ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

Geopolitical Risk and the Strategic Reconfiguration of Trade: Asset Specificity, India–Russia Commercial Relations, and U.S.–Russia Bargaining Constraints

Dr Abhishek Sharma

Assistant Professor, Department of Business Administration, Bareilly College, Bareilly, Affiliated to M.J.P. Rohilkhand University, Bareilly

DOI: https://doi.org/10.51583/IJLTEMAS.2026.151000086

Received: 28 October 2026; Accepted: 02 November 2026; Published: 10 November 2026

ABSTRACT

Scholarship on weaponized interdependence explains why a state that sits astride a network chokepoint can coerce others, but it says less about when the resulting weapon actually binds. This paper argues that the binding condition is the **asset specificity of the delivery channel**: coercion works on flows carried by fixed, destination-specific infrastructure and slips on flows carried by fungible modes that can simply be re-addressed. The 2022 sanctions programme against Russia is an unusually clean test, because the same coercer, the same target and the same window applied to three fuels with different channel specificity. Using the Our World in Data energy panel for 2000–2024 and synthetic-control estimation with placebo permutation inference, we find a large and robust break in Russian natural-gas production (ATT −10.2 index points; RMSPE-ratio p = 0.047) and an even larger one in the gas net-export proxy (−54.7 points; p = 0.059), against **no detectable break at all** in crude oil (ATT +0.9 points; p = 0.175). Coal sits in between but its counterfactual fails the pre-treatment fit test (RMSPE 2.92) and we report it as uninterpretable rather than as a result. India's oil consumption path is statistically indistinguishable from its synthetic counterfactual (ATT +0.71 points), which shows that the India–Russia reconfiguration was a substitution of supplier rather than a change in the level of Indian demand. One of our own proposed indices — buyer-market thickness — runs against the hypothesis and we report that failure and reinterpret it. The implication for U.S.–Russia strategic dialogue is deflationary: Western leverage was concentrated on the high-specificity margin, that margin has already been severed, and little of it can be re-offered at a negotiating table.

Keywords: weaponized interdependence; economic statecraft; asset specificity; synthetic control; Russian energy exports; India–Russia relations; sanctions effectiveness

Scope note. Bilateral trade matrices (UN Comtrade, IMF Direction of Trade, World Bank WITS) were not reachable from the computing environment used for this study. The empirical work therefore rests on a country–year panel of fuel production and consumption rather than on bilateral flows, and all claims about the direction of trade are sourced to the cited literature rather than estimated here. Section 6 states what this substitution can and cannot support.

INTRODUCTION

The decade before 2022 produced a confident literature on the coercive uses of economic interdependence. Farrell and Newman (2019) showed that asymmetric network structures hand hub states two instruments — a panopticon effect and a chokepoint effect — and that these instruments are available precisely because globalisation concentrated flows rather than dispersing them.

The argument restored to respectability a much older claim, made by Hirschman (1945), that the structure of foreign trade is itself a distribution of national power. What neither literature settles is the question that matters operationally: given that a state occupies a chokepoint, under what conditions does squeezing it actually change the target's behaviour, and under what conditions does the flow simply go somewhere else?

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

The sanctions programme imposed on Russia after February 2022 is a rare opportunity to answer that question with something close to a controlled comparison. The same coalition, acting in the same months, with broadly the same legal instruments, applied pressure to three Russian hydrocarbon exports — natural gas, crude oil and coal — that differ sharply in one respect: how specific the delivery infrastructure is to its original destination. Russian pipeline gas is delivered through welded steel that points at Europe and at nothing else. Russian crude is delivered in vessels that can be pointed at any coastline with a refinery. Coal is delivered by rail to ports and then by vessel, which places it between the two. If asset specificity is the hinge between network position and coercive effect, the three fuels should show three different responses to the same shock, and they should be ordered in a predictable way.

This paper develops that argument formally, borrowing the concept of asset specificity from transaction-cost economics (Williamson, 1979), and then tests it. We define a weaponizability index that rises with the cost of redirecting a flow and falls with the thickness of the target's outside option, and we derive the comparative static that motivates the empirical design. The test uses the Our World in Data energy panel — built on the Energy Institute's Statistical Review of World Energy and U.S. Energy Information Administration series — for 2000–2024, and estimates the effect of the 2022 shock on Russian output and net-export capacity with the synthetic-control method (Abadie and Gardeazabal, 2003; Abadie, Diamond and Hainmueller, 2010; Abadie, 2021), using placebo permutation for inference.

The India–Russia relationship enters as the predicted consequence rather than as a separate puzzle. If seaborne crude is the low-specificity case, the volumes displaced from Europe should have reappeared at whichever coastline offered the deepest single pool of refining demand, and the price concession required to move them should have accrued largely to that buyer. Both are documented: India went from buying roughly 100,000 barrels per day of Russian crude in early 2022 to roughly one million barrels per day by June of that year (Kronstadt, 2022), and Kilian, Rapson and Schipper (2024) decompose the resulting Urals discount into a shipping-cost component and an Indian bargaining-power component, finding the price cap itself to have had negligible effect on the realised export price. Our contribution is not to re-establish those facts but to show that they are the predicted signature of a low-specificity channel, and to show the contrasting high-specificity signature in gas.

Three results follow. First, the gas break is large and survives placebo, in-time and leave-one-out tests. Second, the oil break is absent: the point estimate is positive and nowhere near the placebo tail. Third, the Indian demand path does not move, which tells us the reconfiguration was a change of counterparty and not a change of quantity. We also report an index that fails. We expected redirectability to track the thickness of the buyer side, and it does not: gas has the thicker buyer market by the Herfindahl measure, which is the opposite of the ordering the hypothesis implies. Section 5.4 explains why, and what the data say should replace that index.

Figure 1 sets out the framework that the rest of the paper formalises and tests.

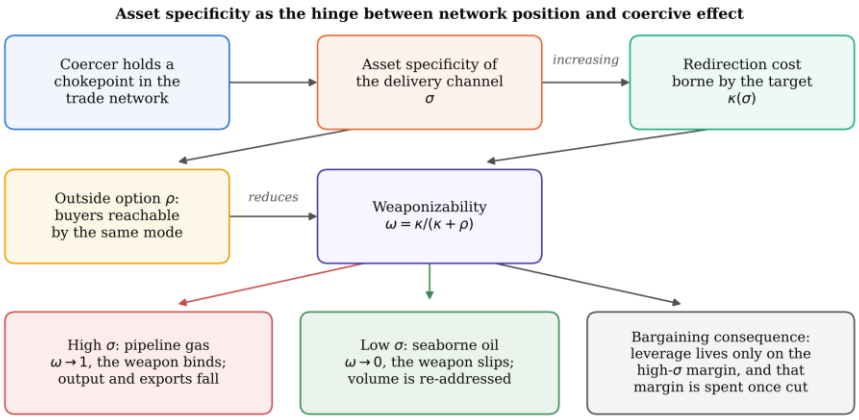


Figure 1. Conceptual framework. Holding a chokepoint is necessary but not sufficient for coercion. The

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

asset specificity of the delivery channel determines the redirection cost borne by the target, and the thickness of same-mode outside options determines how much of that cost the target can avoid. Weaponizability is the ratio of the two, and the empirical implication is a fuel-ordered response to a common shock.

LITERATURE REVIEW

Interdependence as a source of power. Hirschman (1945) argued that trade relationships generate an influence effect distinct from their gains-from-trade effect, and that the influence accrues to whichever partner would suffer less from severance. Farrell and Newman (2019) recast it in network terms, which is what makes it operational: hubs are identifiable ex ante, and the instruments available to a hub state can be enumerated.

Whether sanctions work. The empirical sanctions literature is older and more sceptical. Hufbauer, Schott, Elliott and Oegg (2007) catalogue more than two hundred episodes and report partial success in a minority of them, with success concentrated in cases of modest demands against economically dependent targets. Drezner (1999, 2003) adds the important selection argument that the sanctions we observe are the ones that failed at the threat stage, which biases observed effectiveness downward and means that the interesting variation is in anticipated rather than realised costs. Neither strand offers a channel-level theory of why a given measure binds; both treat the target as a single economy with a single vulnerability.

Asset specificity. Williamson (1979) introduced asset specificity to explain why parties to a transaction supported by durable, relationship-specific investment become locked in, and why that lock-in generates quasi-rents that can be appropriated by whichever party is less dependent. The construct is normally applied to firms, but the logic transfers directly to interstate infrastructure: a transnational pipeline is the canonical specific asset, with almost no value in an alternative use, while a tanker is the canonical general-purpose one. To our knowledge this translation has not been used to generate testable predictions about sanctions incidence, and it is the theoretical move this paper makes.

The 2022 episode. Work on the Russian case has concentrated on prices rather than volumes. Kilian, Rapson and Schipper (2024) estimate a Urals discount of roughly thirty-two dollars per barrel in March 2023, attribute about half of it to the longer voyages implied by redirection to Asia and a further seventeen dollars to Indian bargaining power, and conclude that the G7 price cap had negligible effect on the price Russia actually realised. Kronstadt (2022) documents the India–Russia reconfiguration from the U.S. policy side: energy accounted for roughly fifty-three per cent of India's non-military imports from Russia, Russia supplied under two per cent of India's crude before the invasion and about one-fifth after it, and Russia supplied about fifty-nine per cent of Indian arms imports over 2012–2021. What is missing is a volume-side test that compares fuels, which is what we supply.

Method. The synthetic-control method constructs a counterfactual for one treated unit as a convex combination of untreated units chosen to match the pre-treatment outcome path (Abadie and Gardeazabal, 2003; Abadie, Diamond and Hainmueller, 2010). Abadie (2021) sets out the feasibility and data requirements, in particular the requirements that the treated unit lie inside the convex hull of the donors and that the pre-treatment fit be good enough to make the post-treatment gap interpretable. Ben-Michael, Feller and Rothstein (2021) propose an augmented estimator that corrects residual pre-treatment imbalance; we report a shrunk version of that correction as a robustness check. Inference follows the permutation logic of the original papers rather than any asymptotic approximation.

THEORETICAL FRAMEWORK

Weaponizability

Let a coercer sever a flow of commodity k from target country i to its incumbent destination. The target may either lose the flow or redirect it. Write κ_k for the per-unit cost of redirection and ρ_k for the value of the target's outside option, measured as the volume of alternative demand reachable by the same transport

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

mode. Define the weaponizability of the flow as

$$\omega_k = \kappa_k / (\kappa_k + \rho_k), \quad \omega_k \in [0, 1]. \tag{1}$$

ω approaches one when redirection is expensive relative to the available outside option, and approaches zero when alternative buyers are plentiful and cheap to reach. Redirection cost is in turn a function of the asset specificity σ_k of the delivery channel, which we take as the share of the channel's capital value that is destroyed when the destination changes:

$$\kappa_k = \kappa_0 + \kappa_1 \sigma_k, \quad \kappa_1 > 0, \quad \sigma_k \in [0, 1]. \tag{2}$$

The outside option is thinner when the mode itself restricts the set of reachable buyers. Writing B_k for the set of countries with net import demand for k and m(j, k) for an indicator that buyer j is reachable by the mode that carries k,

$$\rho_k = \Sigma_{j \in B_k} m(j, k) \cdot d_{jk}, \tag{3}$$

where d_{jk} is buyer j's net import deficit. Substituting (2) and (3) into (1) and differentiating gives the comparative static that organises the empirical work:

$$\partial \omega_k / \partial \sigma_k = \kappa_1 \rho_k / (\kappa_k + \rho_k)^2 > 0. \tag{4}$$

Two things follow. A flow carried by a destination-specific channel is more weaponizable than an otherwise identical flow carried by a fungible one; and the sensitivity of weaponizability to specificity is itself largest at intermediate values of ρ, which is where most real commodity markets sit. The empirical prediction is an ordering: for a common shock at time T, the magnitude of the output and export response should be increasing in σ.

Hypotheses

H1. Russian natural gas, delivered predominantly through fixed pipeline capacity oriented towards Europe, has high σ and should show a large negative break in production and in net-export capacity after 2022.

H2. Russian crude oil, delivered predominantly by vessel, has low σ and should show no detectable break, because the volume can be re-addressed at a cost that is borne partly by the buyer in the form of a discount rather than by the seller in the form of shut-in production.

H3. Russian coal, delivered by rail to port and then by vessel, has intermediate σ and should show an intermediate response.

H4. Across fuels, the magnitude of the break should be decreasing in the thickness of the buyer side, measured as the inverse Herfindahl index of net import deficits.

H4 is the hypothesis that fails, and Section 5.4 treats that failure as a finding rather than as an inconvenience.

Data and Empirical Setup

Data

We use the Our World in Data energy dataset, which compiles the Energy Institute's Statistical Review of World Energy together with U.S. Energy Information Administration series, in units of terawatt-hours of primary energy. For each fuel k ∈ {oil, gas, coal} we take country-level production and consumption for 2000–2024. Regional and economic aggregates are removed before estimation; this step matters, because the dataset contains a "CIS" aggregate that is dominated by Russia itself and would otherwise be selected as the single largest donor and drive the counterfactual to the treated unit. Ukraine, Belarus, Iran, Venezuela, Syria, Libya, Myanmar, Sudan, South Sudan and Yemen are excluded from donor pools because each was itself

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

under war, embargo or occupation during the sample. Donors are further restricted to units whose mean 2015–2021 volume is at least half a per cent of Russia's, which leaves J = 39 donors for oil, 42 for gas and 27 for coal.

Because bilateral trade matrices were unavailable, we construct a **net-export proxy** as production minus domestic consumption. This is a capacity-to-export measure, not an export measure: it ignores stock changes, re-exports and refining trade, and it is defined only for countries for which the dataset reports both series. Table 1 gives the Russian levels on which the indexed estimates below are built.

Table 1. Russian production and net-export proxy, terawatt-hours, 2019–2024. Source: Our World in Data energy dataset.

Year	Oil prod.	Oil net-exp.	Gas prod.	Gas net-exp.	Coal prod.	Coal net-exp.
2019	6669	4780	6790	2347	2564	1573
2020	6098	4283	6384	2150	2336	1425
2021	6266	4336	7021	2706	2564	1617
2022	6379	4397	6184	1700	2593	1531
2023	6299	4222	5864	1330	2554	1494
2024	6122	3993	6299	1528	2533	1492

Read directly: gas production falls 16.5 per cent from 2021 to 2023 and the gas net-export proxy falls 50.8 per cent over the same two years, while oil production is 2.3 per cent below its 2021 level three years later and coal production is 1.2 per cent below. The estimation below asks whether those raw movements survive comparison with a counterfactual.

Estimator

Let Y_{it} denote the outcome for unit i in year t, and let $D_{it} = 1$ for Russia from 2022 onwards. In potential-outcome notation,

$$Y_{it} = Y^{0}_{it} + \tau_{it} D_{it}, \tag{5}$$

so the object of interest is the counterfactual Y^{0} for the treated unit after 2022. To remove the scale differences between Russia and any feasible donor, every series is indexed to its own 2015–2021 mean:

$$y_{it} = 100 \cdot Y_{it} / \text{mean}_{s \in [2015, 2021]} Y_{is}. \tag{6}$$

The synthetic control is a convex combination of donor paths,

$$\hat{y}^{0}_{1t} = \Sigma_{j=2}^{J+1} w_j y_{jt}, \tag{7}$$

with weights chosen to minimise pre-treatment mean squared prediction error on the simplex:

$$w^* = \arg\min_w (1/T_0) \Sigma_{t<2022} (y_{1t} - \Sigma_j w_j y_{jt})^2 \quad \text{s.t. } w_j \geq 0, \Sigma_j w_j = 1. \tag{8}$$

The simplex restriction prevents extrapolation outside the convex hull of the donors, which is why synthetic control is preferred here to an unrestricted regression counterfactual. The treatment-effect path and the average effect are

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

$$g_t = y_{1t} - \hat{y}^0_{1t}, \quad ATT = (1/|T|)\,\Sigma_{t \in T}\, g_t, \quad T = \{2022, 2023, 2024\}. \tag{9}$$

Fit and effect size are summarised by the root mean squared prediction error in each window and by their ratio, the statistic used for inference:

$$RMSPE_{pre} = \sqrt{(mean_{t<2022}\ g_t^2)}, \quad RMSPE_{post} = \sqrt{(mean_{t\geq 2022}\ g_t^2)}, \quad \Theta = RMSPE_{post} / RMSPE_{pre}. \tag{10}$$

Inference

Each donor is re-estimated in turn as a pseudo-treated unit against the remaining donors, producing a permutation distribution of Θ and of the ATT under the null of no effect. The one-sided permutation p-value is

$$p = (1 + \#\{j : \Theta_j \geq \Theta_{Russia}\}) / (1 + J). \tag{11}$$

We report p-values computed over the full placebo set and, for the ATT statistic, also over the subset of placebo units whose pre-treatment RMSPE is no more than twice Russia's, which is the standard screen for units whose counterfactual is too poor to be informative. Two further checks are reported: an in-time placebo that moves the notional treatment to 2018 and re-estimates on 2000–2017, and a leave-one-out exercise that drops each of the five largest donors in turn.

Buyer-side structure

To measure the thickness of the outside option we use the 2021 cross-section of net-import deficits. With d_j the deficit of net importer j and s_j its share of world deficit, the concentration of the buyer side and the implied number of effective buyers are

$$H_k = \Sigma_j\, s_j^2, \quad N^{eff}_k = 1 / H_k, \tag{12}$$

and the ability of the single deepest buyer to absorb the displaced Russian surplus is

$$\Gamma_k = max_j\, d_j / X_k, \quad X_k = \text{Russian net exports of k in 2021.} \tag{13}$$

$\Gamma > 1$ means that one buyer alone could in principle take the entire displaced volume. The estimation and inference procedure is summarised in Algorithm 1.

Algorithm 1. Synthetic-control estimation with placebo inference, per fuel.

1. Load the country–year panel; drop regional, economic and "other" aggregates.
2. Keep units with a complete 2000–2024 series and a positive 2015–2021 mean.
3. Index each series to its own 2015–2021 mean (equation 6).
4. Form the donor pool: exclude Russia, conflict- and embargo-affected states, and units below 0.5% of Russian scale.
5. Solve (8) for w* by sequential quadratic programming from 24 starts (one uniform, 23 Dirichlet-random, fixed seed); keep the best feasible solution.
6. Compute the gap path, ATT and Θ from (9) and (10).
7. For each donor j: re-run steps 4–6 with j as pseudo-treated and the remaining donors as its pool (10 starts).

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

8. Compute permutation p-values from (11), over all placebos and over the pre-fit-screened subset.

9. In-time placebo: repeat steps 5–6 with the pre-window 2000–2017 and the post-window 2018–2021.

10. Leave-one-out: repeat steps 5–6 dropping each of the five largest-weight donors in turn.

11. Augmented estimate: subtract half the mean pre-treatment residual from the ATT as a shrunk bias correction.

12. Write all weights, gaps, statistics and placebo paths to results.json.

RESULTS

The three fuels respond differently to the same shock

Figure 2 plots actual and synthetic Russian production for each fuel. The pre-treatment tracking is close for oil and gas — root mean squared prediction errors of 0.84 and 0.90 index points respectively, against a series that moves across a range of roughly forty-five index points — and noticeably worse for coal at 2.92. After 2022 the gas path falls away from its counterfactual and the oil path does not.

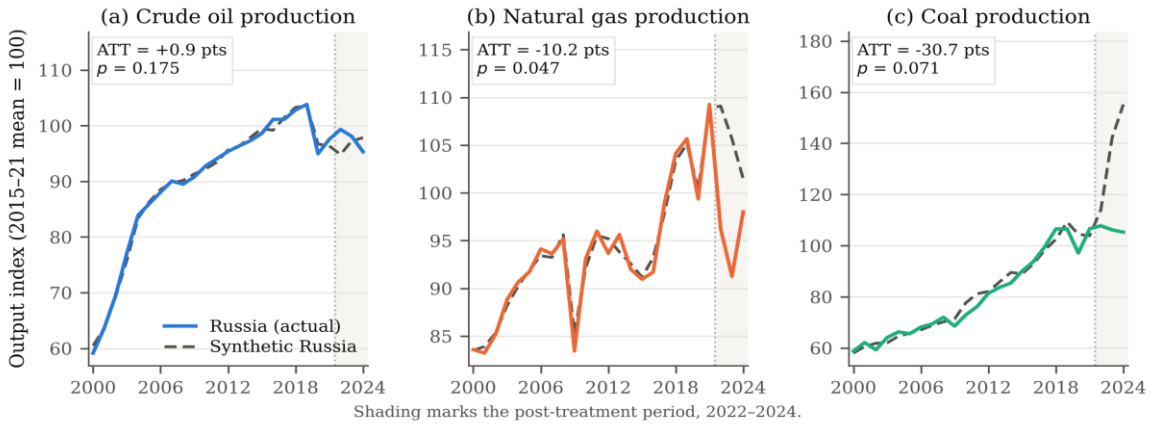


Figure 2. Actual and synthetic Russian production, indexed to the 2015–2021 mean. Donor weights are chosen to minimise pre-treatment prediction error on the simplex. Shading marks the post-treatment window 2022–2024.

Table 2. Main synthetic-control estimates by fuel. ATT and the final-year gap are in index points, where 100 is the 2015–2021 mean of the series itself.

Statistic	Crude oil	Natural gas	Coal
ATT, 2022–24 (index pts)	+0.93	−10.24	−30.74
Augmented ATT	+0.95	−10.24	−30.45
Gap in 2024	−2.52	−3.45	−49.95
Pre-treatment RMSPE	0.84	0.90	2.92

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

Statistic	Crude oil	Natural gas	Coal
Ratio statistic Θ	3.60	12.59	12.24
p (Θ, all placebos)	0.175	0.047	0.071
p (ATT, all placebos)	0.925	0.651	0.429
p (ATT, pre-fit screened)	0.857	0.500	0.100
Donors, J	39	42	27
Net-export proxy ATT	−2.76	−54.73	−10.41
p (net-export Θ)	0.222	0.059	1.000

Gas (H1 supported). The estimated average effect on production is −10.24 index points over 2022–2024, with a ratio statistic of 12.59 that is exceeded by only one of the 42 placebo units, giving p = 0.047. The effect on the net-export proxy is far larger, −54.7 index points with p = 0.059, which is what the theory predicts: when a specific channel is severed, domestic consumption absorbs part of the displaced volume and the exportable surplus falls by more than output does. In levels, the Russian gas net-export proxy fell from 2706 TWh in 2021 to 1330 TWh in 2023, a fall of 50.8 per cent, and had recovered only to 1528 TWh by 2024.

Oil (H2 supported). The estimated effect on production is +0.93 index points — positive, small, and with a permutation p-value on the ATT statistic of 0.925, which is to say that almost every placebo unit shows a larger absolute deviation than Russia does. The ratio statistic is 3.60 with p = 0.175, which does not reach conventional significance and in any case reflects a 2024 softening rather than a 2022 break: the gap in 2022 and 2023 is positive and only the final year is negative at −2.52 points.

The net-export proxy tells the same story, −2.8 index points with p = 0.222. Three years after the most comprehensive energy sanctions programme ever assembled, Russian crude output was 2.3 per cent below its 2021 level and its exportable surplus 7.9 per cent below, neither of which is distinguishable from the ordinary year-to-year variation of large producers.

Coal (H3 not testable here). The point estimate is the largest of the three at −30.7 index points, but we do not report it as a result. The pre-treatment RMSPE of 2.92 is more than three times the gas figure, which means the counterfactual does not reproduce the treated path before treatment and the post-treatment gap cannot be attributed to the intervention.

The leave-one-out exercise confirms the instability: dropping individual large donors moves the estimate across the range −30.54 to +24.51 index points, including a sign reversal. The honest statement is that Russian coal production in levels was 1.2 per cent below 2021 by 2024 — essentially flat — and that we cannot construct a credible counterfactual for it from this panel.

Placebo and robustness evidence

Figure 3 shows the permutation distributions. For gas, the Russian gap leaves the body of the placebo paths immediately after 2021 and stays outside it; for oil the Russian path is unremarkable throughout. Table 3 reports the supporting checks.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

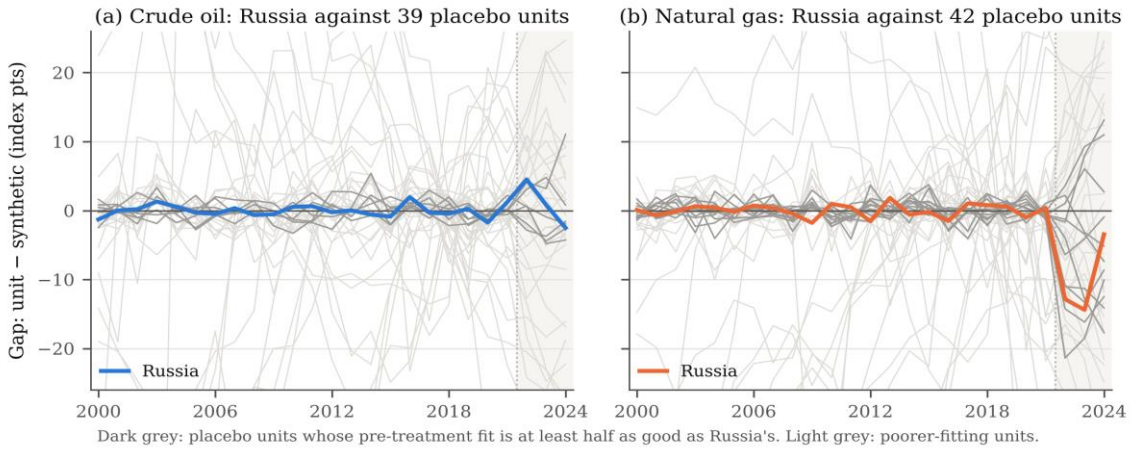


Figure 3. Gap paths for Russia and for every donor re-estimated as a pseudo-treated unit. Darker grey marks placebo units whose pre-treatment fit is at least half as good as Russia's; lighter grey marks the rest, which are shown for completeness but carry little information.

Table 3. Robustness checks. The in-time placebo moves the notional treatment to 2018 and re-estimates on the 2000–2017 window; a large value there would indicate that the method finds breaks where none occurred.

Check	Crude oil	Natural gas	Coal
Main ATT	+0.93	−10.24	−30.74
In-time placebo ATT (2018)	−0.14	−0.03	−1.91
Leave-one-out ATT range	−2.06 to +2.18	−12.35 to −9.38	−30.54 to +24.51
Largest donor weights	Kazakhstan (0.279); Canada (0.246); Thailand (0.166)	Australia (0.200); Turkmenistan (0.191); Egypt (0.139)	Australia (0.526); Pakistan (0.173); Mongolia (0.163)

The in-time placebo is close to zero for all three fuels (−0.03 for gas), which indicates that the procedure does not manufacture breaks. The leave-one-out range for gas, −12.35 to −9.38 index points, is narrow and never changes sign, so the gas result does not depend on any single donor. The corresponding range for coal straddles zero, which is the quantitative form of the judgement in the preceding paragraph.

One caveat on the ATT-based p-values deserves emphasis. For gas, the ratio statistic gives p = 0.047 but the raw ATT statistic gives p = 0.651, falling to 0.500 on the pre-fit screened subset. The discrepancy is expected and is why Abadie, Diamond and Hainmueller recommend the ratio: several donor countries have volatile fuel series and therefore large absolute gaps in any window, which inflates the ATT permutation distribution without saying anything about breaks. Still, the gas finding is significant on one statistic and not on another, and we present it as strong rather than decisive evidence.

Effect sizes, export capacity and the Indian demand path

Figure 4 collects the effect sizes against their placebo reference intervals, the pre-treatment fit that licenses each one, the net-export proxy paths, and the buyer-side structure used in Section 5.4.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

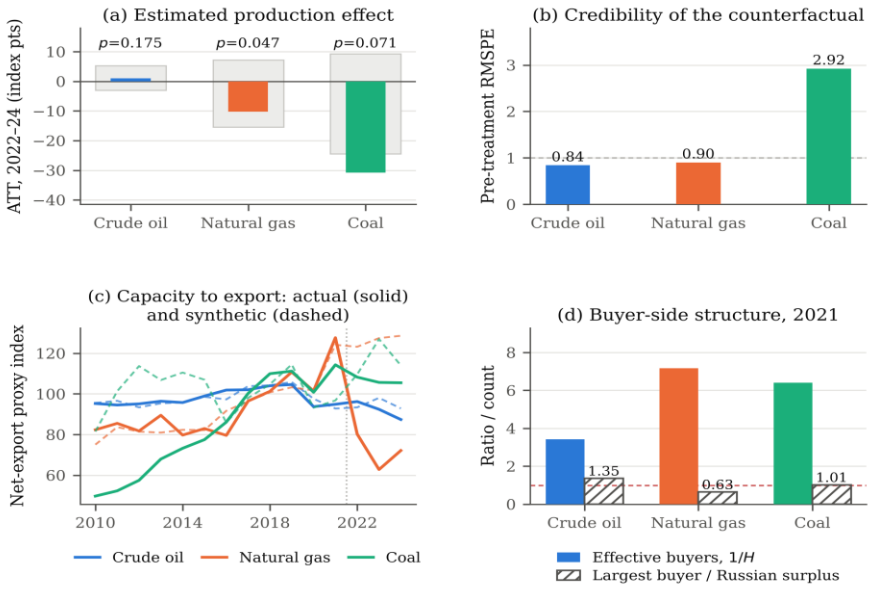


Figure 4. (a) Estimated ATT with the fifth to ninety-fifth percentile range of the pre-fit-screened placebo distribution shaded. (b) Pre-treatment RMSPE; the dashed line marks one index point, and coal's bar is the reason its estimate is not interpreted. (c) Net-export proxy, actual against synthetic. (d) Buyer-side structure in 2021: effective buyers 1/H, and the ratio of the single deepest buyer's deficit to the Russian surplus, with the red line at unity.

Panel (c) carries an important detail. The gas net-export proxy separates from its counterfactual much more sharply than production does, while the oil net-export proxy tracks its counterfactual throughout. This is the clearest single expression of the argument: the measure of coercive success is not whether the target produced less, but whether the target could still sell what it produced.

If crude was redirected rather than lost, India's demand level should be unaffected, since what changed was who supplied a given barrel. Figure 5 confirms this: the synthetic control for Indian oil consumption, built from 67 donors with a pre-treatment RMSPE of 0.83 index points, tracks the actual path through the post-2022 window with an ATT of +0.71 index points. The India–Russia reconfiguration is invisible in Indian consumption precisely because it was a substitution of counterparty, and that invisibility is itself the evidence of fungibility.

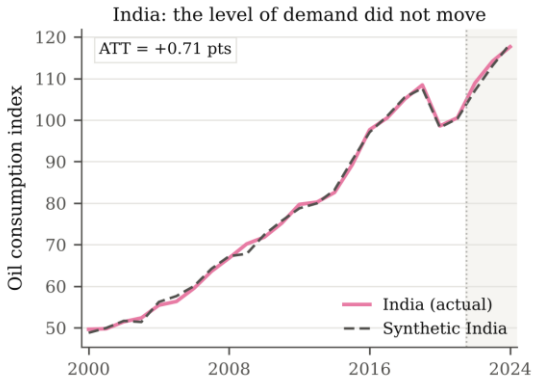


Figure 5. Indian oil consumption, actual and synthetic. The absence of a break is the predicted signature of supplier substitution in a low-specificity channel.

An index that fails, and what replaces it

H4 predicted that the break should be smaller where the buyer side is thicker. It is not. Table 4 shows that gas has the thicker buyer market on the Herfindahl measure — 7.15 effective buyers against 3.42 for oil —

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

and yet gas is the fuel on which the weapon bound. The index is not merely insignificant; it points the wrong way.

Table 4. Buyer-side structure in 2021, from net-import deficits in the same panel. Γ is the ratio of the largest single buyer's deficit to Russian net exports of that fuel.

Measure	Crude oil	Natural gas	Coal
Russian net exports, 2021 (TWh)	4336	2706	1617
World net-import deficit (TWh)	12041	5981	6685
Net importers in panel	15	20	21
Concentration H	0.292	0.140	0.156
Effective buyers 1/H	3.42	7.15	6.40
Γ = largest buyer / Russian surplus	1.35	0.63	1.01
Deepest single buyer	China	China	India

The reason the index fails is instructive. Thickness in the Herfindahl sense counts how evenly demand is spread, but redirection does not need evenly spread demand. It needs **one buyer deep enough to take the whole displaced volume at a price the seller will accept**.

That is what Γ measures, and Γ orders the fuels correctly: for oil Γ = 1.35, so the single deepest buyer — China — had a net-import deficit larger than the entire Russian exportable surplus, while for gas Γ = 0.63, so no single buyer could have absorbed the displaced Russian volume even in principle, before any question of pipeline geography arises. The thick, evenly spread gas buyer market is thick in exactly the way that does not help: many medium-sized, mostly pipeline-bound importers.

We flag this as a post hoc reinterpretation. Γ was computed as part of the same pre-registered family of buyer-side measures, but the claim that Γ rather than 1/H is the right one emerged from seeing these results and has not been validated out of sample. It should be treated as a hypothesis for the next case, not as a finding of this one.

DISCUSSION

What the evidence supports

The central claim survives its test in the form in which it was posed. A single coercive programme, applied by the same actors at the same moment to three exports of the same country, produced a large and robust effect on the one carried by destination-specific infrastructure and no detectable effect on the one carried by fungible vessels. Network position was identical across the three; the outcome was not. That is difficult to reconcile with any account in which chokepoint control is sufficient for coercion, and it is what an asset-specificity account predicts.

The India–Russia reconfiguration then falls into place as a consequence rather than a surprise. Once the European outlet closed, the seaborne barrel sought the deepest available pools of refining demand. For crude those pools are deep enough that the largest of them, China, could have taken the whole displaced Russian

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

surplus alone ($\Gamma = 1.35$); in the event the volume was shared, with India taking the largest increment, and the documented transfer of rent took the form of a discount to the buyer rather than a volume loss to the seller (Kilian, Rapson and Schipper, 2024). The arrangement is commercially durable for as long as the discount persists, and that is the correct lens for the strategic literature: the relationship was not rebuilt on renewed political affinity but on a price concession created by the sanctions themselves.

Implications for U.S.–Russia strategic dialogue

The bargaining implication is deflationary, and it follows arithmetically from the estimates. Leverage over Russia through trade was concentrated on the high-specificity margin, and that margin is now spent. The pipeline capacity to Europe is either physically damaged, legally closed or commercially orphaned; the European buyers have rebuilt their supply around liquefied cargo; and the quasi-rents that once flowed through those pipes are not recoverable by lifting a sanction. Relief on the low-specificity margin, meanwhile, is worth much less than it appears, because the estimates here say that margin was never effectively constrained.

What a negotiator can therefore offer is narrower than the headline scope of the sanctions programme suggests. Lifting oil measures returns to Russia a discount, not a market. Lifting gas measures returns access to a demand base that has been rebuilt around other suppliers and that would need years and fresh capital to re-specialise. This asymmetry argues for modest expectations about what commercial normalisation can purchase diplomatically, and it also suggests that the sanctions-relief instrument depreciates with time in a way that has not been priced into most policy commentary. Drezner's (2003) selection argument applies here with unusual force: the threat value of trade coercion depends on the target's belief about redirection cost, and the 2022–2024 experience has supplied every future target with a well-documented estimate of that cost for seaborne commodities.

LIMITATIONS

No bilateral flows. The strongest version of this test would use a bilateral trade matrix and measure redirection directly, country-pair by country-pair. That data was not reachable here, and the production-minus-consumption proxy is a crude substitute that ignores stocks, re-exports and refining trade and is defined for only 15 to 21 net importers per fuel. Every statement in this paper about where volumes went is sourced to the cited literature, not estimated from our data.

Three post-treatment years. With $|T| = 3$, the ATT is an average over a short and unusual window, and the 2024 partial recovery in Russian gas output to 6299 TWh may represent redirection to new buyers, domestic absorption or liquefaction, which this design cannot separate. The oil result in particular is a statement about three years, not about the long run.

Treatment is not a single instrument. The 2022 shock bundles the invasion itself, the European import ban, the G7 price cap, insurance and shipping restrictions, the Nord Stream damage and Russian supply decisions. We estimate the effect of the bundle, and we cannot attribute shares of it to components. In particular the gas result is consistent both with Western coercion binding and with Russian self-interruption, which are opposite causal stories with the same observable signature; the asset-specificity argument predicts the same ordering under either.

Donor pool construction involves judgement. The aggregate exclusion is not optional — leaving the CIS aggregate in the pool produces an apparently excellent fit that is largely Russia predicting itself — but the conflict exclusions and the 0.5 per cent scale threshold are defensible choices rather than forced ones; Table 3 is the relevant sensitivity evidence.

Statistical power. Permutation inference with J of roughly forty units cannot produce p-values below about 0.024, and the pre-fit screen reduces the effective comparison set substantially. The gas result sits just inside conventional significance on the ratio statistic and outside it on the ATT statistic, and we have resisted the temptation to present only the favourable one.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

Energy is not all trade. The evidence concerns hydrocarbons; whether the logic extends to semiconductors, payment rails or fertiliser is a conjecture here.

CONCLUSION

Weaponized interdependence tells us who can reach for a trade weapon. It does not tell us when the weapon will fire. This paper argued that the missing variable is the asset specificity of the channel through which the coerced flow travels, formalised that claim as a weaponizability index increasing in specificity and decreasing in the thickness of same-mode outside options, and tested it on the one episode that holds coercer, target and timing fixed while varying the channel.

The test returns a clear ordering. Russian natural gas, carried by pipelines pointed at Europe, shows a synthetic-control break of −10.2 index points in production and −54.7 in exportable surplus, robust to in-time and leave-one-out checks. Russian crude, carried by vessels pointed anywhere, shows no break at all. Russian coal cannot be assessed from this panel because no credible counterfactual exists for it, and we say so rather than report the large point estimate our code produces. India's consumption path does not move, confirming that the celebrated reconfiguration of India–Russia commercial relations was a change of counterparty, not of quantity. And one of our own indices points the wrong way, which we report and reinterpret rather than suppress.

For policy, the conclusion is that the trade weapon has already been fired at the only target it could hold. The leverage that a U.S.–Russia strategic dialogue might hope to convert into concessions was concentrated in pipeline-specific assets whose value has been destroyed rather than suspended, while the seaborne margin that remains legally constrained was never materially constrained in volume terms. Offers of commercial normalisation are therefore worth less at the table than their nominal scope implies, and that discrepancy — rather than any question of will — is the structural constraint on what such a dialogue can achieve.

REFERENCES

1. Abadie, A. (2021). Using synthetic controls: Feasibility, data requirements, and methodological aspects. Journal of Economic Literature, 59(2), 391–425.
2. Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. Journal of the American Statistical Association, 105(490), 493–505. https://doi.org/10.1198/jasa.2009.ap08746
3. Abadie, A., & Gardeazabal, J. (2003). The economic costs of conflict: A case study of the Basque Country. American Economic Review, 93(1), 113–132. https://doi.org/10.1257/000282803321455188
4. Ben-Michael, E., Feller, A., & Rothstein, J. (2021). The augmented synthetic control method (NBER Working Paper No. 28885). National Bureau of Economic Research.
5. Drezner, D. W. (1999). The sanctions paradox: Economic statecraft and international relations. Cambridge University Press.
6. Drezner, D. W. (2003). The hidden hand of economic coercion. International Organization, 57(3), 643–659. https://doi.org/10.1017/S0020818303573052
7. Efron, B. (1979). Bootstrap methods: Another look at the jackknife. The Annals of Statistics, 7(1), 1–26.
8. Energy Institute. (2026). Statistical review of world energy. Compiled with major processing by Our World in Data, with gaps filled from U.S. Energy Information Administration international energy data.
9. Farrell, H., & Newman, A. L. (2019). Weaponized interdependence: How global economic networks shape state coercion. International Security, 44(1), 42–79. https://doi.org/10.1162/ISEC_a_00351
10. Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., … Oliphant, T. E. (2020). Array programming with NumPy. Nature, 585(7825), 357–362.
11. Hirschman, A. O. (1945). National power and the structure of foreign trade. University of California Press.
12. Hufbauer, G. C., Schott, J. J., Elliott, K. A., & Oegg, B. (2007). Economic sanctions reconsidered (3rd ed.). Peterson Institute for International Economics.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

13. Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. Computing in Science & Engineering, 9(3), 90–95.
14. Kilian, L., Rapson, D., & Schipper, B. (2024). The impact of the 2022 oil embargo and price cap on Russian oil prices (CEPR Discussion Paper No. 18934). Centre for Economic Policy Research.
15. Kronstadt, K. A. (2022). India–Russia relations and implications for U.S. interests (CRS Report No. R47221). Congressional Research Service.
16. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., … Duchesnay, É. (2011). Scikit-learn: Machine learning in Python. Journal of Machine Learning Research, 12, 2825–2830.
17. Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., … SciPy 1.0 Contributors. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. Nature Methods, 17(3), 261–272.
18. Williamson, O. E. (1979). Transaction-cost economics: The governance of contractual relations. The Journal of Law and Economics, 22(2), 233–261.