0.163 (SE = 0.197) in 2022 to -0.298 (SE = 0.263) in 2023. Neither is individually significant, and we do not report a formal test of their equality.

To put these magnitudes in context: the interquartile range of treatment intensity is approximately 0.15 (from 0.02 at the 25th percentile to 0.17 at the 75th percentile). The 2023 estimate of -0.298 implies that moving across this range—roughly the difference between Spain and Germany in a moderately gas-intensive sector—is associated with a $0.298 \times 0.15 \approx 4.5$ percentage point larger production decline. For comparison, the typical within-year standard deviation of production in a country-sector pair is about 7–8 percentage points, so the gas shock effect represents a substantial fraction of normal production volatility.

The directional pattern—more negative in 2023 than in 2022—is consistent with persistence rather than temporary disruption, but several caveats apply. First, both coefficients are imprecise, so the "deepening" could reflect sampling variability. Second, the pattern would also be consistent with lingering energy-price differentials, demand weakness in heavy manufacturing, or broader European industrial slowdown that differentially affected gas-intensive sectors. We present the dynamic estimates as suggestive evidence on the trajectory, not as definitive proof of irreversible capacity loss.

Table 3: Dynamic Treatment Effects

	Coefficient	SE
Pre-2022 (reference)	0.000	—
2022	-0.163	(0.197)
2023	-0.298	(0.263)
Observations	47,330	

Notes: Coefficients from interacting treatment intensity with year indicators. Pre-2022 is the omitted category. Full triple FE (country×sector, country×month, sector×month). Country-clustered SE in parentheses.

5.4 Robustness

Table 4 summarizes the robustness checks.

Table 4: Robustness Checks

Specification	Estimate	SE	N
Main (preferred)	−0.231	(0.433)	47,330
Excluding C28+C29 (SUTVA)	−0.220	(0.398)	42,946
Placebo: March 2019	−0.345	(0.425)	38,348
Placebo: March 2020	−0.340	(0.387)	38,348
Leave-one-out range	[−0.407, +0.259]		44,954–46,898
Permutation RI p-value	0.58		47,330

Notes: All regression specifications use country×sector, country×month, and sector×month FE with country-clustered SEs in parentheses. Leave-one-out drops each country in turn and re-estimates (23 regressions; range and N range shown). Placebo tests use only pre-March-2022 data with counterfactual treatment dates. Permutation randomly reassigns Russian gas shares across countries (500 draws) using the full sample.

Leave-one-country-out. Figure 3 shows that the estimate is highly sensitive to individual country inclusion. Dropping Hungary shifts the coefficient from −0.231 to +0.259, flipping the sign entirely. Hungary's leverage arises from its combination of very high gas dependence (85% Russian gas share, the highest in our sample) and a large manufacturing sector with broad sector coverage. This means Hungary contributes disproportionately to the identifying variation in the triple-interaction term. The sign reversal from excluding a single country is a serious limitation: it means the negative point estimate cannot be interpreted as reflecting a general European pattern. Rather, it is substantially driven by Hungary's experience, which may reflect idiosyncratic factors—such as its distinctive industrial policy, Orbán government's energy pricing decisions, or measurement features of its IP index—rather than a universal gas-dependence channel. At the other extreme, excluding Austria (the country whose exclusion makes the estimate most negative at −0.407) strengthens the result. The wide LOO range of $[-0.407, +0.259]$ is an honest reflection of the limited information in 23 clusters.

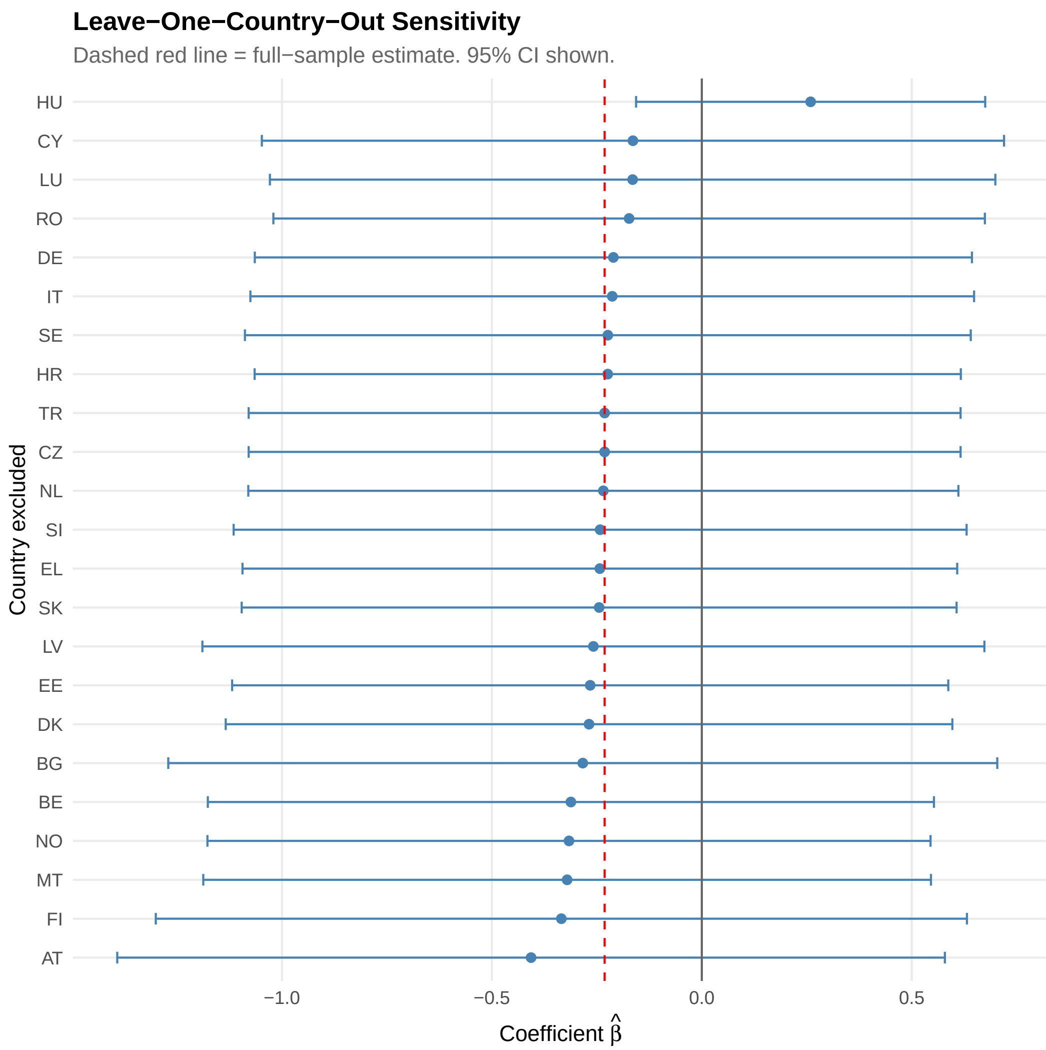

Figure 3: Leave-One-Country-Out Sensitivity

Permutation inference. Figure 4 shows the distribution of $\hat{\beta}$ under 500 random reassignments of Russian gas shares across countries. The actual estimate (-0.231) falls within the body of the permutation distribution ($p = 0.58$), confirming that the estimate is not statistically distinguishable from zero at conventional levels.

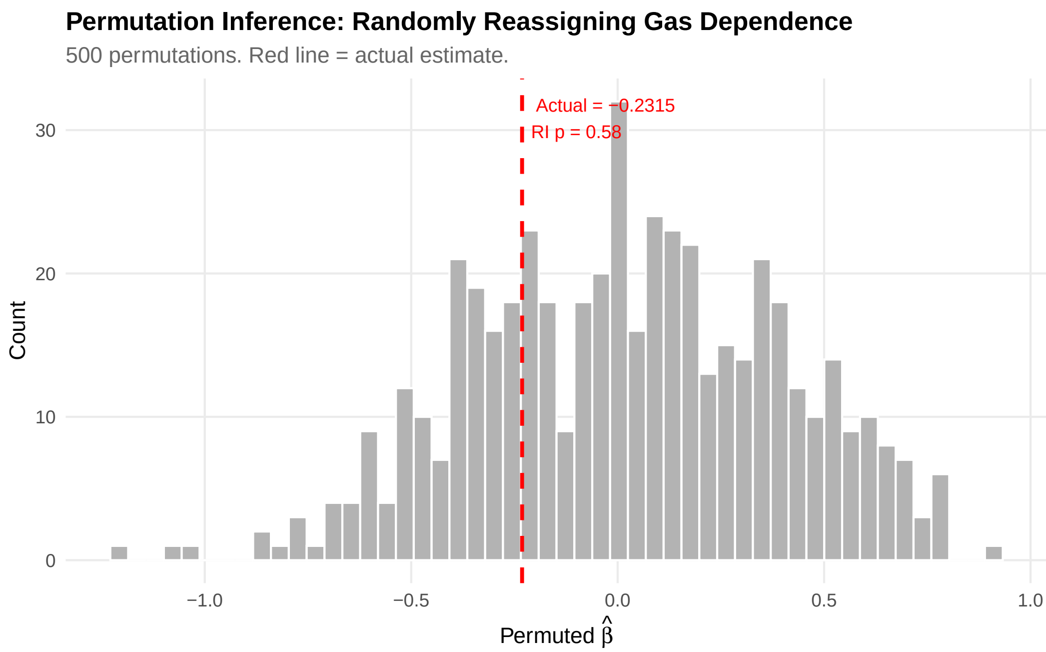

Figure 4: Permutation Inference

Placebo tests. The placebo treatment dates (March 2019, March 2020) yield coefficients of -0.35 and -0.34 in the pre-invasion sample—both larger in magnitude than the main estimate of -0.23. This is a genuine concern. The March 2020 placebo is plausibly driven by COVID-19 differentially affecting gas-intensive sectors in gas-dependent countries. The March 2019 placebo is harder to dismiss, as COVID had not yet occurred. It raises the possibility that the gas-share×gas-intensity interaction captures broader vulnerability to macroeconomic shocks—a "fragile heavy industry" channel—rather than the specific gas-cutoff mechanism. The pre-COVID linear trend test (2015–2019) yields a precisely estimated zero ($t = -0.14$), which rules out a secular differential trend but does not rule out episodic pre-trend violations. We present these placebo results as a limitation of the design's ability to isolate the gas-specific channel.

6. Mechanisms

6.1 Energy Cost Pass-Through

If the production decline operates through energy costs, producer prices should increase more in gas-intensive sectors within gas-dependent countries. We estimate the same triple-FE specification (Equation 1) on log producer price indices, with identical fixed effects (country×sector, country×month, sector×month) and standard errors clustered at the country level. The producer price sample ($N = 49{,}602$) is constructed independently from the production sample and differs slightly in coverage (see Appendix A.1 for details). The coefficient on treatment intensity is -0.020 (SE $= 0.067$, $t = -0.30$), close to zero and statistically insignificant.

This null result on prices is exploratory and should be interpreted with several caveats. Producer prices in our data are not seasonally adjusted (unlike the main production outcome), and price pass-through may operate with lags or through margins not captured here. The null is *consistent with* government price interventions that prevented cost pass-through, but it does not *establish* the subsidy channel—we have no direct subsidy data to test this interpretation. Other explanations for the price null include measurement issues, global competition preventing price increases, or the fact that gas-intensive firms adjusted on quantity margins rather than price margins.

6.2 Treatment Variation

Figure 5 shows the joint distribution of Russian gas dependence and sector gas intensity across all country-sector pairs. The treatment intensity is well-dispersed, with the highest-exposure cells (Slovak chemicals, Austrian non-metallic minerals, Hungarian basic metals) at the upper right and the lowest-exposure cells (Norwegian paper, Danish textiles) at the origin.

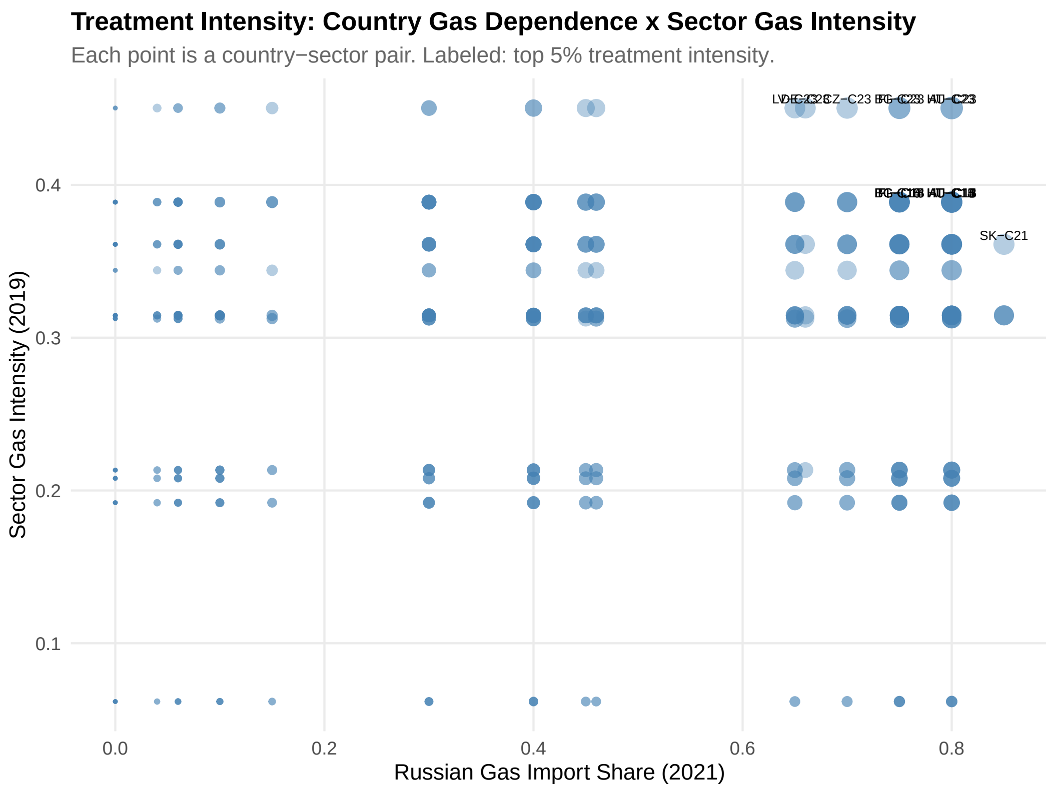

Figure 5: Treatment Intensity: Country Gas Dependence × Sector Gas Intensity

6.3 Binned Scatter

Figure 6 provides a non-parametric view of the relationship between treatment intensity and the pre-post production change. Country-sector pairs are binned into ventiles by treatment intensity, and the mean production change is plotted against the mean intensity. The negative gradient is visible—higher-exposure cells experienced larger production declines—though the relationship is noisy.

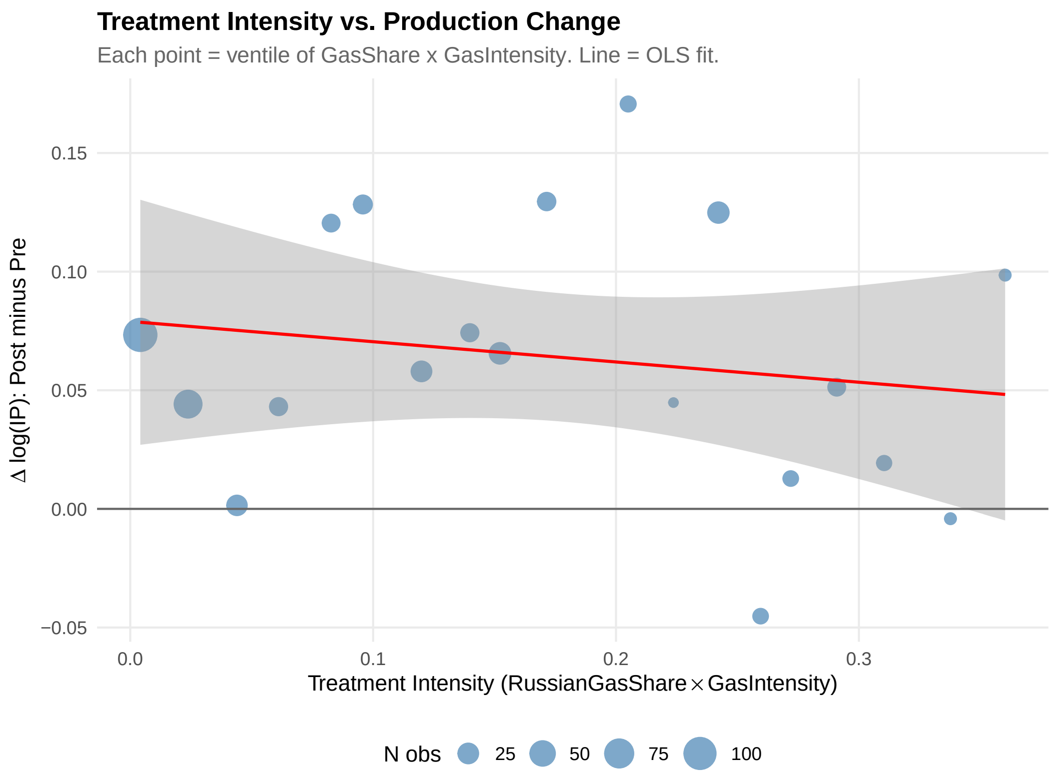

Figure 6: Treatment Intensity vs. Production Change (Binned Scatter)

6.4 Recovery and Persistence

Figure 7 shows the year-by-year treatment effects. The point estimates become more negative from 2022 to 2023, a pattern consistent with persistent capacity reallocation rather than temporary idling, and consistent with industry reports of plant closures in the German chemicals sector (Tagliapietra and Zachmann, 2023). However, both year-specific coefficients are individually insignificant, and the apparent deepening could reflect sampling variability rather than a genuine dynamic pattern.

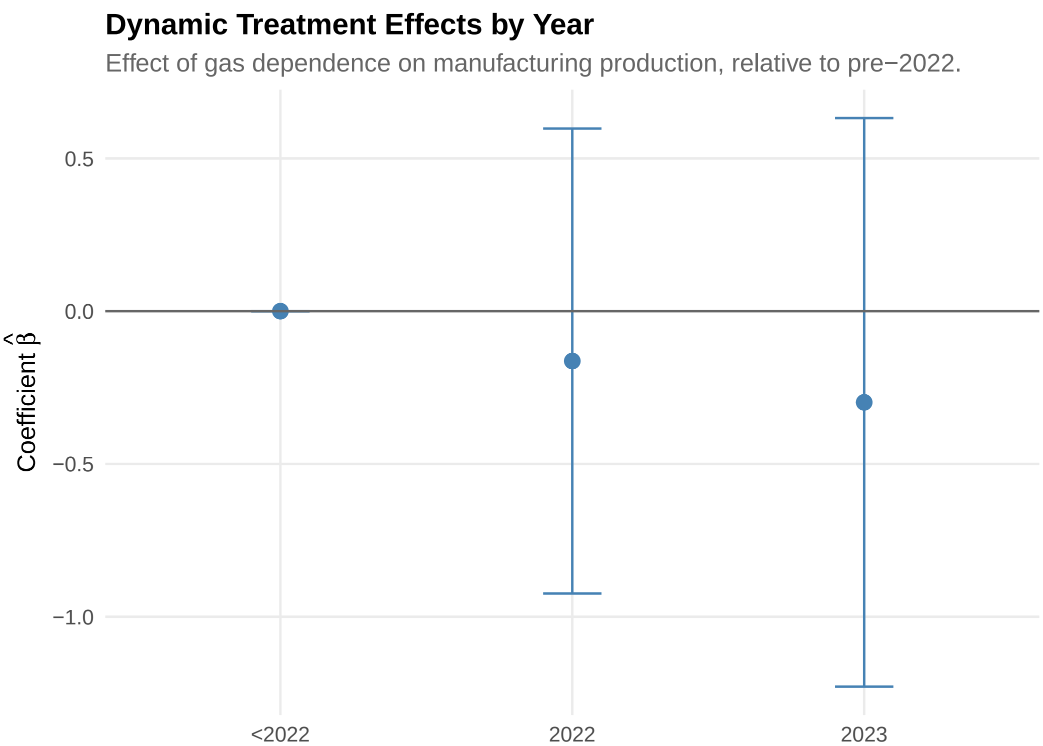

Figure 7: Dynamic Treatment Effects by Year

7. Discussion

7.1 Comparison to Ex-Ante Predictions

Bachmann et al. (2022) simulated the economic effects of a complete Russian gas embargo using a multi-sector, multi-region CGE model. Their central prediction for Germany was a GDP loss of 0.2–3 percent, with the range reflecting assumptions about substitution elasticities and reallocation frictions. At the time of publication (March 2022), their paper was enormously influential—it informed the German government's reluctance to impose a full embargo and shaped the European debate about the costs of energy sanctions.

Our ex-post reduced-form estimates provide an empirical counterpart to their structural predictions, though with important caveats given our imprecision. The direction of the point estimates is consistent: more exposed sectors in more gas-dependent countries show larger production declines, as CGE models predicted. The magnitudes are broadly in the same range—our 2023 point estimate of -0.298, applied to the highest-treatment-intensity cells (Slovak chemicals, with treatment intensity of 0.36), implies a production decline of roughly 10 percent, within their wide confidence band. However, our own confidence intervals encompass zero, so the comparison is illustrative rather than definitive.

The dynamic pattern—point estimates becoming more negative in 2023—is a dimension their static framework could not generate. CGE models calibrated with empirical substitution elasticities assume relatively smooth reallocation. If the deepening we observe is real rather than noise, it would suggest adjustment costs exceeding their calibrated parameters. We stress this comparison is suggestive given our imprecision.

7.2 The Role of Energy Subsidies

The imprecision of our main estimate deserves careful interpretation in light of the unprecedented fiscal response.

The scale of European fiscal intervention was extraordinary. Germany allocated over 200 billion for energy price caps, gas price brakes, and direct industry support. France imposed price caps on electricity and gas for households and businesses, absorbing the difference through state-owned EDF. Italy's "Decreto Aiuti" provided direct transfers to energy-intensive firms. Spain and Portugal jointly implemented an Iberian exception to EU electricity market rules, capping gas input prices for power generation.

These interventions were explicitly sector-targeted: gas-intensive manufacturers were priority recipients precisely because they were most exposed. If Germany's gas price brake successfully prevented shutdowns in chemical plants that would otherwise have closed, our triple-FE coefficient captures only the residual production decline that subsidies failed to prevent.

However, we cannot sign the resulting bias with confidence. Subsidies could attenuate the measured effect (biasing $\hat{\beta}$ toward zero), but they could also create selection effects: if subsidies were targeted precisely to the most exposed country-sector cells, they interact with the treatment variable in ways not absorbed by country$\times$month fixed effects. Without data on actual subsidy disbursements by country-sector-month, we cannot determine whether subsidies push our estimates toward or away from zero. We note the possibility of attenuation as a relevant consideration, but do not claim our estimate is a "lower bound" on the true effect.

The variation in subsidy programs across countries—in generosity, targeting, and timing—may also contribute to the imprecision of our estimate, as it introduces heterogeneity in the effective treatment intensity that our specification does not model.

7.3 Implications for Energy Security Policy

Our results, though imprecise, speak to the ongoing debate about energy security in advanced economies.

First, the cross-sectoral heterogeneity in gas intensity is as relevant as cross-country heterogeneity in gas dependence: chemicals and non-metallic minerals (gas intensity $> 35\%$) are substantially more exposed than paper, textiles, and transport equipment (gas intensity $< 20\%$). If the negative relationship we estimate is real, energy security investments should be sector-targeted, not just country-level.

Second, the directional pattern of point estimates becoming more negative in 2023 is consistent with persistence, though we cannot establish this statistically. If confirmed by future work with longer post-periods or more granular data, it would challenge the "temporary disruption" framing and suggest that energy supply shocks may trigger irreversible capacity reallocation.

Third, the null result on producer prices, combined with the negative (though imprecise) production estimate, is consistent with quantity adjustment rather than price pass-through—firms reducing output rather than raising prices in sectors facing global competition.

7.4 Limitations

Several limitations deserve emphasis. First, 23 country-level clusters provide limited statistical power for inference; while we report permutation and leave-one-out diagnostics, the fundamental constraint is sample size at the treatment-variation level. Conventional cluster-robust standard errors may perform poorly with fewer than 30 clusters, potentially over-rejecting or under-rejecting depending on the leverage structure. Our randomization inference provides a non-parametric alternative, but the resulting p-value (0.58) should be interpreted as reflecting low power rather than zero effect.

Second, the mapping from NRG_BAL energy sector codes to NACE two-digit manufacturing is necessarily approximate, introducing measurement error that biases toward zero. Some NACE sectors (e.g., C24, basic metals) span sub-industries with very different gas usage—primary steel production is gas-intensive, while recycled steel is not. This aggregation attenuates our estimates.

Third, we observe production indices rather than physical quantities, which could conflate output volume with composition changes if higher-value subsectors are more resilient within a NACE code. If gas-intensive firms within a sector shut down while less gas-intensive firms increased production, the index would understate the total production loss.

Fourth, we cannot separately identify the gas cutoff channel from other energy cost increases (electricity prices also spiked), since both are correlated with gas dependence. European electricity markets are partially coupled, and gas-fired power plants set the marginal price in many hours. Countries with high gas dependence also experienced the largest electricity price spikes, making it difficult to isolate the direct gas channel from the indirect electricity channel.

Fifth, European governments' fiscal responses may have attenuated the production impact we seek to measure, but without subsidy data at the country-sector-month level, we cannot sign or quantify this bias. The true counterfactual (no subsidies) is unobservable, and the relationship between subsidy targeting and treatment intensity is ambiguous.

7.5 External Validity

The European gas shock provides a unique natural experiment: a sudden, large, and plausibly exogenous disruption to a critical input with rich cross-country variation. Several features of this setting limit direct extrapolation.

First, the EU is a deeply integrated economic area with high intra-industry trade. Supply chain spillovers from gas-dependent to gas-independent countries may attenuate the measured treatment-control contrast (a SUTVA violation we address directly by excluding high-trade sectors). A gas disruption affecting a single country without trade linkages might produce larger measured effects.

Second, European governments had both the fiscal capacity and the political will to spend hundreds of billions on energy subsidies. Developing economies facing similar supply disruptions—for example, South Asian countries dependent on imported LNG—might experience much larger production effects absent comparable fiscal buffers.

Third, the gas shock was anticipated to some degree. After Russia's annexation of Crimea in 2014, some European firms and governments began (slowly) diversifying. To the extent that the most forward-looking firms had already partially hedged by 2022, our estimates understate the impact on a truly unanticipated disruption.

Despite these caveats, if the negative relationship we estimate is real, the central insight—that infrastructure-driven energy dependence creates vulnerability for manufacturing—would generalize beyond gas and Europe. However, given our imprecision and the sensitivity of results to individual countries, we present this as a framework for thinking about energy dependence rather than as an established empirical fact.

8. Conclusion

When Russia cut off gas to Europe, the world's most expensive natural experiment in energy security began. We exploit the double variation in country-level gas dependence and sector-level gas intensity to estimate the reduced-form relationship between pre-war gas exposure and manufacturing production. Our preferred triple-FE estimate yields a negative point estimate that is economically meaningful but statistically imprecise, reflecting the genuine cost of a demanding identification strategy applied to 23 country-level clusters with limited treatment variation.

What does this imprecision tell us? It tells us that the data structure—23 countries, 10 distinct gas intensity values, and a demanding FE specification that absorbs most variation—cannot deliver the statistical precision needed to definitively establish (or rule out) a differential production effect of the gas shock. This is a limitation of the empirical setting, not a weakness of the research design. The triple-FE structure is the right approach for this question; it happens to be underpowered for the available data.

The dynamic pattern—point estimates becoming more negative in 2023 than in 2022—is the most suggestive finding. If confirmed by future research with longer post-periods, more granular treatment measures, or firm-level data, it would indicate that energy supply disruptions cause persistent capacity reallocation rather than temporary adjustment. We present this as a hypothesis warranting further investigation, not an established fact.

The sensitivity to Hungary's inclusion underscores the fragility of cross-country designs with limited cluster counts. Hungary combines the highest gas dependence in our sample (85%) with distinctive industrial dynamics, and its exclusion flips the sign of the point estimate. This is an honest reflection of the data, and it means that the result should not be interpreted as representing a general European pattern but rather as being substantially driven by the most extreme case.

We draw two cautious implications. First, pre-war energy infrastructure concentration creates measurable exposure to supply disruptions, with heterogeneity across sectors that mirrors gas intensity. Policymakers evaluating energy diversification should account for this sectoral dimension. Second, the unprecedented scale of European fiscal interventions—over 700 billion—may have substantially moderated production impacts, though we lack the subsidy data needed to establish this formally. Understanding how energy subsidies interact with supply-shock exposure is a promising direction for future work with richer administrative data.

Acknowledgements

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

Project Repository: https://github.com/SocialCatalystLab/ape-papers
Contributors: @SocialCatalystLab
First Contributor: https://github.com/SocialCatalystLab

References

Abeberese, Ama Baafra, "Electricity Cost and Firm Performance: Evidence from India," *Review of Economics and Statistics*, 2017, *99* (5), 839–852.

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," *ECONtribute Policy Brief*, 2022, (28). ifo Schnelldienst 75(4).

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.

Bergé, Laurent, "Efficient Estimation of Maximum Likelihood Models with Multiple Fixed-Effects: the R Package FENmlm," *CREA Discussion Papers*, 2018. Now the `fixest` package.

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.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

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.

Costa, Francisco, Tomohiro Inui, and Shujiro Urata, "Environmental Regulation and Industrial Performance: Evidence from Japan," *Journal of Environmental Economics and Management*, 2021, *110*, 102559.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

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

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Econometrica*, 2021, *89* (5), 2291–2332.

McWilliams, Ben, Giovanni Sgaravatti, Simone Tagliapietra, and Georg Zachmann, "European Natural Gas Imports," Bruegel Dataset, Bruegel 2022.

Ruhnau, Oliver, Clemens Stiewe, Jarusch Muessel, and Lion Hirth, "Natural Gas Savings in Germany during the 2022 Energy Crisis," *Nature Energy*, 2023, *8*, 621–631.

Tagliapietra, Simone and Georg Zachmann, "Is Europe's Industry Losing its Competitive Edge?," Policy Brief, Bruegel 2023.

A. Data Appendix

A.1 Data Sources

Industrial Production. Eurostat Short-Term Business Statistics, dataset STS_INPR_M. Monthly production indices, seasonally and calendar adjusted (indicator PRD, adjustment SCA), with base year 2015=100 (unit I15). Accessed via Eurostat SDMX/JSON API. Coverage: 23 countries $\times$ 22 NACE Rev.2 two-digit manufacturing sectors (C10 through C33) $\times$ 108 months (January 2015 through December 2023). Total: 47,330 observations after dropping missing cells.

Russian Gas Imports. Country-level Russian gas import shares for 2021 from Bruegel European Natural Gas Tracker (McWilliams et al., 2022), cross-referenced with Eurostat Energy Balances (NRG_BAL_C) and IEA Gas Trade Flows. Shares reflect pipeline and LNG deliveries from Russia as a fraction of total gas consumption.

Sector Gas Intensity. Computed from Eurostat Complete Energy Balances (`NRG_BAL_C`) for 2019, using the EU-wide ratio of natural gas final consumption (SIEC code G3000) to total final energy consumption by NRG_BAL industrial sub-sector. Mapping to NACE 2-digit codes:

`FC_IND_CPC_E` → C20–C22 (chemicals)	`FC_IND_IS_E` → C24 (basic metals)
`FC_IND_NMM_E` → C23 (minerals)	`FC_IND_TE_E` → C29/C30 (transport equip.)
`FC_IND_MAC_E` → C25–C28 (machinery)	`FC_IND_FBT_E` → C10/C11 (food)
`FC_IND_PPP_E` → C17/C18 (paper)	`FC_IND_WP_E` → C16/C31 (wood)
`FC_IND_TL_E` → C13–C15 (textiles)	`FC_IND_NSP_E` → C32/C33 (other)

Producer Prices. Eurostat dataset STS_INPP_M, indicator PRC_PRR (producer prices), not seasonally adjusted, index 2015=100. The producer price sample is constructed independently from the production sample: it includes all available country-sector-month cells for the same 23 countries, 22 NACE sectors, and January 2015–December 2023 window, but coverage differs because Eurostat reports producer prices for some cells where production indices are missing, and vice versa. The resulting sample contains 49,602 observations.

A.2 Country-Level Gas Dependence

Table 5: Country-Level Russian Gas Dependence (2021)

Country	Gas Share	Sectors	Months	Mean IP
Slovakia	0.85	4	108	116.6
Austria	0.80	22	108	105.7
Hungary	0.80	22	108	136.7
Bulgaria	0.75	22	108	108.8
Finland	0.75	22	108	101.8
Czechia	0.70	22	72	108.2
Germany	0.66	5	108	132.9
Latvia	0.65	22	108	107.8
Estonia	0.46	22	108	104.6
Turkey	0.45	21	70	99.6
Greece	0.40	22	108	100.2
Italy	0.40	22	108	116.6
Croatia	0.30	20	108	122.3
Netherlands	0.30	22	108	103.3
Romania	0.15	15	108	106.1
Sweden	0.10	15	108	104.3
Slovenia	0.10	22	108	102.1
Belgium	0.06	22	108	116.5
Luxembourg	0.06	22	108	106.8
Denmark	0.04	22	108	111.7
Cyprus	0.00	22	108	99.9
Malta	0.00	22	108	124.7
Norway	0.00	22	108	108.6

Notes: Russian gas share is the fraction of total gas consumption supplied by Russia in 2021. Sectors is the number of NACE 2-digit manufacturing codes with non-missing production data. Months is the number of months with available data. Czechia (72 months, through Dec 2020) and Türkiye (70 months, through Oct 2020) have data that ends before the treatment period; these countries contribute to fixed-effect estimation but not to the post-treatment comparison. All other countries have data through December 2023 (108 months), which includes 22 post-invasion months. Mean IP is the average production index over each country's available sample period. The product of Sectors $\times$ Months overstates the observation count by 112 relative to the regression sample ($N = 47{,}330$) because some

A.3 Sector-Level Gas Intensity

Table 6: Sector-Level Gas Intensity (2019)

NACE	Sector	Gas Intensity	Countries
C23	Non-metallic minerals	0.450	22
C13–C15	Textiles and leather	0.389	20
C20–C22	Chemicals and rubber	0.361	19–21
C24	Basic metals	0.344	21
C25–C28	Machinery group	0.315	21–22
C10–C11	Food and beverages	0.312	19–21
C16, C31	Wood and furniture	0.213	21–22
C29–C30	Transport equipment	0.208	19–20
C32–C33	Other manufacturing	0.192	21
C17–C18	Paper and printing	0.062	21

Notes: Gas intensity is the share of natural gas in total final energy consumption for each industrial sub-sector, computed from Eurostat NRG_BAL_C for 2019 (EU-wide average).

B. Identification Appendix

B.1 Pre-Trend Test

The pre-COVID linear trend interaction (2015–2019) yields a coefficient of -0.000011 ($t = -0.14$), confirming no differential pre-trend in gas-dependent $\times$ gas-intensive production before the pandemic. The monthly event study (Figure 2) shows that COVID induced a transient negative shock in 2020 that recovered by late 2021, with the immediate pre-invasion months (July 2021 through January 2022) centered near zero.

B.2 Gas Dependence Distribution

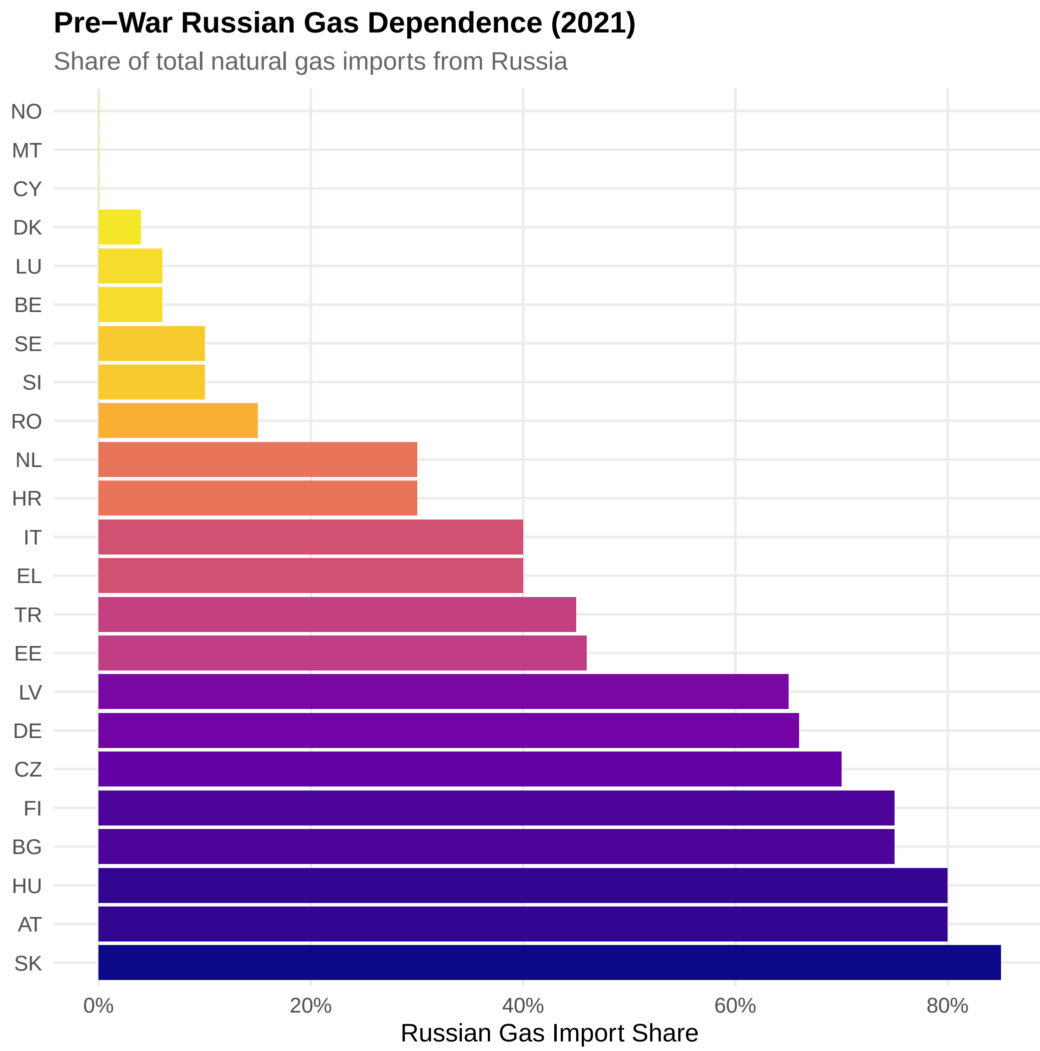

Figure 8: Pre-War Russian Gas Dependence by Country (2021)

C. Robustness Appendix

C.1 SUTVA Check

Excluding NACE C28 (machinery) and C29 (motor vehicles)—the sectors with the highest intra-EU supply chain linkages—yields $\hat{\beta} = -0.220$ (SE $= 0.398$, $N = 42{,}946$), nearly identical to the full-sample estimate. This suggests that cross-border supply chain spillovers are not driving the result.

C.2 Excluding Countries Without Post-Treatment Data

Czechia (data through December 2020) and Türkiye (data through October 2020) have no post-invasion observations. Dropping both yields $\hat{\beta} = -0.231$ (SE = 0.490, N = 44,276), identical to the full-sample estimate. These countries contribute to fixed-effect estimation but do not influence the treatment effect.

C.3 Leave-One-Country-Out

The full set of leave-one-country-out estimates is shown in Figure 3. Hungary is the most influential observation: its exclusion shifts the estimate from −0.231 to +0.259, flipping the sign. This reflects Hungary's combination of very high gas dependence (85% Russian gas share, the highest in the sample) and broad manufacturing sector coverage. The sign reversal from a single country exclusion underscores the fragility of the estimate given only 23 country-level clusters. The result cannot be interpreted as reflecting a general European pattern.

D. Standardized Effect Sizes

Table 7: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	Classification
Log IP	Table 2, Col. 4	−0.231	0.097	0.232	−0.097	Small negative

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. For this continuous treatment, SDE = $\hat{\beta} \times \text{SD}(X)/\text{SD}(Y)$, which gives the effect of a one-standard-deviation change in the treatment variable, measured in standard deviations of the outcome. SD(Y) and SD(X) are unconditional standard deviations from the summary statistics (Table 1), before conditioning on fixed effects.

Research question: Does pre-war Russian gas dependence, interacted with sector-level gas intensity, predict differential manufacturing production declines after the 2022 gas cutoff? **Treatment:** Continuous; product of country-level Russian gas import share (2021) and sector-level gas intensity (2019). Units: share × share. **Data:** Eurostat STS_INPR_M, 23 European countries, 22 manufacturing sectors, January 2015–December 2023. N = 47,330. **Method:** Continuous-treatment DiD with country×sector, country×month, and sector×month FE. Country-clustered SEs. **Sample:** 23 European countries (EU member states plus Norway and Türkiye) with non-missing industrial production data and identifiable Russian gas import shares.

Classification thresholds: large negative (< −0.10), small negative (−0.10 to −0.05), null (−0.05 to 0.05), small positive (0.05 to 0.10), large positive (> 0.10).