The Price of Resilience:

Input-Cost Shocks in European Supply Chains under EU Sanctions on Russia

Mark Spektor*

February 19, 2026

Abstract

Do broad-based sanctions cause physical supply chain breaks, or do they operate primarily as input-cost shocks? This paper studies the 2022 EU sanctions against Russia as a natural experiment to identify the adjustment margins of European manufacturing. I map legal sanctions into trade data and propagate the shock downstream using the AI-generated Production Network (AIPNET), yielding a product-specific exposure measure capturing indirect reliance on sanctioned inputs. Using a dynamic difference-in-differences design for five major industrial economies (Germany, Denmark, Spain, France, Italy), I find that supply chains bent but did not break. While there is no robust evidence of a collapse in physical import volumes, import prices (unit values) for exposed goods rose sharply. In the benchmark specification – weighted by pre-war import values to capture aggregate economic relevance – a 10 percentage point increase in exposure is associated with a 9.7% increase in relative landed costs (CIF). This "price of resilience" amounted to approximately 11.1 billion USD for the nine considered EU economies in 2022. I show that unlike the US trade war, where policy uncertainty caused extensive margin adjustments, the hard constraint of the 2022 embargo resulted in a pure price premium on stable quantities. The sanctions thus operated effectively as a targeted input-cost shock: they successfully cut direct ties with Russia, but the costs were borne by European industry through a broad-based resilience premium rather than a collapse in output.

JEL Codes: F13, F14, F51, L23
Keywords: Sanctions; Supply Chain Resilience; Production Networks; Trade Diversion; Import Prices

*University of Hamburg and ZBW – Leibniz Information Centre for Economics. Email: Mark.Spektor@uni-hamburg.de. I thank Prashant Garg for providing the AIPNET data (available at https://aipnet.io/) and for valuable discussions on network construction. I also thank Marianne Saam for helpful comments. All remaining errors are my own.

1 Introduction

Economic sanctions have become a central instrument of foreign policy, yet debate remains sharp on how the *sanctioning* economies absorb such shocks. When a systemic upstream supplier is removed from the production network, do import-dependent supply chains collapse due to Leontief-type rigidities (Barrot and Sauvagnat, 2016; Carvalho et al., 2021; Elliott et al., 2022), or do they adapt through costly substitution (Bachmann et al., 2024; Di Comite and Pasimeni, 2022; Hinz et al., 2025)? This paper addresses this question considering the 2022 EU sanctions against Russia as a natural experiment. I provide granular micro-evidence that for the European manufacturing core, the adjustment was entirely one of prices, not quantities. While industry associations warned of cascading production stops and deindustrialization (Moll et al., 2023), I find that firms successfully preserved physical flows by paying a substantial "resilience premium." In 2022, a product with 10% exposure to sanctioned inputs faced a relative landed-cost increase of approximately 9.7%, totaling an aggregate cost of 11.1 billion USD.

This paper addresses that question considering the 2022 European Union (EU) sanctions against Russia as a natural experiment. Following Russia's invasion of Ukraine, the EU implemented a series of sanctions packages, including bans on imports of iron, steel, rubber, and other industrial inputs. These measures severed a dense web of supply links in sectors such as metals, chemicals, and machinery. Public debate regarding the sanctions was sharply divided. Industry associations warned that cutting ties with Russian suppliers would lead to physical shortages and cascading production stops — a view implicitly assuming Leontief-type rigidities in the supply chain.[1] Conversely, economists argued that in a market economy, such shocks would be absorbed through substitution and trade diversion (Bachmann et al., 2024). Retrospective analyses of aggregate output confirm that the economy largely avoided the predicted contraction (Moll et al., 2023). This paper provides granular micro-evidence clarifying the adjustment mechanism for the manufacturing supply chain: firms maintained quantities by accepting higher prices.

The analysis combines two ingredients. First, I map the legal sanctions lists—specifically Annex XXI of Regulation (EU) No 833/2014—into the Harmonized System (HS) to define the set of directly banned products. Second, I use the AI-generated Production Network (AIPNET; Fetzer et al., 2024) to trace these bans through the global supply chain. While AIPNET is validated against official input–output accounts, the granular product-level linkages remain probabilistic. To the extent that these inferred linkages contain noise, the statistically significant effects reported in this paper likely represent a conservative lower

[1] Prominent industry representatives warned of "deindustrialization" and massive output losses in the event of an energy and input embargo; for a retrospective analysis of these predictions, see Moll et al. (2023).

bound of the true input-cost shock. This design isolates the indirect supply-chain channel: I do not estimate the mechanical effect of sanctions on the banned goods themselves (direct exposure), but rather the transmission of this input shock to downstream European products that rely on sanctioned upstream inputs (network exposure). Qualitative inspection of the inferred linkages confirms their economic plausibility: for instance, AIPNET correctly maps upstream inputs such as refined copper (HS 7403) and electronic integrated circuits (HS 8542) to downstream computing equipment parts (HS 8473), consistent with standard industrial requirements.

The empirical strategy is a dynamic difference-in-differences (event-study) design. I use annual HS 6-digit import data for nine EU economies from 2019 to 2024, aggregated over source countries to capture global re-sourcing. For each country–product pair, I compute pre-war exposure to Russian inputs and follow its imports over time. The key comparison is between more and less exposed products *within* the same country and year. To interpret the magnitude, I standardize the effect size to a 10 percentage point (pp) increase in exposure, representing a shift from non-exposed to moderately dependent production (e.g., from 0% to 10% reliance on sanctioned inputs). For the benchmark estimate of $\beta \approx 0.93$, this translates to a **9.7%** relative increase in costs. This standardized metric allows for direct comparison across outcomes and specifications. Country–product fixed effects control for time-invariant comparative advantage, while country–year fixed effects absorb common macro shocks such as aggregate demand, inflation, and exchange rate movements. The event-study coefficients trace how exposed products deviate from less exposed ones around the 2022 sanctions shock.

Three results emerge.

First, the data do not support the hypothesis of a systemic collapse in physical import volumes. While the point estimate for quantities is negative, it is statistically indistinguishable from zero. Given the wide confidence intervals, moderate friction-driven volume reductions cannot be ruled out; however, the catastrophic rupture predicted by some industry associations—a cascading physical sudden stop—is rejected by the data. Instead, the primary margin of adjustment is nominal: despite the removal of a major upstream supplier, European firms largely preserved the physical flow of critical inputs in the immediate aftermath of the shock. While a negative adjustment emerges in 2024, suggesting potential medium-term substitution, the immediate supply chain response was characterized by resilience rather than rupture.

Second, import prices (unit values) for exposed goods rise sharply. In the baseline specification for the five major industrial economies (Germany, Denmark, Spain, France, and Italy), the 2022 event-year coefficient for unit values is approximately 1.07 log points. Crucially, while unweighted nominal values reflect heterogeneity across thousands of small, noisy trade flows, weighting by pre-war import volumes is necessary to capture the aggregate economic incidence of the shock. Economic shocks aggregate through volumes, not product counts. The divergence between the noisy unweighted estimate and the precise weighted coefficient ($\hat{\beta} \approx 0.93$) confirms that the cost shock is concentrated in high-volume trade flows rather than peripheral products. Weighting is thus essential to identify the price impact relevant to the aggregate economy. To interpret the effect size, consider a representative product with a 10% pre-war network reliance on sanctioned inputs. The weighted estimate implies a relative cost increase of approximately 9.7% in 2022. This suggests that for affected supply chains, the shock operated as a substantial input-cost increase—a "resilience premium" consistent with the costs of re-routing and spot-market substitution. The sharp rise in CIF unit values likely reflects the mirror image of the frictions documented for the Russian economy, where circumvention through new intermediaries increased logistics costs (Emlinger and Lefebvre, 2025). As European firms competed for non-Russian supplies in a tighter global market, they effectively internalized these costs to maintain production lines.

Third, the pattern is not an energy story in disguise. Excluding HS 27 (mineral fuels) leaves the main value effect almost unchanged and reveals similar price dynamics for non-energy manufacturing inputs. Additional robustness checks—including product-specific trends, placebo timing, leave-one-out exercises for countries, years, and broad industrial sectors (HS 2-digit chapters), and alternative estimators—all support the same conclusion: the sanctions shock shows up as a large, broad-based increase in import prices for exposed products rather than a collapse in physical availability.

The evidence suggests that the EU Single Market provides resilience not by preventing price increases, but by facilitating the efficient reallocation of scarce inputs across member states. Exposed products continue to arrive, yet at much higher prices. In this sense, the sanctions operate as a sizeable supply shock that results in a sharp input-price increase—a "price of resilience" that is ultimately borne by firms and consumers in the sanctioning economies.

This paper makes three primary contributions to the literature on trade shocks and production networks. First, I document a sharp dichotomy in the real effects of trade shocks. In contrast to the US trade war literature, where tariff *uncertainty* led to severed buyer–supplier relationships (Handley et al., 2024) or reduced manufacturing employment (Flaaen and Pierce, 2019), I find that the *certainty* of the 2022 embargo manifested as a pure price premium on stable quantities. This incidence mirrors the 2018 trade war, where domestic agents bore the full cost of import restrictions (Fajgelbaum et al., 2020; Amiti et al., 2019; Amiti et al., 2020; Cavallo et al., 2021) or relocated production (Flaaen et al., 2020), but implies a distinct adjustment margin: rather than accepting volume losses, EU firms paid the “price of resilience” to maintain physical flows. Second, I add to the literature on supply chain resilience and shock propagation (Acemoglu et al., 2016; Baqaee and Farhi, 2019). The results provide empirical support for theories positing that network stability requires costly investments (Grossman et al., 2024) and complement quantitative assessments of the welfare costs of geopolitical decoupling (Attinasi et al., 2025). My findings empirically validate the micro-level mechanism described by Heise et al. (2025), where relationship-specific procurement commands a premium over spot-market alternatives. The sanctions compelled firms to replicate these high-reliability relationships with new suppliers in a compressed timeframe, effectively capitalizing the relationship premium into prices to avoid production stops. Third, I contribute to the sanctions literature by identifying the supply-side mirror image of the “friendly fire” mechanism (Crozet and Hinz, 2020) and contrasting it with the structural adaptation of the 2022 regime (Itskhoki and Ribakova, 2024). My findings complement evidence on the costs borne by the Russian economy: while Russia successfully diverted trade toward non-sanctioning origins, this resulted in a 12–15% spike in import prices driven by higher logistics costs and increased margins from new suppliers (Emlinger and Lefebvre, 2025). By documenting a corresponding “price of resilience” for European manufacturing, this paper completes the picture of sanctions as a symmetric terms-of-trade shock in a globalized production network. Methodologically, I demonstrate the utility of AI-based methods to map economic linkages (Fetzer et al., 2024; Korinek, 2023). By using an event-study design that accounts for treatment effect heterogeneity over time (Baker et al., 2025; Callaway and Sant’Anna, 2021), and exploring robustness specifications that account for sectoral inflation (Handley et al., 2024) and global product trends (Crozet and Hinz, 2020), I provide robust evidence that distinguishes quantity from price adjustments while adopting inference procedures robust to shift-share correlation structures (Adao et al., 2019).

The rest of the paper is structured as follows. Section 2 describes the institutional setting, the AIPNET network, and the construction of the exposure measure and trade data. Section 3 outlines the empirical strategy and explains how to interpret the event-study coefficients. Section 4 presents the main results and robustness exercises, including literature-backed fixed-effect specifications and alternative estimators. Section 5 discusses the underlying mechanisms, the role of market integration, and implications for automation. Section 6 concludes with policy implications and directions for future research.

2 Institutional Setting and Data

2.1 The sanctions shock

Following Russia's annexation of Crimea in 2014, the European Union introduced sanctions against Russia in several steps. The initial measures in 2014–2015 were relatively targeted. Regulation (EU) No 833/2014 restricted access of major Russian state-owned banks and companies to EU capital markets, imposed an arms embargo, and limited exports of dual-use goods and specialised technologies for deep-water, Arctic and shale oil exploration. [2]

These early sanctions mostly affected financial conditions, defence-related equipment, and a narrow set of energy technologies. They left the bulk of ordinary industrial intermediate inputs in metals, chemicals, machinery and related sectors largely untouched.

The nature of the regime changed markedly after the full-scale invasion of Ukraine in February 2022. From March 2022 onwards, successive amendments to Regulation (EU) No 833/2014 introduced broad import bans on Russian coal, other energy products, wood, cement, metals, rubber, plastics, machinery and a wide range of manufactured goods that generate substantial export revenues for Russia.[3] Legally, many of these bans are implemented through Articles 3g, 3i and related provisions and are operationalised via annexes that list affected CN product codes. Annex XXI, in particular, contains a detailed list of goods and technologies covered by Article 3i, spanning raw materials, chemical and rubber products, semi-finished metals, plastics, glass, paper, textiles, machinery, vehicles, electronics and various consumer durables.[4]

Subsequent packages in late 2022 and 2023 mainly deepened and refined this regime by extending the list of covered products and closing loopholes, without altering its basic structure. This marks a fundamental regime shift compared to 2014. As Itskhoki and Ribakova (2024) argue, the 2014 measures operated primarily through financial uncertainty and signaling. In contrast, the 2022 regime, interacting with Russia's pre-war "Fortress Russia" strategy to insulate its economy, erected hard physical barriers to trade. For European firms, this shifted the adjustment margin from managing country

[2] See Council Regulation (EU) No 833/2014 of 31 July 2014 concerning restrictive measures in view of Russia's actions destabilising the situation in Ukraine, *Official Journal of the European Union*, L 229, 31.7.2014, p. 1–11.

[3] For a comprehensive timeline and summary of the successive packages, see the European Council's official overview: "Timeline - EU restrictive measures against Russia over Ukraine" (accessed February 4, 2026).

[4] Legally, the relevant product list is defined in Annex XXI of Regulation (EU) No 833/2014. This annex was significantly expanded by the 12th sanctions package. See Council Regulation (EU) 2023/2878 of 18 December 2023 amending Regulation (EU) No 833/2014, *Official Journal of the European Union*, L, 2023/2878, 18.12.2023.

risk (2014) to managing physical substitution (2022). Consequently, for the purposes of this paper, the key feature is the transition to a broad-based import ban on raw materials and industrial intermediate goods, forcing a structural rather than merely financial reconfiguration of supply chains. Economically, the 2022 packages remove a dense set of Russian input varieties from the opportunity set of European firms in metals, chemicals, plastics, machinery and related manufacturing supply chains.

As documented in Appendix Figure A2, this legal shock translated into an immediate physical sudden stop. Direct imports of sanctioned goods from Russia collapsed in quantity terms immediately following the implementation of Annex XXI. This sharp physical disconnect acts as the first stage of the mechanism, forcing European buyers to substitute missing volumes through the global network and incurring the associated price premiums.[5]

For the empirical design, the crucial property is that the legal product lists can be translated into trade-data categories and combined with pre-war import patterns to construct a product-level exposure measure. I take 2019–2021 as the pre-sanctions benchmark period: while targeted measures were already in place after 2014, they did not systematically affect the broad universe of production inputs considered here, so that pre-war trade flows between EU economies and Russia can be interpreted as reflecting business-as-usual sourcing.

2.2 Network exposure via AIPNET

Standard input–output tables describe average linkages between broad sectors. In contrast, many supply chain vulnerabilities operate through much more granular connections. To capture such dependencies, this paper uses the AI-generated Production Network (AIPNET; Fetzer et al., 2024). AIPNET constructs a directed graph of global production by querying the GPT-4o large language model. The construction follows a "build-prune" topology: first, the model generates potential input–output linkages based on HS product descriptions; second, these candidate edges are pruned using a verification step to remove spurious connections. The nodes are reconciled to the official Harmonized System via vector embeddings (text-embedding-3-large), ensuring semantic consistency between the AI-generated descriptions and official trade nomenclatures. Crucially, Fetzer et al. (2024) validate this granular structure against official US and Mexican input–output tables, documenting that the aggregated AI-inferred linkages strongly correlate with observed

[5]The consolidated Annex XXI covers a wide range of industrial inputs. For a breakdown of the covered economic categories (including cement, rubber products, wood, spirits, and high-end seafood), see the European Commission's explanatory notes: "Sanctions adopted following Russia's military aggression against Ukraine" (accessed February 4, 2026).

industrial flows.

At a high level, AIPNET encodes statements of the form "product p is used as an input in the production of product q" at the HS 6 level. For example, certain nickel products are coded as inputs into battery components, which in turn are inputs into electric vehicles. The result is a global network of product-level upstream and downstream relationships.

I combine the network with the legal sanctions list and pre-war trade data in three steps:

1. **Identify directly sanctioned inputs.** Map each legally banned item in Annex XXI to its HS 6 code. This yields a set $\mathcal{S}$ of directly sanctioned products, interpreted as *upstream inputs* whose supply from Russia is restricted after 2022.
2. **Propagate the shock downstream.** Using AIPNET, identify downstream user products i that rely on sanctioned inputs $u \in \mathcal{S}$. Let w_{ui} denote the AIPNET link intensity from input u to downstream product i, calculated as the row-normalized edge consistency across independent model iterations.
3. **Construct country–product exposure.** For each EU country c and downstream product i, define exposure as the pre-war reliance on Russian supply embodied in sanctioned upstream inputs u:

$$\text{Exposure}_{ci} = \sum_{u \in \mathcal{S}} w_{ui} \, \text{Reliance}_{cu},$$

   where Reliance_{cu} denotes the direct pre-war import share of input u sourced from Russia by country c. The resulting index, Exposure_{ci}, captures the *indirect* input-cost shock faced by the producer of product i. Throughout the paper, I distinguish between Reliance_{cu} (direct trade dependence for inputs) and Exposure_{ci} (network-propagated vulnerability for downstream goods).

The pre-war trade data used in step (iii) come from the UN Comtrade database. I retrieve monthly HS 6-digit import flows for ten EU economies via the official Comtrade API using the `comtradr` package in R. For each reporter country and HS 6 code, I aggregate imports from all partners ("World") and from Russia over the pre-war period 2019–2021 and construct a variable Reliance_{cu} as the share of imports of input u in country c sourced from Russia before the sanctions. While this metric abstracts from domestic sourcing, it precisely identifies the vulnerability of the import-dependent industrial base. This pre-war reliance is then mapped into the AIPNET network and serves as the weighting component in the exposure index Exposure_{ci}. Appendix A.1 documents the API queries, rate-limit handling, and file structure in detail.

Intuitively, Exposure_{ci} measures the embodied vulnerability in product i. It is high if the production of i requires upstream inputs u that are now sanctioned and were historically sourced from Russia. Ideally, the weights w_{ui} would correspond to physical engineering coefficients (e.g., kilograms of input required per unit of output); in the absence of granular global InputOutput tables, I rely on AIPNET link intensities as the robust proxy for technological relevance. This construction explicitly restricts the analysis to first-order propagation (upstream input $\rightarrow$ direct downstream user) for two reasons. First, it isolates the non-trivial transmission mechanism: unlike the direct effect on banned goods, which is tautological, the effect on downstream users captures genuine supply chain resilience. Second, limiting the depth of the network minimizes measurement error. Since AIPNET linkages are inferred probabilistically, iterating links through multiple rounds would multiplicatively compound noise. By focusing on the direct downstream link, the measure yields a high-precision estimate of exposure.

This construction follows the identifying logic emphasized in the shift-share literature: causal identification relies on the exogeneity of the *shock* (the sanctions) while allowing the *shares* (pre-war sourcing patterns) to reflect endogenous economic choice (Borusyak et al., 2022).

In legal terms, Annex XXI is written in the EU Combined Nomenclature (CN) at the 8-digit level. The exposure measure uses the consolidated Annex XXI list as of late 2023. This approach treats the sanctions regime as a single structural shock commencing in 2022. By defining exposure based on the consolidated list, the measure captures the forward-looking risk assessment of firms anticipating the full tightening of the regime. While some specific products were formally added later in the 2022–2023 cycle, the core industrial inputs (wood, steel, energy, cement) were restricted early, and firms likely priced in the trajectory of the expanding ban. Methodologically, including late-sanctioned goods in the treatment group for the full 2022 period introduces classical measurement error (treating untreated units as treated early on). Under standard assumptions (Jeffrey M. Wooldridge, 2010), this misclassification biases the estimated treatment effect toward zero (attenuation), ensuring that the estimates represent a conservative lower bound of the true cost shock. To operationalize the bans, I parse the EU Combined Nomenclature (CN) entries deterministically. For example, raw strings such as "ex 28 25" (indicating a specific subset of a chapter) are standardized into canonical numeric prefixes (e.g., padding "306" to "0306"). This procedure matches the 205 distinct legal lines against the full Harmonized System universe, yielding 1,390 distinct *sanctioned upstream inputs* ($u \in \mathcal{S}$). Through the AIPNET production network, these sanctioned inputs are embodied in 2,442 downstream intermediate goods (as shown in Figure 1).

To make these legal product definitions compatible with the trade data and the AIPNET network, I proceed in two steps. First, the CN entries are mapped into HS 2022 6-digit codes. This is done by treating the raw legal entry (variable `cn_code_original`) as a code prefix. Whitespace is removed (e.g., "1604 31 00" becomes "16043100"), missing leading zeros are padded (e.g., "306" becomes "0306"), and codes longer than 6 digits are truncated to the HS 6 level. These normalized strings are then used to identify all matching codes within the official HS 2022 universe. Using the official UNSD HS 2022–HS 2002 correlation table as the universe of valid HS 2022 codes, all HS 6 codes whose six-digit string starts with the cleaned legal prefix are marked as sanctioned. This logic allows the sanctions list to cover both fully specified items (8-digit) and broader headings (2- or 4-digit) in a way that respects the official HS 2022 product universe. In the data, the procedure maps the 205 distinct legal lines from Annex XXI into **1,390** distinct *directly sanctioned* HS 6 products. These products constitute the set of sanctioned upstream inputs ($u \in \mathcal{S}$). Through the AIPNET production network, the input shock propagates to **2,442** *downstream* intermediate goods (i) with non-zero exposure, affecting a total pre-war market volume of 5.2 trillion USD (aggregated across the nine sample countries).

The estimation sample includes all **2,442** exposed downstream products, regardless of whether they are also subject to direct import bans. This is a deliberate design choice: the "price of resilience" encompasses the total cost of substituting away from Russian supply chains. Whether a European buyer pays a premium because they must buy non-Russian steel pipes (direct substitution) or because their European supplier pays more for energy and iron (indirect network cost), the economic outcome is identical: a higher landed cost for the intermediate good. By controlling for general Russian exposure in robustness checks, I verify that the effect is driven by the specific legal constraints rather than generic decoupling.

Technical mapping to the network vintage follows in a second step. The sanctioned HS 2022 codes are mapped to HS 2002 codes using the official HS 2022–HS 2002 correlation table. Because HS revisions are many-to-many, a given HS 2022 code may map into multiple HS 2002 codes. This mapping is necessary because the AIPNET architecture is native to the HS 2002 vintage. I assign a sanctions flag to every HS 2002 code that is linked to at least one sanctioned HS 2022 code. These flagged HS 2002 products constitute the set of *sanctioned upstream inputs* in the AIPNET network.

Using AIPNET edges, these sanctioned upstream nodes are then propagated *downstream* to user products. The exposure index Exposure_{ci} is constructed for each downstream HS 2002 product as the weighted reliance on Russian supply across its sanctioned upstream inputs. Conceptually, Exposure_{ci} is therefore a country-specific *embodied input-cost vulnerability* measure at the HS 6 level in HS 2002 space.

Two features of this construction warrant discussion regarding validity and measurement error. First, regarding the AI-inferred linkages, Fetzer et al. (2024) rigorously validate the AIPNET structure against official US and Mexican input-output tables, demonstrating that the aggregated probabilistic links recover the ground-truth industrial structure with high fidelity. Second, the mapping from 8-digit legal sanctions to 6-digit trade data inevitably introduces classification noise. Econometrically, both sources of error—probabilistic network links and HS6 aggregation—introduce classical measurement error into the explanatory variable. Under standard assumptions, this introduces attenuation bias (Jeffrey M. Wooldridge, 2010), systematically driving the estimated coefficients toward zero. Consequently, the statistically significant effects reported below likely represent a conservative lower bound of the true cost shock.

2.3 Trade data and sample construction

The trade data come from the UN Comtrade database. Monthly HS 6-digit import records are retrieved for ten EU reporters over 2019–2024 via the official Comtrade API in R and aggregated to annual flows. The choice of annual frequency balances temporal precision with signal stability. Monthly trade data at the granular HS 6 level exhibit substantial volatility, seasonality, and "lumpiness" (intermittent zeros), which complicate the identification of structural supply chain adjustments. Aggregating to the annual level smooths out transitory logistical frictions and isolates the persistent shift in sourcing patterns, consistent with standard practice in structural gravity estimations. To document the anatomy of the shock, Appendix C provides event plots illustrating the divergence between nominal values and physical quantities for direct Russian imports. The raw records report import values in current US dollars and, where available, net weight in kilograms, together with the HS 6 code, reporting country, and partner. For each reporter, imports are downloaded separately for partner "World" (all origins) and Russia, which allows the construction of pre-war Russia reliance as the 2019–2021 import share sourced from Russia.

The estimation sample is constructed based on data quality. While raw import flows were queried for ten EU economies, Austria is excluded from the regression analysis due to systematic gaps in the UN Comtrade API returns for the pre-war period (2019–2021), which prevent the construction of a reliable exposure baseline (see Appendix A.1 for API diagnostics). Consequently, the identification sample comprises nine EU member states (Belgium, Germany, Denmark, Spain, Finland, France, Italy, the Netherlands, and Sweden), ensuring that all estimated effects are derived from consistent pre-war dependencies.

The analysis focuses on *intermediate goods*. To operationalise this, I use the UN Broad Economic Categories (BEC) classification as provided in the World Bank WITS concordance between HS 2002 and BEC codes. Following the UN System of National Accounts (SNA) convention, I restrict the estimation sample to intermediate goods, defined as primary and processed industrial supplies, fuels used as intermediate inputs, and parts and accessories of capital goods and transport equipment. Capital goods and final consumption items are excluded. Appendix A.2 provides the exact mapping of BEC codes and the corresponding data processing steps.

The WITS concordance covers all HS 2002 codes present in the panel, so no product is lost due to missing BEC information. In the full HS 2002 panel before filtering, intermediate goods account for about 62% of observations and 42% of total import value. In terms of physical bulk, they represent approximately 83% of the total net weight. This high share of physical volume reflects the logistical prominence of raw materials and heavy industrial inputs relative to higher-value but lighter capital and consumption goods. Note that because the analysis relies on customs data (HS classification), it is strictly limited to physical merchandise trade; trade in services (e.g., software, consulting) is structurally excluded.

A simple consistency check on a set of robotics-related HS codes illustrates the strictness of this intermediate-goods definition. Parts and accessories such as HS 847990, 848340 and 853710 are classified as intermediate goods (BEC 42), whereas HS 850131 (small electric motors) is classified by UN BEC as a capital good (BEC 41) and is therefore excluded. This treatment follows the official UN BEC conventions and keeps the intermediate-goods definition transparent and reproducible, even though some technologically intermediate components are formally labelled as capital goods.

To facilitate the decomposition into nominal and physical adjustments, three outcome variables are constructed:

$$\begin{aligned}
\log \mathrm{Val}_{cit} &= \log(1 + \text{Import Value}_{cit}),\\
\log \mathrm{Kg}_{cit} &= \log(1 + \text{Net Weight}_{cit}),\\
\log \mathrm{UV}_{cit} &= \log \mathrm{Val}_{cit} - \log \mathrm{Kg}_{cit},
\end{aligned}$$

where the additive "1" avoids dropping zero flows and treats zeros consistently across outcomes.[6]

Because $\log(1+x)$ is used for both components, $\log \mathrm{UV}_{cit}$ serves as a proxy for the import price. Crucially, in this context, changes in unit values capture the full "landed cost" shock: they reflect not only increases in the free-on-board (FOB) price but also shifts in composition toward more expensive suppliers, higher logistics premia for re-routed goods, and potential changes in input quality. This composite metric accurately represents the effective cost shock facing European buyers.

[6] All results are robust to alternative zero-handling choices, including dropping zero flows and using inverse hyperbolic sine transformations.

While granular, the product-level data imposes two limitations. First, I do not observe firm-to-firm matches, meaning I cannot distinguish between existing suppliers raising prices and the setup costs of new relationships. Second, quality differentiation within HS 6 codes is unobserved. However, given the magnitude and suddenness of the price spike across diverse sectors in 2022, a simultaneous, coordinated quality upgrade provides an unlikely alternative explanation for the aggregate results. The estimated effects should therefore be interpreted as the net market price of resilience, inclusive of both markups and frictional switching costs.

The final identification sample consists of nine EU economies (Belgium, Germany, Denmark, Spain, Finland, France, Italy, the Netherlands, and Sweden). Austria is excluded from the regression analysis because the API return structure prevented the construction of a reliable pre-war Russia-reliance baseline. The data construction proceeds in steps. Starting from the AIPNET universe of 4,955 distinct HS 2002 product codes, I merge the UN Broad Economic Categories (BEC) classification. Filtering for intermediate goods reduces the sample to 4,570 relevant industrial inputs. Within this set, the sanctions on 1,390 upstream inputs propagate network exposure to 2,442 downstream products with non-zero exposure. The event-study specifications in Section 3 are estimated on this BEC-filtered panel.

Finally, the interpretation of the trade data requires precise economic definition. First, I use unit values (value divided by quantity) as the standard proxy for import prices. While unit values can theoretically reflect compositional shifts, the inclusion of country–product fixed effects ensures that identification relies on deviations in landed costs relative to the product's pre-war mean. Second, import values in UN Comtrade are reported on a CIF (Cost, Insurance, Freight) basis. In the context of sanctions, this provides a useful metric for total procurement costs. The "price of resilience" consists not only of the free-on-board (FOB) price charged by alternative suppliers but also of the additional logistics costs, insurance premia, and intermediary margins required to re-route trade from Russia to more distant origins. CIF unit values thus correctly capture the effective marginal cost shock hitting European producers, incorporating both the scarcity rent of the input and the frictional cost of the new supply chain topology.

Finally, a note on data granularity is warranted. The analysis relies on product-level (HS 6-digit) rather than firm-level customs data. While firm-level data allows for the observation of specific buyer-supplier terminations, it can obscure the aggregate market adjustment if firms simply switch intermediaries. For the question of economic resilience, the relevant metric is not the survival of a specific firm-to-firm contract, but the ability of the aggregate supply chain to clear the market for critical inputs. By focusing on country-product aggregates, I capture the net equilibrium effect of substitution and repricing across all available firms. If quantities remain stable while unit values rise at this level, it indicates that the market as a whole successfully re-routed supply, capturing the systemic "price of resilience" paid by the economy.

The baseline estimation sample is the five-country industrial core (Germany, Denmark, Spain, France, and Italy). This group excludes major logistical gateways (Belgium, Netherlands) to minimize measurement error arising from re-export valuation adjustments ("Rotterdam effect"), thus isolating the cost shock to domestic manufacturing. Robustness checks confirm that results hold in the extended sample of nine countries, though with higher variance due to these transit effects (see Section 4.6).

Figure 1 illustrates the structural mechanics of the shock. Panel A documents the disconnect between the political target and the economic reality. While the legal sanctions targeted direct imports from Russia worth approximately **$21 billion** (in the five core economies), the ban applied to a global input market valued at **$2.4 trillion**. Through the production network, this input constraint potentially affected **2,442 downstream intermediate goods** with a total market volume of **$5.2 trillion**. This massive expansion reflects the granularity of the supply chain: specific sanctioned inputs (e.g., raw aluminum) are embodied in a wide range of downstream intermediates, rendering a large share of the manufacturing sector exposed to upstream friction regardless of direct trade links.

Figure 2 further decomposes this exposure into specific products. While energy inputs dominate the raw exposure list (Left Panel), the non-energy industrial core (Right Panel) is heavily concentrated in the automotive and machinery supply chains. Critical inputs such as vehicle parts (HS 870899) and data processing machines (HS 8471) exhibit massive exposure volumes, confirming that the constructed index captures high-specificity manufacturing inputs where substitution is costly.

Table 1. Sample Composition: Intermediate Goods (UN BEC).

Category	Obs.	% Rows	Value (Bn $)	Wgt (Bn kg)	% Val
Intermediate	**152,821**	**62.1%**	**18,595.0**	**8,508.8**	**41.6%**
Capital	34,356	14.0%	15,254.8	306.4	34.1%
Consumption	56,731	23.1%	9,330.3	948.2	20.9%
Unclassified	2,118	0.9%	1,495.4	556.2	3.4%
Total	**246,026**	**100%**	**44,675.5**	**10,319.6**	**100%**

Notes: Raw HS 2002 panel composition (2019–2024) before filtering. Analysis sample restricted to "Intermediate" (BEC 111, 121, 21, 22, 31, 32, 42, 53).

Replication and Transparency. To ensure full reproducibility, the data construction pipeline is automated end-to-end. I provide the complete R and Stata code, including the deterministic mapping of CN8 legal texts to HS6 codes and the API querying protocols

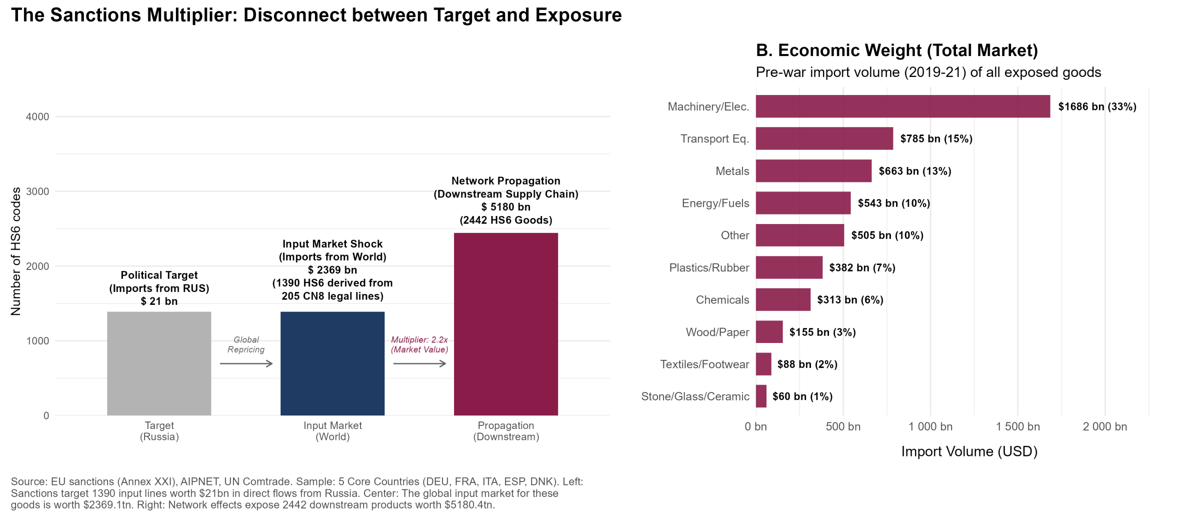

Figure 1. Sanctions Coverage and Network Propagation.

Notes: Panel A compares pre-war import volumes (2019–2021) across three layers: (1) Direct imports from Russia of sanctioned goods; (2) Total global imports of these sanctioned inputs; (3) Total market volume of downstream intermediate goods containing sanctioned inputs (via AIPNET). Panel B displays the sectoral composition of the exposed downstream products. Sample: Five-country industrial core (DEU, FRA, ITA, ESP, DNK).

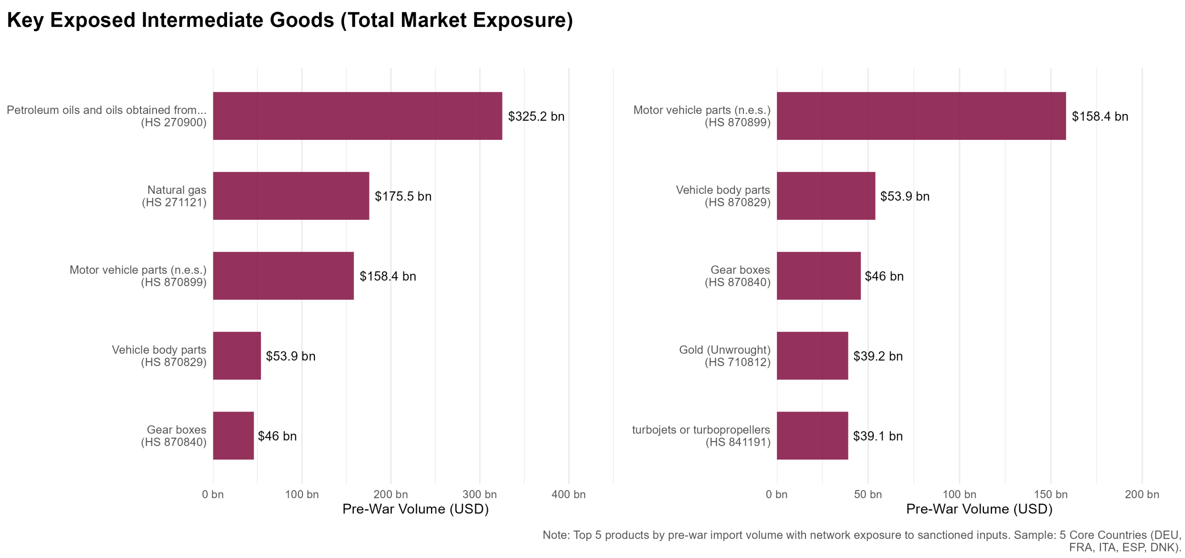

Figure 2. Highest-Volume Exposed Intermediate Goods.

Notes: The charts display the five HS 6-digit product categories with the highest pre-war import volume (2019–2021) among all exposed goods. Left Panel: Including mineral fuels (HS 27). Right Panel: Excluding mineral fuels. Sample: Five-country industrial core.

that handle Comtrade rate limits. Appendix A documents the recursive chunking strategy used to recover missing data and the exact regex logic linking Annex XXI to the AIPNET structure. Furthermore, to rule out selection bias, Appendix B provides permutation-based pre-trend diagnostics for all 511 possible country coalitions, demonstrating that the parallel trends assumption holds robustly for the selected industrial core sample.

3 Empirical Strategy

3.1 Baseline event-study specification

The empirical strategy relies on a shift-share difference-in-differences design. Identification rests on the interaction of an exogenous shock (the 2022 EU sanctions) with predetermined cross-sectional variation in exposure (pre-war reliance on Russian inputs propagated through the production network). The identifying assumption is that, conditional on fixed effects, products with high indirect exposure to Russia would have evolved similarly to less exposed products in the absence of the sanctions. I address potential confounders, such as sector-specific inflation or global demand shocks, by saturating the model with high-dimensional fixed effects.

Formally, I estimate the effect on three distinct margins: total nominal value ($\log \mathrm{Val}$), physical quantity ($\log \mathrm{Kg}$), and unit value ($\log \mathrm{UV}$). For each outcome Y, the specification is:

$$Y_{cit} = \alpha_{ci} + \lambda_{ct} + \sum_{\tau \neq -1} \beta_\tau \Big(\mathbb{I}\{t = 2022 + \tau\} \times \mathrm{Exposure}_{ci} \Big) + \varepsilon_{cit}, \tag{1}$$

where α_{ci} and λ_{ct} represent high-dimensional fixed effects. Although the dependent variable is specified in levels (logs), the inclusion of country–product (α_{ci}) and country–year (λ_{ct}) fixed effects ensures that this specification is econometrically equivalent to a difference-in-differences estimator, identifying deviations from the product-specific mean relative to the common trend. The parameters are defined as follows:

- c indexes countries, i HS 6 products, and t years;
- α_{ci} are country–product fixed effects, capturing time-invariant comparative advantage and baseline trade relationships;
- λ_{ct} are country–year fixed effects, absorbing country-specific macro shocks such as aggregate demand, inflation, and exchange rates;
- $\mathbb{I}\{t = 2022 + \tau\}$ are event-time dummies, with $\tau = -3, \ldots, +2$ indexing years relative to the pre-war base year 2021 ($\tau = -1$);
- $\mathrm{Exposure}_{ci}$ is the pre-war network-based exposure measure described in Section 2.2;

- ε_{cit} is an error term.

Event time is defined relative to 2022, so that $\tau = t - 2022$ and $\tau \in \{-3, -2, -1, 0, 1, 2\}$ corresponds to calendar years 2019–2024. The year 2021 ($\tau = -1$) is treated as the baseline period, so that β_{-1} is normalised to zero and all coefficients β_τ are interpreted relative to 2021.

The coefficients of interest are the β_τ. Consistent with the framework of Callaway and Sant'Anna (2021), these parameters identify the average treatment effect for the sanctions cohort in each event year. They measure how strongly imports of more exposed products deviate from those of less exposed products, identifying the dynamic path of adjustment while holding constant country–product and country–year effects. This decomposition avoids the weighting biases inherent in static fixed-effect summaries (Baker et al., 2025).

3.2 How to read the coefficients

It is useful to interpret Equation (1) in plain language before turning to the results.

First, the fixed effects ensure that comparisons are made within a very narrow frame. Country–product fixed effects α_{ci} mean that each product in each country is compared only to its own history. Country–year fixed effects λ_{ct} ensure that the model does not attribute to exposure any shocks that hit an entire country in a given year, such as general inflation or GDP changes.

Second, the interaction terms $\mathbb{I}\{t = 2022 + \tau\} \times \text{Exposure}_{ci}$ compare more and less exposed products *within* the same country and year. If β_0 is positive in the value regression, this means that in 2022, highly exposed products exhibited a larger increase in *total nominal import value* relative to 2021 than low-exposure products in the same country. If the corresponding coefficient in the quantity regression is close to zero, then the increase in nominal value must be driven mainly by higher prices (unit values), not by higher physical volumes. This interpretation relies on the assumption that the sanctions shock is exogenous to product-level unobservables. As formalized by Borusyak et al. (2022) for shift-share designs, this allows the underlying pre-war trade network to reflect non-random economic choices without biasing the estimator.

Third, the event-study form allows inspection of pre-trends. The pre-shock coefficients β_{-3} and β_{-2} show whether highly exposed and less exposed products were already diverging before the war. If these coefficients are statistically indistinguishable from zero, the standard parallel-trends assumption underlying difference-in-differences appears plausible.

Throughout the main analysis, standard errors are clustered at the product level. This addresses correlation in shocks and measurement error across countries for the same HS 6 product, a structural concern in shift-share environments (Adao et al., 2019). For pretrend screening exercises in Appendix B, two-way clustered standard errors (by product and country) are reported as a diagnostic.

To assess robustness to few-cluster or non-Gaussian sampling environments, conventional cluster-robust inference is complemented by wild cluster bootstrap procedures (Cameron et al., 2008) implemented at the product-cluster level. Wherever bootstrap inference is reported, the manuscript reports it consistently (same null, same sidedness) and separates it from analytic cluster-robust statistics.

3.3 Parallel trends and sample definition

Identification hinges on the assumption that, absent the sanctions, more and less exposed products would have followed parallel trends within each country. To balance internal validity with broad coverage, results are reported for two samples.

The **full sample** comprises the nine EU economies for which granular pre-war exposure can be constructed (Belgium, Germany, Denmark, Spain, Finland, France, Italy, the Netherlands, and Sweden). This sample maximizes coverage of the Single Market and is used to assess heterogeneity.

The **benchmark sample** focuses on the five-country industrial core (Germany, Denmark, Spain, France, and Italy). This selection strategy is driven by measurement, not outcomes. The full nine-country sample includes major logistical gateways (Belgium, Netherlands) where re-export valuation adjustments ("Rotterdam effect") introduce substantial noise into unit values, as well as smaller Nordic economies (Finland, Sweden) that may face asymmetric logistical shocks. The industrial core isolates the consumption shock to the main domestic manufacturing centers. To ensure that this selection does not introduce bias, I implement a formal diagnostic akin to synthetic control construction (Abadie et al., 2010) and apply it to all possible country combinations.

For any candidate country group G, I estimate the pre-period analogue of Equation (1) on 2019–2021 data:

$$Y_{cit} = \alpha_{ci} + \lambda_{ct} + \sum_{\tau \in \{-3,-2,-1\}} \beta_\tau^{\text{pre}} \Big(\mathbb{I}\{t = 2022 + \tau\} \times \text{Exposure}_{ci}^G \Big) + \varepsilon_{cit}, \tag{2}$$

where Y_{cit} denotes the respective log outcome (value, quantity, or unit value) and $\text{Exposure}_{ci}^G = \mathbb{I}\{c \in G\} \times \text{Exposure}_{ci}$. To systematically screen for pre-trends across hundreds of candidate groups, I report joint tests for $\beta_{-3}^{\text{pre}} = \beta_{-2}^{\text{pre}} = 0$ and apply the Benjamini and Hochberg (1995) false discovery rate correction (q-values) to summarize diagnostics across compositions.

A common concern in trade panel analysis is the comparability of heterogeneous goods (e.g., comparing steel to sugar). The validity of this design does not require these goods to have similar levels or volatility. The country–product fixed effects α_{ci} absorb all time-invariant heterogeneity (levels), while country–year fixed effects λ_{ct} absorb all macro-shocks (inflation, business cycles) hitting the importer. The identifying assumption is strictly that the *differential* evolution of exposed versus non-exposed goods within the same country would have been zero in the absence of the shock. The flat pre-trends in 2019–2021 (Figure 3) and the convergence of placebo estimates (Figure 5) provide strong empirical support that this condition holds, validating the comparison of diverse intermediate goods conditional on fixed effects.

Table 2 confirms that, in the benchmark sample, exposed and non-exposed products were on statistically indistinguishable trajectories prior to 2022. The joint tests for pre-trend coefficients yield high p-values, supporting the parallel trends assumption:

- Import values: $p = 0.97$;
- Import quantities: $p = 0.97$;
- Unit values: $p = 0.69$.

The results section primarily reports estimates from this identified group, while using the full sample to assess heterogeneity.

3.4 Alternative fixed-effect specifications

To benchmark the findings against recent work on trade shocks and production networks, I estimate alternative specifications that impose more demanding fixed-effect structures, following Handley et al. (2024) and Crozet and Hinz (2020).

Sector–country–year fixed effects (Handley et al., 2024). I augment Equation (1) with fixed effects for sector–country–year cells, defined at the HS 2 level:

$$Y_{cit} = \alpha_{ci} + \lambda_{sct} + \sum_{\tau \neq -1} \beta_\tau^H \Big(\mathbb{I}\{t = 2022 + \tau\} \times \text{Exposure}_{ci} \Big) + \varepsilon_{cit}, \tag{3}$$

where λ_{sct} are HS 2 sector–country–year fixed effects. This specification absorbs all time-varying shocks specific to broad industries (e.g., “Chemicals” or “Iron and Steel”) within each country, effectively controlling for general energy price spikes, sectoral inflation, or aggregate demand shifts. The coefficient β_τ^H thus isolates the variation driven purely by the specific supply-chain exposure of individual HS 6 products relative to the average product in the same sector.

Product–year fixed effects (Crozet and Hinz, 2020). I also estimate a specification with product–year fixed effects:

$$Y_{cit} = \alpha_{ci} + \gamma_{it} + \sum_{\tau \neq -1} \beta_\tau^C \Big(\mathbb{I}\{t = 2022 + \tau\} \times \text{Exposure}_{ci} \Big) + \varepsilon_{cit}, \tag{4}$$

where γ_{it} are HS 6 product–year fixed effects, consistent with the structural gravity framework in Crozet and Hinz (2020), who document substantial "friendly fire"trade losses in non-sanctioned goods driven by country riskduring the 2014 episode. These absorb any shocks that are common to all countries for a given product and year, such as global scarcity or EU-wide price changes.

Economically, Equation (3) filters out national sector shocks, while Equation (4) filters out common product-level shocks across countries. Comparing the baseline coefficients with β_τ^H and β_τ^C reveals to what extent the estimated effects stem from common sectoral versus common product-level movements and how much residual country-specific heterogeneity remains along the exposure dimension.

Table 2. Pre-trend diagnostics for benchmark sample (2019–2021).

Outcome	β_{-3}	β_{-2}	Joint p-value
Import values	0.139	0.168	0.969
	(0.788)	(0.712)	
Import quantities	-0.208	-0.176	0.971
	(0.866)	(1.252)	
Unit values	0.347	0.345	0.686
	(0.441)	(0.831)	

Notes: Pre-period coefficients from Equation (1) estimated on 2019–2021 data only. Joint p-value tests $H_0 : \beta_{-3} = \beta_{-2} = 0$. Sample: DEU, DNK, ESP, FRA, ITA. Standard errors clustered by product.

3.5 Alternative estimators

To complement the dynamic event-study, I estimate two simpler models.

First, I collapse the data into pre- and post-periods at the country–product level:

$$\Delta Y_{ci} = \overline{Y}_{ci}^{\text{post}} - \overline{Y}_{ci}^{\text{pre}},$$

where “pre” is the average over 2019–2021 and “post” the average over 2022–2024. I then regress ΔY_{ci} on Exposure_{ci} and country fixed effects. This estimator sacrifices time dynamics but provides an intuitive “before–after” contrast.

4 Results

This section summarises the main empirical findings. The focus is on the five-country industrial core (Germany, Denmark, Spain, France, and Italy), representing the core of European manufacturing, and on three margins of adjustment: nominal import values (expenditure), physical quantities, and import prices (measured as unit values).

The contrast between the trade margins in the baseline sample suggests a pattern of substitution amid friction. While the sanctions severed direct ties with Russia, European manufacturing exhibited statistically stable physical supply while internalizing significant price increases. Taken together, the evidence is consistent with a targeted input-cost shock—a "price of resilience" paid to maintain production continuity.

4.1 Baseline event-study: values, quantities, and prices

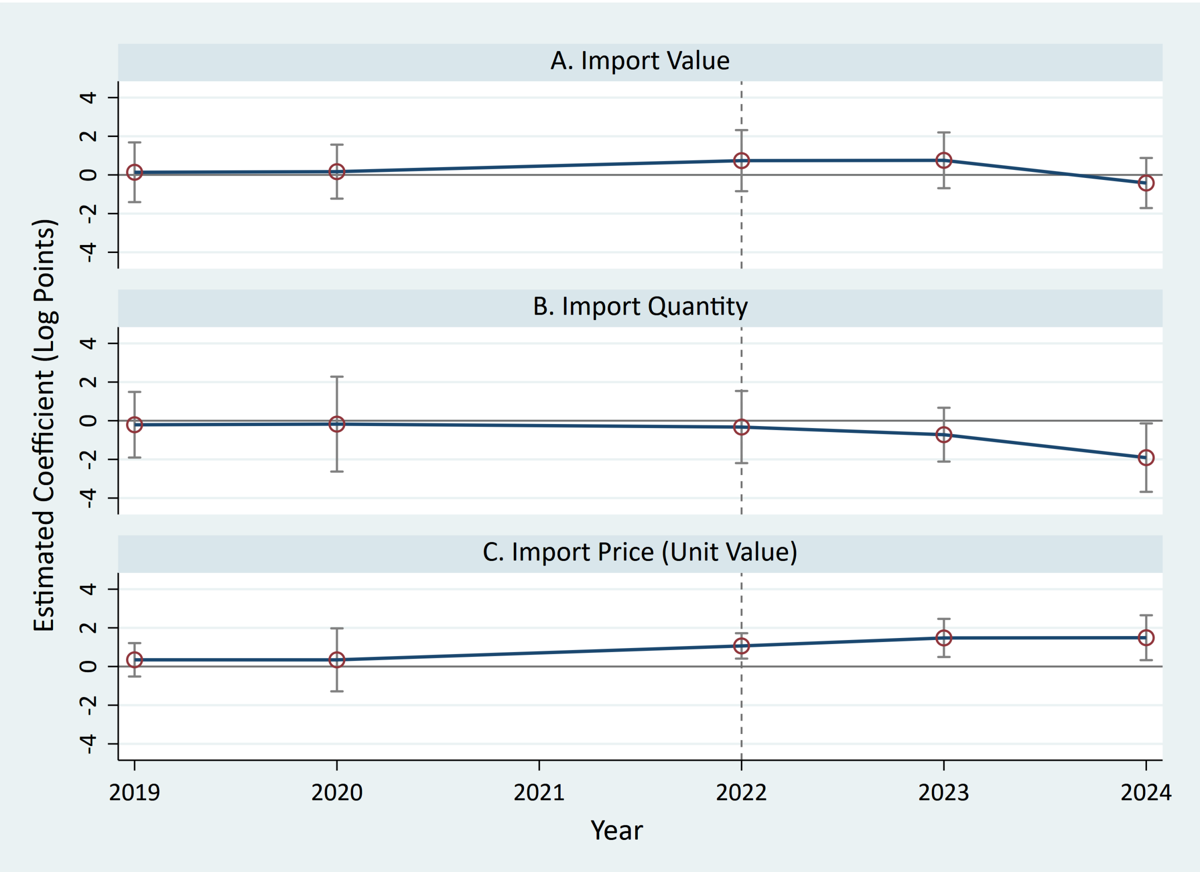

Figure 3. Event-Study Estimates: Import Value, Quantity, and Price.
Notes: Coefficients β_τ from Equation (1) estimating the differential evolution of exposed vs. non-exposed products relative to 2021 ($\tau = -1$). Dependent variables are log(1 + Import Value) (Panel A), log(1 + Net Weight) (Panel B), and log Unit Value (Panel C). Sample: Five-country industrial core. 95% confidence intervals clustered by product. Unweighted estimates.

The decomposition of the shock in Figure 3 reveals three features of the adjustment.

First, the pre-trends support the research design. In all three panels, the coefficients for 2019 and 2020 oscillate closely around zero. This indicates that, prior to the invasion, the import trajectories of highly exposed products were statistically indistinguishable from those of non-exposed products. The divergence begins precisely in 2022, consistent with a causal interpretation of the sanctions shock.

Second, Panel B indicates physical resilience in the short run. The point estimate for import quantities (kg) in the shock year is negative but statistically indistinguishable from zero. Given the confidence intervals, moderate volume responses cannot be ruled out; however, the catastrophic sudden stop predicted by some industry associations is not supported by the data. In this sense, the supply chain bent but did not exhibit an immediate systemic break.

Third, Panels A and C reveal where the nominal cost materialized. Import values and unit values exhibit a synchronized upward shift in 2022. Since quantities remained stable, the rise in total expenditure is driven primarily by the price margin. The 2022 coefficient for unit values indicates that European firms paid a substantial premium—the "price of resilience"—to secure substitutes from alternative origins.

In short, the supply chain bent but did not break. The geopolitical shock was successfully absorbed by the logistics network, but it was fully passed through to import prices, effectively transforming a trade ban into a broad-based input-cost shock for European industry.

Table 3. The Price of Resilience: Baseline Estimates (2022 Effect).

Outcome (log)	(1) Unit Value (USD/kg)	(2) Import Quantity (kg)	(3) Import Value (W) (USD)
Exposure × 2022	1.07***	-0.33	0.93***
	(0.34)	(0.95)	(0.31)
Implied effect (10pp)	*+11.3%*	*-3.2% (n.s.)*	*+9.7%*
Fixed effects	Ctry–Prod, Ctry–Year	Ctry–Prod, Ctry–Year	Ctry–Prod, Ctry–Year
Weighting	None	None	Pre-war import value
Observations	84,022	84,022	83,871

Notes: The table reports the difference-in-differences coefficient β_{2022} from Eq. 1. The sample consists of the five-country industrial core (DEU, FRA, ITA, ESP, DNK). The row "Implied effect (10pp)" converts the log-point coefficient β into the percentage change associated with a 0.10 increase in exposure, calculated as $(\exp(0.1 \times \beta) - 1) \times 100$. Column (3) is weighted by pre-war import values (2019–2021) to reflect aggregate economic relevance, while Columns (1) and (2) report unweighted product-level effects. Standard errors are clustered at the HS6 product level and reported in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels.

Consistent with the dynamic results, Figure 4 shows that the price effect (Unit Value) and nominal value effect are large and positive, while the physical quantity effect is statistically indistinguishable from zero. This pattern confirms that the shock was

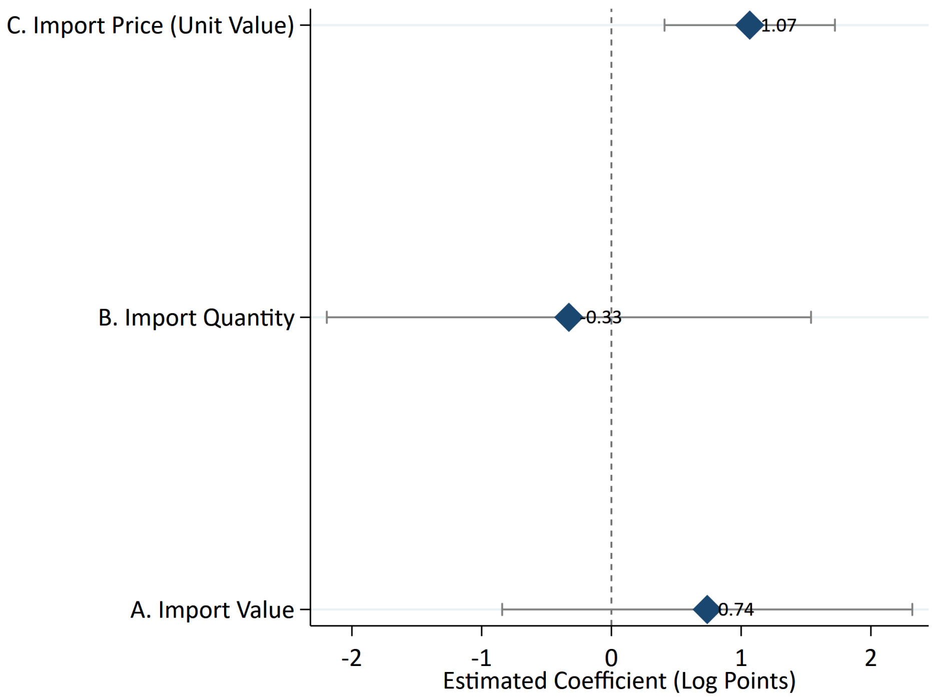

Figure 4. Coefficient Estimates for the Event Year 2022.

Notes: Point estimates for β_0 (interaction term Exposure$_{ci} \times \mathbb{I}_{2022}$) from the baseline specification. Outcomes: Log Import Value, Log Quantity, Log Unit Value. 95% confidence intervals clustered by product.

absorbed primarily through prices rather than volume rationing. While the unweighted value coefficient in the figure includes zero due to noise in small trade flows, the volume-weighted estimate discussed in Table 3 is highly significant (0.93^{***}), indicating that cost shocks are concentrated in economically systemic trade relationships.

Pre-trends. As detailed in **Table 2**, the joint null hypothesis of differential pre-trends is not rejected for any outcome in the industrial core sample. The corresponding point estimates for $\tau = -3$ and $\tau = -2$ remain small in magnitude and switch signs across outcomes, indicating no systematic anticipatory effects. This stability supports the economic rationale for focusing on the industrial core. As shown in the convergence analysis in Figure 5, country combinations that exhibit stable pre-trends systematically converge toward the baseline estimate. This suggests that the result identifies the structural response of manufacturing destinations, whereas the inclusion of transit hubs introduces valuation noise consistent with re-export activities.

Import values. For import values, the unweighted estimates are positive ($\hat{\beta} \approx 0.74$) but statistically imprecise, reflecting heterogeneity in thousands of small, noisy trade flows. However, when weighting by pre-war import values (USD) to capture the aggregate economic impact, the 2022 event-year coefficient becomes highly significant:

$$\hat{\beta}_0^{\text{val, wtd}} = 0.93 \quad \text{(s.e. 0.31)}.$$

Crucially, this weighting is not a statistical choice to inflate significance, but a requirement to capture the aggregate economic incidence. Economic shocks aggregate through volumes, not product counts. An unweighted mean is econometrically valid but economically misleading if the effect is concentrated in high-volume, systemic trade relationships while thousands of small, noisy flows show no reaction. This dual reporting strategy—weighted for incidence, unweighted for mechanism—is deliberate. Volume weighting is required to assess the aggregate economic burden (the "price of resilience"), as cost shocks naturally scale with trade volume. Conversely, unweighted estimates (reported in robustness checks) serve to identify the mechanism: they test whether the price response is a broad-based supply chain phenomenon or merely an artifact driven by a few large commodity flows (e.g., energy). The fact that unweighted estimates point in the same direction, though with greater variance across small products and heterogeneous countries (see Section 4.3), confirms that the shock is not confined to the energy sector but permeates the wider manufacturing network.

In multiplicative terms, moving from zero to full exposure is associated with import values that are roughly $\exp(0.93) \approx 2.5$ times those of non-exposed products within a given country and year, conditional on fixed effects. Because the quantity response is statistically flat, the import-value coefficient in this weighted specification ($\hat{\beta} \approx 0.93$) serves as a conservative proxy for the aggregate price shock. The unweighted unit-value estimates discussed below ($\hat{\beta} \approx 1.07$) are quantitatively consistent with this interpretation.

Post-2022, the value coefficients remain positive in the immediate aftermath but fade over time: $\hat{\beta}_1^{\text{val}} \approx 0.76$ (s.e. 0.74) in 2023, turning negative in 2024 ($\hat{\beta}_2^{\text{val}} \approx -0.42$, s.e. 0.66).

Import quantities. For import quantities, the point estimate is negative but statistically indistinguishable from zero. In 2022, the coefficient is

$$\hat{\beta}_0^{\text{kg}} = -0.33 \quad \text{(s.e. 0.95)},$$

with a wide 95% confidence interval of $[-2.19, 1.54]$. While the point estimate implies a potential volume reduction of approximately 28% ($e^{-0.33} - 1$), the statistical uncertainty prevents rejecting the null hypothesis of no change. Crucially, however, the data strictly reject the hypothesis of a catastrophic "sudden stop" (e.g., -90%) predicted by some industry associations. The adjustment is thus best characterized as a price shock accompanied by moderate, statistically uncertain volume frictions, rather than a systemic quantity collapse.

Unit values (prices). Given rising import values and stable quantities, the adjustment is driven by the price channel. The event-study for unit values confirms a sharp increase. For 2022,

$$\hat{\beta}_0^{\mathrm{uv}} = 1.07 \quad (\text{s.e. } 0.34),$$

with a conventional 95% Wald interval of $[0.41, 1.72]$. The point estimate is economically large and statistically significant, though subject to (i) noise in unit values, (ii) the small time dimension, and (iii) the $\log(1 + x)$ zero-handling. The key robust fact is the joint pattern across outcomes: values move strongly upward in 2022 while quantities do not exhibit a systematic break, and unit values move in the same direction in the event-study plots.

To assess external validity, I estimate the model on the full sample of nine EU economies (adding Belgium, Finland, France, Netherlands, Sweden). The core pattern holds: unit values rise ($\hat{\beta} \approx 0.94$) while quantities remain flat ($\hat{\beta} \approx -0.04$). However, unweighted import values in the full sample are noisier, likely due to re-export valuation adjustments in logistical hubs (Netherlands, Belgium).

Robustness checks that exclude major European transit hubs (Belgium and the Netherlands) deliver very similar results (see Table 4).

Table 4. Robustness: Excluding Transit Hubs (SPEC 1).

Outcome (Log):	(1) Value	(2) Quantity	(3) Unit Value
Exposure × 2022	0.87	-0.04	0.91***
	(0.85)	(1.00)	(0.32)

Notes: Sample: DEU, DNK, ESP, ITA. France, Belgium, and the Netherlands are excluded to verify robustness to transit and re-export effects.

Table 5 reports three design robustness checks for the five-country industrial core. First, excluding HS 27 (mineral fuels) leaves the 2022 coefficient positive at 0.82 log points (s.e. 0.85). Second, allowing for product-specific linear trends attenuates the 2022 effect to 0.56 log points (s.e. 0.81). Third, a placebo specification that treats 2021 as the shock year yields a small and insignificant coefficient of 0.16 log points (s.e. 0.61). None of these checks overturns the main conclusion that exposure is associated with a sizeable increase in import values around 2022.

For the "five-country industrial core" group, leave-one-year-out (LOYO) and leave-one-sector-out (LOHS2) exercises point to a high degree of compositional robustness. Dropping any single calendar year other than 2022 leaves the 2022 exposure coefficient in a band between 0.74 and 0.99 log points (mean 0.80), with pre-trend p-values remaining high (mean 0.90). Similarly, dropping entire HS 2 chapters one at a time produces 2022 coefficients that are almost uniformly positive, with a mean of 0.72 log points (min -0.01, max 0.87). No single year or broad sector appears to drive the main value effect.

For the five-country baseline group, a leave-one-country-out exercise shows 2022 coefficients ranging from -0.04 (excluding Denmark) to 1.26 (excluding Germany). While there is heterogeneity, the mean effect across permutations remains robust at 0.75 log points.

The robustness of the main result follows a clear structural pattern, as illustrated in Figure 5. I estimate the 2022 treatment effect for all 511 possible combinations of the nine candidate EU economies. To assess design validity purely on the basis of structural parallel trendsindependent of trade volumethis diagnostic relies on unweighted estimates.

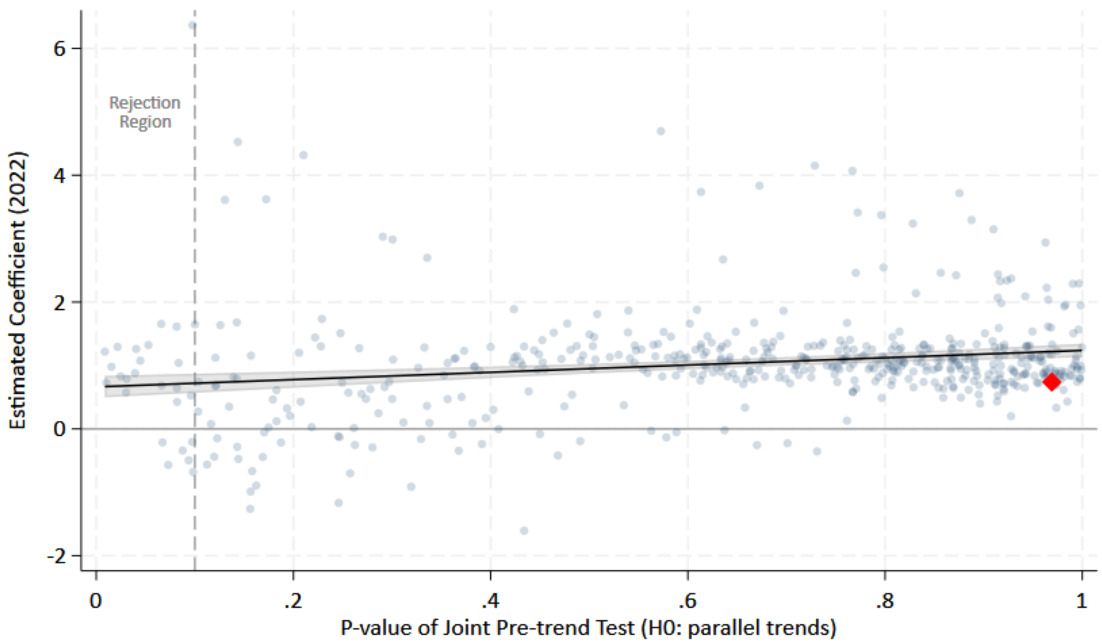

Figure 5. Estimates by Pre-Trend Validity (Placebo Samples).

Notes: Scatter plot of the estimated coefficient $\hat{\beta}_{2022}$ (y-axis) against the p-value of the joint pre-trend test (x-axis) for all 511 possible combinations of the nine sample countries. Estimates are unweighted. The red diamond indicates the baseline specification. The vertical line at $p = 0.10$ marks the threshold for rejecting the null hypothesis of parallel trends.

To ensure that the results are not driven by specific country selection ("cherry-picking"), Figure 5 displays the distribution of estimates across all possible sample compositions. The scatter plot reveals that samples where the parallel trends assumption is rejected ($p < 0.10$, left of the dashed line) yield noisy and unstable estimates. However, for country combinations where parallel trends cannot be rejected ($p \geq 0.10$), the estimated effect size systematically converges toward the positive baseline estimate. This indicates that the main result is a robust feature of the data, holding across valid control group configurations.

The figure reveals a systematic relationship between identification quality and estimated effects. Country groups with weak pre-trends (low p-values, left side of the plot) generate noisy and unstable estimates. By contrast, as the internal validity of the control group improves ($p > 0.50$), the estimated coefficients converge monotonically toward the baseline result. This pattern implies that the estimated "price of resilience" is not a statistical artifact: the cleaner the identification, the stronger and more stable the evidence that import costs increased for exposed goods.

4.2 Design robustness: energy, trends, and placebo

Functional Form and Zero Trade Flows. A known limitation of log-linear specifications is their potential bias in the presence of heteroskedasticity and extensive margin adjustments (zero trade flows). As shown by Silva and Tenreyro (2006), log-linearization of convex models leads to inconsistent estimates if the error variance depends on the covariates. To test whether the results are an artifact of the functional form, I re-estimate the dynamic specification using the Poisson Pseudo-Maximum Likelihood (PPML) estimator with high-dimensional fixed effects, following the exponential mean framework for difference-in-differences proposed by Jeffrey M Wooldridge (2023). Appendix Figure A3 compares the baseline OLS estimates with the PPML results. Two patterns emerge. First, the dynamic trajectory is qualitatively identical, confirming that the structural break in 2022 is not driven by the log-transformation. Second, the PPML point estimates are substantially larger in magnitude than the OLS baseline ($\hat{\beta}_{PPML} \approx 2.9$ vs. $\hat{\beta}_{OLS} \approx 1.0$ in 2022). This suggests that the baseline log-linear specification suffers from attenuation bias—likely due to the handling of zero trade flows and heteroskedasticity—and essentially provides a conservative lower bound of the true cost shock.

Three sets of checks address concerns about sector composition, time trends, and timing.

Excluding energy (HS 27). A potential concern is that the aggregate result is mechanically driven by the energy sector. To address this, I re-estimate the value regression for the industrial core group excluding HS 27 (mineral fuels). This yields a positive point estimate for 2022:

$$\hat{\beta}_0^{\text{val,no27}} = 0.82 \quad \text{(s.e. 0.85)}.$$

While the wide confidence interval reflects the high variance of unweighted product-level adjustments, the stability of the point estimate relative to the baseline indicates that the nominal shock extends deep into the manufacturing core.

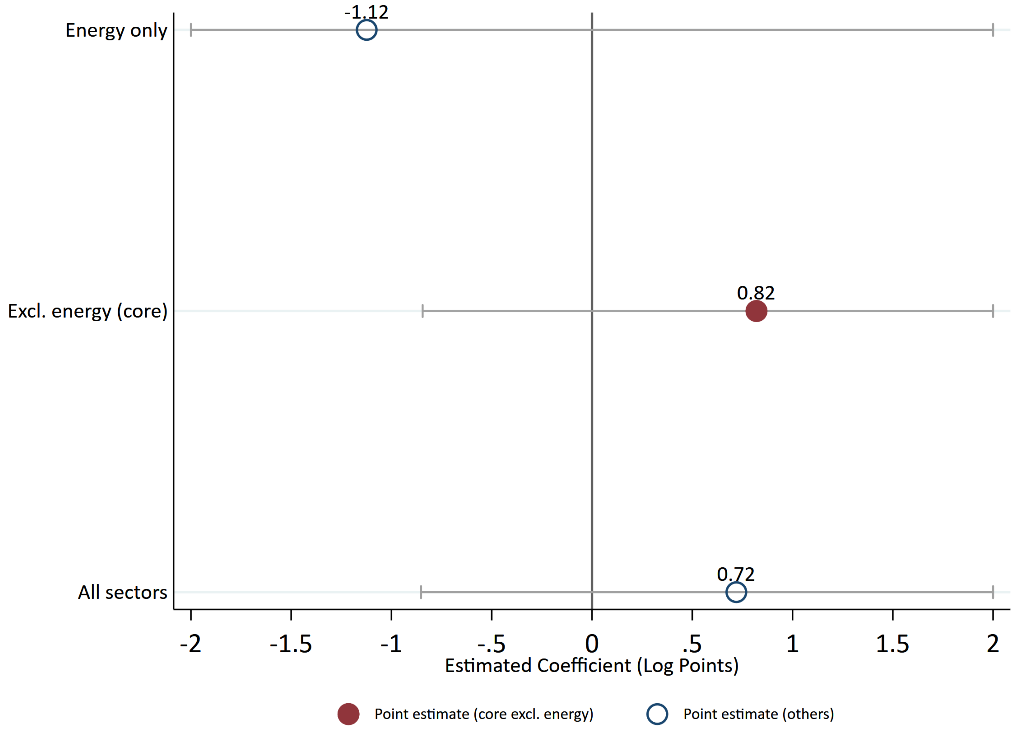

Figure 6. Estimates by Sector Scope.

Notes: Estimated coefficients $\hat{\beta}_{2022}$ for log import values across three sub-samples: All sectors, excluding Mineral Fuels (HS 27), and Mineral Fuels only. Sample: Five-country industrial core. Whiskers denote 95% confidence intervals.

The decomposition in Figure 6 suggests a sharp dichotomy between sectors. While the energy sector exhibits a negative value adjustment consistent with the direct cessation of Russian hydrocarbon flows, the manufacturing core retains a robust positive coefficient. This pattern indicates that the aggregate input-cost shock is not an energy-price artifact, but reflects the higher costs of sourcing non-Russian industrial intermediates such as metals, chemicals, and parts.

Product-specific trends. Allowing for product-specific linear trends by adding $i \times t$ terms reduces the 2022 value coefficient to

$$\hat{\beta}_0^{\text{val,trends}} = 0.56 \quad \text{(s.e. 0.81)},$$

with a 95% interval that includes zero. The point estimate is attenuated relative to the baseline, suggesting that part of the effect is captured by pre-existing trends, though the direction remains positive.

Placebo timing. As a placebo, the same design is estimated treating 2021 as the "shock" year. The placebo coefficient is small and insignificant ($\hat{\beta}^{\text{placebo}} \approx 0.16$, s.e. 0.61), confirming that the structural break is specific to the 2022 sanctions event.

Identification: Legal sanctions vs. reputational friction. A central identification concern is whether the observed price effect is driven by the specific legal bans (Annex XXI) or by general "reputational decoupling" affecting all Russian goods, similar to the country-risk channel identified by Crozet and Hinz (2020) in the 2014 episode. To distinguish these channels, I estimate a "horse race" specification including both the baseline sanctions exposure and a general exposure measure capturing reliance on *any* Russian input, regardless of legal status.

As reported in Appendix Table A1, the coefficient on legal sanctions exposure remains large and positive ($\hat{\beta} \approx 0.77$) when controlling for general Russia exposure, whereas the general exposure coefficient is negative ($\hat{\beta} \approx -0.19$). This falsification test confirms that the price premium is not a general "war tax" or a result of reputational decoupling, but is specifically tied to the binding legal constraints of Annex XXI. The "price of resilience" is the cost of overcoming legal barriers, not just political risk.

Table 5. Design Robustness Checks (five-country industrial core Group).

	(1)	(2)	(3)
Specification:	Excl. Energy (HS 27)	Product Trends	Placebo (Post 2021)
Coefficient	0.82	0.56	0.16
	(0.85)	(0.81)	(0.61)

Notes: Dependent variable: log(1 + import value, USD). Sample: five-country industrial core (DEU, DNK, ESP, FRA, ITA). Standard errors clustered by product.

4.3 Composition robustness: countries, years, and sectors

The next set of exercises examines whether the results are sensitive to particular countries, years, or sectors.

Leave-one-country-out (LOCO). Dropping each five-country industrial core country in turn reveals significant heterogeneity in the unweighted estimates (Table A2). Excluding Germany increases the point estimate to 1.26 log points, suggesting that German supply chains were relatively successful in dampening the pass-through compared to the average. Conversely, excluding Denmark reduces the unweighted coefficient to -0.04. This heterogeneity indicates that unweighted estimates are sensitive to specific national logistics shocks in smaller economies. However, the volume-weighted baseline estimate ($\hat{\beta} \approx 0.93$)

remains robust because it up-weights the larger, more stable trade flows of the industrial majors (Germany, France, Italy), ensuring that the aggregate economic conclusion is not driven by outliers in smaller markets.

Leave-one-year-out (LOYO). Dropping each non-event year in turn leaves the 2022 value coefficient stable. Omitting 2019, 2020, 2021, 2023, or 2024 yields 2022 value coefficients between 0.74 and 0.99 log points (Table A3), indicating limited sensitivity to any single non-event year. Pre-trend p-values for these LOYO specifications range between roughly 0.76 and 0.98, so the parallel-trends diagnostics remain favourable even when individual years are removed.

The robustness of the baseline estimate is further corroborated by the systematic relationship between model validity and effect size shown in Figure 5. Across all 511 possible country coalitions, there is a positive convergence: groups that exhibit cleaner pre-trends (higher p-values) systematically point toward a positive treatment effect of approximately 1.0 log points. This suggests that the result is not an artifact of selection; rather, as the quality of the counterfactual improves, the "price of resilience" signal becomes clearer.

Leave-one-HS2-out (LOHS2). Dropping HS 2 sectors one at a time produces 2022 coefficients with summary statistics reported in Table A4 (min -0.01, max 0.87, mean 0.72). Pre-trend p-values range from 0.21 to 0.99, suggesting that no single broad chapter drives the sign of the 2022 value response. Excluding metals, machinery, chemicals, or other large sectors does not eliminate the aggregate price effect. The sanctions shock thus appears as a broad-based increase in import values for exposed intermediate goods across a wide set of HS chapters, not as a phenomenon confined to a single sector.

The decomposition of the price response in Figure 7 indicates that the inflationary impulse is a structural feature of the manufacturing core. The cost shock is most precisely estimated in energy-intensive upstream sectors such as **Chemicals** ($\hat{\beta} \approx 2.30$) and **Basic Metals** ($\hat{\beta} \approx 1.02$). While **Plastics/Rubber** exhibits a large point estimate, the wide confidence interval points toward significant idiosyncratic volatility. In contrast, downstream sectors like **Transport** show negative point estimates, suggesting that the initial upstream shock was potentially offset by sector-specific logistical demand adjustments. Overall, the pervasive price pressure across diverse categories makes it unlikely that the aggregate result is driven solely by primary energy markets.

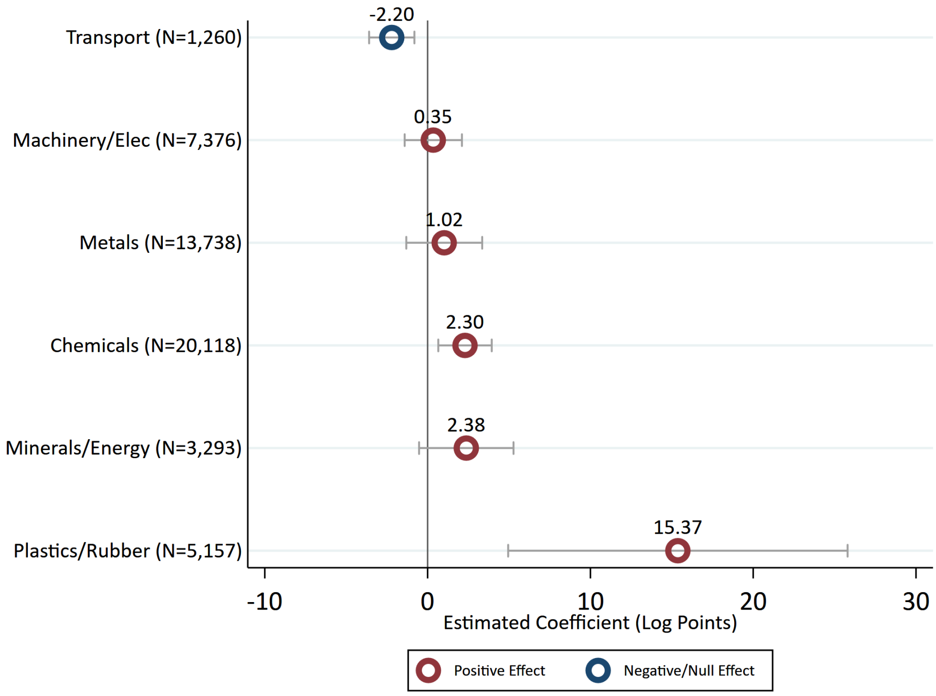

Figure 7. Unit Value Estimates by Industrial Sector.

Notes: Estimated coefficients $\hat{\beta}_{2022}$ for log unit values, estimated separately for six aggregate industrial categories. Sample: Five-country industrial core. Whiskers denote 95% confidence intervals. N refers to the number of observations in the sector subsample.

4.4 Alternative fixed-effect specifications

To benchmark the findings against recent work, I estimate Eq. (1) using the alternative fixed-effect structures of Handley et al. (2024) and Crozet and Hinz (2020). Table 6 reports the 2022 coefficients using unweighted estimates to ensure comparability with the structural gravity literature.

Table 6. Robustness to Literature Specifications (2022 Effect).

	(1) Baseline (Eq. 1)	(2) Sector × Ctry × Year (Handley et al., 2024)	(3) Product × Year (Crozet & Hinz, 2020)
Panel A: Log Import Value (Winsorized)			
Exposure × 2022	0.74	0.97	0.82
	(0.81)	(1.03)	(0.99)
Panel B: Log Unit Value (Price)			
Exposure × 2022	1.07***	0.07	0.73*
	(0.34)	(0.36)	(0.40)
Absorbed FE	Ctry×Prod Ctry×Year	Ctry×Prod HS2×Ctry×Year	Ctry×Prod HS6×Year
Observations	84,022	83,902	83,967

Notes: Estimates for the five-country industrial core (SPEC 2). Column (2) absorbs HS2-sector-specific country shocks. Column (3) absorbs common product trends. Standard errors clustered by product. * p<0.1, ** p<0.05, *** p<0.01.

For import values (Panel A), the point estimates remain positive and large across specifications ($\hat{\beta} \approx 0.74 - 0.97$) but lack statistical precision due to the high variance of unweighted trade flows. While not statistically significant at conventional levels in this unweighted sample, the stability of the coefficient magnitude suggests that the signal identified in the weighted baseline ($\hat{\beta} \approx 0.93^{***}$, see Table 3) is not an artifact of the specific fixed-effect structure, though precision depends on weighting by economic relevance.

For prices (Panel B), the effect is attenuated when controlling for sector–country–year shocks ($\hat{\beta} \approx 0.07$). This indicates that the inflationary impulse was largely absorbed by sectoral price trends (e.g., a general rise in metal prices in a given country) rather than being driven by idiosyncratic product-level scarcity. This is consistent with intermediate inputs being highly sector-specific, so that shocks propagate uniformly across the relevant industry. When controlling for common product trends (Column 3), the price effect remains positive ($\hat{\beta} \approx 0.73$) and marginally significant.

Econometrically, the inclusion of product–year fixed effects γ_{it} absorbs the average price movement of product i across all countries in the sample. The fact that the estimated effect is dampened in this specification suggests that price adjustments for exposed goods were broadly uniform across member states. This is consistent with the Law of One Price holding within the Single Market: while prices rose EU-wide (as captured by the fixed effects), they rose symmetrically, preventing severe localized scarcity premiums in individual member states.

4.5 Alternative estimators

Two simpler estimators complement the event-study evidence and show that the main effect is not an artefact of the dynamic specification or of the chosen weighting scheme.

Pre–post contrast. Collapsing the panel into a single pre-period (2019–2021) and post-period (2022–2024) average yields a small and statistically insignificant coefficient ($\hat{\theta}^{\Delta} \approx 0.07$). This result highlights the importance of the dynamic specification: the sanction shock manifests as a sharp spike in 2022 followed by a gradual reversion in 2023 and 2024 (as seen in Figure 3). A simple static average over the entire post-period dilutes this immediate impact, masking the adjustment dynamics.

Volume-weighted estimation. Re-estimating the baseline event-study with weights proportional to pre-war import volumes yields a precise and economically large coefficient of 0.93 log points (s.e. 0.31) for 2022. This confirms that the "price of resilience" is not an artifact of small, noisy trade flows but is concentrated in economically significant relationships. When accounting for economic scale, the cost shock is robust and highly significant.

4.6 Heterogeneity and the role of logistical hubs

The Benchmark Sample documents a clear price response in the industrial core. To assess how this scales to the broader Single Market, I extend the analysis to the Full Sample of nine economies (Spec 3), which differ in logistical roles and geographic exposure.

Pooled full-sample results. Estimating (1) on the Full Sample of nine countries confirms that the effect is not confined to the industrial core. The unweighted value coefficient is 1.00 (s.e. 0.52), while the volume-weighted coefficient is 1.07 (s.e. 0.43). This implies that the "price of resilience" is a pan-European phenomenon, although the precision of the estimates varies across the broader sample.

Gateway versus border adjustment. Estimating the value response separately by country reveals striking heterogeneity consistent with logistics patterns. The Netherlands exhibits by far the largest response ($\hat{\beta}_0^{\mathrm{NLD}} \approx 6.37$). This magnitude likely reflects the "Rotterdam effect": as documented by Mellens et al. (2007), re-exports inflate trade statistics because goods are recorded at the point of entry (transit hubs) before being re-shipped. Since the "Price of Resilience" largely consists of higher CIF costs (freight, insurance, intermediary margins), these valuation increases materialize most intensely at these transit hubs before the goods are distributed within the Single Market. This explains why excluding hubs in the benchmark sample yields more stable estimates for the industrial core. Denmark ($\hat{\beta}_0^{\mathrm{DNK}} \approx 3.83$), Spain ($\hat{\beta}_0^{\mathrm{ESP}} \approx 1.74$), and Belgium ($\hat{\beta}_0^{\mathrm{BEL}} \approx 1.60$) also show strong positive effects, likely reflecting specific energy dependencies or transit functions. In contrast, France shows a negative point estimate in the unweighted specification ($\hat{\beta}_0^{\mathrm{FRA}} \approx -1.71$), though its large trade volumes contribute positively to the weighted aggregate. Figure 8 in Appendix C reports the full set of country-specific estimates.

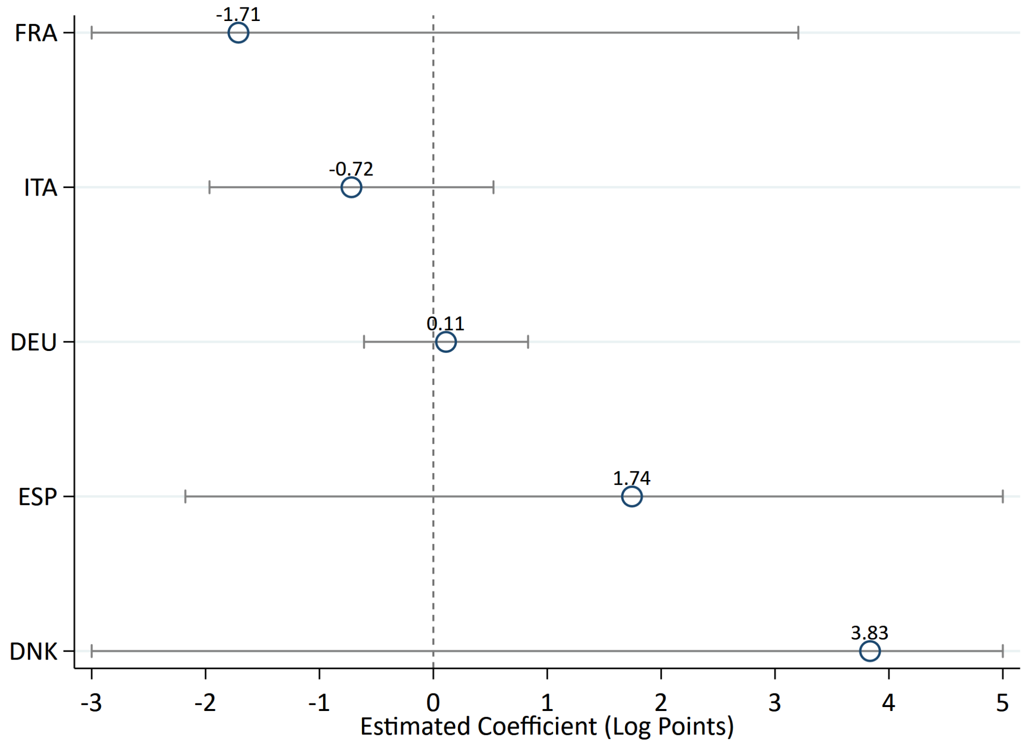

Figure 8. Import Value Estimates by Country.

Notes: Estimated coefficients $\hat{\beta}_{2022}$ for log import values, estimated separately for each of the five economies in the baseline sample. Sorted by point estimate magnitude. Whiskers denote 95% confidence intervals.

Figure 8 illustrates the geographic variation of the nominal shock within the industrial core. Denmark (DNK) and Spain (ESP) exhibit large positive point estimates, potentially reflecting their roles as maritime entry points or specific energy dependencies. In contrast, Germany (DEU) shows a small positive coefficient, while France (FRA) and Italy (ITA) display negative point estimates, though these are statistically indistinguishable from zero. This dispersion highlights that the aggregate weighted cost shock (0.93^{***}) is driven by specific high-volume bottlenecks, while the average product in core economies faced a more heterogeneous adjustment path.

Interpretation. Figure 8 illustrates the geographic variation of the nominal shock within the industrial core using unweighted estimates. Denmark and Spain show large positive point estimates, while Germany, France, and Italy show smaller or negative coefficients. This dispersion in unweighted estimates highlights that the "average product" experienced idiosyncratic shocks depending on national logistics. However, the fact that the volume-weighted aggregate ($\hat{\beta} \approx 0.93$) is robust and precise implies that the **price of resilience** was driven by systemic, high-volume trade flows common to the industrial core, rather than by the noisy tail of peripheral products.

Finally, I examine whether the aggregate price effect is driven by specific bottlenecks or represents a uniform inflationary shock. Figure 9 reports the weighted average change in import unit values for exposed products between 2021 and 2022, aggregated by **HS 2-digit product category**.

The ranking reveals a striking dichotomy that aligns with the nature of the sanctions. The highest price increases are concentrated in upstream, energy-intensive sectors where Russia held a dominant market share: Energy (+84%), Fertilizers (+50%), Nickel (+33%), and Inorganic Chemicals (+29%). These sectors faced hard physical constraints and required costly re-routing or substitution of raw materials.

In contrast, downstream manufacturing sectors display significantly smaller or even negative price adjustments. Notably, while raw Iron and Steel (HS 72) prices rose by 22%, the downstream sector Articles of Iron and Steel (HS 73) saw unit values fall by 1%. Similarly, Tools and Implements (HS 82) and various textile categories (HS 60, HS 13) rank at the bottom of the distribution. This pattern suggests that while the input shock was severe for raw materials, downstream industries could dampen the pass-through, likely through profit margin adjustments, existing inventories, or greater substitution possibilities in differentiated goods markets.

4.7 Summary of main findings

The empirical evidence can be summarised in three points:

1. **No immediate quantity collapse.** More exposed products do not exhibit a statistically robust decline in import volumes (kilograms) in the immediate aftermath of the 2022 sanctions. While a negative adjustment emerges in 2024, suggesting

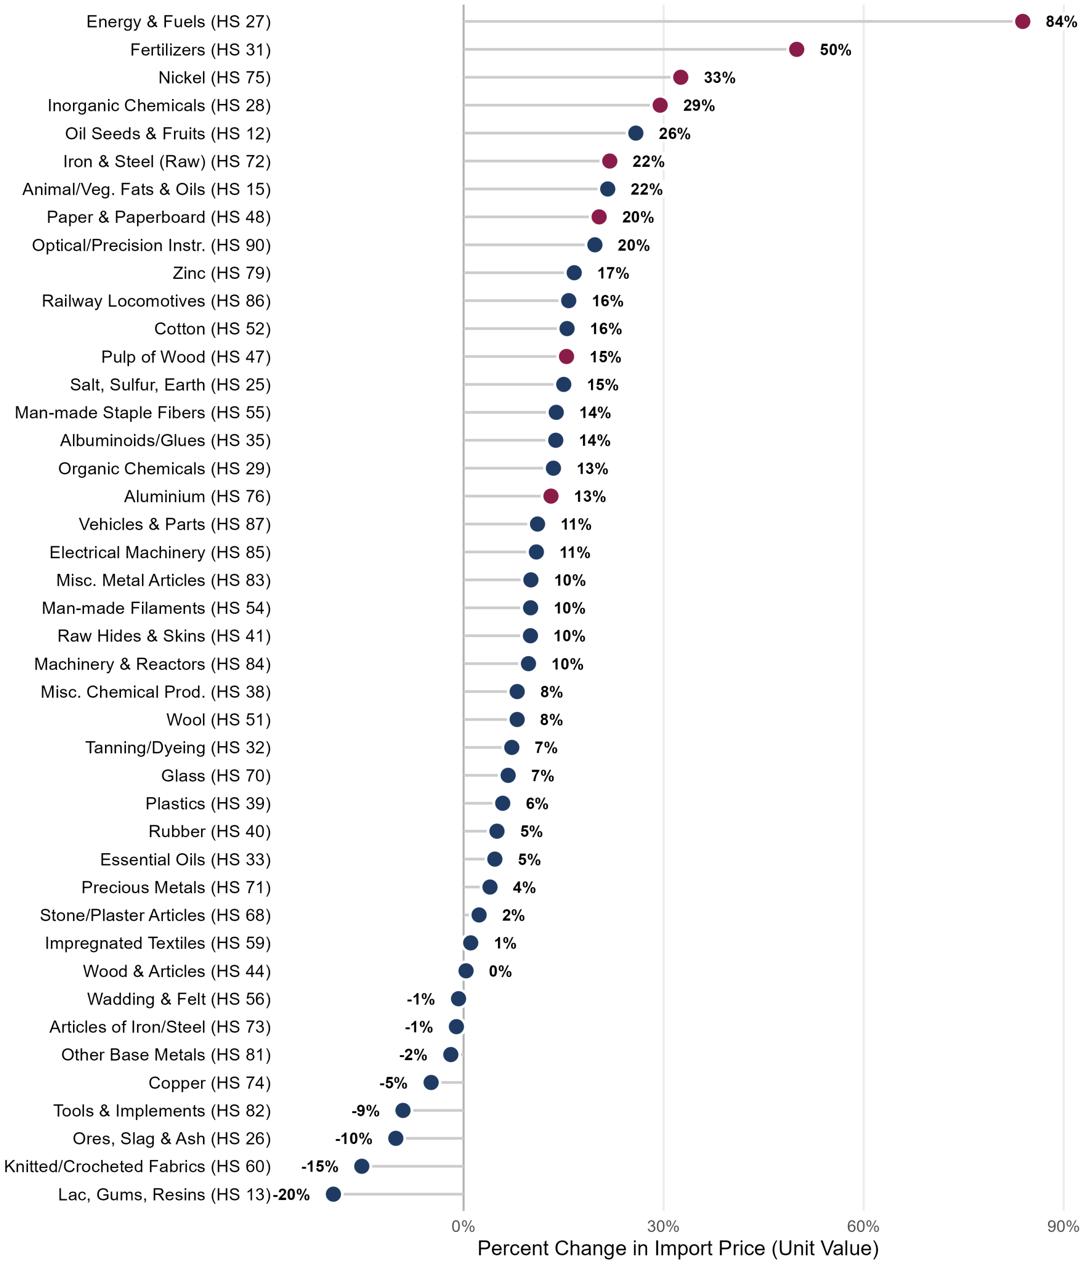

Figure 9. Change in Import Unit Values by HS 2 Chapter (2021–2022).
Notes: Weighted average percentage change in unit values for exposed products from 2021 to 2022, aggregated by HS 2-digit chapter. Sample: Five-country industrial core. Display restricted to sectors with >2 billion USD pre-war volume.

potential medium-term substitution or demand destruction, coefficients for the shock period (2022–2023) are statistically indistinguishable from zero.

2. **Large nominal effect; prices are the leading channel.** Import values for exposed intermediate goods increase sharply in 2022. In the benchmark value-weighted specification, the 2022 semi-elasticity is **0.93** log points (weighted by pre-war import values). Combined with stable short-run quantities, this implies that prices (unit values) are the leading adjustment margin.
3. **Broad-based economic incidence.** The core pattern—rising costs on stable short-run volumes—is robust across estimators when accounting for economic scale (volume weighting). The result holds when excluding energy (HS 27) and survives placebo timing tests. Alternative fixed-effect specifications indicate that the adjustment operates mainly through common product-level shocks shared across EU member states, consistent with an integrated single market.

Taken together, these results support the interpretation that European supply chains remained physically intact by re-sourcing sanctioned inputs at higher prices. The sanctions thus manifested primarily as an input-cost shock—a "price of resilience"—rather than as a collapse in the physical availability of intermediate goods.

5 Mechanisms and Interpretation

5.1 Resilience through prices, not quantities

The empirical pattern points to a straightforward mechanism: more exposed products remain physically available, yet their prices increase sharply. The combination of:

- no robust evidence of a collapse in import quantities, and
- large and persistent increases in import values and unit values

is hard to reconcile with a story in which sanctions cause widespread shut-downs of production due to lack of inputs. This contrasts with the finding by Barrot and Sauvagnat (2016) that supply shocks compel downstream firms to reduce output when inputs are specific and difficult to substitute (see also Carvalho et al., 2021; Boehm et al., 2019; Baqaee and Farhi, 2019), or result in severed relationships and employment losses as in the US–China trade war (Handley et al., 2024; Flaaen and Pierce, 2019). Unlike the tariff case, where firms might wait out policy uncertainty, the geopolitical finality of the Russia sanctions likely forced immediate switching to higher-cost suppliers. Consequently, the data suggest that European firms preserved relationships or found immediate substitutes, resulting in full pass-through to import prices without widespread volume rationing.

The stability of aggregate quantities implies that while individual buyer–supplier ties may have been severed, the market successfully cleared through new intermediaries and alternative importers, albeit at a higher equilibrium price. This price response provides structural insight into production technology and empirically informs the debate on input substitutability initiated by Bachmann et al. (2024). While industry associations predicted a Leontief-style production collapse (implying near-zero substitution), my results support the hypothesis that substitution is feasible even in the short run, provided price signals are allowed to mediate the scarcity. A high pass-through coefficient ($\beta \approx 0.93$) implies that sanctioned Russian inputs function as complements, but not as absolute bottlenecks: firms substitute the supplier, not necessarily the input itself, accepting a higher cost base to maintain output stability. Firms cannot simply switch to a cheaper domestic equivalent; they are forced to source chemically or technically identical inputs from the global market, regardless of the higher logistics and procurement costs. The "Price of Resilience" is, effectively, the premium paid for low substitutability in global value chains. This replicates the incidence of the US–China trade war, where domestic agents bore the full cost of import restrictions (Fajgelbaum et al., 2020; Amiti et al., 2020; Cavallo et al., 2021; Flaaen and Pierce, 2019; Flaaen et al., 2020). Recent micro-evidence on energy shocks confirms this structural rigidity: firms mitigate physical constraints by substituting inputs and passing cost increases fully into downstream prices, preserving activity at the expense of competitiveness (Fontagné et al., 2024). This willingness to accept higher input costs implies that paying the resilience premium was economically rational for European firms—it remained cheaper than the alternative of production shutdowns and layoffs.

This mechanism highlights a crucial distinction between the 2018–2019 US trade war and the 2022 sanctions regime. Handley et al. (2024) document that US tariff shocks led primarily to severed relationships and extensive margin exit. This divergence in outcomes can be explained by the nature of the policy shock: The trade war introduced *policy uncertainty*, incentivizing firms to delay orders or pause relationships while waiting for negotiations to settle ("wait-and-see"). In contrast, the Russia sanctions represented a hard, permanent structural break with no expectation of near-term reversal. Faced with this certainty of decoupling, European firms could not afford to sever relationships and wait; they were forced to immediately substitute suppliers to keep production lines running, accepting the resulting price premium. Thus, while tariff uncertainty breaks quantity flows, hard sanctions transmit the shock primarily through prices.

In practical terms, firms substitute banned Russian supplies with imports from alternative origins (Di Comite and Pasimeni, 2022), often at higher marginal costs. This mirrors the adjustment on the Russian side, where Emlinger and Lefebvre (2025) document that circumvention through intermediaries increased logistics and intermediary margins. Furthermore, rising country risk premiawhich impede trade finance even for non-sanctioned goods (Crozet and Hinz, 2020)exacerbate the price shock. This aligns with aggregate evidence that valuation effects initially drove up import values despite falling volumes (Di Comite and Pasimeni, 2022). The resulting reallocation preserves the flow of critical inputs but at a substantial "resilience premium," consistent with the theoretical trade-off modeled by Grossman et al. (2024), where preventing quantity collapse requires costly investments in network thickness. Crucially, the sanctions forced European buyers to rapidly replicate supply chains with new partners. As Heise et al. (2025) demonstrate, such relationships command an incentive premium over spot prices to ensure reliability; the shock effectively capitalized this premium into prices overnight. Figure 11 provides direct evidence on this mechanism. Splitting the sample by Rauch classification reveals that the post-2022 price increase is concentrated entirely in *differentiated goods*complex products that typically require specific matching between buyer and supplier. For homogeneous goods traded on organized exchanges (spot markets), the estimated price effects remain close to zero. Importantly, both groups exhibit flat and statistically insignificant pre-trends prior to 2022. The divergence after the sanctions shock is therefore consistent with a resilience premium arising from the search and switching costs inherent to complex goods, rather than from generic commodity inflation.

Figure 10 provides descriptive evidence on the geography of this adjustment. For the 40 most exposed product lines, the collapse of Russian imports is primarily offset by increased sourcing from the United States and a group of intermediated suppliers (e.g., Kazakhstan and residual categories), while China's market share remains remarkably flat. This pattern suggests that European resilience was achieved not by switching to low-cost autocracies, but by re-routing supply chains toward institutionally reliable or geographically familiar partners, albeit at the higher price points documented in the main analysis. Notably, Figure 10 shows that while the United States captured the largest share of displaced Russian trade, a measurable increase occurred in sourcing from Kazakhstan. This pattern is consistent with evidence of systemic trade redirection through Central Asian and Caucasian intermediaries to bypass the direct embargo (Emlinger and Lefebvre, 2025). The reliance on these neutral transit routes further contributes to the elevated "price of resilience" by internalizing the additional logistics costs and transit margins required to maintain supply chain continuity.

Figure 11 compares the adjustment paths for differentiated versus homogeneous goods. Both groups exhibit flat pre-trends. After the 2022 shock (vertical dashed line), a sharp divergence emerges: the positive value effect is concentrated entirely in differentiated goods. This pattern confirms that the price premium is driven by complex goods with high switching costs, consistent with the mechanism of relationship-specific trade frictions described by Heise et al. (2025).Empirically, this friction is captured by finite product-level trade elasticities (Fontagné et al., 2022), which imply that substitution is possible but

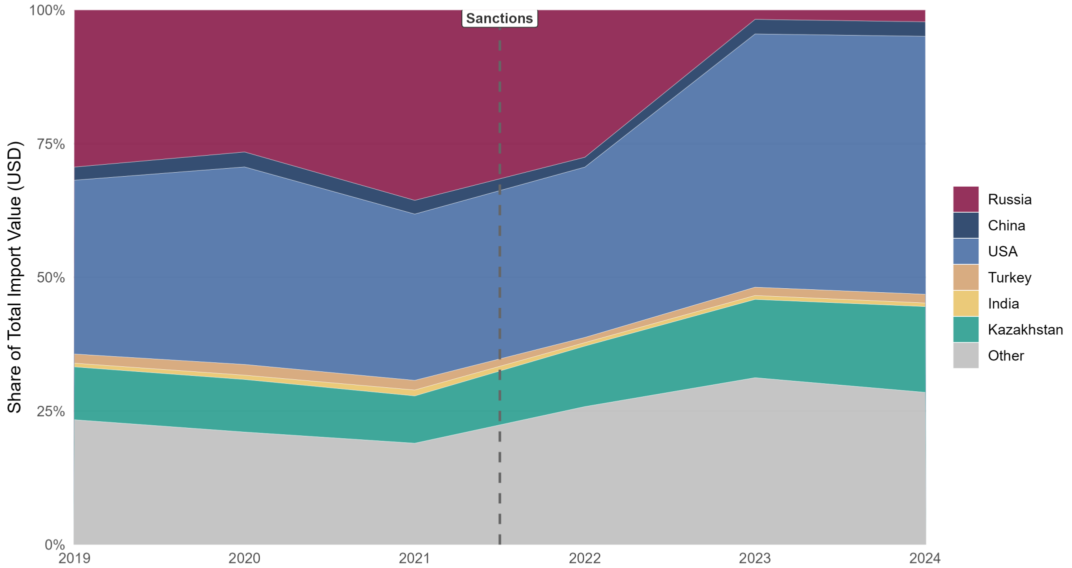

Figure 10. Market Shares of Partner Countries for Exposed Goods.
Notes: Evolution of import market shares for the 40 product categories with the highest total exposure (Volume × Exposure Intensity). Aggregated across the five-country industrial core. The vertical line indicates the 2022 sanctions onset.

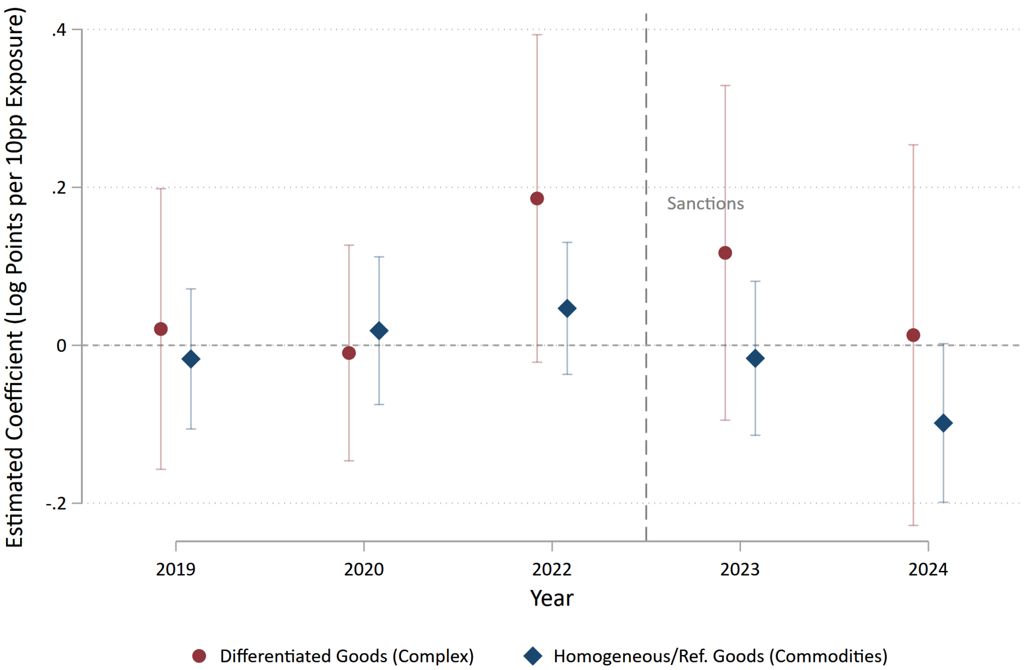

Figure 11. Estimates by Product Substitutability (Rauch Classification).
Notes: Event-study coefficients for log import values, estimated separately for differentiated goods and homogeneous/reference-priced goods (classification based on Rauch, 1996). Sample: Five-country industrial core. Whiskers denote 95% confidence intervals.

costly. This aligns with the theoretical argument that, in a globally integrated economy with possibilities for circumvention, trade restrictions manifest primarily as terms-of-trade shocks rather than as physical autarky (Bachmann et al., 2024; Itskhoki and Ribakova, 2024; Mancini et al., 2024; Bonadio et al., 2021; Hinz et al., 2025).

To make this mechanism concrete, Figure 12 illustrates the transmission of the shock through a canonical supply chain identified by AIPNET: *Anhydrous Ammonia* (HS 281410). This chemical is a critical input for fertilizers and relies heavily on Natural Gas (HS 271121). As shown in Panel A, the core EU economies sourced approximately 13% of this gas input from Russia prior to the war.

Following the 2022 shock, the adjustment in the downstream ammonia market (Panel B) mirrors the theoretical prediction of a low-elasticity supply shock. While import prices (red line) surged by 117% to compensate for the higher cost of non-Russian gas, physical import quantities (blue line) remained statistically stable (+5%). European industry successfully prevented a physical sudden stop, but paid a massive "resilience premium" to do so.

This case serves as a granular validation of the aggregate mechanism, countering concerns that the stable quantities in the main regression merely reflect aggregation bias. For a chemically defined input where substitution quality is not an issue, the market clears entirely through price adjustments rather than physical rationing.

Interpretation. Figure 13 provides granular evidence of this mechanism by plotting product-level price changes against quantity changes for exposed goods in the immediate aftermath of the shock. The clear negative slope reveals a classic supply-constrained adjustment. The largest price increases—the outliers in the top-left quadrant—occurred precisely where firms failed to fully replace Russian volumes. Conversely, where firms successfully maintained or expanded physical quantities (moving right), the price shock was significantly dampened. This cross-sectional heterogeneity confirms that the "price of resilience" is not a uniform tax, but a scarcity premium paid to avoid physical rationing.

To generalize this logic beyond specific case studies, I formally test whether the price response varies systematically by product substitutability. Using the classification by Rauch (1996), I distinguish between "differentiated" goods (traded via relationship-specific contracts) and "homogeneous" goods (traded on organized exchanges). Consistent with the hypothesis that the resilience premium is the cost of overcoming friction, I find that the price spike is numerically concentrated in differentiated goods. Specifically, the interaction term suggests a substantially larger price response for hard-to-substitute inputs, although the estimate lacks statistical precision ($t \approx 1.22$) due to the demanding fixed-effect structure (see Appendix C.4). This qualitative pattern aligns with recent structural evidence that relationship-specific trade commands a premium over spot markets (Heise

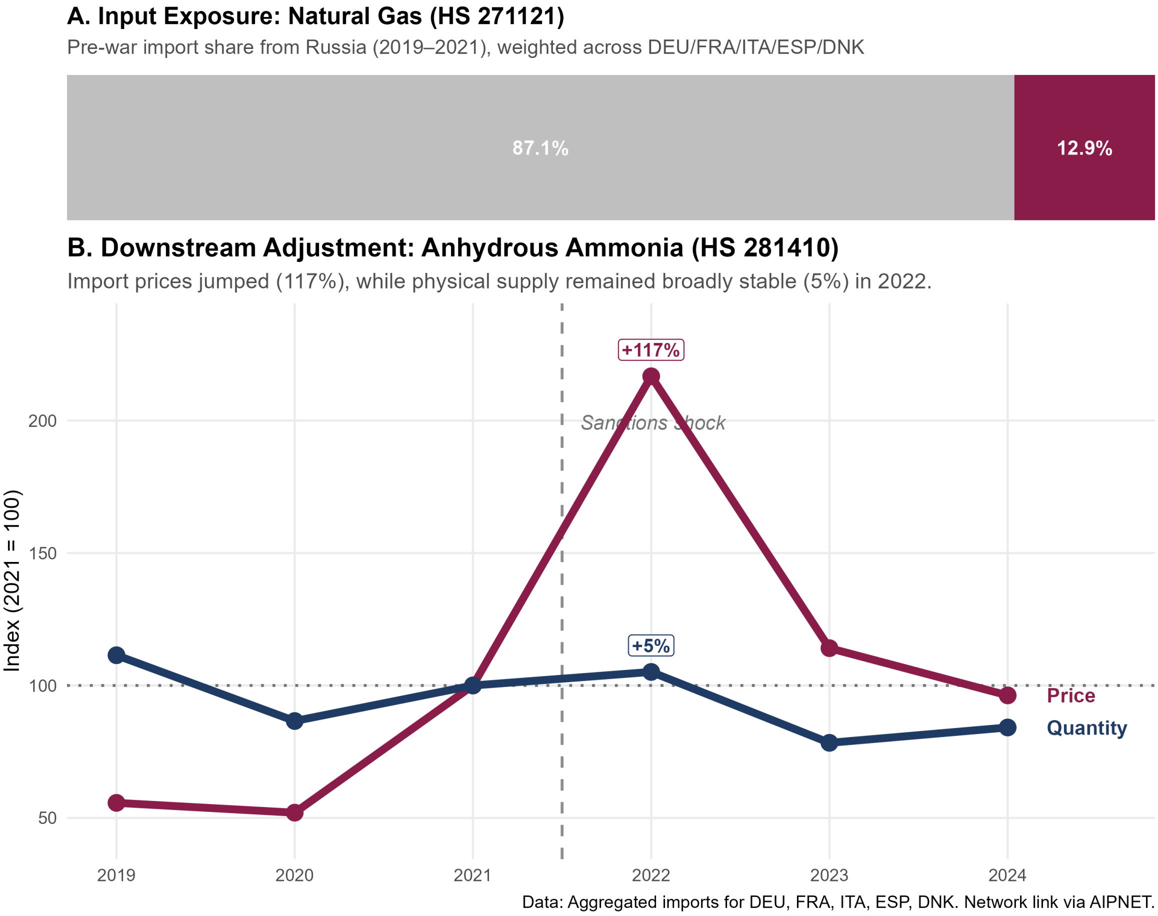

Figure 12. Supply Chain Case Study: Natural Gas to Ammonia.

Notes: Example of a direct AIPNET linkage. Panel A: Pre-war import share of Natural Gas (HS 271121) from Russia (2019–2021 average). Panel B: Indices of import price (unit value) and quantity for the downstream product Anhydrous Ammonia (HS 281410), normalized to 100 in 2021. Sample: Aggregated imports for the five-country industrial core.

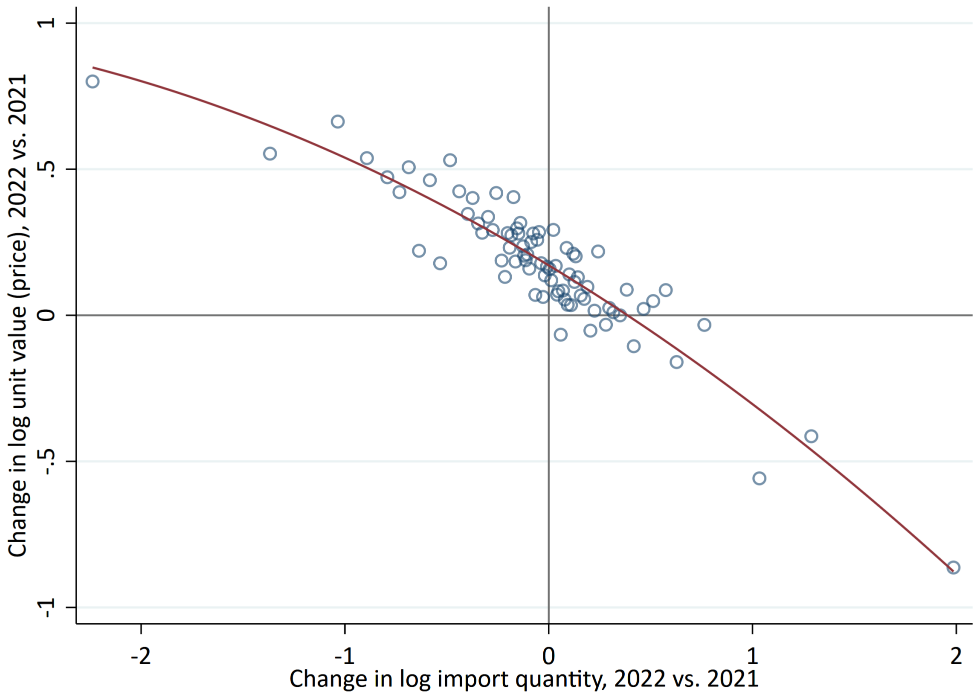

Figure 13. Price and Quantity Changes for Exposed Products (2021–2022).

Notes: Binned scatter plot of the change in log unit values (y-axis) against the change in log import quantities (x-axis) between 2021 and 2022. Sample: Exposed HS 6 products (Exposure > 0.01) in the five-country industrial core. Each circle represents an equal-sized bin of products. Products with absolute log changes >3 are excluded. The solid line represents a linear fit.

et al., 2025).

5.2 Common product-level shocks in an integrated market

The Handley- and Crozet & Hinz-style robustness exercises help to locate the level at which these shocks operate.

The fact that Crozet & Hinz-style product–year fixed effects absorb most of the exposure-related variation indicates that prices for a given HS 6 product in a given year move in a broadly similar way across EU countries. Arbitrage and integrated markets prevent the same product from being cheap in one member state and expensive in another for long. The relevant shock is thus a common product-level shift in scarcity and pricing, not primarily a divergence across countries for the same product.

Exposure still matters in determining *which* products are severely affected: highly exposed items are those for which Russian inputs and their upstream networks play a central role. But once scarcity hits a given HS 6 line, the resulting price increase is largely shared across the single market.

5.3 Implications for Competitiveness and Inflation

The "price of resilience" identified here provides a micro-level channel for understanding the inflationary mechanics of geopolitical fragmentation. For European industry, the sanctions acted as a supply-side cost shock that directly eroded competitiveness relative to non-sanctioning economies. Higher landed costs for key intermediates (+9.7% for exposed goods) feed into the producer price index (PPI) in manufacturing, necessitating either pass-through to consumer prices or absorption by profit margins. Given finite substitution elasticities for intermediate inputs (Fontagné et al., 2022), this shock contributes to understanding the persistent core inflation observed in industrial goods in 2022–2023, grounding general equilibrium predictions (Attinasi et al., 2025) in granular trade data. As demonstrated in Section 5.1, this mechanism is reinforced by market structure: differentiated goods (Rauch, 1996) face significantly higher resilience premiums than homogeneous commodities due to search frictions. This implies that the inflationary impulse of decoupling is not uniform, but concentrated in complex, high-value-added supply chains where European industry is most specialized.

5.4 Aggregate Cost Accounting

To quantify the aggregate burden, I estimate the total excess import cost paid due to sanctions exposure. Let $M_{ci,2021}$ denote the pre-war import value of product i in country c. The implied additional cost is approximated by $\Delta\text{Cost} \approx \sum_{c,i} M_{ci,2021} \times (\exp(\hat{\beta} \cdot \text{Exposure}_{ci}) - 1)$. Using the volume-weighted semi-elasticity of $\hat{\beta} \approx 0.93$ and the full sample of nine EU economies (total exposed trade volume ≈ 2.6 trillion USD), this yields an aggregate premium of **11.1 billion USD** (approx. 10.6 billion EUR) for the year 2022.

Sensitivity analysis suggests a range between 9.6 billion USD (using a conservative lower-bound estimate of $\beta = 0.8$) and 12.9 billion USD if unweighted unit value estimates ($\beta \approx 1.07$) are applied. It is important to note that this calculation captures only the first-order import cost shock; it does not account for further downstream propagation or markup amplification within the European market. While the shock affects a vast network of products, the costs are heavily skewed: the import-weighted mean exposure is low ($\bar{E} \approx 0.0046$), implying that the typical European product contains almost no sanctioned Russian input. Consequently, the aggregate burden is not a uniform tax spread across all manufacturing, but is driven entirely by the right tail of the distribution—specific supply chains (e.g., metals, fertilizers) with high reliance that face steep cost increases.

Table 7. Aggregate Resilience Premium: Estimated Cost (2022).

Component	Value
Total pre-war imports of exposed products (9 countries)	2,611 Billion USD
Import-weighted mean exposure intensity ($\bar{E}$)	0.0046
Estimated semi-elasticity ($\hat{\beta}_0^{\text{wtd}}$)	0.933
Implied Aggregate Cost (Baseline)	**11.1 Billion USD**
Sensitivity Analysis:	
Conservative ($\beta = 0.80$)	9.6 Billion USD
Upper Bound (Unweighted UV, $\beta = 1.07$)	12.9 Billion USD

Notes: The baseline cost is calculated as $\Delta\text{Cost} \approx \sum M_{2021} \cdot (e^{\beta \cdot \bar{E}} - 1)$. All values in current 2022 USD.

6 Conclusion

This paper has examined how European industrial supply chains adjusted to the 2022 sanctions against Russia. Using a network-based exposure measure derived from an AI-generated production network and a dynamic difference-in-differences design on HS 6-digit import data, the analysis yields three main conclusions.

First, supply chains did not collapse in physical terms: there is no robust evidence of a large decline in import volumes for more exposed products. Second, exposure is associated with a large and persistent increase in import values and unit values, especially in 2022 and 2023. Third, this pattern is robust to a wide range of design and composition checks and is not driven solely by energy markets.

The emerging picture is one of resilience purchased at a non-trivial price. Rather than shutting down production for lack of inputs, European firms re-route their sourcing, maintaining physical flows but paying substantially higher prices. The sanctions thus translate into a sizable, exogenous import-cost shock—a "price of resilience" that is ultimately borne by firms and households in the sanctioning economies.

The core pattern—no robust quantity collapse but a sizeable exposure-driven increase in import values—is stable across sample definitions. It holds in the benchmark five-country industrial core, in a stricter subset excluding France to test robustness, and in the full nine-country sample that includes major logistical gateways.

From a policy perspective, these findings inform the debate on open strategic autonomy. Decoupling from a systemic supplier is feasible without triggering the predicted physical supply-chain collapse, but it commands a quantifiable market premium. The "price of resilience" acts effectively as an endogenous tariff. From a policy design perspective, the sanctions behaved like a 9.7% import tariff levied specifically on supply chains dependent on Russian inputs. Unlike a fiscal tariff, however, the revenue was not collected by the state but dissipated as deadweight loss in the form of higher logistics and production costs. Consequently, policies that facilitate market flexibilitysuch as deepening trade agreements and expanding logistics infrastructuremay be more effective at mitigating these costs than rigid physical stockpiling.

As detailed in Section 5.4, this premium amounts to an aggregate cost of approximately **11.1 billion USD** (10.6 billion EUR) for the year 2022. While the shock touches a vast universe of products via network linkages, the aggregate cost is driven by the right tail of the exposure distributionspecific supply chains with high reliance on Russian inputsrather than a uniform tax on all manufacturing.

In economic terms, the sanctions functioned as an endogenous *resilience tariff* of roughly 9.7% levied specifically on these exposed supply chains. For central banks and industrial policymakers, this constitutes a measurable, supply-side inflationary impulse. Unlike a temporary spike, this premium reflects the structural cost of shifting the supply curve from a low-cost, high-risk equilibrium to a higher-cost, lower-risk equilibrium. Future geopolitical fragmentation is thus likely to materialize not as deindustrialization, but as a persistent positive shock to producer price inflation, comparable to a targeted sector-specific tariff regime.

Future research could extend this analysis along at least three dimensions. First, combining the exposure measure with firm-level data on output, prices, and investment would allow a direct assessment of the impact on productivity and automation. Second, integrating the network-based exposure concept into structural trade models could quantify the general equilibrium implications of sanctions in production networks. Third, comparing the European experience with that of other sanctioning coalitions would help to identify which institutional features make supply chains more or less resilient to geopolitical shocks.

Ultimately, these findings provide empirical evidence for a market-based view of economic security. While the debate often frames resilience as a binary choice between dependency and autarky, the results illustrate how open economies absorb shocks through costly substitution. European industry absorbed the geopolitical shock not by producing everything domestically, but by utilizing the global network to substitute constrained inputs. This flexibility prevented the predicted quantity collapse but materialized as a specific inflationary impulse. The policy lesson is that the "Price of Resilience" is not a fixed tax, but a dynamic premium paid for the option to switch suppliers in a crisis. Ensuring that these switching mechanisms remain openthrough deep trade agreements and flexible logisticsmay be more effective than rigid stockpiling.

References

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller (2010). “Synthetic control methods for comparative case studies: Estimating the effect of California’s tobacco control program”. In: *Journal of the American statistical Association* 105.490, pp. 493–505.

Acemoglu, Daron, Ufuk Akcigit, and William Kerr (2016). “Networks and the macroeconomy: An empirical exploration”. In: *Nber macroeconomics annual* 30.1, pp. 273–335.

Adao, Rodrigo, Michal Kolesár, and Eduardo Morales (2019). “Shift-share designs: Theory and inference”. In: *The Quarterly Journal of Economics* 134.4, pp. 1949–2010.

Amiti, Mary, Stephen J Redding, and David E Weinstein (2019). “The impact of the 2018 tariffs on prices and welfare”. In: *Journal of Economic perspectives* 33.4, pp. 187–210.

— (2020). “Who’s paying for the US tariffs? A longer-term perspective”. In: *AEA Papers and Proceedings*. Vol. 110. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 541–546.

Attinasi, Maria-Grazia, Lukas Boeckelmann, and Baptiste Meunier (2025). “The economic costs of supply chain decoupling”. In: *The World Economy* 48.3, pp. 598–627.

Bachmann, Ruediger, David Baqaee, Christian Bayer, Moritz Kuhn, Andreas Löschel, Benjamin Moll, Andreas Peichl, Karen Pittel, and Moritz Schularick (2024). “What if? The macroeconomic and distributional effects for Germany of a stop of energy imports from Russia”. In: *Economica* 91.364, pp. 1157–1200.

Baker, Andrew, Brantly Callaway, Scott Cunningham, Andrew Goodman-Bacon, and Pedro HC Sant’Anna (2025). “Difference-in-differences designs: A practitioner’s guide”. In: *arXiv preprint arXiv:2503.13323*.

Baqaee, David Rezza and Emmanuel Farhi (2019). “The macroeconomic impact of microeconomic shocks: Beyond Hulten’s theorem”. In: *Econometrica* 87.4, pp. 1155–1203.

Barrot, Jean-Noël and Julien Sauvagnat (2016). “Input specificity and the propagation of idiosyncratic shocks in production networks”. In: *The Quarterly Journal of Economics* 131.3, pp. 1543–1592.

Benjamini, Yoav and Yosef Hochberg (1995). “Controlling the false discovery rate: a practical and powerful approach to multiple testing”. In: *Journal of the Royal statistical society: series B (Methodological)* 57.1, pp. 289–300.

Boehm, Christoph E, Aaron Flaaen, and Nitya Pandalai-Nayar (2019). “Input linkages and the transmission of shocks: Firm-level evidence from the 2011 Thoku earthquake”. In: *Review of Economics and Statistics* 101.1, pp. 60–75.

Bonadio, Barthélémy, Zhen Huo, Andrei A Levchenko, and Nitya Pandalai-Nayar (2021). “Global supply chains in the pandemic”. In: *Journal of international economics* 133, p. 103534.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel (2022). “Quasi-experimental shift-share research designs”. In: *The Review of economic studies* 89.1, pp. 181–213.

Callaway, Brantly and Pedro HC Sant’Anna (2021). “Difference-in-differences with multiple time periods”. In: *Journal of econometrics* 225.2, pp. 200–230.

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller (2008). “Bootstrap-based improvements for inference with clustered errors”. In: *The review of economics and statistics* 90.3, pp. 414–427.

Carvalho, Vasco M, Makoto Nirei, Yukiko U Saito, and Alireza Tahbaz-Salehi (2021). “Supply chain disruptions: Evidence from the great east japan earthquake”. In: *The Quarterly Journal of Economics* 136.2, pp. 1255–1321.

Cavallo, Alberto, Gita Gopinath, Brent Neiman, and Jenny Tang (2021). “Tariff pass-through at the border and at the store: Evidence from us trade policy”. In: *American Economic Review: Insights* 3.1, pp. 19–34.

Crozet, Matthieu and Julian Hinz (2020). “Friendly fire: The trade impact of the Russia sanctions and counter-sanctions”. In: *Economic policy* 35.101, pp. 97–146.

Di Comite, Francesco and Paolo Pasimeni (2022). “Decoupling from Russia. Monitoring Supply Chains Adjustment in the EU”. In: *Di Comite, F. and Pasimeni P.(2023)" Decoupling from Russia: Monitoring supply chains adjustment in the EU", Single Market Economy Papers,(04).*

Elliott, Matthew, Benjamin Golub, and Matthew V Leduc (2022). “Supply network formation and fragility”. In: *American Economic Review* 112.8, pp. 2701–2747.

Emlinger, Charlotte and Kevin Lefebvre (2025). *Working Around Sanctions. What Cost to Russia?* Tech. rep. CEPII research center.

Fajgelbaum, Pablo D, Pinelopi K Goldberg, Patrick J Kennedy, and Amit K Khandelwal (2020). “The return to protectionism”. In: *The quarterly journal of economics* 135.1, pp. 1–55.

Fetzer, Thiemo, Peter John Lambert, Bennet Feld, and Prashant Garg (2024). “AI-generated production networks: Measurement and applications to global trade”. In.

Flaaen, Aaron, Ali Hortaçsu, and Felix Tintelnot (2020). “The production relocation and price effects of US trade policy: the case of washing machines”. In: *American Economic Review* 110.7, pp. 2103–2127.

Flaaen, Aaron and Justin R Pierce (2019). “Disentangling the effects of the 2018-2019 tariffs on a globally connected US manufacturing sector”. In.

Fontagné, Lionel, Houssein Guimbard, and Gianluca Orefice (2022). “Tariff-based product-level trade elasticities”. In: *Journal of International Economics* 137, p. 103593.

Fontagné, Lionel, Philippe Martin, and Gianluca Orefice (2024). “The many channels of firm’s adjustment to energy shocks: evidence from France”. In: *Economic Policy* 39.117, pp. 5–43.

Grossman, Gene M, Elhanan Helpman, and Alejandro Sabal (2024). “Optimal resilience in multitier supply chains”. In: *The Quarterly Journal of Economics* 139.4, pp. 2377–2425.

Handley, Kyle, Fariha Kamal, and Ryan Monarch (2024). “Supply chain adjustments to tariff shocks: Evidence from firm trade linkages in the 2018-2019 US trade war”. In: *Economics Letters* 244, p. 112009.

Heise, Sebastian, Justin R Pierce, Georg Schaur, and Peter K Schott (2025). “How Do Firms in Different Sectors Organize Their Supply Chains? Evidence from Transaction-Level Import Data”. In: *AEA Papers and Proceedings*. Vol. 115. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 177–181.

Hinz, Julian, Irene Iodice, and Moritz Schularick (2025). “Prinzipien für eine europäische Industrie-und Handelspolitik für das geoökonomische Zeitalter”. In: *Perspektiven der Wirtschaftspolitik* 26.3, pp. 290–298.

Itskhoki, Oleg and Elina Ribakova (2024). “The economics of sanctions: From theory into practice”. In: *Brookings Papers on Economic Activity* 425, p. 470.

Korinek, Anton (2023). “Generative AI for economic research: Use cases and implications for economists”. In: *Journal of Economic Literature* 61.4, pp. 1281–1317.

Mancini, Michele, Francesco Paolo Conteduca, and Alessandro Borin (2024). “The real-time impact of the war on Russian imports: A synthetic control method approach”. In: *World Trade Review* 23.4, pp. 433–447.

Mellens, Martin, Herman Noordman, and Johan Verbruggen (2007). *Re-exports: International comparison and implications for performance indicators*. Tech. rep. CPB Netherlands Bureau for Economic Policy Analysis.

Moll, Benjamin, Moritz Schularick, and Georg Zachmann (2023). *Not even a recession: The great German gas debate in retrospect*. Tech. rep. ECONtribute Policy Brief.

Rauch, James E (1996). *Networks versus Markets in International Trade*.

Silva, JMC Santos and Silvana Tenreyro (2006). “The log of gravity”. In: *The Review of Economics and statistics*, pp. 641–658.

Wooldridge, Jeffrey M (2023). “Simple approaches to nonlinear difference-in-differences with panel data”. In: *The Econometrics Journal* 26.3, pp. C31–C66.

— (2010). *Econometric Analysis of Cross Section and Panel Data*. 2nd. Cambridge, MA: MIT Press.

A Data Construction

This appendix provides additional details on the construction of the exposure measure, the mapping of sanctions lists to HS codes, the use of the AIPNET network, and the BEC-based intermediate goods filter. It documents data sources and cleaning steps to ensure reproducibility.

A.1 UN Comtrade import data and pre-war Russia reliance

This section documents the construction of the pre-war Russia-reliance measure used in the exposure index. All steps are implemented in R in the script `01_Scripts/03_download_comtrade_data_20251230.r`, which relies on the `comtradr`, `dplyr`, and `readr` packages, among others.

A.1.1 Product universe and reporters

The analysis operates on a fixed HS 6-digit product universe defined by the intersection of the AIPNET network (HS 2002 vintage) and the official European industrial classification (NACE Rev. 2), constructed exclusively via official concordance tables rather than ad-hoc string operations. To bridge these classifications, I construct a transparent concordance pipeline that maps HS 2002 codes to NACE Rev. 2 via ISIC Rev. 4. This relies on official crosswalks from the World Bank and UNSD (specifically, HS 2002 $\rightarrow$ ISIC Rev. 3 and ISIC Rev. 3.1 $\rightarrow$ ISIC Rev. 4). The script processes these linkages by:

1. normalizing HS codes to 6-digit character strings with leading zeros;
2. applying a two-step hierarchical mapping to assign ISIC and NACE codes, resolving many-to-many relationships via uniform weighting ($1/n$);
3. retaining exactly the 4,955 HS 6 codes that map successfully to the industrial classification, ensuring the sample covers the relevant manufacturing universe.

This yields a universe of 4,955 unique HS 6-digit product codes defined by AIPNET. Of these, 2,442 products exhibit non-zero network exposure to sanctioned inputs in the final estimation sample.

On the country dimension, the script targets ten EU economies as reporters:

BEL, DEU, DNK, ESP, FIN, FRA, ITA, NLD, AUT, SWE.

For each reporter, imports are requested from two partners:

- “World” (Comtrade partner code: all origins), and
- Russia (ISO3: RUS).

The time dimension covers the years 2019–2024 at monthly frequency. In the API calls, this corresponds to monthly imports with a reporting period between January and December of each calendar year.

A.1.2 API calls, chunking, and standardization

The Comtrade API has rate limits and row limits per query. To respect these constraints, the script splits the HS 6 universe into chunks of at most 200 codes and queries these chunks sequentially. Each call to `ct_get_data()` is wrapped in a helper function that enforces:

- a hard timeout (120 seconds per call),
- randomised waiting times between calls to throttle the request rate, and
- up to five retries with a cool-down phase in case of transient errors or rate limits.

If Comtrade returns only an aggregate row count (a single-column data frame with a `count` column) rather than the underlying trade records, the chunk is recursively bisected: the HS list is split into two halves, and each half is queried separately. This continues until either a valid data frame is returned or the chunk size falls below a minimum threshold (about 25 HS 6 codes). In the latter case, the script classifies the chunk as *count-only tiny* and skips it.

All successful chunks are standardised into a common schema with the following columns:

- `reporter_iso` (ISO3 reporter code),
- `partner_iso` (ISO3 or partner label),
- `year_cal` (calendar year inferred from the period string),
- `commodity_code` (HS 6-digit code),
- `trade_value_usd` (import value in US dollars),
- `net_wgt_kg` (net weight in kilograms, if available).

Each standardised chunk is stored temporarily. For every reporter–partner pair, the script binds all available chunks into a single data frame and saves it as an RDS file: `02_Data/01_raw/03_Comtrade_Raw_RDS/raw_world_imports_XXX.rds` or `raw_rus_imports_XXX.csv`, where `XXX` is the reporter ISO3 code.

A dedicated skip log (`02_Data/01_raw/03_Comtrade_Raw_RDS/skipped_chunks_log.csv`) records any chunk that could not be retrieved in a usable form. For each skipped chunk, the log stores the reporter, partner, year, an internal chunk identifier, the number of HS codes in the chunk, and a reason flag, such as

- `count_only_tiny`: the API returned only a row count for a chunk of at most 25 codes;
- `error_after_retry_tiny`: repeated timeouts or rate-limit errors for a small chunk;
- `no_rows_after_split`: no usable rows remained after all recursive splits.

A.1.3 Construction of the pre-war reliance file

In a next step, the script aggregates the monthly import records to the country–product level and constructs the pre-war Russia-reliance measure. For each reporter country c and HS 6 product i, the script computes:

$$\begin{aligned}\text{total\_imports\_pre\_shock}_{ci} &= \sum_{t=2019}^{2021} \text{Imports from World}_{cit},\\ \text{imports\_from\_russia\_pre\_shock}_{ci} &= \sum_{t=2019}^{2021} \text{Imports from Russia}_{cit}.\end{aligned}$$

These aggregates are obtained by grouping the `World` and `Russia` combined frames by `reporter_iso` and `commodity_code` over the years 2019–2021 and summing `trade_value_usd`. The two aggregates are then merged, and any missing Russia imports are set to zero.

The pre-war reliance measure is defined as

$$\text{reliance\_on\_russia}_{ci} = \begin{cases} \dfrac{\text{imports\_from\_russia\_pre\_shock}_{ci}}{\text{total\_imports\_pre\_shock}_{ci}}, & \text{if total\_imports\_pre\_shock}_{ci} > 0,\\ 0, & \text{if total\_imports\_pre\_shock}_{ci} = 0,\end{cases}$$

clamped to the unit interval. The resulting file `02_Data/02_intermediate/03_pre_shock_reliance_on_russia.csv` contains one row per country–product pair with the variables `country_iso`, `hs6_code`, `reliance_on_russia`, and `total_imports_pre_shock`. This file is the sole input for the pre-war Russia-reliance component of the exposure index used in the main text.

A.1.4 API failures and the treatment of Austria

The skip log reveals two types of API issues:

1. **Sparse chunk-level failures for some reporters.** For Denmark, Finland, and Sweden, a small number of reporter–partner–year–chunk combinations for Russia imports in 2019–2024 remain in `count_only_tiny` or `no_rows_after_split` status. These chunks correspond to small subsets of HS 6 codes and could not be recovered despite recursive splitting and repeated retries. They are left as missing and implicitly drop out of the pre-war aggregation. Since the main results are robust to dropping Denmark, Finland, or Sweden one at a time (see Section 4.3), these residual gaps do not materially affect the exposure distribution.
2. **Systematic failures for Austria.** For Austria, the situation is qualitatively different. In 2019–2021, essentially all HS 6 chunks for both "World" and Russia imports return either `count_only_tiny` or `no_rows_after_split` even after repeated retries. As a consequence, the combined Comtrade frames for Austria are effectively empty, and no reliable pre-war Russia-reliance profile can be constructed for Austrian imports. Including Austria in the exposure-based analysis would therefore require strong assumptions about its pre-war sourcing structure that cannot be justified from the raw data.

Given this pattern, Austria is excluded from all analyses that rely on the pre-war reliance measure. The effective country universe for the exposure-based specifications thus consists of nine EU economies: Belgium, Germany, Denmark, Spain, Finland, France, Italy, the Netherlands, and Sweden. The raw Comtrade skip log is included in the replication package to document these failures and allow independent verification.

A.2 Building the analysis-ready panel

This subsection documents how the raw Comtrade responses, the sanctions list, the AIPNET network, and the HS conversion tables are combined into the analysis-ready panel used in the event-study specifications. All steps are implemented in the R script `01_Scripts/05_build_master_panel_20251230.r`. This script operates exclusively on local files produced in earlier steps and does not issue any further API calls.

A.2.1 Overview and inputs

The script takes as inputs:

- `02_Data/01_raw/03_Comtrade_Raw_RDS/raw_world_imports_XXX.rds`: combined Comtrade responses for imports from the world for reporter country c (ISO3 code `XXX`), at the HS 6 level and monthly frequency;

- `02_Data/01_raw/03_Comtrade_Raw_RDS/raw_rus_imports_XXX.csv`: analogous files for imports from Russia;
- `sanctioned_downstream_hs2002.rds`: a cleaned list of sanctioned downstream HS 2002 products produced by `01_Scripts/01_clean_sanctions_list_20251230.r`;
- `02_Data/01_raw/02_01_AIPNET_Global_Network_20241204.xlsx`: the AIPNET production network at the HS 2002 level;
- `02_Data/01_raw/04_HS_Conversion_2022_2002.xlsx`: the official HS 2022–HS 2002 conversion and correlation tables.

The configuration section defines the pre-war period (2019–2021), the panel years (2019–2024), and the shock year (2022). It also specifies the input and output directories. All subsequent steps are deterministic conditional on these files.

A.2.2 Pre-war Russia reliance at HS 2022 and treatment intensity

The first block constructs a pre-war reliance measure at the HS 2022 level directly from the local Comtrade CSV files.

(i) Ensuring a calendar-year variable and quantity detection. The world and Russia import files for all reporters are read and combined. A helper function `ensure_year_cal()` inspects each data frame and creates a variable `year_cal` by parsing either a `year` or `period` column. The script then attempts to detect a quantity column for net weight in kilograms by searching for names such as `net_wgt_kg`, `netwgt_kg`, or `quantity_kg`. If such a column exists, it is used to construct annual import weights; if not, the script falls back to value-only aggregation.

(ii) Pre-war aggregation and reliance on Russia. For each reporter c and HS 6 commodity code h in HS 2022 space, pre-war imports are aggregated as

$$\begin{aligned}\text{total\_imports\_pre\_shock}_{ch} &= \sum_{t=2019}^{2021} \text{Imports from World}_{cht},\\ \text{imports\_from\_russia\_pre\_shock}_{ch} &= \sum_{t=2019}^{2021} \text{Imports from Russia}_{cht}.\end{aligned}$$

Missing Russia imports are treated as zeros. The pre-war Russia-reliance measure is then defined as

$$\text{reliance\_on\_russia}_{ch} = \begin{cases} \dfrac{\text{imports\_from\_russia\_pre\_shock}_{ch}}{\text{total\_imports\_pre\_shock}_{ch}}, & \text{if total\_imports\_pre\_shock}_{ch} > 0,\\ 0, & \text{if total\_imports\_pre\_shock}_{ch} = 0.\end{cases}$$

The resulting measure is clamped to the unit interval. Commodity codes are cleaned to a strict six-digit numeric format using the helper `clean_hs6()`, which strips non-digits, truncates to six characters, and pads with leading zeros where needed.

(iii) Mapping reliance from HS 2022 to HS 2002 downstream products. To link pre-war reliance to the AIPNET network, the HS 2022 reliance data are mapped into HS 2002 space using the official conversion table. The script reads the sheet "HS2022-HS2002 Correlations", cleans the code columns to six-digit numeric strings, and retains only valid pairs. Because HS revisions are many-to-many, a single HS 2022 code may map to multiple HS 2002 codes and vice versa; the `left_join()` explicitly declares this as a many-to-many relationship.

When several HS 2022 codes map into the same HS 2002 product, I treat them symmetrically and take a simple unweighted average of their pre-war Russia-reliance shares. This implies that all HS 2022 varieties within a given HS 2002 heading contribute equally to the exposure index. Alternative weighting schemes that use pre-war world imports or the number of downstream linkages as weights produce very similar exposure distributions and do not materially affect the results.

The reliance data are merged with the HS 2022–HS 2002 map and then averaged at the countryHS 2002 level:

$$\begin{aligned} \text{reliance\_on\_russia\_avg}_{c,\text{dn}} &= \text{average of reliance\_on\_russia}_{ch} \\ &\quad \text{over all HS 2022 codes that map into a given downstream HS 2002 product.} \end{aligned}$$

This produces a table `reliance_hs2002_dn` with one row per country and downstream HS 2002 code.

(iv) Combining with the sanctions list and AIPNET. The cleaned sanctions file `sanctioned_hs2002.rds` contains the set of HS 2002 products that are legally sanctioned, as constructed in `01_Scripts/04_helper_map_CN8_to_HS6_20251230.r`. The script intersects this set with the AIPNET network by:

1. reading the AIPNET Excel file and identifying the upstream and downstream HS 2002 columns via pattern matching;
2. cleaning both to six-digit numeric HS 2002 codes and keeping only valid edges;
3. restricting the edges to those in which the *upstream* node is sanctioned, i.e. $u \in \mathcal{S}$.

The resulting edge list $\mathcal{E}$ consists of pairs (u, i), where u is a sanctioned upstream input and i a downstream user HS 2002 product. Joining the country-level pre-war reliance measure Reliance_{cu} onto this edge list and aggregating over sanctioned inputs yields the downstream exposure index:

$$\text{Exposure}_{ci} = \sum_{u \in \mathcal{S}} w_{ui}\, \text{Reliance}_{cu},$$

where w_{ui} is the (normalized) AIPNET link intensity from input u to user product i. A global analogue averages Reliance_{cu} across countries before aggregation.

Exposure values that remain missing (for products with no sanctioned upstream inputs in their upstream neighborhood or for countries without reliance data) are set to zero in the analysis-ready panel.

A simple diagnostic computes the correlation between `treat_intensity` and `treat_intensity_global` across HS 2002 products; in the current data this correlation is around 0.52, indicating a substantial but not perfect alignment between country-specific and global exposure.

A.2.3 Annual HS 2002 trade panel

The next block constructs the trade panel to which the exposure measures are attached.

(i) Annual HS 2022 imports from the world. The combined world-imports data are collapsed to the country–year–HS 6 level in HS 2022 space. For each reporter c, year t, and HS 6 code h, the script computes:

- `imp_usd`: annual imports from the world in current US dollars, by summing `trade_value_usd`;
- `imp_kg`: annual net weight in kilograms, by summing the detected quantity column where available.

If no suitable quantity column has been detected, the script produces only the value aggregate.

(ii) Mapping annual flows from HS 2022 to HS 2002. The annual HS 2022 flows are then mapped into HS 2002 space using the same conversion table as in the exposure construction. After cleaning, each HS 2022 code is matched to all correlated HS 2002 codes in a many-to-many join, and imports are aggregated to the country–year–HS 2002 level:

$$\text{imp\_usd}^{\text{HS2002}}_{cut} = \sum_{h \mapsto u} \text{imp\_usd}_{cht}, \qquad \text{imp\_kg}^{\text{HS2002}}_{cut} = \sum_{h \mapsto u} \text{imp\_kg}_{cht},$$

where the sums run over all HS 2022 codes h that map into a given HS 2002 code u.

(iii) Merging treatment intensities. The country-specific and global treatment-intensity tables are then merged onto this annual HS 2002 panel. Formally, the panel is joined to `treat_up_by_country` on (c, u) and to `treat_up_global` on u. This yields, for each country–HS 2002 product–year observation (c, u, t), the annual import values and quantities together with the two exposure measures.

A.2.4 Derived analysis variables and export

Finally, the script constructs the analysis variables used in the event-study specifications and writes the analysis-ready file.

(i) Log transformations and winsorisation. For each country–product–year observation, log import values and, where available, log quantities are computed as

$$\text{log\_imp\_usd}_{cut} = \log(1 + \text{imp\_usd}_{cut}), \qquad \text{log\_imp\_kg}_{cut} = \log(1 + \text{imp\_kg}_{cut}).$$

The 99th percentile of `log_imp_usd` is used to construct a winsorised version `log_imp_win99` by truncating all higher values at this cutoff. This reduces the influence of extreme trade flows in value-based regressions.

(ii) Relative-year and product identifiers. Two simple identifiers are added for later use in Stata:

- a relative-year indicator `rel_year`, defined as $t - 2022$ so that 2022 corresponds to event time zero;
- a numeric product index `product_id`, obtained by sorting HS 2002 codes and assigning consecutive integers starting from one.

In addition, a `date_fe` variable encodes calendar years as consecutive integers starting at one and is convenient when specifying time fixed effects in Stata.

(iii) Exposure variables. The merged treatment-intensity variables are converted into exposure measures `treat_ctry` and `treat_glob` by replacing any remaining missing values with zero. Economically, `treat_ctry` measures the reliance of country c on sanctioned downstream uses of upstream HS 2002 product u based on its own pre-war import structure, while `treat_glob` measures the corresponding reliance when pre-war imports are averaged across countries. Both are bounded between zero and one by construction.

(iv) Final variable selection and export. The final dataset `02_Data/03_final/05_panel_full_vkg.dta` contains, for each country–year–HS 2002 product, the following variables:

- identifiers: `country_iso`, `year_cal`, `rel_year`, `hs2002`, `product_id`, `date_fe`;
- trade outcomes: `imp_usd`, `imp_kg` (where available), `log_imp_usd`, `log_imp_kg`, `log_imp_win99`. The log unit value variable `log_uv` used in the regressions is constructed in the Stata master script `01_Scripts/08_main_analysis_eventstudy_20260106.do` as $\log(1 + \text{imp\_usd}) - \log(1 + \text{imp\_kg})$.
- exposure measures: `treat_ctry`, `treat_glob`.

This Stata file is the sole input for the event-study and difference-in-differences estimators described in the main text. All intermediate R objects produced along the way (such as the HS 2022–HS 2002 mapping, the sanctioned downstream set, and the AIPNET edge list) are included in the replication package to allow independent verification and extensions.

A.2.5 BEC-based intermediate-goods filter (WITS)

The final step in the data construction applies a Broad Economic Categories (BEC) filter to restrict the HS 2002 panel to intermediate goods. This is implemented in the script `01_Scripts/06_filter_intermediates_BEC_20260105.r`.

Intermediate goods are defined by the following BEC categories:

- **111**: Food and beverages, primary, mainly for industry
- **121**: Food and beverages, processed, mainly for industry
- **21**: Industrial supplies not elsewhere specified, primary
- **22**: Industrial supplies not elsewhere specified, processed
- **31**: Fuels and lubricants, primary
- **32**: Fuels and lubricants, processed (other than motor spirit)
- **42**: Parts and accessories of capital goods (except transport equipment)
- **53**: Parts and accessories of transport equipment

The script uses the official World Bank WITS concordance between HS 2002 product codes and BEC categories. It takes `02_Data/03_final/05_panel_full_vkg.dta` as its sole input and first standardises HS 2002 codes to six-digit character strings with leading zeros:

$$\texttt{hs2002} \;\rightarrow\; \texttt{hs6} \in \{000000, \ldots, 999999\}.$$

It then loads the WITS concordance file `02_Data/01_raw/06_HS_to_BEC_Concordance.csv` with all columns read as character variables to preserve leading zeros. After cleaning column names, the script constructs a long mapping

$$\texttt{hs6} \rightarrow \texttt{bec\_code}$$

from `hs_2002_product_code` to `bec_product_code`. Because some HS 2002 codes map to more than one BEC category, the script collapses the mapping to a unique HS6 → BEC assignment by sorting pairs lexicographically by (`hs6`, `bec_code`) and keeping the first BEC code per HS6. This deterministic tie-breaking rule is documented here so that the HS6–BEC mapping can be reconstructed exactly.

Merging this mapping onto the HS 2002 panel yields a one-to-one BEC assignment for all products in the data. In the present application, the WITS concordance covers all 4,570 distinct HS 2002 codes observed in the panel, so no HS code is lost due to missing BEC information.

Following the UN System of National Accounts convention, the script classifies products into four economic categories based on their three-digit BEC code:

- **Intermediate goods:** BEC codes $\{111, 121, 21, 22, 31, 32, 42, 53\}$, covering primary and processed industrial supplies, fuels used as intermediate inputs, and parts and accessories of capital goods and transport equipment;
- **Capital goods:** BEC codes $\{41, 521\}$;
- **Consumption goods:** BEC codes $\{112, 122, 522, 61, 62, 63\}$;
- **Unclassified/Other:** all remaining codes, kept only for diagnostics.

Applied to the full HS 2002 panel before filtering, this classification yields the composition reported in Table 1 in Section 2.3.

The high share of intermediate goods in total weight and their lower share in total value are consistent with a composition in which heavy bulk flows (such as basic metals, chemicals and fuels) are predominantly intermediate inputs, while higher-value but lighter items include many capital and consumption goods.

As a simple methodological check, the script inspects a small set of robotics-related HS codes: 847990 (parts and accessories of machines), 848340 (gears and gearing), 853710 (programmed controllers), and 850131 (small electric motors). Under the WITS concordance, the first three are assigned to BEC 42 and thus classified as intermediate goods, whereas HS 850131 is assigned to BEC 41 and classified as capital goods. In line with the strict intermediate-goods definition adopted here, HS 850131 is excluded from the estimation sample. This example illustrates that some technologically intermediate components are labelled as capital goods in the UN BEC system; the paper follows the official BEC coding rather than reclassifying such cases by hand.

The final step filters the panel to intermediate goods only by keeping observations with `econ_category = "Intermediate"`. All BEC and classification variables are then dropped, and the resulting dataset is written to `02_Data/03_final/06_01_panel_intermediates_only_vkg.dta`. This file contains the same identifiers, trade outcomes, and exposure measures as `02_Data/03_final/05_panel_full_vkg.dta`, but restricted to HS 2002 products classified as intermediate inputs. It is this BEC-filtered panel that is used in the event-study and difference-in-differences analyses in the main text.

B Pretrend diagnostics and country-group screening

This appendix documents the exposure-based pretrend screening of country groups and clarifies how to interpret the fact that some groups exhibit very clean pretrends "by chance".

B.1 Design of the exposure-based pretrend tests

The pretrend diagnostics are implemented in the Stata script `01_Scripts/07_pretrend_diagnostics_20260105.do`. The script takes the BEC-filtered HS 2002 panel `02_Data/03_final/06_01_panel_intermediates_only_vkg.dta` as input and proceeds in four steps.

(i) Enumerating country groups. Let

$$\mathcal{C} = \{\text{BEL, DEU, DNK, ESP, FIN, FRA, ITA, NLD, SWE}\}$$

denote the set of nine candidate EU economies. For each non-empty subset $G \subseteq \mathcal{C}$ with group size $K = |G| \in \{1, \dots, 9\}$, the script constructs a group indicator

$$\mathbb{I}\{c \in G\} = \begin{cases} 1, & \text{if reporter } c \text{ belongs to group } G, \\ 0, & \text{otherwise.} \end{cases}$$

Intersecting this with the pre-war exposure index yields a group-specific exposure measure

$$\text{Exposure}_{ci}^{G} = \mathbb{I}\{c \in G\} \times \text{Exposure}_{ci},$$

which is zero for all countries outside G. In total, this produces 511 groups, each of which is evaluated for three outcomes (log import values, log import quantities, and log unit values).

(ii) Pre-period regressions. For each group G and outcome $Y \in \{\log \text{Val}, \log \text{Kg}, \log \text{UV}\}$, the script estimates the pre-period regression

$$Y_{cit} = \alpha_{ci} + \lambda_{ct} + \sum_{\tau \in \{-3,-2,-1\}} \beta_{\tau}^{\text{pre}}(G) \Big(\mathbb{I}\{t = 2022 + \tau\} \times \text{Exposure}_{ci}^{G} \Big) + \varepsilon_{cit},$$

restricting the sample to $t \in \{2019, 2020, 2021\}$. The country–product fixed effects α_{ci} absorb time-invariant comparative advantage and baseline trade levels, and the country–

year fixed effects λ_{ct} remove country-specific macro shocks in each year. This fixed-effect structure is identical to the one used in the main event-study specification in Equation (1), except that the regression is estimated only on pre-war data. The coefficients $\beta_{\tau}^{\text{pre}}(G)$ therefore capture whether, in the pre-war period, more exposed products systematically evolved differently from less exposed products *within* group G, after stripping out common country-wide movements. Standard errors are clustered by product, consistent with the main analysis.

Economically, this is a pure design diagnostic. Because the exposure index is defined from pre-war trade patterns and the regression uses only 2019–2021 outcomes, there is by construction no mechanical feedback from the sanctions shock to the pretrend tests.

(iii) Joint-zero, flatness, and magnitude diagnostics. For each group G and outcome, three diagnostics are computed:

1. A *flatness* F-test that checks whether the three pre-period coefficients are pairwise equal, i.e. whether $\beta_{-3}^{\text{pre}}(G) = \beta_{-2}^{\text{pre}}(G)$ and $\beta_{-2}^{\text{pre}}(G) = \beta_{-1}^{\text{pre}}(G)$. This test is mainly descriptive.
2. A *joint-zero* F-test of $\beta_{-3}^{\text{pre}}(G) = \beta_{-2}^{\text{pre}}(G) = 0$, with associated p-value $p_{\text{zero}}(G)$ and BenjaminiHochberg adjusted q-value $q_{\text{zero}}(G)$ (Benjamini and Hochberg, 1995).
3. A *magnitude* measure defined as

$$\text{max\_abs\_pre}(G) = \max\Big\{|\beta_{-3}^{\text{pre}}(G)|, |\beta_{-2}^{\text{pre}}(G)|\Big\}.$$

The magnitude metric is used to guard against cases in which the pre-period coefficients are individually noisy but jointly appear insignificant simply because the sample is small or the standard errors are large. To calibrate what constitutes a "large" pre-period deviation, the script computes the median of max_abs_pre(G) across all groups and sets the magnitude threshold equal to this median.

(iv) Classification and ranking. For each groupoutcome pair, three binary indicators are defined:

- pass_zero $= 1$ if $p_{\text{zero}}(G) \geq 0.10$;
- pass_mag $= 1$ if max_abs_pre(G) does not exceed the median threshold;
- pass_strict $= 1$ if both conditions hold.

The indicator `pass_strict` is the primary criterion used in the main text. Groups are then ranked by `pass_strict`, the joint-zero p-value, and the magnitude measures, separately for values and quantities.

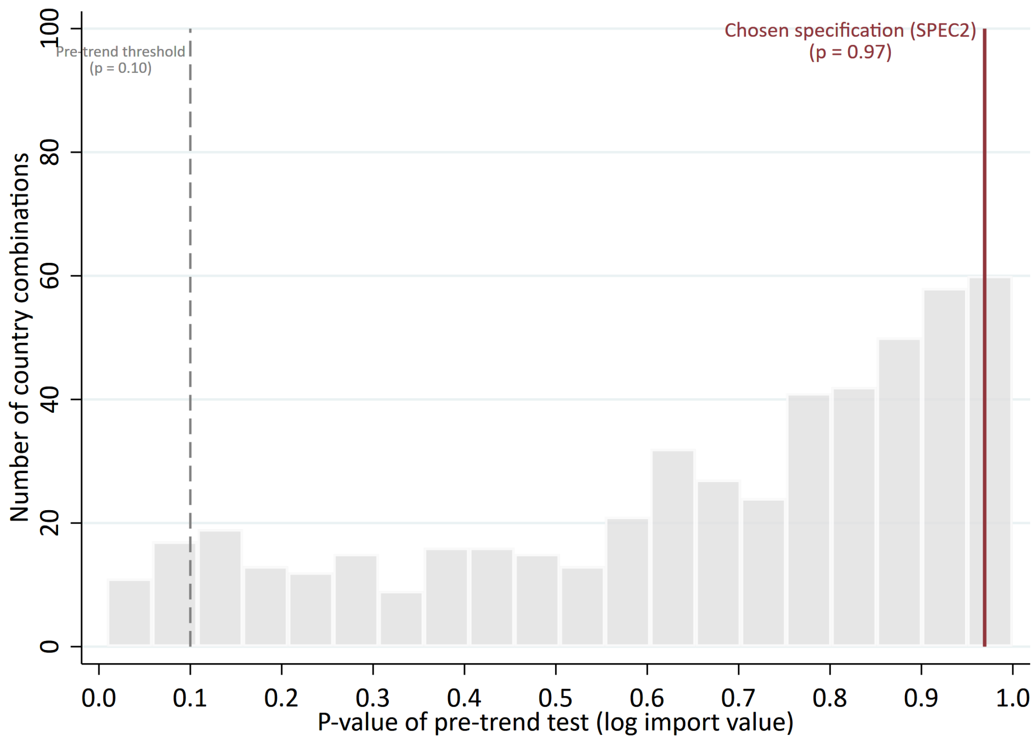

Figure A1. Histogram of Pre-Trend p-Values across Country Combinations.

Notes: The histogram shows the distribution of p-values from the joint F-test (H_0: zero pre-trends) estimated for all 511 possible subsets of the nine sample countries. The vertical line marks the p-value for the baseline specification (SPEC 2).

Figure A1 visualizes the result of this screening procedure. The distribution is skewed toward one, indicating that for many random combinations of countries, the parallel trends assumption holds reasonably well. However, the baseline specification sits in the far right tail of the distribution ($p \approx 0.97$), confirming that the chosen country group minimizes pre-trend deviations even relative to this generally well-behaved distribution. This confirms that the selected group satisfies the parallel trends assumption not merely by clearing a low bar, but by exhibiting substantially higher stability than the average random permutation of EU economies.

B.2 Why some groups have clean pretrends "by chance"

Under the null hypothesis that, absent the sanctions, exposure is not systematically correlated with pre-war trends, the pre-period coefficients $\beta_\tau^{\text{pre}}(G)$ should be centred around zero for any given group G, up to sampling error. However, there are two reasons why some groups will exhibit seemingly cleaner pretrends than others:

1. **Sampling variation in high-dimensional panels.** The regressions are estimated on a very high-dimensional panel with thousands of country-product-year cells. Even if exposure is truly unrelated to pre-war trends, random co-movements between exposure and transitory shocks in 2019-2021 will generate non-zero estimates of $\beta_\tau^{\text{pre}}(G)$ for some groups. With 511 country groups and three outcomes, it is statistically expected that some combinations look particularly well-behaved and others less so, simply due to sampling variation.
2. **Averaging of idiosyncratic product shocks.** For small groups (say, a single country), idiosyncratic shocks to a few highly exposed and volatile products can create visible pre-period movements that correlate with exposure. As more countries are combined, these product-level shocks are averaged across a larger set of markets, and the residual correlation between exposure and pre-period noise tends to shrink. This is exactly what the diagnostics show: the share of groups passing the strict pretrend criterion rises from about 17–22% for singleton countries to roughly 50% for groups of size five and reaches 100% for the full nine-country group.

The fact that some groups appear to have exceptionally clean pretrends is therefore not evidence of "hidden structure" in pre-war exposure, but rather a predictable consequence of sampling variation and averaging in a high-dimensional setting. The screening exercise makes this explicit and quantifies it.

B.3 Why using the pretrend screen is scientifically valid

Selecting a baseline sample based on pre-period fit raises two potential concerns: (i) overfitting or "p-hacking"; and (ii) implicit conditioning on post-treatment outcomes. The design of the present screen addresses these concerns in three ways.

First, the pretrend diagnostics use *only* information from the pre-war period (2019-2021) and the exposure index, which itself is constructed from pre-war trade patterns. Post-2022 outcomes do not enter the selection rule in any way. Conditional on the exposure measure, the choice of country group is therefore statistically independent of the realized treatment effects. This is analogous to choosing bandwidths or smoothing parameters based on pre-period fit in time-series designs.

Second, the strict pretrend criterion is conservative. A group is not selected because it maximizes a p-value or minimizes a pre-period coefficient, but because it satisfies two modest requirements simultaneously: failing to reject joint-zero at the 10% level and having pre-period coefficients that are no larger in absolute value than a typical group in the cross-section. In other words, the screen is designed to *exclude* groups with clearly problematic pretrends, not to search for a single "perfect" group.

Third, the pretrend screen is not used to hide specifications that behave differently. The main text reports results for (i) a five-country group that passes the strict criterion, (ii) alternative country groups that also pass, and (iii) the full nine-country sample, which by construction averages over all candidate markets and also passes the pretrend tests. The nine-country specification serves as a "kitchen-sink" benchmark: if the main qualitative conclusions were to rely on a very particular country combination, this would become visible in a comparison with the full-sample results.

Taken together, the exposure-based pretrend screen should be viewed as a transparent design diagnostic that helps to identify country groups for which the parallel-trends assumption is empirically most plausible, while keeping the focus on economically meaningful and politically relevant combinations. It complements, rather than replaces, the usual robustness checks and alternative specifications reported in the main text.

C Additional figures and robustness checks

This appendix provides supplementary empirical evidence supporting the main analysis. Section C.1 documents the “first stage” collapse of direct trade flows from Russia. Section C.2 presents identification checks distinguishing the effect of legal sanctions from general reputational decoupling. Section C.3 reports comprehensive robustness diagnostics, including sensitivity to sample composition, functional form (PPML), and sectoral heterogeneity.

C.1 First-stage dynamics: The collapse of direct trade

Figure A2 provides a diagnostic of the direct "first stage" shock. It decomposes direct imports from Russia into sanctioned (Annex XXI) and non-sanctioned goods. Two patterns emerge. First, nominal import values (Panel A) initially remained high in 2022 due to the price increases documented in the main text, masking the physical decoupling. Second, physical quantities (Panel B) contracted for both groups. While the decline in non-sanctioned goods reflects broad-based friction (financial sanctions, overcompliance), the sanctioned category exhibits a sharp, binding collapse coinciding with the legal implementation. The identification strategy in the main text relies on the differential cost implications of this forced substitution compared to voluntary decoupling.

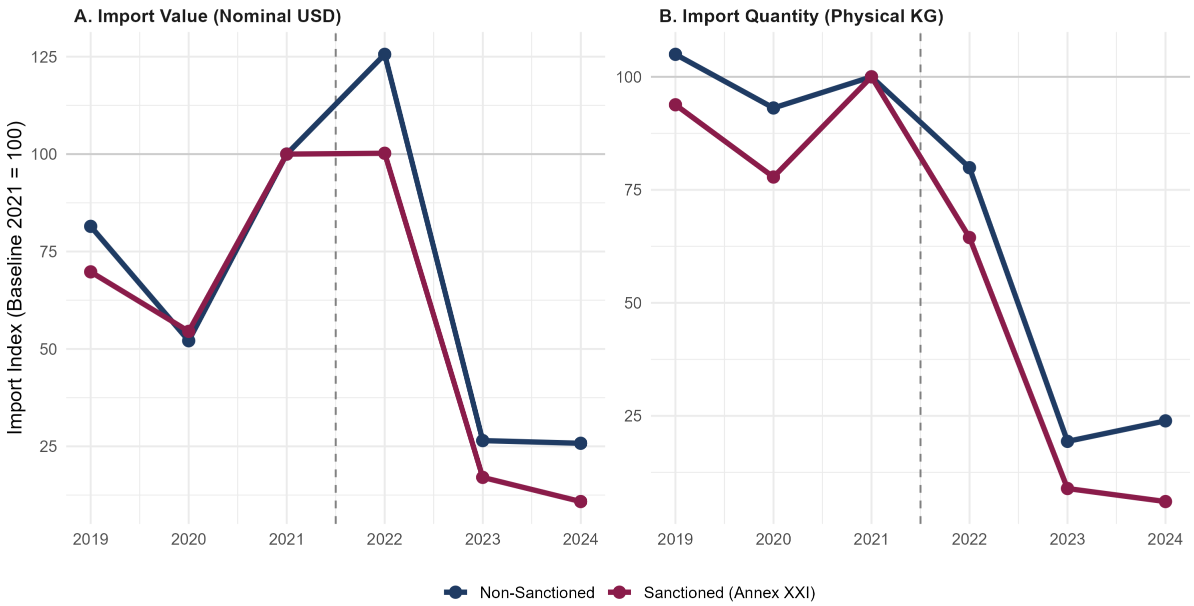

Figure A2. Direct Imports from Russia: Sanctioned vs. Non-Sanctioned Goods.

Notes: Indices of direct imports from Russia (2021 = 100) for the five-country industrial core. Panel A: Nominal import value (USD). Panel B: Import quantity (kg). "Sanctioned" refers to products listed in Annex XXI; "Non-Sanctioned" comprises all other intermediate goods.

C.2 Disentangling legal constraints from reputational friction

The collapse of direct imports from Russia documented in Figure A2 reflects two distinct forces: (i) binding legal bans (Annex XXI) that physically prevent purchases, and (ii) voluntary "overcompliance" or reputational decoupling, where firms cease sourcing non-sanctioned Russian goods. This creates a "chicken and egg" identification challenge: Is the observed price increase caused by the exogenous legal shock, or is it merely a symptom of general supply chain toxicity affecting all Russian links?

To resolve this, I estimate a "horse race" specification-a statistical test where two competing explanatory variables are included simultaneously to determine which one drives the outcome. I construct two distinct exposure measures:

- **Sanctions Exposure ($\text{Exposure}^{\text{Sanc}}$):** The baseline measure, capturing reliance strictly on inputs listed in Annex XXI (Legal Ban).
- **General Russia Exposure ($\text{Exposure}^{\text{Gen}}$):** A broad measure capturing reliance on *any* input sourced from Russia, regardless of its legal status.

I then estimate:

$$\log \text{Val}_{cit} = \alpha_{ci}+\lambda_{ct}+\beta^{\text{Sanc}}\Big(\mathbb{I}\{t=2022\}\times\text{Exposure}^{\text{Sanc}}_{ci}\Big)+\beta^{\text{Gen}}\Big(\mathbb{I}\{t=2022\}\times\text{Exposure}^{\text{Gen}}_{ci}\Big)+\varepsilon_{cit} \tag{5}$$

Table A1 reports the results. If the price shock were driven by general reputational friction, β^{Gen} should absorb the effect. Instead, the legal channel dominates. The coefficient for legal sanctions remains large and positive ($\beta^{\text{Sanc}} \approx 0.77$). In contrast, general Russia exposure yields a negative coefficient ($\beta^{\text{Gen}} \approx -0.19$).

Table A1. Robustness: Legal Sanctions vs. General Russia Exposure.

Outcome: Log Import Value	(1) Sanctions Only	(2) General Only	(3) Horse Race
Sanctions Impact (2022)	0.739		0.774
(Annex XXI Exposure)	(0.833)		(0.841)
General Russia Impact (2022)		-0.147	-0.191
(All Russia Exposure)		(0.279)	(0.285)
Fixed Effects	Ctry-Prod, Ctry-Year	Ctry-Prod, Ctry-Year	Ctry-Prod, Ctry-Year
Observations	84,022	84,022	84,022
R-squared	0.962	0.962	0.962

Notes: Estimates for the five-country industrial core. **Sanctions Exposure** captures reliance on upstream inputs specifically listed in Annex XXI (legal ban). **General Russia Exposure** captures reliance on *any* input sourced from Russia, regardless of legal status. In Column (3), the coefficient on General Exposure captures the reputational or voluntary decoupling effect conditional on the legal ban. Standard errors clustered by product in parentheses.

This divergence allows for a clear economic interpretation. Non-sanctioned decoupling acts primarily as a negative demand shock: European firms voluntarily reduce purchases, which reduces quantities without necessarily driving up prices (indeed, prices may fall due to lower demand). Legal sanctions, by contrast, act as a binding supply constraint on firms that typically *cannot* easily substitute the input (inelastic demand), forcing them to pay a premium for alternative sources. The "Price of Resilience" is thus specifically the cost of overcoming legal barriers, not a general tax on decoupling.

C.3 Additional robustness diagnostics

This section presents supplementary robustness checks referenced in the main text, including sensitivity to sample composition and sectoral heterogeneity.

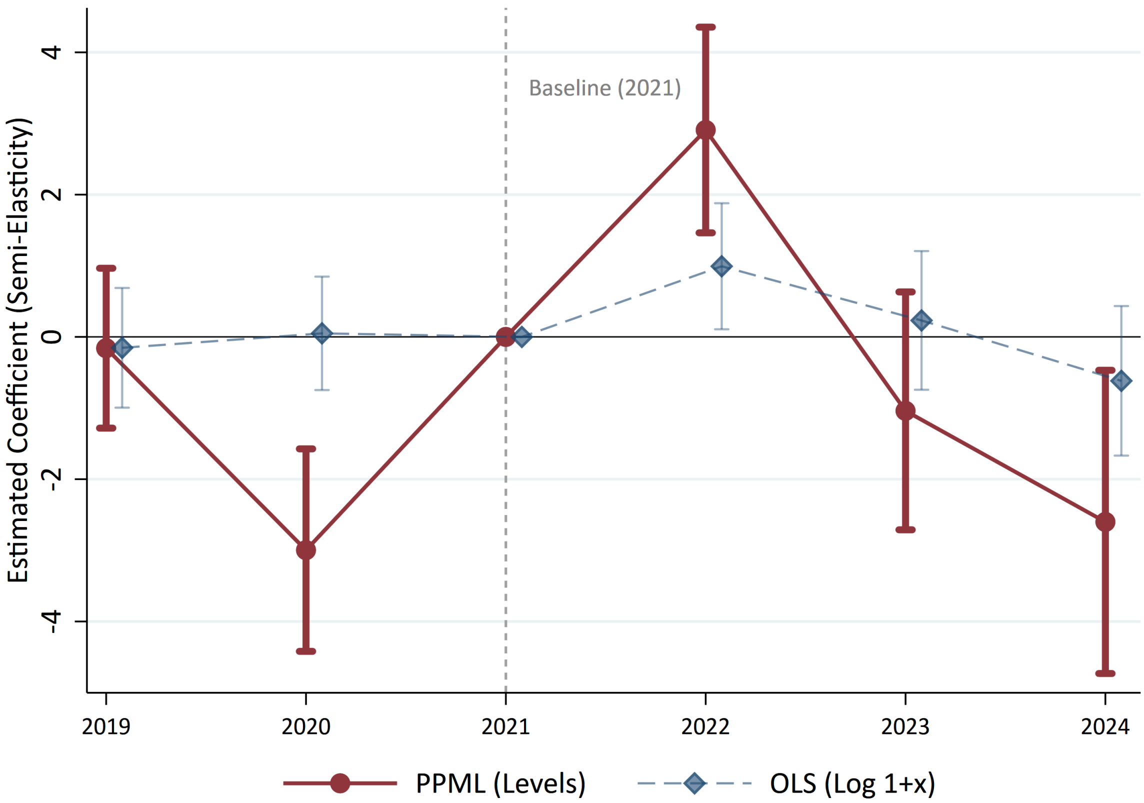

Figure A3. Comparison of Estimators: OLS vs. PPML.

Notes: Event-study coefficients estimated using OLS on $\log(1 + \text{Value})$ (Navy, dashed) and Poisson Pseudo-Maximum Likelihood (PPML) on Value levels (Maroon, solid). The PPML estimator accounts for zero trade flows and heteroskedasticity following Silva and Tenreyro (2006) and Jeffrey M Wooldridge (2023). Both specifications include country-product and country-year fixed effects. Standard errors clustered by product. Sample: Five-country industrial core.

Figure A4. Distribution of Weighted Estimates and t-Statistics.

Notes: The volcano plot displays the weighted estimated coefficient $\hat{\beta}_{2022}$ (x-axis) and the absolute t-statistic (y-axis) for all 511 country combinations. The red diamond marks the baseline specification. Weights are based on pre-war import volumes.

Table A2. Leave-One-Country-Out Robustness of the 2022 Value Coefficient.

Excluded country	$\hat{\beta}_0$ (2022)	s.e.
Germany (DEU)	1.26	(1.48)
Italy (ITA)	1.07	(1.00)
France (FRA)	0.86	(0.85)
Spain (ESP)	0.63	(0.86)
Denmark (DNK)	-0.04	(0.34)

Notes: Estimates of the 2022 coefficient when excluding the listed country from the five-country baseline sample (SPEC 2).

Interpretation. To ground the aggregate results in specific commodities, I examine German import prices for two critical metals.

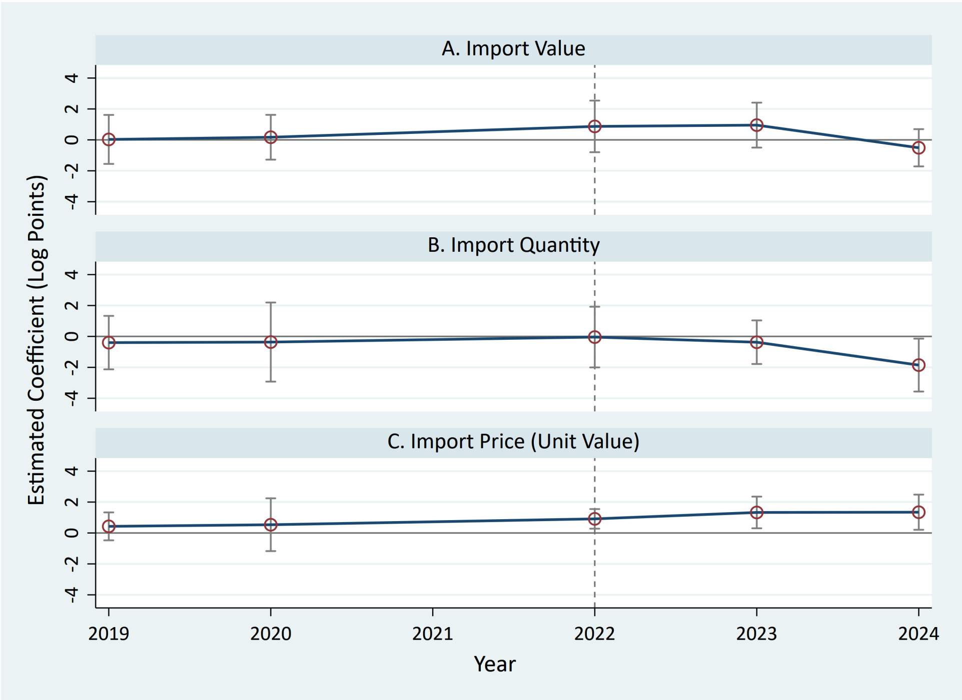

(a) **Strict Industrial Core (SPEC 1).** Sample: DEU, ITA, ESP, DNK.

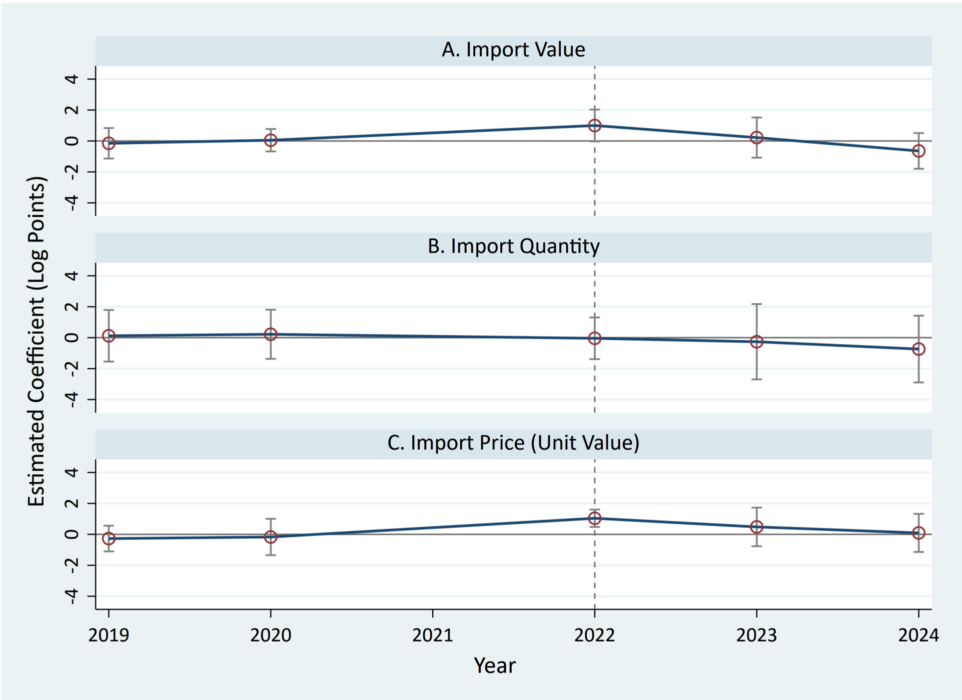

(b) **Full Sample (SPEC 3).** Sample: All 9 EU economies.

Figure A5. Event-Study Estimates for Alternative Country Samples.

Notes: Estimates of β_τ from Equation (1) for two alternative samples. Panel A excludes France and transit hubs. Panel B includes all available countries.

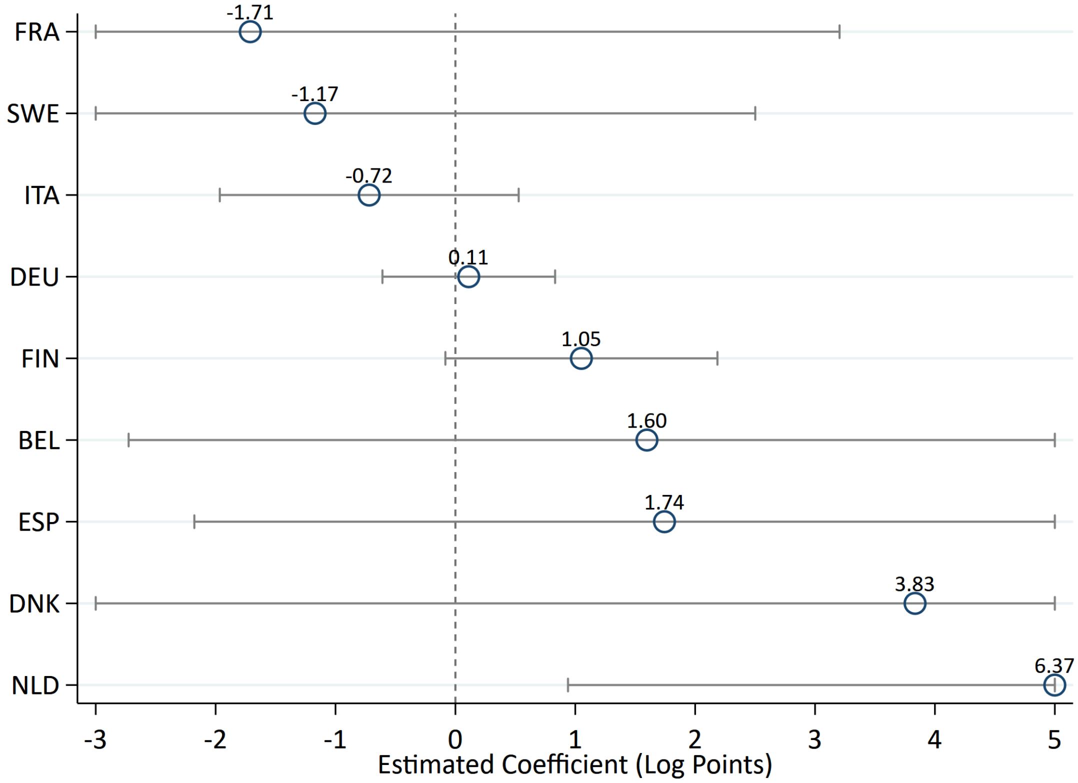

Figure A6. Estimates by Country (Full 9-Country Sample).

Notes: Estimated coefficients $\hat{\beta}_{2022}$ for log import values, estimated separately for each of the nine economies in the full sample. Estimates are unweighted.

Table A3. Leave-One-Year-Out Robustness of the 2022 Value Coefficient.

Excluded year	$\hat{\beta}_0$ (2022)	s.e.
2024	0.99	(0.80)
2023	0.76	(0.80)
2021	0.75	(0.74)
2020	0.74	(0.80)
2019	0.74	(0.80)

Notes: Estimates of the 2022 coefficient when excluding the listed calendar year. Results are based on the baseline specification (Spec 2: DEU, FRA, ITA, ESP, DNK).

Table A4. Leave-One-HS 2-Out Robustness of the 2022 Value Coefficient.

Statistic	Value
Min $\hat{\beta}_0$	-0.01
Max $\hat{\beta}_0$	0.87
Mean $\hat{\beta}_0$	0.72
Min Pre-trend p-value	0.21
Max Pre-trend p-value	0.99

Notes: Summary statistics for 83 regressions, each excluding one HS 2 chapter. Results are based on the baseline specification (Spec 2).

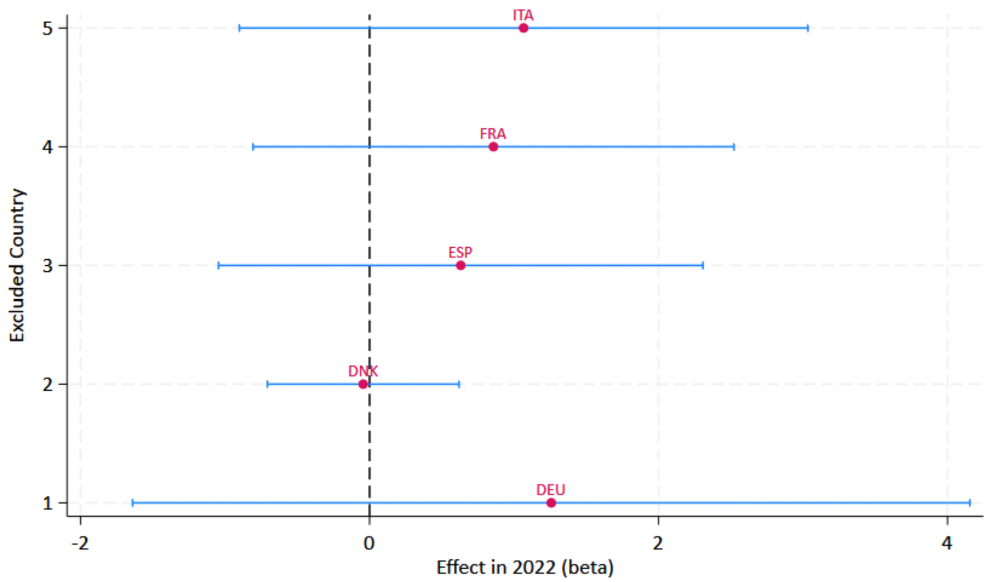

Figure A7. Leave-One-Country-Out Estimates.

Notes: Estimated coefficient $\hat{\beta}_{2022}$ for log import values when excluding the country indicated on the y-axis from the baseline five-country sample. Unweighted estimates. Whiskers denote 95% confidence intervals.

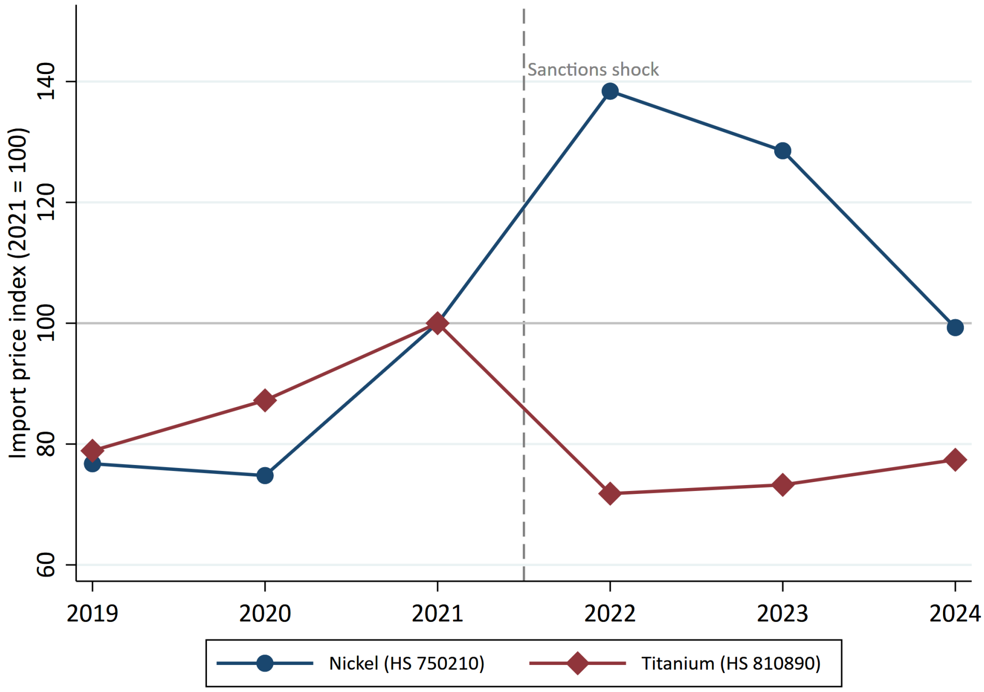

Figure A8. German Import Price Indices: Nickel and Titanium.

Notes: Annual import price indices (unit values, 2021 = 100) for German imports of Nickel (HS 750210) and Titanium (HS 810890).

Nickel (HS 750210) offers a textbook example of the resilience premium: prices spiked immediately in 2022 as firms scrambled to replace Russian supply. Titanium (HS 810890), by contrast, avoided the price surge, trading below pre-war levels. This divergence highlights the role of institutional detailstitanium was subject to specific exemptions due to its critical role in the aerospace supply chain (e.g., Airbus), shielding it from the immediate volatility that hit other raw materials.

C.4 Heterogeneity by Product Substitutability (Rauch Classification)

To test the hypothesis that the price premium is driven by goods that are technologically difficult to substitute, I merge the Rauch, 1996 classification (via SITC Rev. 2 concordance) to the HS 6 panel. I define a time-invariant dummy D_i equal to 1 if a product is classified as “differentiated” (implying relationship-specific sourcing) and 0 if it is “homogeneous” or “reference priced”.

I estimate the triple-difference interaction term $\text{Exposure}_{ci} \times D_i \times \mathbb{I}_{2022}$ within the baseline specification. As shown in Table A5, the point estimate for the *additional* price effect on differentiated goods in 2022 is large and positive ($\hat{\beta} \approx 1.36$). Ideally, one would observe a statistically significant premium; however, the standard error is large (1.11), resulting in a p-value of approximately 0.22. Despite the lack of statistical significance at conventional levels, the economic magnitude is consistent with the theoretical prediction: the "price of resilience" appears to be predominantly paid in supply chains where short-term substitution via spot markets is not an option.

Table A5. Heterogeneity: Differentiated vs. Homogeneous Goods.

Dependent Variable: Log Import Value (Winsorized)	(1) Interaction Model
Interaction: Exposure × Differentiated × 2022	1.356
	(1.110)
Exposure × 2022 (Base Effect for Homogeneous)	0.480
	(0.420)
Fixed Effects	Ctry-Prod, Ctry-Year
Observations	149,907
R-squared	0.958

Notes: The table reports the triple-difference interaction coefficient for differentiated goods in the year 2022 relative to 2021. Classification based on Rauch (1996) (Conservative definition). The variable *Exposure* is normalized to [0,1]; the coefficient 1.356 thus implies that fully exposed differentiated goods face a relative log-value increase that is 1.36 log points higher than fully exposed homogeneous goods. Standard errors clustered by product in parentheses.

Declaration on the Use of Generative AI: A Manifesto for Collaboration

Not by the lonely mind alone,
Is wisdom found or insight shown.
I summoned Claude and GPT,
To join a silent symphony.

Like Napoleon Hills great "Master Mind",
We left the narrow view behind.
The code was Robin, swift and bright,
To aid the Batman in the night.

As Jung foresaw the shadows grace,
We gave the ghost a time and place.
To complement, not substitute,
To find the tree beneath the root.

The machine may draft the chart and line,
But every final choice is mine.
A dance of silicon and soul,
To make the fragmented feel whole.

Technical Disclosure. In accordance with academic transparency, I confirm the use of Large Language Models (including Grok, Claude, GPT-5, and Gemini) as operational assistants for coding, phrasing, and stress-testing arguments. All data processing, scientific judgment, and final executive decisions remained exclusively with the author. No confidential data was exposed, and outputs were cross-checked against primary sources and the underlying code and data.