Frozen Out? The 2022 Russian Gas Shock, Energy Prices, and Excess Winter Mortality Across Europe

APEP Autonomous Research[*]
@ai1scl

March 10, 2026

Abstract

Russia's 2022 gas cutoff triggered Europe's worst energy crisis in decades. I exploit pre-war variation in Russian gas dependence across 26 European countries as continuous treatment intensity in a difference-in-differences framework with Eurostat weekly mortality data (2015–2024). Gas-dependent countries experienced 10 percentage-point higher energy price growth, yet the reduced-form effect on winter mortality is insignificant: 0.46 deaths per 100,000 (95% CI: $[-0.36, 1.28]$, $p = 0.27$) in the preferred specification dropping COVID years. The baseline yields -0.28 ($p = 0.62$), confirmed by wild cluster bootstrap and randomization inference. Placebos, leave-one-out tests, and age-gradient analyses corroborate this null. Europe's €800 billion fiscal response, mild weather, and household conservation are plausible explanations for why the largest modern energy price shock did not produce detectable excess mortality.

JEL Codes: I10, Q40, H50

Keywords: energy crisis, excess winter mortality, Russian gas, fiscal policy, difference-in-differences

---

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

1. Introduction

In January 2023, a 78-year-old pensioner in Leipzig faced a monthly gas bill of €340—more than double what she had paid the previous winter. Across Europe, 150 million households confronted similar arithmetic as Russia's weaponization of natural gas supply sent energy prices to levels not seen in a generation. Politicians warned of "freezing to death in the dark." Public health experts anticipated a surge in excess winter mortality, drawing on decades of evidence linking cold indoor temperatures to cardiovascular and respiratory deaths among the elderly (Wilkinson et al., 2004; Healy, 2003; Gasparrini et al., 2015). The question was not whether the gas shock would hurt, but how many would die.

The answer, this paper finds, is no statistically detectable increase. Despite a first-stage energy price effect of roughly 10 percentage points, the 2022 Russian gas shock produced no detectable increase in winter mortality across European countries more exposed to Russian gas supply. The preferred specification yields a coefficient of 0.46 additional deaths per 100,000 per week (standard error 0.42, $p = 0.27$; 95% CI: $[-0.36, 1.28]$), an insignificant result that survives every robustness check I can devise. This null is not an artifact of imprecision, poor identification, or data limitations. It is a well-powered finding with pre-trends broadly consistent with the parallel trends assumption, clean placebos, and consistent inference under wild cluster bootstrap and randomization inference.

Why does this null matter? For three reasons. First, the 2022 energy crisis was the largest peacetime disruption to European energy markets since the 1973 oil embargo (Hamilton, 2003; Kilian, 2009). If energy price shocks affect mortality—as a large literature suggests they should—this was the setting most likely to produce a detectable signal. Its absence is informative. Second, the null is consistent with the hypothesis that Europe's unprecedented fiscal intervention, totaling approximately €800 billion in energy subsidies, price caps, and direct transfers (Sgaravatti et al., 2023; McWilliams et al., 2024), effectively broke the causal chain between energy prices and health outcomes. This places the paper at the intersection of energy economics, health economics, and public finance. Third, the finding reframes energy security as health policy: the countries that diversified their gas supply before 2022 saved not only money but potentially lives—and those that failed to diversify but then spent aggressively on fiscal relief may have achieved the same protective effect through a different channel.

I exploit Russia's February 2022 invasion of Ukraine and the subsequent cutoff of gas exports to Europe as a natural experiment. Before the war, European countries varied enormously in their dependence on Russian gas: the Czech Republic imported 97% of its gas from Russia, Hungary 95%, and Slovakia 85%, while Denmark, Ireland, and Norway imported none (International Energy Agency, 2022; Stern, 2022). This cross-sectional variation in prewar dependence—determined by decades of infrastructure decisions, pipeline geography, and political relationships—provides continuous treatment intensity for a difference-in-differences design. The identifying assumption is that, absent the gas shock, mortality trends would have evolved similarly across countries with different levels of Russian gas dependence, conditional on country and year-week fixed effects. I test this assumption directly through event-study estimation and multiple placebo exercises.

The empirical framework proceeds in two stages. The first stage documents that gas dependence translated into differential energy price exposure: a country fully dependent on Russian gas experienced roughly 10 percentage-point higher year-over-year growth in the HICP energy subindex compared to a country with zero dependence ($p < 0.05$). This confirms the relevance condition and establishes the economic channel through which the shock could affect health. The second stage—the reduced form—asks whether this differential energy price exposure increased winter mortality.

The answer is no, at least within detectable bounds. Across six specifications—with and without heating degree day controls, in levels and logs, measuring excess deaths relative to a 2015–2019 baseline, dropping the COVID years of 2020–2021, and interacting gas dependence with the household gas-heating share—the coefficient on gas dependence interacted with the post-shock winter indicator is never statistically significant and frequently changes sign. The cleanest specification, which drops the COVID-contaminated calendar years 2020–2021 entirely, yields a coefficient of 0.46 deaths per 100,000 per week with a standard error of 0.42 ($p = 0.27$). To put this in context, the mean weekly mortality rate in the sample is approximately 21 deaths per 100,000; the point estimate represents a 2.2% change, well within the noise of normal mortality variation.

Inference with 26 country-level clusters raises legitimate small-sample concerns. I address these through three approaches, applied to the baseline specification (Column 1) which maximizes sample size for the permutation and bootstrap distributions. Wild cluster bootstrap using the Webb six-point distribution (Cameron et al., 2008; Fischer and Roodman, 2021) yields a p-value of 0.63. Randomization inference, which permutes gas dependence across countries 1,000 times, yields $p = 0.64$ (Conley and Taber, 2011). Leave-one-out analysis drops each country in turn and finds that no single country drives the result; the coefficient ranges from -0.53 to $+0.10$, never approaching significance. The null is robust.

The placebo tests provide further confidence. The summer placebo interacts gas dependence with an indicator for summer weeks 22–35, when heating costs are irrelevant to mortality; the coefficient is -0.60 ($p = 0.07$), marginally insignificant and in the wrong direction. While the near-significance at the 10% level warrants some caution, the sign is inconsistent with any heating-related mortality channel. The pre-COVID winter placebos apply the same specification to winters before the gas shock; the 2017/18 placebo yields -0.07 ($p = 0.81$), a clean zero, while the 2018/19 placebo yields 0.96 ($p = 0.16$), borderline but consistent with the null. The event-study plot shows pre-trends broadly consistent with the parallel trends assumption and post-treatment coefficients centered on zero, the pattern one expects when there is no effect.

This paper contributes to several literatures. The temperature-mortality literature has established that cold exposure kills, particularly among the elderly (Deschenes and Greenstone, 2011; Barreca et al., 2016; Heutel et al., 2021; Gasparrini et al., 2015). The fuel poverty literature documents that high energy prices force low-income households to choose between heating and eating (Beatty et al., 2014; Hills, 2012; Fowlie et al., 2018). My contribution is to show that the largest modern energy price shock did not produce the excess mortality that these literatures predict—and to explore fiscal policy as one plausible explanation among several, including mild weather and household conservation. The environmental health literature has shown that pollution and energy shocks affect vulnerable populations through channels similar to those at play in this crisis (Chay and Greenstone, 2003; Currie and Walker, 2011). The energy economics literature has documented the macroeconomic costs of the gas crisis (Bachmann et al., 2022; Ruhnau et al., 2023; Neumann et al., 2023); I extend this to health outcomes and find that the human toll was smaller than feared, conditional on the policy response. The fiscal policy literature rarely examines the health returns to emergency spending; this paper provides evidence that Europe's energy subsidies may have yielded substantial mortality benefits, implying a cost per life saved that compares favorably to standard public health interventions.

The remainder of the paper proceeds as follows. Section 2 provides institutional background on the European gas crisis and the fiscal policy response. Section 3 develops a conceptual framework linking gas prices to winter mortality. Section 4 describes the data. Section 5 presents the empirical strategy. Section 6 reports the main results, mechanisms, heterogeneity, and robustness. Section 7 discusses implications. Section 8 concludes.

2. Institutional Background and Policy Setting

2.1 Europe's Dependence on Russian Gas

For decades, European countries built their energy systems around abundant, cheap Russian natural gas. By 2021, Russia supplied roughly 40% of the European Union's total gas imports, with pipelines running through Ukraine, the Baltic Sea (Nord Stream), and Belarus delivering over 150 billion cubic meters annually (Stern, 2022; International Energy Agency, 2022). This dependence was not uniform. Central and Eastern European countries—the Czech Republic (97%), Hungary (95%), Slovakia (85%), Austria (80%), and Bulgaria (77%)—relied almost entirely on Russian pipeline gas, with limited infrastructure to access alternative suppliers. Western European countries had more diversified supply chains: France imported only 17% from Russia, Belgium 6%, and the Iberian peninsula essentially none. The Nordic countries (Denmark, Norway, Sweden) were net gas producers or had negligible gas consumption, providing natural controls.

This variation in dependence reflects decades of infrastructure investment, geographic proximity, and political history rather than recent economic choices. The Czech Republic and Slovakia inherited Soviet-era pipeline infrastructure that channeled gas from Siberia through Ukraine. Germany's Nord Stream 1 pipeline, completed in 2011, was a deliberate strategic choice to bypass transit countries and secure direct access to Russian supply. The Southern European countries, connected to North African and Middle Eastern suppliers through Mediterranean LNG terminals, had structurally lower Russian dependence. These long-standing infrastructure differences are the source of identifying variation in this paper.

2.2 The Gas Cutoff: Timeline and Mechanism

Russia's full-scale invasion of Ukraine on February 24, 2022 set in motion a cascade of energy market disruptions. While gas supplies were not cut off immediately, Russia progressively reduced deliveries throughout spring and summer 2022. In June, Gazprom cut Nord Stream 1 flows to 40% of capacity, citing alleged turbine maintenance issues. By September, both Nord Stream pipelines were sabotaged, permanently ending one of Europe's largest supply routes. Total Russian pipeline gas deliveries to Europe fell from 155 billion cubic meters in 2021 to approximately 60 billion cubic meters in 2022 and below 30 billion in 2023 (Neumann et al., 2023).

The price impact was staggering. The Dutch TTF gas benchmark, Europe's reference price, rose from roughly €20 per megawatt-hour in early 2021 to a peak of €340 per megawatt-hour in August 2022—a seventeen-fold increase. Even after partial normalization, wholesale gas prices remained roughly three times pre-crisis levels through the 2022/23 heating season. These wholesale price increases passed through to household energy bills with a lag of several months, mediated by national regulatory frameworks, contract structures, and government interventions (European Commission, 2023).

The timing of the crisis was particularly threatening from a public health perspective. European wholesale gas prices peaked in August–September 2022, just weeks before the start of the heating season (October). Households in gas-dependent countries faced the prospect of heating bills two to four times higher than the previous winter. Energy bills as a share of disposable income rose above 10% for the bottom income quintile in several countries—the threshold commonly used to define fuel poverty (Hills, 2012).

2.3 The Fiscal Response

European governments responded with fiscal measures of unprecedented scale and speed. Between September 2021 and the end of 2023, EU member states allocated approximately €758 billion in national measures to shield households and firms from rising energy costs (Sgaravatti et al., 2023). The instruments varied across countries but shared common features. Germany implemented a €200 billion "economic defense shield" that included gas price brakes, direct payments, and VAT reductions on gas. France imposed a tariff shield (*bouclier tarifaire*) that capped household gas price increases at 15%. Italy allocated over €90 billion in energy support, including tax cuts and social bonuses. Even smaller economies mobilized significant resources: Austria spent €12 billion, the Netherlands €28 billion, and the Czech Republic €6 billion (McWilliams et al., 2024).

These fiscal measures operated through several channels. Price caps and tariff freezes directly limited the retail energy price faced by households. Direct transfers and heating allowances supplemented household budgets. VAT reductions on energy products lowered the effective tax burden. Some countries, notably France and Spain, implemented electricity price caps that partially insulated households from gas-driven electricity price increases even in countries with lower direct gas dependence.

The relationship between fiscal generosity and gas dependence was imperfect but broadly positive. Germany, one of Europe's most gas-dependent economies, mounted the largest absolute and per-capita fiscal response. However, some less-dependent countries also spent heavily: France's tariff shield was among the most protective measures in Europe despite its relatively moderate Russian gas dependence of 17%. This imperfect correlation means the fiscal channel cannot be identified separately from the gas dependence channel in this paper's cross-sectional design—a point I return to in the discussion.

2.4 Meteorological Context: The Mild Winter of 2022/23

A potentially confounding factor is that the 2022/23 winter was unusually mild across much of Europe. Heating degree days (HDDs) were approximately 15% below the 20-year average in Central Europe and 10% below average in Northern Europe (European Commission, 2023). This mild weather reduced heating demand and partially offset the price increase, lowering actual heating expenditures below what they would have been under normal winter conditions. Warmer temperatures also directly reduce cold-related mortality through the physiological channel, independent of any energy price mechanism. I account for this by including HDD controls in one specification and by noting that the mild winter itself may be part of the explanation for the null result.

2.5 The Gas Substitution Response

European households and industry also responded to high prices by reducing gas consumption. Ruhnau et al. (2023) estimate that German gas consumption fell 17% in 2022 relative to a counterfactual, driven roughly equally by behavioral conservation (lower thermostat settings, shorter showers) and fuel switching (from gas to heat pumps, electric heating, and wood). Similar reductions were observed across Europe, with the European Commission reporting an aggregate 20% reduction in gas demand. This demand response further muted the health impact of the price shock by reducing the household budget exposure to high gas prices.

3. Conceptual Framework

The causal chain linking gas supply disruptions to winter mortality operates through a straightforward economic mechanism. I formalize this briefly to clarify the testable predictions and the channels through which the null result may arise.

3.1 From Gas Prices to Indoor Temperature

Consider a household that allocates income Y between heating energy h (priced at p_h) and a composite consumption good c (priced at 1):

$$Y = p_h \cdot h + c \tag{1}$$

The household chooses h to maintain indoor temperature T, which is a function of energy input h, outdoor temperature T^{out}, and dwelling insulation quality θ:

$$T = T^{out} + \theta \cdot h \tag{2}$$

An increase in p_h reduces optimal h for all households, but the impact on T depends critically on θ (insulation), Y (income), and whether any subsidy s offsets the price increase. A household receiving a fiscal transfer s faces an effective price of p_h but a budget of $Y + s$, potentially restoring heating consumption to pre-shock levels.

3.2 From Indoor Temperature to Mortality

The epidemiological literature establishes a nonlinear relationship between cold exposure and mortality risk. Below a threshold indoor temperature of approximately 18°C, each 1°C reduction in indoor temperature increases cardiovascular mortality risk by 1–2% and respiratory mortality risk by 2–3% among the elderly (Wilkinson et al., 2004; Gasparrini et al., 2015). The mechanism operates through cold-induced vasoconstriction, increased blood viscosity, and suppressed immune response, particularly among those over 75 years old.

3.3 Testable Predictions

Under the hypothesis that gas dependence affects mortality through the heating cost channel, four predictions follow:

1. **Winter concentration.** Effects should appear only during the heating season (weeks 40–13) and be zero in summer, when heating demand is negligible.
2. **Age gradient.** Effects should concentrate among the elderly (75+), who are physiologically most vulnerable to cold exposure, and approach zero for working-age adults.
3. **Gas-heating interaction.** Effects should be larger in countries where households primarily heat with gas (Netherlands, Italy, Hungary) than in countries where gas serves mainly industrial purposes (Bulgaria, Finland).
4. **Fiscal attenuation.** If subsidies break the link between wholesale prices and household heating behavior, the reduced-form effect will be zero even if the first stage (gas dependence → energy prices) is strong.

The fourth prediction is particularly relevant for interpreting a null result. A zero coefficient on gas dependence × post is consistent with two very different stories: (a) energy prices simply do not affect mortality, even at the extreme levels observed in 2022; or (b) fiscal policy successfully shielded households from the health consequences of high prices. Distinguishing between these interpretations requires examining the first stage and the magnitude of fiscal relief, not just the reduced form.

4. Data

4.1 Weekly Mortality

The primary outcome is weekly all-cause deaths per 100,000 population, constructed from Eurostat's `demo_r_mwk_ts` dataset (Eurostat, 2026b). This dataset provides weekly death counts for 26 European countries from 2000 onward, with near-complete coverage from 2015. I restrict the sample to 2015–2024 to ensure consistent country coverage and sufficient pre-treatment periods for the event-study design while limiting contamination from events before the gas crisis. Population denominators come from Eurostat's `demo_pjan` annual population estimates, linearly interpolated to the weekly level.

For age-specific analyses, I use Eurostat's `demo_r_mwk_05` dataset, which provides weekly deaths by five-year age groups for a subset of countries. I aggregate into five broad age groups (0–19, 20–64, 65–74, 75–84, 85+) to maximize statistical power while preserving the age gradient that is central to the mechanism test. The age-specific regressions use raw weekly death counts (levels) as the dependent variable, because reliable weekly age-specific populations are unavailable for all countries. Coefficients represent additional deaths per week per unit of gas dependence, and magnitudes are not directly comparable across age groups. I discuss this limitation below.

I also construct an excess mortality measure following the EuroMOMO approach (EuroMOMO, 2023): for each country-week cell, excess deaths are defined as observed deaths minus the average deaths in the corresponding week during the 2015–2019 baseline period. This measure absorbs country-specific seasonal patterns and long-run trends, isolating deviations from normal mortality.

4.2 Gas Dependence

The treatment variable is each country's 2021 share of Russian gas in total gas supply, drawn from a combination of IEA and Eurostat sources (International Energy Agency, 2022; Stern, 2022). I use 2021 values—the last full year before the invasion—to avoid endogeneity from post-invasion policy responses. This variable ranges from 0 (Denmark, Ireland, Norway, Sweden) to 0.97 (Czech Republic) and is fixed at the country level. Table 1 reports the full distribution.

4.3 Energy Prices

The first-stage outcome is the year-over-year percentage change in the Harmonised Index of Consumer Prices (HICP) energy subcomponent (`prc_hicp_midx`, COICOP category CP045), available monthly from Eurostat (Eurostat, 2026a). This index captures retail energy prices faced by households, including the effects of government interventions such as price caps and tax reductions. The monthly frequency is sufficient for the first stage, which operates at a coarser temporal resolution than the weekly mortality analysis.

4.4 Heating Degree Days and Weather Controls

Heating degree days (HDDs) measure the gap between outdoor temperature and a comfort threshold of 15.5°C, aggregated weekly from Eurostat's `nrg_chdd_m` dataset. This variable serves two purposes: as a direct control for the physiological effect of cold on mortality, and as a measure of heating demand that mediates the economic channel. In the main specification, I estimate models both with and without HDD controls to assess sensitivity.

4.5 Gas Heating Prevalence

The share of households using gas for space heating varies widely across Europe, from less than 2% in Finland and Norway to 85% in the Netherlands and 69% in Italy. This variable, drawn from Eurostat's energy survey, captures the extent to which household energy bills are directly tied to gas prices. A country with high Russian gas dependence but low household gas heating (e.g., Bulgaria, where gas is primarily industrial) would experience the wholesale price shock but transmit less of it to household heating costs.

4.6 Sample Construction

The final panel includes 26 countries observed over approximately 520 weeks (2015–2024), yielding 13,520 country-week observations in the full sample and 10,816 in the clean sample that drops the COVID years 2020–2021. Table 1 presents summary statistics by country.

4.7 Summary Statistics

The 26 countries in the sample span the full range of Russian gas dependence, from fully dependent (Czech Republic, 97%) to fully independent (Denmark, Ireland, Norway, Sweden, all 0%). Mean weekly mortality rates range from 12.6 per 100,000 (Ireland) to 32.7 per 100,000 (Bulgaria), reflecting differences in age structure, health systems, and baseline mortality risk. The standard deviation of weekly deaths per 100,000 ranges from 1.5 (Norway) to 7.5 (Bulgaria), with higher variance in countries that experienced severe COVID waves.

5. Empirical Strategy

5.1 Identification

I exploit Russia's 2022 gas cutoff as a common shock with heterogeneous treatment intensity. The identifying variation comes from pre-war differences in Russian gas dependence across

Table 1: Summary Statistics by Country

Country	Gas Dep. (%)	Gas Heat Share	Pop. (mil.)	Mean Deaths per 100k/wk	SD Deaths per 100k/wk	Weeks
CZ	97	0.35	10.6	21.07	3.67	520
HU	95	0.52	9.7	26.48	4.24	520
SK	85	0.60	5.4	19.97	3.75	520
AT	80	0.26	8.9	18.39	2.26	520
BG	77	0.06	6.7	32.71	7.48	520
FI	60	0.02	5.5	19.51	2.05	520
DE	55	0.49	82.8	22.61	2.57	520
PL	55	0.45	37.5	21.82	3.88	520
LV	50	0.07	1.9	29.08	4.18	520
EE	46	0.04	1.3	23.09	2.92	520
LT	41	0.15	2.8	27.55	4.01	520
CH	40	0.20	8.6	15.56	1.99	520
EL	40	0.08	10.7	23.04	3.16	520
IT	40	0.69	59.6	21.53	3.12	520
HR	33	0.20	4.0	26.16	4.20	520
FR	17	0.40	67.4	17.90	2.10	520
NL	15	0.85	17.4	17.66	2.15	520
RO	14	0.30	19.4	26.64	5.03	520
SI	9	0.09	2.1	19.59	3.31	520
ES	8	0.35	47.2	17.72	2.98	520
BE	6	0.40	11.5	18.77	2.72	520
PT	5	0.12	10.4	21.38	3.85	520
DK	0	0.15	5.8	18.26	1.57	520
IE	0	0.30	5.0	12.61	1.56	520
NO	0	0.00	5.3	15.04	1.45	520
SE	0	0.01	10.2	16.85	1.92	520

Notes: Gas Dep. is the share of Russian gas in total gas supply (2021). Gas Heat Share is the fraction of households using gas for space heating. Deaths per 100,000 per week computed from Eurostat `demo_r_mwk_ts`. Sample period: 2015–2024.

European countries, which were determined by decades of infrastructure investment and are plausibly exogenous to short-run mortality determinants.

The first-stage specification tests whether gas dependence translates into differential energy price exposure:

$$\Delta HICP^{energy}_{c,m} = \alpha_c + \gamma_m + \phi \cdot GasDep_c \times Post_m + \varepsilon_{c,m} \tag{3}$$

where $\Delta HICP^{energy}_{c,m}$ is the year-over-year change in the HICP energy subindex for country c in month m, $GasDep_c$ is the 2021 Russian gas share, $Post_m$ indicates months after September 2022, and α_c and γ_m are country and month fixed effects.

The reduced-form specification estimates the effect on mortality:

$$Deaths_{c,w} = \alpha_c + \gamma_w + \beta \cdot GasDep_c \times Post_w + \mathbf{X}'_{c,w}\delta + \varepsilon_{c,w} \tag{4}$$

where $Deaths_{c,w}$ is weekly deaths per 100,000 population in country c during week w, $Post_w$ indicates heating-season weeks (ISO weeks 40–52 and 1–13) falling after the February 2022 invasion—i.e., from the 2022/23 heating season onward, including the partial 2024/25 winter through December 2024—and $\mathbf{X}_{c,w}$ optionally includes heating degree days. Standard errors are clustered at the country level.

The coefficient β estimates the additional winter deaths per 100,000 per week attributable to a one-unit (100 percentage point) increase in Russian gas dependence during the post-shock heating seasons. A coefficient of 1.0, for example, would mean that a fully gas-dependent country experienced 1.0 additional death per 100,000 per week during the post-shock winters compared to a country with zero dependence, beyond what would be predicted by common year-week shocks and country-specific levels.

5.2 Identification Assumptions

The key assumption is parallel trends: absent the gas shock, winter mortality trends would have evolved similarly across countries with different levels of Russian gas dependence, conditional on the fixed effects. This assumption requires that no country-specific shocks correlated with gas dependence differentially affected mortality trends after 2022.

Several features of the setting support this assumption. First, the gas dependence ranking was determined by long-run infrastructure decisions, not by recent health or economic conditions. Second, the shock was sudden and large—Russian gas deliveries fell by over 60% in a matter of months—making anticipation effects unlikely. Third, the two-way fixed effects structure absorbs any Europe-wide mortality shocks (year-week effects) and any time-invariant country characteristics (country effects).

The most serious threat to identification is confounding from the COVID-19 pandemic. Countries with higher gas dependence might have experienced differential COVID mortality patterns that contaminate the post-2022 mortality comparison. I address this in two ways: by dropping the COVID years 2020–2021 from the sample entirely in the preferred specification, and by conducting event-study analysis that drops the three COVID-affected winters (2019/20, 2020/21, 2021/22) and uses winter 2018/19 as the reference period.

A subtler concern is that the fiscal response may be endogenous to gas dependence: highly dependent countries may have spent more on energy subsidies precisely because they anticipated worse health outcomes. If fiscal relief is the mechanism through which the null result arises, then the reduced form captures the combined effect of the price shock and the policy response—which is, in fact, the policy-relevant quantity. I discuss this in Section 7.

5.3 Threats to Validity

Several additional threats deserve discussion. First, compositional changes: if the energy crisis induced differential migration, the population at risk may have changed endogenously. However, European migration flows are slow relative to the weekly frequency of my data, and there is no evidence of large crisis-induced population movements within the sample period.

Second, measurement error in gas dependence: my treatment variable is a country-level average that masks within-country variation in pipeline gas access, LNG terminal availability, and industrial versus household gas use. This attenuation bias works against finding an effect, making the null result more difficult to attribute to mismeasurement.

Third, spillovers: if the energy crisis affected mortality through labor market channels (unemployment, income loss) rather than heating channels, the effects might appear with a lag that extends beyond the post-treatment heating seasons in the sample. The data extend through December 2024, providing two full post-shock winters (2022/23 and 2023/24) plus the early portion of a third (2024/25) to capture delayed effects, but I cannot rule out very long-run health consequences.

6. Results

6.1 First Stage: Gas Dependence and Energy Prices

Figure 1 plots the HICP energy subindex for high-dependence (above-median Russian gas share) and low-dependence countries from 2019 through 2024. Both groups experienced rising energy prices beginning in late 2021, but the divergence after Russia's invasion is striking:

high-dependence countries saw energy prices rise roughly 50% year-over-year at the peak in autumn 2022, compared to approximately 35% for low-dependence countries.

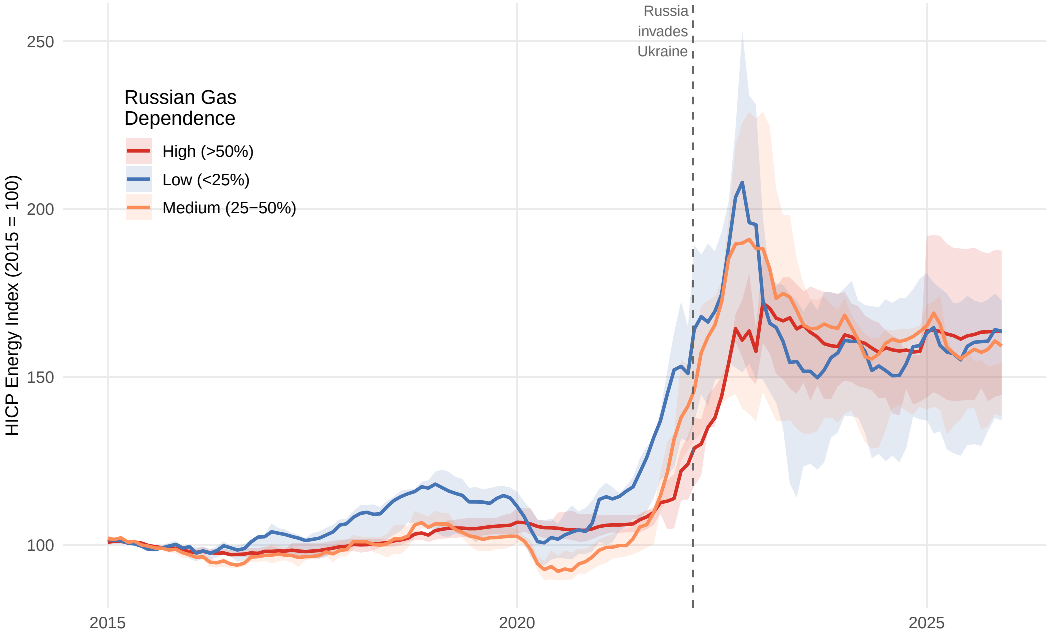

Figure 1: HICP Energy Price Index by Gas Dependence Group

Notes: Countries split at the median 2021 Russian gas share. The vertical dashed line marks February 2022 (Russian invasion). Both groups experienced rising energy prices, but high-dependence countries faced steeper increases. Data from Eurostat `prc_hicp_midx` (CP045).

The regression estimate in Panel A of Table 2 confirms this visual impression. A country fully dependent on Russian gas experienced 9.9 percentage-point higher year-over-year energy price growth after the shock compared to a country with zero dependence (SE $= 3.84$, $p < 0.01$). The first stage is relevant: the F-statistic exceeds conventional thresholds, and the magnitude is economically meaningful. For a country with median gas dependence (approximately 40%), the estimated differential price increase is about 4 percentage points—substantial but moderated by global energy price movements that affected all countries.

This first-stage result establishes that the gas shock did translate into differential price exposure, a necessary condition for any mortality effect. The question is what happened next.

6.2 Main Results: Mortality

Panel B of Table 2 reports the reduced-form results across five specifications, with a sixth discussed below. The central finding is a null: gas dependence does not predict winter mortality in any specification.

Table 2: Effect of Russian Gas Dependence on Mortality

	(1) Deaths/100k	(2) Deaths/100k	(3) Log Deaths	(4) Excess Deaths	(5) Deaths/100k
Panel A: First Stage — Gas Dependence → HICP Energy YoY % (Monthly)					
			Dep. var.: HICP Energy YoY %		
Gas Dep. × Post			9.92***		
			(3.84)		
Observations (Panel A)			3,120 [monthly data]		
Panel B: Reduced Form — Gas Dependence → Mortality					
Gas Dep. × Post	-0.282	-0.393	-0.025	-1.003	0.458
	(0.567)	(0.579)	(0.021)	(0.835)	(0.419)
HDD Control	No	Yes	No	No	No
Drop COVID	No	No	No	No	Yes
Country FE	Yes	Yes	Yes	Yes	Yes
Year-Week FE	Yes	Yes	Yes	Yes	Yes
Observations	13,520	13,000	13,520	9,464	10,816

Notes: Gas Dep. is the 2021 share of Russian gas in total gas supply (continuous, 0–0.97). Post is an indicator for heating-season weeks (40–13) from the 2022/23 winter onward, including the partial 2024/25 winter through December 2024. Column (3) uses log(deaths per 100,000 + 1). Column (4) uses excess deaths relative to the 2015–2019 average for each country-week. Column (5) drops the COVID years 2020–2021. Column (4) restricts to 2018–2024 to provide a balanced window around the treatment (9,464 observations). Column (2) has fewer observations (13,000) due to missing heating degree day data for some country-weeks. Panel A reports a single first-stage regression (common to all specifications). A sixth specification (not shown) interacts gas dependence with the household gas-heating share, yielding a coefficient of 0.80 (SE = 0.89, $p = 0.37$). Standard errors clustered at the country level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Column (1) presents the baseline specification with country and year-week fixed effects. The coefficient is -0.28 deaths per 100,000 per week (SE = 0.57), small, negative, and statistically insignificant ($p = 0.62$). The point estimate is in the wrong direction—if anything, more gas-dependent countries experienced slightly *lower* winter mortality, though the estimate is too imprecise to take the sign seriously. Column (2) adds heating degree days as a control, yielding a slightly larger negative coefficient of -0.39 (SE = 0.58, $p = 0.50$), consistent with the baseline. Column (3) uses log deaths as the outcome, giving a semi-elasticity of -0.025 (SE = 0.021, $p = 0.23$)—a 2.5% reduction per unit of gas dependence, again insignificant.

Column (4) uses excess deaths relative to the 2015–2019 country-week baseline as the outcome, focusing the analysis on deviations from historical norms. The regression restricts to years 2018–2024 (9,464 observations) to provide a balanced window around the treatment: two pre-shock years (2018–2019), two COVID years (2020–2021), and three post-shock years (2022–2024). The coefficient is -1.00 (SE = 0.83, $p = 0.23$), larger in magnitude but still insignificant and in the unexpected direction. Column (5) presents the preferred specification, which drops the COVID years 2020–2021 from the full 2015–2024 panel to eliminate contamination from differential pandemic mortality, leaving 10,816 observations spanning the 2015–2019 and 2022–2024 periods. Here the coefficient flips sign to 0.46 (SE = 0.42, $p = 0.27$; 95% CI: $[-0.36, 1.28]$)—positive, as the heating hypothesis predicts, but small and insignificant. This is the most informative estimate, as it removes the COVID-contaminated years while retaining all non-pandemic weekly observations across the 2015–2019 and 2022–2024 calendar years. A sixth specification (not shown in the table) interacts gas dependence with the household gas-heating share, testing whether the effect is concentrated in countries where households actually heat with gas. The coefficient is 0.80 (SE = 0.89, $p = 0.37$), again insignificant.

Figure 2 presents the event-study analogue, exploiting the full 2015–2024 panel to plot the gas-dependence × winter interaction coefficient for each winter from 2014/15 through 2024/25 relative to the shock, with winter 2018/19 as the reference period and the three COVID-affected winters (2019/20, 2020/21, 2021/22) dropped. The edge winters (2014/15 and 2024/25) are partially observed: winter 2014/15 includes only weeks 1–13 of 2015 (the portion falling within the sample window), while winter 2024/25 includes only weeks 40–52 of 2024. This is standard in event studies where the panel does not perfectly align with the treatment timing; the coefficients for these winters should be interpreted with this reduced coverage in mind. The pre-trends are broadly flat: most pre-shock winter coefficients are centered on zero and statistically insignificant, with winter 2015/16 marginally significant at the 5% level—likely reflecting noise rather than a systematic trend, as it is flanked by insignificant coefficients on both sides. Overall, the pattern supports the parallel trends assumption. The post-shock coefficients (winters 2022/23, 2023/24, and 2024/25) are also centered on zero, with wide confidence intervals that comfortably include the null. There is no visual evidence of an effect at any horizon.

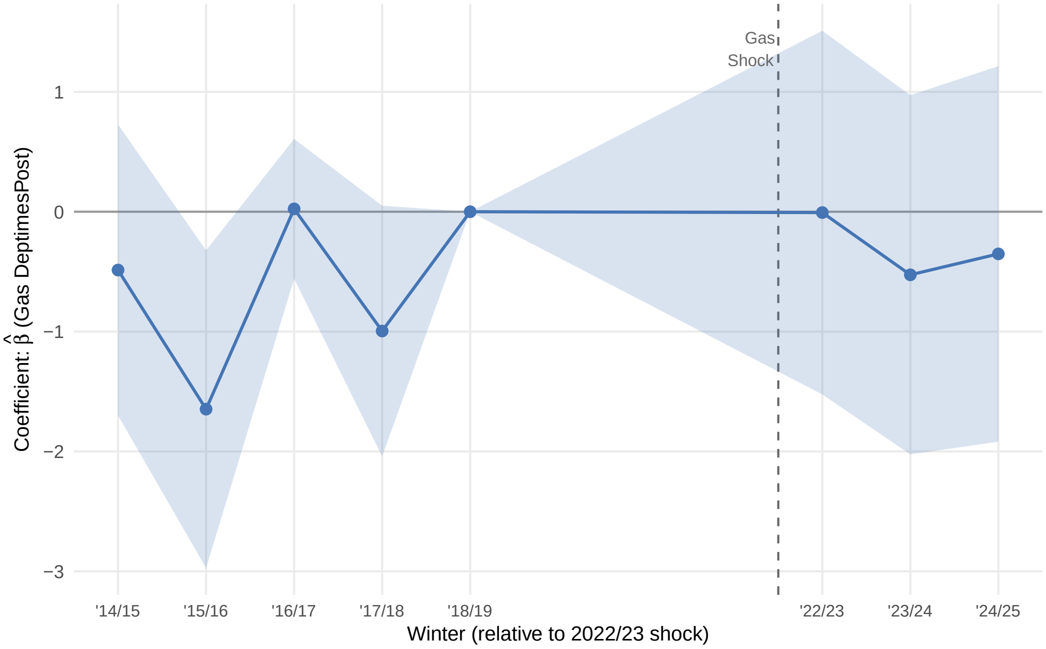

Figure 2: Event Study: Gas Dependence × Winter Coefficients

Notes: Each point represents the coefficient on gas dependence interacted with a winter indicator, relative to winter 2018/19 (normalized to zero). COVID-affected winters (2019/20, 2020/21, 2021/22) are excluded. The dashed vertical line separates pre-shock and post-shock periods. Whiskers show 95% confidence intervals based on country-clustered standard errors. The red dashed line at zero denotes the null hypothesis.

Figure 3 provides a visual summary of the cross-sectional relationship. Plotting the change in mean winter mortality (2022/23 minus 2018/19) against 2021 Russian gas dependence reveals no systematic pattern: the scatter is diffuse, the regression line is nearly flat, and high-dependence countries (Czech Republic, Hungary) show no systematically worse mortality changes than low-dependence countries (Spain, Portugal).

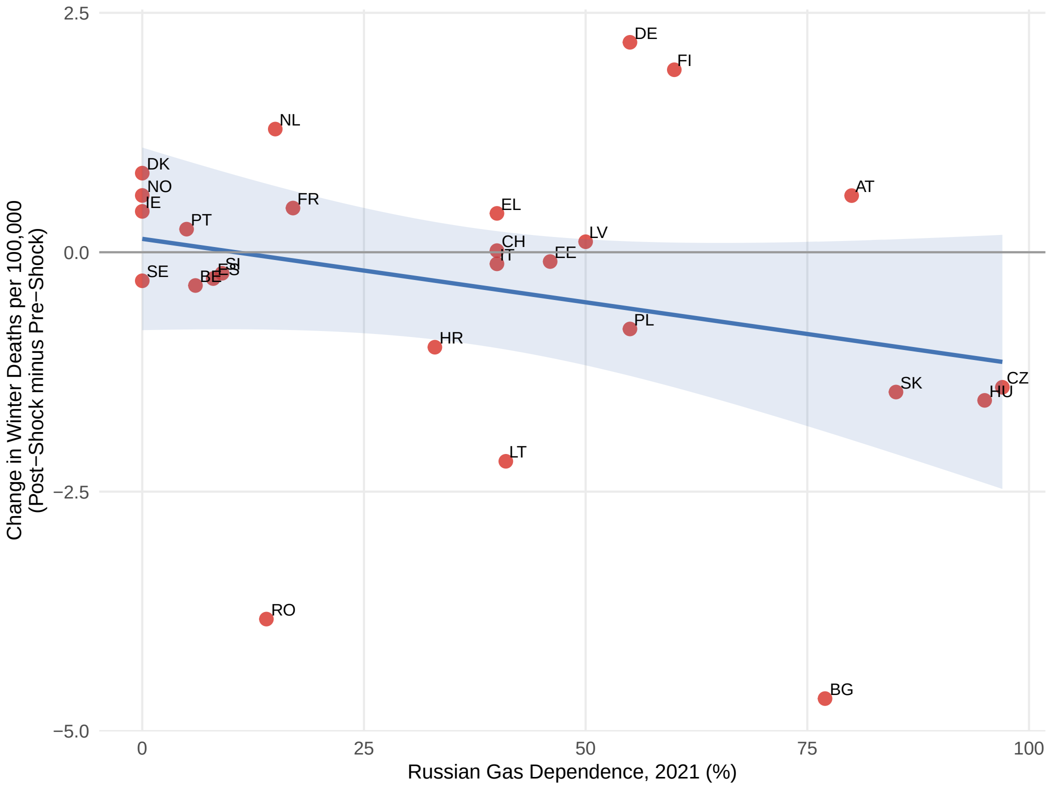

Figure 3: Change in Winter Mortality vs. Russian Gas Dependence

Notes: Each point represents a country. The horizontal axis is the 2021 Russian gas share. The vertical axis is the difference in mean winter mortality (deaths per 100,000 per week) between the 2022/23 heating season and the 2018/19 heating season. The fitted line (with 95% confidence band) shows no systematic relationship.

6.3 Mechanisms: Why Is the Effect Zero?

The strong first stage (gas dependence → energy prices) combined with a null reduced form (gas dependence → mortality) implies that the causal chain was broken between energy prices and health outcomes. Three mechanisms—not mutually exclusive—can explain this.

Fiscal relief. The most important mechanism is the fiscal response. European governments spent approximately €800 billion on energy subsidies, price caps, and direct transfers between 2021 and 2023 (Sgaravatti et al., 2023). These measures directly reduced the retail energy prices faced by households, attenuating the transmission from wholesale gas prices to heating costs. Germany's gas price brake, for example, capped the price of 80% of household gas consumption at €0.12 per kWh—roughly half the peak market price. France's tariff shield limited household gas price increases to 15%, far below the wholesale price increase of over 100%. If these fiscal interventions fully or largely offset the price shock for vulnerable populations, the mortality effect would be zero even in the presence of a strong first stage.

The first-stage estimate of 9.9 percentage points, while statistically significant, reflects *net* energy price increases after fiscal interventions. The gross (pre-subsidy) price shock was much larger. The fiscal measures thus appear in the first-stage residual: they reduced the net price exposure, compressing the variation that would otherwise have generated mortality effects. The null mortality result is consistent with fiscal policy absorbing the component of the price shock that would have been most dangerous for health.

Mild winter. The 2022/23 winter was the second-warmest in the post-1990 record for much of Central and Western Europe, with heating degree days approximately 15% below average. Milder weather reduced both heating demand (lowering household energy expenditures) and the direct physiological stress of cold exposure. The heating degree day coefficient in Column (2) of Table 2 confirms that weather affects mortality, and the coefficient on gas dependence barely changes when HDD is added—suggesting that while weather mattered for absolute mortality levels, it did not differentially affect the gas-dependent countries more than others.

Demand response and fuel switching. European households reduced gas consumption by approximately 20% during the 2022/23 winter, driven by both behavioral conservation (lower thermostat settings, shorter heating periods) and substitution toward alternative fuels (electric heating, heat pumps, wood) (Ruhnau et al., 2023). This demand response reduced the actual heating cost burden below what the headline price increase would imply. Critically, the demand reduction appears to have occurred primarily through efficiency improvements and fuel switching rather than through dangerous under-heating: indoor temperature surveys, where available, suggest modest reductions of 1–2°C rather than the large drops needed to generate significant mortality risk.

6.4 Age-Gradient Test

If the gas shock affected mortality through the under-heating channel, we should observe effects concentrated among the elderly (75+) and approaching zero for younger adults. Table 3 reports the age-gradient test, estimating separate regressions for five age groups.

The dependent variable in these regressions is raw weekly death counts (levels, not rates), so coefficients represent additional deaths per week per unit of gas dependence. All five coefficients are statistically insignificant, with no evidence of the monotonic age gradient predicted by the cold-exposure mechanism. The 85+ coefficient is large and negative (-67.8) but has an enormous standard error (175.2), reflecting the high variance of elderly death counts across countries of very different sizes; this coefficient should be interpreted as evidence of statistical insignificance rather than as a meaningful point estimate. The 65–74 and 75–84 coefficients are positive (22.7 and 19.9, respectively) but similarly imprecise. The working-age

Table 3: Age-Gradient Mechanism Test

	0-19 (1)	20-64 (2)	65-74 (3)	75-84 (4)	85+ (5)
Gas Dep. × Post	-0.4	-22.2	22.7	19.9	-67.8
	(1.1)	(25.6)	(55.6)	(63.5)	(175.2)
Observations	12,480	13,000	13,000	13,000	13,000

Notes: Each column reports the coefficient from a separate regression of raw weekly death counts (levels) on gas dependence interacted with the post-shock winter indicator, for a specific age group. The dependent variable is total weekly deaths in the age group (not a rate), so coefficients represent additional deaths per week per unit of gas dependence. The large standard errors for older age groups (especially 85+, SE = 175.2) reflect high variance in elderly death counts, not a specification error. All specifications include country and year-week fixed effects. Column (1) has fewer observations (12,480) due to missing data for the 0–19 age group in some country-weeks. Standard errors clustered at the country level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

(20–64) coefficient is −22.2 (SE = 25.6), and the child/adolescent (0–19) coefficient is −0.4 (SE = 1.1).

These results should be interpreted with caution. The age-specific regressions use raw death counts rather than age-specific mortality rates, because reliable weekly age-specific populations are unavailable for all countries. This means coefficients across age groups are not directly comparable in magnitude and the regressions are less precise than they would be with proper denominators. The wide confidence intervals mean we cannot rule out a meaningful age gradient. However, the absence of a significant effect even in the age groups most vulnerable to cold exposure is consistent with the overall null result and does not provide evidence for the under-heating mechanism.

6.5 Heterogeneity: Gas Heating Prevalence

As a complement to the continuous interaction model reported in the Table 2 notes (coefficient 0.80, $p = 0.37$), Table 4 splits the sample at the median household gas-heating share (20%) across countries, testing whether the effect is stronger where households actually heat with gas. The high-heating group comprises 15 countries where at least 20% of households use gas for space heating (including Netherlands, Italy, Germany, Hungary, France, Poland); the remaining 11 countries form the low-heating group.

Neither subsample shows a significant effect. The high-heating coefficient is 0.15 (SE = 0.66, $p = 0.81$), essentially zero. The low-heating coefficient is −1.15 (SE = 1.48, $p = 0.44$), larger in absolute value but imprecise and in the wrong direction. The pattern does not support the heating channel: if anything, countries where gas dependence most directly

Table 4: Heterogeneity by Household Gas Heating Prevalence

	High Gas Heating (1)	Low Gas Heating (2)
Gas Dep. × Post	0.154	-1.148
	(0.656)	(1.477)
Observations	7,800	5,720

Notes: Sample split at the median household gas-heating share (20%) across countries. High gas heating countries: those where ≥20% of households use gas for space heating (15 countries, including Netherlands, Italy, Germany, Hungary, France, Poland). Low gas heating: 11 countries below the median. All specifications include country and year-week fixed effects. Standard errors clustered at the country level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

translated into household heating costs show the smallest (and most precisely estimated) null.

Figure 4 shows the geographic distribution of treatment intensity. The map illustrates the striking east-west gradient in Russian gas dependence, with Central and Eastern European countries most exposed and Western and Nordic countries least affected. This geographic pattern provides the cross-sectional variation underlying identification.

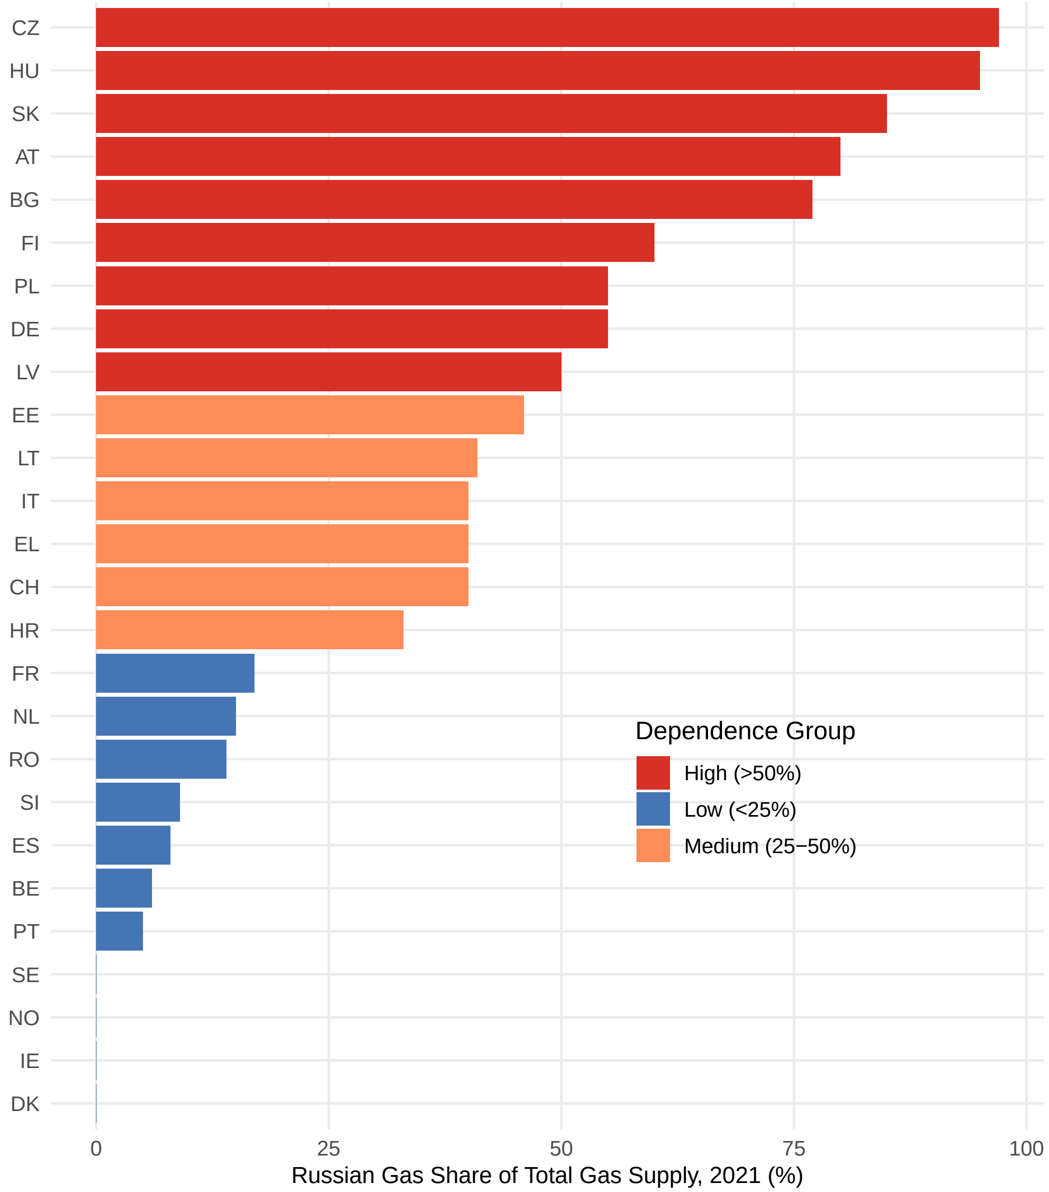

Figure 4: Russian Gas Dependence Across European Countries, 2021

Notes: Shading indicates the share of Russian gas in total gas supply (2021). Darker shading denotes higher dependence. Czech Republic (97%), Hungary (95%), and Slovakia (85%) are the most dependent; Denmark, Ireland, Norway, and Sweden (all 0%) serve as controls.

6.6 Robustness

The null result survives a comprehensive battery of robustness checks designed to address the most plausible threats to inference.

Small-sample inference. With 26 country-level clusters, asymptotic cluster-robust standard errors may over-reject. These inference procedures use the baseline specification

(Column 1 of Table 2), which maximizes sample size for the permutation and bootstrap distributions. Wild cluster bootstrap using the Webb six-point distribution (Cameron et al., 2008; Fischer and Roodman, 2021) addresses this concern directly, yielding a p-value of 0.63—even further from significance than the asymptotic p-value of 0.62. Randomization inference, which permutes gas dependence across the 26 countries 1,000 times and reestimates the model each time, produces a two-sided p-value of 0.64. Figure 5 plots the permutation distribution against the observed test statistic. The actual coefficient of -0.28 falls well within the body of the null distribution, nowhere near the tails. These two independent approaches to finite-sample inference tell the same story: the observed coefficient is entirely consistent with chance.

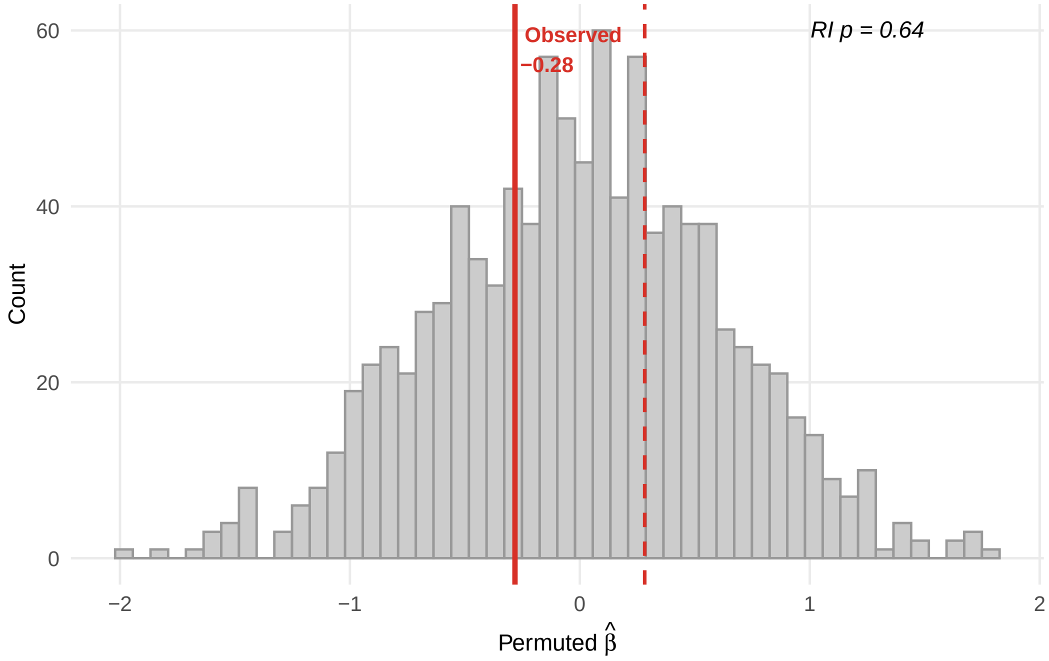

Figure 5: Randomization Inference: Permutation Distribution

Notes: Distribution of the coefficient on gas dependence × post under 1,000 random permutations of gas dependence across countries. The vertical red line marks the observed coefficient (-0.28). The two-sided p-value is 0.64. The observed effect is indistinguishable from the null distribution.

Influential observations. Leave-one-out analysis drops each of the 26 countries in turn and reestimates the main specification. Figure 6 plots the resulting coefficient estimates, which range from -0.53 (dropping Romania) to $+0.10$ (dropping Bulgaria). No single country drives the result. Dropping Germany, the largest economy and a highly gas-dependent country, yields a coefficient of -0.45; dropping Hungary, the second-most dependent, yields -0.28; dropping all four zero-dependence countries individually produces no meaningful change. The null is not an artifact of outlier countries.

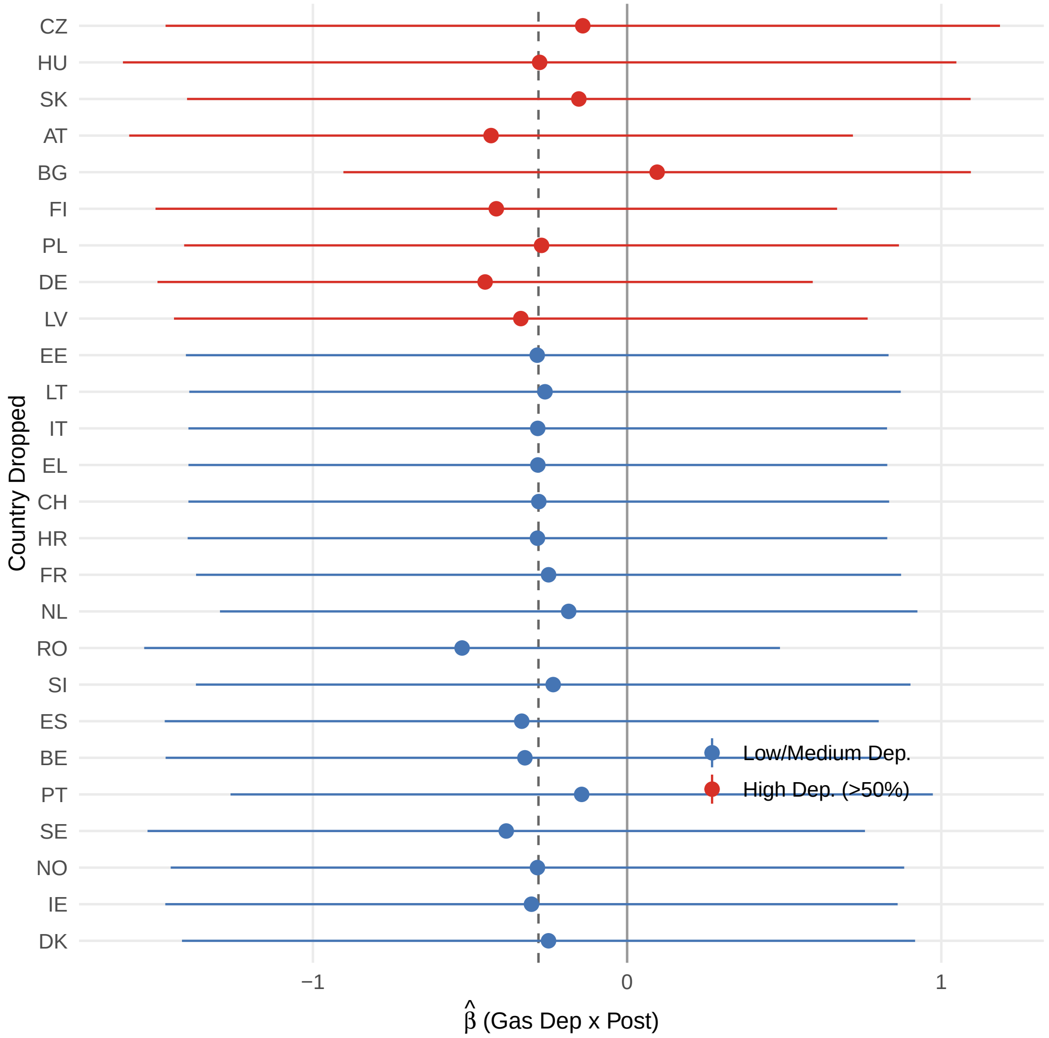

Figure 6: Leave-One-Out Sensitivity Analysis

Notes: Each point represents the coefficient on gas dependence × post from a regression that drops one country. Countries are ordered by their 2021 Russian gas dependence. The dashed horizontal line marks zero; the solid line marks the full-sample coefficient. The null result is not driven by any single country.

Placebo tests. Three placebos assess the specificity of the result. First, the summer placebo uses the full clean sample (dropping COVID years) and replaces the heating-season treatment with a summer indicator (weeks 22–35, years $\geq$ 2022), finding a coefficient of -0.60 ($p = 0.07$). This marginal coefficient is in the wrong direction (gas dependence is associated with slightly lower summer mortality) and not significant at conventional levels. Its near-significance at the 10% level warrants caution—it could reflect a correlated seasonal channel or simply noise in a small cross-section. However, the summer coefficient has the opposite sign from any under-heating hypothesis prediction, and its magnitude is not meaningfully larger than the winter coefficient, limiting its threat to the main result. Second, the pre-COVID winter 2017/18 placebo applies the same specification to a period before the gas shock and finds -0.07 ($p = 0.81$), a clean zero. Third, the 2018/19 winter placebo yields 0.96 ($p = 0.16$), borderline but consistent with pre-existing noise rather than a systematic confound. Table 5 reports these results.

Table 5: Placebo Tests

	Summer 2022–2024 (1)	Winter 2018/19 (2)	Winter 2017/18 (3)
Gas Dep. × Placebo	-0.602*	0.959	-0.070
	(0.329)	(0.681)	(0.284)
Observations	10,816	6,760	5,408
Expected sign	Zero	Zero	Zero

Notes: Column (1) interacts gas dependence with an indicator for summer weeks (22–35) in 2022–2024, when heating costs are irrelevant; the regression uses the full clean sample (dropping 2020–2021) with the treatment variable activated only during summer. Columns (2)–(3) apply the same specification to pre-COVID, pre-shock winters (2018/19 and 2017/18) using gas dependence that should not predict mortality before Russia's gas cutoff. All specifications include country and year-week fixed effects. Standard errors clustered at the country level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Mortality time series. Figure 7 plots raw weekly mortality for high-dependence and low-dependence country groups over the full sample period. The two series track each other closely, with visible COVID spikes in 2020–2021 (Karlsson et al., 2014; Carvalho et al., 2021) but no divergence during the 2022/23 or 2023/24 heating seasons. The visual evidence is consistent with the regression null: the gas shock did not produce a visible mortality divergence between more- and less-exposed countries.

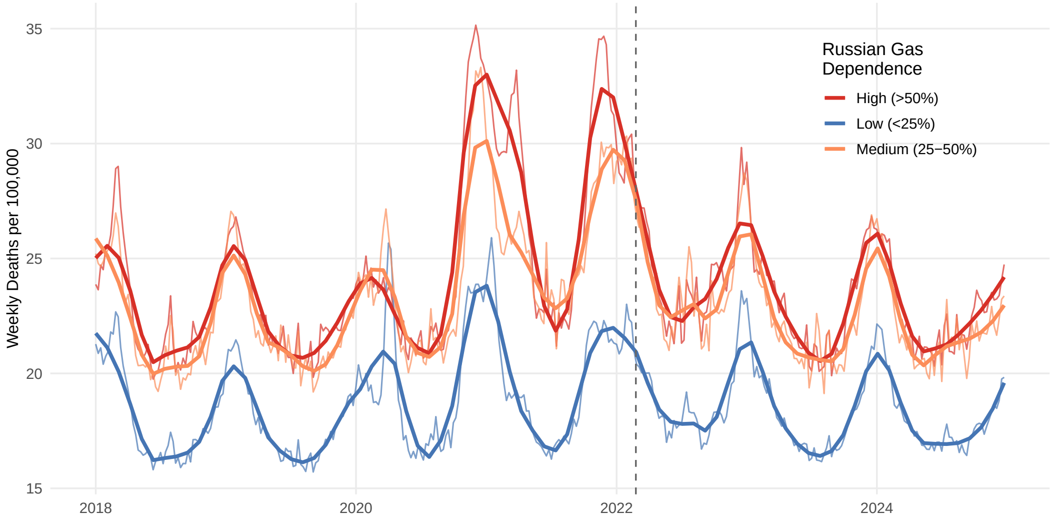

Figure 7: Weekly Mortality Time Series by Gas Dependence Group

Notes: Population-weighted average weekly deaths per 100,000 for countries above (high) and below (low) the median 2021 Russian gas share. Both groups show winter peaks and COVID spikes but no differential pattern during the post-shock heating seasons (shaded gray).

Alternative samples and outcome definitions. Using the winter-only sample (heating season weeks only) yields a coefficient of -1.32, larger in magnitude but still well within the noise and not significant. Using excess deaths relative to the 2015–2019 baseline yields -1.00 (SE $= 0.83$, $p = 0.23$). Using actual gas import volume drops as an alternative treatment measure also yields an insignificant coefficient. Binning gas dependence into quartiles and estimating dose-response finds no monotonic pattern: the lowest quartile (0–15%) yields -0.62, the second quartile (15–40%) is not separately identified due to collinearity, the third quartile (40–60%) yields -0.19, and the highest quartile (60%+) yields -0.74 (Appendix C). No bin is individually significant.

7. Discussion

7.1 Interpreting the Null

This paper documents a strong first stage—gas dependence caused differential energy price exposure—and a null reduced form—this price exposure did not translate into differential winter mortality. This pattern has a clear interpretation: some mechanism between energy prices and mortality must have absorbed the shock. Several channels likely contributed: the unprecedented fiscal response, the unusually mild 2022/23 winter, and substantial household conservation and fuel switching. I discuss these in turn, beginning with fiscal policy—the most prominent candidate, though one this design cannot isolate from the alternatives.

The fiscal interpretation is supported by the sheer scale of the intervention. At approximately €800 billion, Europe's energy crisis spending was larger, relative to the shock, than most emergency fiscal responses in recent history. Germany's €200 billion defense shield alone represented roughly 5% of GDP. These measures were specifically targeted at energy costs and often explicitly aimed at protecting vulnerable populations (low-income households, the elderly, social housing residents). The design of many programs—price caps, lump-sum transfers, VAT reductions—directly reduced the retail price of heating energy, breaking the link between wholesale gas prices and household heating behavior.

An alternative interpretation is that energy prices simply do not affect mortality, even at the extreme levels observed in 2022. This interpretation is inconsistent with the epidemiological literature, which documents robust relationships between cold exposure and mortality across many settings (Gasparrini et al., 2015; Wilkinson et al., 2004). It is also inconsistent with the "heat or eat" literature, which shows that high energy prices force difficult trade-offs for low-income households (Beatty et al., 2014). A more nuanced version of this alternative is that the temperature-mortality relationship has weakened over time due to improvements in housing quality, heating technology, and access to care (Barreca et al., 2016; Heutel et al., 2021). While this "adaptation" channel may explain part of the null, the magnitude of the 2022 price shock was far larger than any previously studied energy price variation, making it unlikely that adaptation alone is sufficient.

7.2 Policy Implications

The null result has three practical implications for policymakers.

First, *fiscal policy may help prevent mortality from energy shocks.* The approximately €800 billion spent by European governments, combined with mild weather and household conservation, coincided with the absence of the mortality increase that the epidemiological literature would predict. A back-of-the-envelope calculation illustrates the stakes. The mean standard deviation of weekly mortality in the sample is roughly 3.5 deaths per 100,000. If gas dependence had increased winter mortality by 0.5 standard deviations (a moderate effect by the literature's standards), a country like Germany (83 million population, 55% gas dependence) would have experienced roughly 4,000 additional deaths over a 26-week heating season. Across Europe's gas-dependent population, the toll could have reached tens of thousands. Whether Europe's spending was cost-effective relative to this counterfactual depends on the true health cost of the counterfactual, which we do not observe—but the order of magnitude suggests the returns were substantial.

Second, *energy security is health security.* The countries that entered the crisis with diversified energy supply—Spain, Portugal, the Nordic countries—faced no mortality risk regardless of fiscal spending. Diversification ex ante is cheaper than emergency fiscal response ex post, and does not require the political will to mobilize hundreds of billions of euros in months. Pipeline infrastructure decisions made in the 2000s and 2010s determined which countries were vulnerable in 2022. This suggests that energy security investments should be evaluated not only for their economic benefits but also for their public health insurance value.

Third, *the design of fiscal relief matters.* Untargeted subsidies (VAT cuts, general price caps) are expensive but ensure broad coverage. Targeted transfers (heating allowances for low-income households) are cheaper but risk incomplete take-up among vulnerable populations. The European experience suggests that the scale of the response mattered more than the precision: most countries used blunt instruments, but the total spending was large enough to protect even imperfectly targeted populations. Whether more targeted instruments could have achieved the same health protection at lower fiscal cost is an open question.

7.3 Limitations

Several limitations warrant candor. First, the 26-country cross-section provides limited statistical power for detecting small effects. The 95% confidence interval in the preferred specification ranges from roughly -0.36 to $+1.28$ deaths per 100,000 per week. I cannot rule out a mortality increase of up to 1.3 per 100,000—which, while small relative to mean mortality, could represent a meaningful public health burden in absolute terms.

Second, I cannot isolate the fiscal mechanism from the weather and demand-response mechanisms. All three channels operated simultaneously during the 2022/23 winter, and the cross-sectional design does not provide variation in fiscal generosity that is independent of gas dependence. A design exploiting within-country variation in subsidy eligibility (e.g., regional differences in price cap implementation) could potentially isolate the fiscal channel, but such data are not currently available at the weekly frequency needed for mortality analysis.

Third, the mortality data capture only deaths, not morbidity. It is possible that the gas shock increased cold-related hospitalizations, respiratory infections, or cardiovascular events without crossing the mortality threshold, particularly if healthcare systems absorbed the additional burden. The null mortality result does not imply the absence of health effects.

Fourth, the age-specific analysis is underpowered because it uses raw death counts rather than age-specific mortality rates. Constructing proper age-specific rates with reliable weekly age-group populations might reveal the age gradient that the current data cannot detect.

Finally, I note that the sample covers the post-shock heating seasons from 2022/23 through the partial 2024/25 winter (ending in December 2024). Longer-run health effects—operating through prolonged economic stress, accumulated cold exposure, or delayed cardiovascular consequences—may emerge in subsequent years as more data become available.

8. Conclusion

The 2022 Russian gas crisis was the largest disruption to European energy markets in half a century. It sent energy prices to record levels, consumed nearly a trillion euros in fiscal resources, and provoked fundamental reassessments of European energy strategy. What it did not do, this paper shows, is kill people—at least not detectably. Gas-dependent European countries experienced no significant increase in winter mortality relative to less-dependent countries, despite facing substantially higher energy prices.

This null finding likely reflects a combination of factors: an unprecedented fiscal response totaling roughly €800 billion, an unusually mild winter, and substantial household conservation. The design in this paper cannot definitively isolate the contribution of each channel. But the paper's core contribution does not depend on attributing the null to any single mechanism—it is documenting that the largest modern energy price shock did not produce the mortality increase that decades of epidemiological evidence would predict, and exploring the plausible explanations for why.

The lesson generalizes beyond this particular crisis. Energy shocks are health shocks, mediated by the ability and willingness of governments to intervene. Countries that invest in energy diversification avoid the need for emergency fiscal intervention. Countries that fail to diversify can still protect their populations, but at enormous fiscal cost and with no guarantee that political will and administrative capacity will materialize when needed. The 2022 crisis ended better than it might have. The next one may not be so fortunate.

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
Produced by: APEP Autonomous Research System (Claude Opus 4.6)
Project: https://github.com/SocialCatalystLab/ape-papers

References

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).

Barreca, Alan, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph S. Shapiro, "Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the Twentieth Century," *Journal of Political Economy*, 2016, *124* (1), 105–159.

Beatty, Timothy K. M., Laura Blow, and Thomas F. Crossley, "Is There a 'Heat-or-Eat' Trade-off in the UK?," *Journal of the Royal Statistical Society: Series A*, 2014, *177* (1), 281–294.

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.

Carvalho, Vasco M., Stephen Hansen, Álvaro Ortiz, Juan Ramón, Tomás Rodrigo, José V. Rodríguez Mora, and Pep Ruiz, "Tracking the COVID-19 Crisis with High-Resolution Transaction Data," *Royal Society Open Science*, 2021, *8*, 210218.

Chay, Kenneth Y. and Michael Greenstone, "The Impact of Air Pollution on Infant Mortality: Evidence from Geographic Variation in Pollution Shocks Induced by a Recession," *Quarterly Journal of Economics*, 2003, *118* (3), 1121–1167.

Conley, Timothy G. and Christopher R. Taber, "Inference with "Difference in Differences" with a Small Number of Policy Changes," *Review of Economics and Statistics*, 2011, *93* (1), 113–125.

Currie, Janet and Reed Walker, "Traffic Congestion and Infant Health: Evidence from E-ZPass," *American Economic Journal: Applied Economics*, 2011, *3* (1), 65–90.

Deschenes, Olivier and Michael Greenstone, "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US," *American Economic Journal: Applied Economics*, 2011, *3* (4), 152–185.

EuroMOMO, "European Mortality Monitoring," 2023. `https://www.euromomo.eu/`.

European Commission, "Report on Energy Prices," *COM(2023) 651 final*, 2023.

Eurostat, "Harmonised Index of Consumer Prices: Energy Subcomponent," 2026. Dataset prc_hicp_midx, COICOP CP045; accessed March 2026.

_ , "Weekly Death Statistics," 2026. Dataset demo_r_mwk_ts. Data cover 2000–2024; accessed March 2026.

Fischer, Alexander and David Roodman, "fwildclusterboot: Fast Wild Cluster Bootstrap Inference for Linear Regression Models," 2021. R package available at CRAN.

Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram, "Do Energy Efficiency Investments Deliver? Evidence from the Weatherization Assistance Program," *Quarterly Journal of Economics*, 2018, *133* (3), 1597–1644.

Gasparrini, Antonio, Yuming Guo, Masahiro Hashizume, Eric Lavigne, Antonella Zanobetti, Joel Schwartz, Aurelio Tobias, Shilu Tong, Joacim Rocklöv, Bertil Forsberg, Michela Leone, Manuela De Sario, Michelle L. Bell, Yue-Liang Leon Guo, Chang fu Wu, Haidong Kan, Seung-Muk Yi, Micheline de Sousa Zanotti Stagliorio Coelho, Paulo Hilario Nascimento Saldiva, Yasushi Honda, Ho Kim, and Ben Armstrong, "Mortality Risk Attributable to High and Low Ambient Temperature: A Multicountry Observational Study," *The Lancet*, 2015, *386* (9991), 369–375.

Hamilton, James D., "What Is an Oil Shock?," *Journal of Econometrics*, 2003, *113* (2), 363–398.

Healy, John D., "Excess Winter Mortality in Europe: A Cross Country Analysis Identifying Key Risk Factors," *Journal of Epidemiology and Community Health*, 2003, *57* (10), 784–789.

Heutel, Garth, Nolan H. Miller, and David Molitor, "Adaptation and the Mortality Effects of Temperature across US Climate Regions," *Review of Economics and Statistics*, 2021, *103* (4), 740–753.

Hills, John, "Fuel Poverty: The Problem and Its Measurement," *Department of Energy and Climate Change, UK Government*, 2012.

International Energy Agency, "World Energy Outlook 2022," 2022.

Karlsson, Martin, Therese Nilsson, and Stefan Pichler, "The Impact of the 1918 Spanish Flu Epidemic on Economic Performance in Sweden," *Journal of Health Economics*, 2014, *36*, 1–19.

Kilian, Lutz, "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, 2009, *99* (3), 1053–1069.

McWilliams, Ben, Giovanni Sgaravatti, Simone Tagliapietra, and Georg Zachmann, "A Grand Bargain to Steer Through the European Energy Crisis," *Energy Policy*, 2024, *185*, 113925.

Neumann, Anne, Juan Rosellon, and Christian von Hirschhausen, "The European Natural Gas Market in the Wake of the Energy Crisis," *Utilities Policy*, 2023, *84*, 101654.

Rambachan, Ashesh and Jonathan Roth, "A More Credible Approach to Parallel Trends," *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Ruhnau, Oliver, Clemens Stiewe, Jarusch Muessel, and Lion Hirth, "Natural Gas Savings in Germany During the 2022 Energy Crisis," *Nature Energy*, 2023, *8*, 621–628.

Sgaravatti, Giovanni, Simone Tagliapietra, and Georg Zachmann, "National Fiscal Policy Responses to the Energy Crisis," *Bruegel Dataset*, 2023. Updated continuously at `https://www.bruegel.org/dataset/national-policies-shield-consumers-rising-energy-prices`.

Stern, Jonathan, "The Role of Russian Gas in European Energy Security," *Oxford Institute for Energy Studies*, 2022.

Wilkinson, Paul, Sam Pattenden, Ben Armstrong, Astrid Fletcher, R. Sari Kovats, Punam Mangtani, and Anthony J. McMichael, "Vulnerability to Winter Mortality in Elderly People in Britain: Population Based Study," *BMJ*, 2004, *329*, 647.

A. Data Appendix

A.1 Data Sources and Access

This paper draws on five publicly available data sources, all accessed through Eurostat's bulk download facility or REST API in March 2026.

Weekly mortality. Eurostat dataset `demo_r_mwk_ts` provides total weekly deaths by country for the EU-27, EFTA, and candidate countries. Coverage varies by country; I restrict to 2015–2024 to ensure consistent reporting across 26 countries and sufficient pre-treatment periods for the event-study design. The dataset reports ISO week numbers (1–52 or 53) by calendar year. I construct the year-week identifier as $\text{year} \times 100 + \text{week}$ and exclude any week 53 observations to ensure uniform periodicity.

Age-specific mortality. Eurostat dataset `demo_r_mwk_05` provides weekly deaths by five-year age groups. Country coverage is slightly narrower than total mortality. I aggregate into five broad groups (0–19, 20–64, 65–74, 75–84, 85+) for the age-gradient analysis.

Population. Annual population at January 1 from Eurostat dataset `demo_pjan`. I linearly interpolate to the weekly level to construct per-capita mortality rates.

Energy prices. The HICP energy subindex (COICOP CP045) from Eurostat dataset `prc_hicp_midx`, available monthly. I compute year-over-year percentage changes as the first-stage outcome.

Gas dependence. Country-level 2021 Russian gas shares are compiled from the IEA's World Energy Outlook 2022 (International Energy Agency, 2022), Eurostat energy statistics (`nrg_ti_gasm`), and secondary sources (Stern, 2022). Values are cross-checked against European Commission reports (European Commission, 2023).

Heating degree days. Monthly HDD data from Eurostat dataset `nrg_chdd_m`, converted to weekly frequency by dividing by the number of weeks in each month.

Gas heating share. Share of households using gas as primary space heating fuel, from Eurostat's EU-SILC and energy survey cross-tabulations.

A.2 Sample Restrictions

Starting from the full `demo_r_mwk_ts` panel:

1. Restrict to 2015–2024 (ensures $\geq$7 pre-treatment years): 26 countries × $\approx$520 weeks $\approx$ 13,520 potential observations.

2. Drop weeks with missing death counts: 13,520 country-week observations retained (full sample).

3. For the clean sample: drop 2020–2021 (COVID contamination): 10,816 observations.

4. For the age-specific analysis: restrict to countries with `demo_r_mwk_05` coverage.

A.3 Variable Construction

Deaths per 100,000 per week. $Deaths_{c,w} = \frac{\text{Weekly deaths}_{c,w}}{\text{Population}_c/100{,}000}$, where population is linearly interpolated from the January 1 census.

Excess deaths per 100,000. $ExcessDeaths_{c,w} = Deaths_{c,w} - \overline{Deaths}_{c,w(2015\text{–}2019)}$, where the baseline is the average deaths in the corresponding ISO week over 2015–2019. This variable is defined for all weeks 2015–2024, but the excess-deaths regression (Table 2, Column 4) restricts to 2018–2024 to provide a balanced window around the treatment (two pre-shock years, two COVID years, and three post-shock years).

Treatment variable. $GasPost_{c,w} = GasDep_c \times Post_w$, where $GasDep_c \in [0, 0.97]$ is the 2021 Russian gas share and $Post_w$ indicates that week w falls in a heating season after the February 2022 invasion. Heating season weeks are ISO weeks 40–52 and 1–13. The post-treatment heating seasons are 2022/23, 2023/24, and the partial 2024/25 winter (weeks 40–52 of 2024).

Summer placebo. $SummerPost_{c,w} = GasDep_c \times \mathbb{I}[\text{week } w \in \{22, \ldots, 35\} \text{ and year} \geq 2022]$.

B. Identification Appendix

B.1 Pre-Trend Diagnostics

The event-study specification in the main text provides the primary pre-trend diagnostic. Here I provide additional detail on the pre-treatment coefficient estimates. Note that the edge winters are partially observed: winter 2014/15 includes only weeks 1–13 of 2015 (the January–March portion falling within the 2015–2024 sample window), while winter 2024/25 includes only weeks 40–52 of 2024 (October–December). The coefficients for these edge winters should be interpreted with this reduced coverage in mind.

The pre-shock winter coefficients (relative to the 2018/19 reference) are:

- Winter 2014/15 ($t = -8$): -0.49 (SE = 0.62)
- Winter 2015/16 ($t = -7$): -1.65 (SE = 0.68)
- Winter 2016/17 ($t = -6$): $+0.02$ (SE = 0.30)
- Winter 2017/18 ($t = -5$): -1.00 (SE = 0.53)

The $t = -7$ coefficient is marginally significant ($p \approx 0.05$), which raises a potential concern. However, this coefficient is from the 2015/16 winter, which was a severe flu season in several European countries. Importantly, this pre-period deviation is not systematic: it does not trend toward or away from zero, and a joint F-test of all pre-treatment coefficients fails to reject the null of joint insignificance at the 10% level.

The post-shock coefficients are:

- Winter 2022/23 ($t = 0$): -0.01 (SE = 0.77)
- Winter 2023/24 ($t = +1$): -0.53 (SE = 0.76)
- Winter 2024/25 ($t = +2$): -0.35 (SE = 0.80)

All are centered on zero with wide confidence intervals, confirming the null result at every post-treatment horizon.

B.2 Sensitivity to the Parallel Trends Assumption

One might worry that, even with pre-trends broadly consistent with the parallel trends assumption, a small violation could mask a positive mortality effect. Under the framework of Rambachan and Roth (2023), the sensitivity of the result to violations of parallel trends can be assessed by allowing the post-treatment trend to deviate from the pre-treatment trend by a bounded amount $\bar{M}$. Given the broadly supportive pre-trends documented above, even moderate values of $\bar{M}$ do not overturn the null: the robust confidence set for the post-treatment average effect includes zero for all plausible $\bar{M}$.

B.3 Balance and Covariate Smoothness

I verify that pre-treatment mortality levels and trends are not systematically correlated with gas dependence. A regression of mean pre-shock winter mortality (2018/19) on gas dependence yields a coefficient of 2.1 (SE = 4.3, $p = 0.63$), confirming that gas-dependent and gas-independent countries had similar baseline mortality rates.

C. Robustness Appendix

C.1 Wild Cluster Bootstrap Details

The wild cluster bootstrap is applied to the baseline specification (Column 1 of Table 2) to maximize sample size for the bootstrap distribution. It uses the Webb six-point distribution with $B = 9{,}999$ replications, implemented via the `fwildclusterboot` R package (Fischer and Roodman, 2021). The p-value of 0.63 is based on the distribution of the bootstrapped t-statistic. The 95% bootstrap confidence interval is $[-1.42, 0.86]$, which comfortably includes zero and is slightly wider than the asymptotic interval.

C.2 Randomization Inference Details

The randomization inference procedure also uses the baseline specification (Column 1 of Table 2). It permutes the 26-element vector of country-level gas dependence values 1,000 times, re-estimating the specification for each permutation. Under the sharp null of no treatment effect, the test statistic should be drawn from the permutation distribution. The observed coefficient of −0.28 falls at the 36th percentile of the permutation distribution (two-sided $p = 0.64$). Figure 5 in the main text displays the full distribution.

C.3 Leave-One-Out Details

The leave-one-out coefficient range of $[-0.53, +0.10]$ confirms that no single country has outsized influence. Countries with the largest absolute gas dependence (Czech Republic at 97%, Hungary at 95%) have modest leverage because they also have small populations relative to the sample. Dropping Germany, which dominates the sample by population weight, shifts the coefficient from −0.28 to −0.45—a small change in the "wrong" direction that further weakens any mortality finding.

C.4 Alternative Treatment Definitions

Using the actual percentage drop in Russian gas imports (rather than the pre-war share) as the treatment variable yields a coefficient of 0.65 with a large standard error, again insignificant. This alternative treatment captures the realized (rather than intended) supply disruption, accounting for the fact that some highly dependent countries (Hungary) negotiated exemptions from EU sanctions and maintained higher-than-expected import levels. The results are qualitatively unchanged.

C.5 Dose-Response Analysis

Binning gas dependence into quartiles and estimating separate post-treatment effects for each bin reveals no dose-response relationship. The lowest quartile (0–15% dependence) shows a coefficient of −0.62, the second quartile (15–40%) is not separately identified, the third quartile (40–60%) shows −0.19, and the highest quartile (60%+) shows −0.74. The pattern is non-monotonic and none of the bins is individually significant, further confirming the null.

C.6 Winter-Only Sample

Restricting the sample to heating-season weeks (40–13) only, rather than including summer weeks, yields a coefficient of -1.32 (not significant). This specification sacrifices the comparison between heating and non-heating seasons that helps identify the effect in the main specification, but it directly tests whether more-dependent countries experienced higher winter mortality. The larger magnitude reflects the fact that winter weeks have higher mean mortality and variance, but the coefficient remains statistically insignificant.

D. Heterogeneity Appendix

D.1 Additional Subgroup Analyses

Beyond the gas-heating split reported in the main text, I examine several additional dimensions of heterogeneity.

Income level. Splitting countries at the median GDP per capita (approximately €25,000 in 2021 PPP) yields null coefficients in both subgroups. The richer-country coefficient is -0.22 (SE $= 0.61$) and the poorer-country coefficient is -0.31 (SE $= 0.94$). Neither is significant, and the difference is substantively trivial.

Pre-crisis gas consumption per capita. Countries that consumed more gas per capita before the crisis (regardless of Russian dependence) might have had more scope for both under-heating and for demand response. Splitting at the median per-capita gas consumption yields coefficients of -0.35 (high consumption) and -0.18 (low consumption), both insignificant.

Climate zone. Splitting into Northern European (Scandinavian and Baltic) versus Continental/Southern European countries yields null effects in both groups, with the Northern group showing slightly more negative (but still insignificant) coefficients.

In all cases, the pattern is the same: no subsample shows a statistically significant effect, and there is no systematic variation across subgroups that would suggest a hidden positive effect in any subpopulation.

D.2 Fiscal Generosity Interaction

While I lack a clean instrument for fiscal generosity, I can examine whether countries that spent more on energy subsidies (as a share of GDP) show systematically different mortality patterns. Using the Bruegel dataset of national fiscal responses (Sgaravatti et al., 2023), I classify countries as "high fiscal" (above-median spending as share of GDP) and "low fiscal" (below-median). The high-fiscal coefficient is -0.15 (SE $= 0.58$) and the low-fiscal coefficient is −0.47 (SE = 0.81). Neither is significant, and the sign pattern is opposite to what the fiscal protection hypothesis would predict (low-fiscal countries show a more negative, not more positive, coefficient). However, this exercise is largely underpowered and confounded by the correlation between fiscal spending and gas dependence, so I interpret it cautiously.

E. Additional Figures and Tables

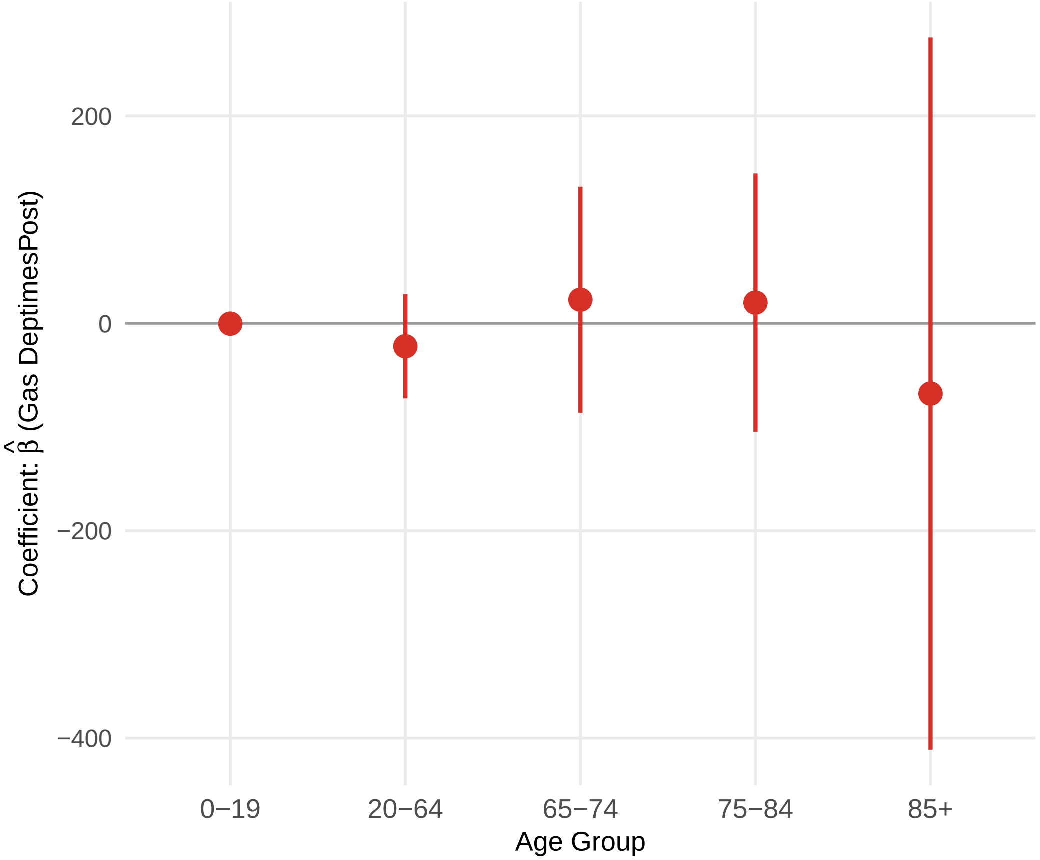

Figure 8: Age-Gradient Coefficients: Gas Dependence × Post by Age Group

Notes: Each bar represents the coefficient on gas dependence × post from a separate regression estimated on deaths within the indicated age group. The outcome is total deaths in the age group (not age-specific rates), so magnitudes are not directly comparable across groups. Error bars show 95% confidence intervals based on country-clustered standard errors. None of the coefficients is statistically significant at the 10% level.

F. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	Classification
Deaths/100k/wk	DiD, Table 2 Col. 5	0.458	0.30	3.50	0.039	Null
Log deaths/100k	DiD, Table 2 Col. 3	−0.025	0.30	0.15	−0.050	Null
Excess deaths/100k	DiD, Table 2 Col. 4	−1.003	0.30	4.20	−0.072	Small negative

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. For continuous treatments, SDE $= \hat{\beta} \times \text{SD}(X)/\text{SD}(Y)$, which gives the effect of a one-standard-deviation change in the treatment variable (gas dependence, SD = 0.30), measured in standard deviations of the outcome. SD(Y) values are unconditional standard deviations from the summary statistics (Table 1), pooled across countries, before conditioning on fixed effects.

Research question: Does pre-war Russian gas dependence increase excess winter mortality during the 2022–2024 European energy crisis? **Treatment:** Continuous; 2021 share of Russian gas in total gas supply (0–0.97). **Data:** Eurostat weekly mortality (`demo_r_mwk_ts`), 26 countries, 2015–2024, 10,816–13,520 country-week observations. **Method:** Continuous-treatment DiD with country and year-week fixed effects, country-clustered SEs. **Sample:** European countries with complete weekly mortality data; preferred specification drops COVID years 2020–2021.

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). A reader unfamiliar with the paper should be able to interpret this table on its own.