The Cost of Rivalry: Geopolitical Fragmentation and Trade during the Cold War*

Timothy Meyer† Nicolas Wesseler‡

Job Market Paper

November 13, 2026

Link to Latest Version Here

Abstract

How costly is geopolitical fragmentation? We use variation in coup success during the Cold War to estimate the causal effect of geopolitical alignment on international trade. An increase in political distance substantially reduces trade. Effects are more than three times larger for relationships involving superpowers and concentrated in more sophisticated products. Formal policy tools like foreign aid and sanctions only explain a small part of the trade reduction. A calibrated trade model suggests that economic incentives reinforced geopolitical blocks during the Cold War. In the present, a substantial share of countries would gain economically by aligning with the rival block.

Keywords: International trade, geopolitical fragmentation, Cold War, trade costs, geoeconomics
JEL Codes: F14, F50, F51, N70

*This document is an output from the research initiative 'Reducing Conflict and Improving Performance in the Economy' (ReCIPE), a programme funded by the Foreign, Commonwealth & Development Office (FCDO), contract reference ReCIPE_LOA_303006_Wesseler. The views expressed are not necessarily those of FCDO. Nicolas Wesseler thanks Keith Head, Ayumu Ken Kikkawa, and Nathan Nunn for their invaluable guidance and support. We thank Rodrigo Adão, John Sturm Becko, Emily Blanchard, Davin Chor, Markus Eberhardt, Jonathan Federle, Jeff Frieden, Andreas Fuchs, Samuel Kortum, Ohyun Kwon, Matteo Maggiori, Scott Orr, Pablo Ottonello, Swapnika Rachapalli, John Ries, and Yoto Yotov for helpful comments and discussions. We are also grateful for comments from participants at the Yale Cowles Summer Conference on International Trade, the GEP/CEPR Postgraduate Conference, the 8th International Conference on European Economics and Politics and The Third Annual Macroeconomic Policy Perspectives Conference at Stanford SIEPR. All errors are our own.

†Columbia University. Email: timothy.meyer@columbia.edu.
‡UBC Sauder School of Business. Email: nicolas.wesseler@sauder.ubc.ca.

1 Introduction

Geopolitical rivalry is reshaping the world economy. The United States has restricted economic ties with China, Europe has sharply reduced its engagement with Russia, and governments around the world face growing pressure to orient their economic relationships around great power competition. These developments have raised concern about the future of global trade in a world of geopolitical fragmentation (Fernández-Villaverde et al., 2024; Gopinath, 2024; Lagarde, 2023).

This paper provides evidence on the effect of political alignment on international trade using the Cold War as a historical laboratory. The Cold War divided the world into rival geopolitical spheres centered on the United States and the Soviet Union, with a large group of Non-Aligned countries navigating between the two superpowers. Countries' positions within this international order were far from fixed, however, as changes in government often brought shifts in political alignment. We exploit variation in coup success during the Cold War that is plausibly exogenous to trade trends to identify the effect of political realignment on international trade. We find that political distance generates substantial barriers to international trade and embed our estimates in a quantitative trade model to study the welfare consequences of geopolitical fragmentation.

Studying the effects of political alignment on international trade is difficult for two reasons. First, reverse causality may arise because countries that trade extensively develop closer political relationships. Second, common economic or geopolitical shocks may jointly shape both trade and political alignment. We therefore focus on political realignments following coups. Coups were frequent during the Cold War, averaging about five successful coups per year, and offer several distinct advantages for identifying the trade effects of political realignment. The success of an attempt plausibly depends on tactical factors and, at times, sheer luck that is unrelated to bilateral trade trajectories. Comparing successful with failed attempts therefore helps isolate the trade effects of the resulting political realignment. Coups are also difficult to predict systematically and produce clearly dated changes in government, unlike revolutions or civil wars that can unfold over extended periods. This reduces the scope for anticipatory trade responses and allow us to trace how trade evolves before and after a realignment. Their occurrence across regime types further allows us to study realignments beyond countries with competitive elections. Our estimates are robust to excluding coup attempts involving superpower intervention.

We employ two complementary identification strategies. Our first strategy exploits variation across a coup country's trading partners, comparing how trade evolves with partners that experience a larger coup-induced change in political distance to partners that experience a smaller one. This design identifies how the *magnitude* of political realignment affects trade and delivers an elasticity that can be mapped directly into models of international trade under decoupling scenarios (e.g., Bonadio et al., 2025). The key identifying assumption is that, conditional on exporter-year, importer-year, and dyad-by-coup fixed effects, trading partners that experience larger versus smaller changes in political distance would otherwise have followed similar bilateral trade trajectories.

Political realignment systematically reshapes bilateral trade. In our baseline analysis, we combine historical data on successful and failed coups from Bjørnskov and Rode (2020) with a measure

of political distance based on voting at the United Nations. We find that following a coup, trade systematically falls with partners from which the new government becomes more politically distant. The effect emerges gradually and persists over time. A one-standard deviation increase in political distance reduces bilateral trade by roughly 15 percent in the medium run, these effects are four times larger for trade with superpowers. We compare our estimates with those from a standard gravity regression that models trade costs as a function of political distance, the standard approach in prior work. Our causal estimates imply smaller trade costs associated with political distance for trade with Non-Aligned countries, but larger costs for trade with superpowers. Crucially, when we study failed coups as a placebo, we find no effect on international trade.

Our second strategy provides an additional test of the causal interpretation by conditioning directly on the occurrence of a coup attempt. In our sample, 51.7 percent of attempts succeed. Success depends on factors such as coordination among plotters, the loyalty of military units, timing, and luck. For example, during Nigeria's failed 1976 coup, a visit from a fellow officer seeking permission to name his newborn son after Olusegun Obasanjo, the country's deputy head of state, delayed Obasanjo's departure for work, causing the plotters to mistake another officer's car for his and allowing him to escape their ambush (Are, 2020). We use failed attempts directly as a control group, comparing the realignment of international trade after successful and failed coups. This strategy requires us to estimate how a country's geopolitical alignment would have evolved had its failed coup attempt succeeded. We use a large language model to classify the geopolitical orientation of the coup leadership for every attempt as pro-U.S., pro-Soviet, or unclear. Identification now requires that, among coup attempts, whether the attempt succeeds is unrelated to the dyad-specific trade trends that would otherwise have occurred.

The success-versus-failure design delivers the same qualitative pattern as our main specification. Among coup attempts that would move a country closer to one side of the geopolitical divide, successful coups produce substantially larger changes in trade with partners on the opposite side than failed coups. This pattern differs by partner type and depends on the direction of realignment. The effect is concentrated among bloc partners (NATO or Warsaw Pact members) when a coup is predicted to increase political distance.

Because the two strategies draw on different sources of variation, the size of the political shift in one case, whether the attempt succeeds in the other, their agreement is reassuring. We verify that our results are robust to identification and measurement concerns. One concern is that, during the Cold War, coups were orchestrated by superpowers (Berger et al., 2013b). We search systematically for superpower support for coups. Even under expansive definitions of superpower support, only 15% of coups received such outside support. Dropping these yields quantitatively similar estimates. We also validate that our results are robust to alternative measures of political alignment constructed from geopolitical events (Fan et al., 2025a) and to an alternative source of variation, natural deaths of country leaders (Jones and Olken, 2005), among other robustness checks. We further treat the post-Cold War period as an out-of-sample test of this relationship. Political realignments after 1990 show little comparable effect on trade, suggesting that the bipolar

structure of the Cold War amplified the trade effects of political realignment.

The estimated effect is economically meaningful. Translating our preferred coup-identified coefficient into an ad valorem tariff equivalent implies that a one-standard-deviation increase in political distance raises bilateral trade costs by approximately 3.1 percent on average. The effect is heterogeneous across dyads, with a tariff equivalent of 3.8 percent for relationships involving at least one bloc member and rising up to 11.5 percent for those involving a superpower. This magnitude is comparable to the trade effects of interstate war. Glick and Taylor (2010) find that trade between newly warring countries remains about 42 percent below its counterfactual level five years after the war ends, a similar order of magnitude to the decline in trade we estimate per one-standard-deviation increase in political distance between a country and a superpower. Political distance therefore acts as a particularly important barrier to trade precisely where geopolitical rivalry is most consequential.

Why does political alignment affect international trade? We provide evidence on the underlying mechanisms using more granular trade data that we construct for a subset of countries and on formal policy actions taken by governments.

We digitize product-level bilateral trade data from Soviet statistical yearbooks and construct similar data for the U.S. to study heterogeneity across products. The trade adjustment is concentrated in machinery and transport equipment, as well as primary goods on the Soviet side, and in chemicals and manufactures and, more weakly, machinery and transport equipment on the US side. Mineral fuels respond only weakly for either superpower. A finer product-level decomposition for the United States points to a common underlying pattern across several independent classifications of traded goods. The decline is smallest for standardized, commodity-like goods and largest for goods that depend on durable relationships or sophisticated production, consistent with trust and relationship-specificity being part of the channel.

We also study to what extent the fragmentation in international trade is driven by formal government policy actions. We consider foreign aid and sanctions, the two most important instruments of foreign economic policy during the Cold War (Baldwin, 1985). Using data on U.S. and Soviet aid, we document that countries that shift away from a superpower following a coup receive less foreign assistance and are more likely to face economic sanctions (though the sanction probability remains low). Combining these with estimates of the effect of aid and sanctions on trade, a back-of-the-envelope calculation suggests that changes in foreign aid could account for roughly 14 percent of the trade response, whereas formal trade sanctions account for less than 1 percent. We illustrate that trade can respond even without formal policy actions in a case study. We use the deterioration in U.S.–Turkish relations following the 1964 Johnson Letter in which the U.S. warned Turkey that it may not be defended by NATO against Soviet retaliation if it proceeded with a planned military intervention in Cyprus. Using a synthetic control approach, we show that U.S.–Turkey trade subsequently diverged substantially from its synthetic counterfactual (it remained flat in an era of rising global integration), even absent any sanctions, formal trade policy or changes in aid.

Finally, we quantify the welfare consequences of these trade costs. We embed the estimated

relationship between political distance and bilateral trade in a quantitative model of international trade and conduct general-equilibrium counterfactuals. We first study the costs of geopolitical fragmentation during the Cold War for world trade and welfare by considering a counterfactual in which we neutralize the trade barriers induced by geopolitical alignment. The resulting costs were substantial and unevenly distributed across countries. Non-Aligned countries bore the largest burden, reflecting the economic cost of political distance from the large Western market. For Eastern-bloc countries, the Iron Curtain accounts for most of the welfare loss, whereas for Western countries the bulk of the costs arise from political frictions within the Western bloc itself.

The structural framework also allows us to ask whether countries' geopolitical allegiances coincided with their economic interests. For each country, we compare the welfare consequences of full alignment with the United States and with the Soviet Union. During the Cold War, political and economic incentives were closely aligned for members of NATO and the Warsaw Pact, while the strongest tension between political positioning and material interest was concentrated in the Non-Aligned world. Using modern data, we then apply the same framework to the contemporary U.S.–China rivalry. The economic stakes of choosing between rival powers are smaller today than during the Cold War because economic ties across geopolitical blocs remain much deeper. Yet political and economic incentives are less consistently aligned: several countries currently aligned with one bloc would gain economically from closer alignment with the other.

Taken together, this paper makes four contributions. First, we provide causal evidence that geopolitical alignment affects international trade. By exploiting abrupt political realignments generated by coups, we move beyond the cross-sectional correlation between political proximity and economic exchange. Second, we show that the effects of geopolitical alignment are highly heterogeneous and shed light on the channels through which they operate. Third, we quantify the welfare consequences of geopolitical fragmentation, both as an aggregate feature of the Cold War and as a source of tension between countries' political allegiances and their economic interests. Fourth, we construct new datasets for studying Cold War political economy, digitized product-level Soviet trade flows, and an LLM-based classification of superpower backing across the full universe of Cold War coup attempts.

This paper contributes to three strands of the literature, work on the relationship between politics and international trade, on geoeconomics and on the economics of leadership transitions.

Politics and International Trade. A substantial literature studies the two-way relationship between international trade and political relations. Early work documents that countries trade more with political allies and less with adversaries (Dixon and Moon, 1993; Pollins, 1989). More recently, work has expanded to cover alternative forms of relationships, including visits by the Dalai Lama, military alliances and war (Fan et al., 2025a; Fuchs and Klann, 2013; Glick and Taylor, 2010; Hinz, 2023; Khalil et al., 2025; Neri-Laine, 2026; Niemann, 2025).[1] Using modern data, Kempf et al. (2025) and Ayyagari et al. (2025) extend these findings to the firm level. For the Cold War period,

[1] A related literature studies the effect of trade on international politics, including alignment, democracy and war (Feng et al., 2026; Kleinman et al., 2024; Martin et al., 2008; Tabellini and Magistretti, 2024).

Berger et al. (2013b) show that CIA interventions increase trade with the U.S. Closest to our setting, Gopinath et al. (2025) study the trade patterns between Cold War blocks and compare them to the present, while Campos et al. (2024) study the trade and welfare consequences of the Iron Curtain for Post-War Europe. Relative to these papers, we provide a new identification strategy to study the effect of geopolitics on trade and highlight implications for non-aligned countries.

Geoeconomics and Economic Statecraft. Our paper also contributes to a growing literature on how states use economic relationships to pursue geopolitical objectives. These questions have a long history in international relations (Baldwin, 1985; Hirschman, 1945) and have recently been taken up again by economists. Much of recent work in geoeconomics studies how countries can strategically use economic relationships for geopolitical goals (Becko and O'Connor, 2024; Broner et al., 2025; Clayton et al., 2025, 2026a,b; Kooi, 2024; Levchenko et al., 2026; Liu et al., 2026; Liu and Yang, 2025; Madsen and Prat, 2026; Meyer and Wesseler, 2026; Thoenig, 2024). A related strand uses structural trade models to evaluate fragmentation scenarios directly, specifying geopolitical blocs and imposing counterfactual increases in trade barriers between them, from fragmentation in strategic sectors (Bolhuis et al., 2023; Cerdeiro et al., 2021) to broader increases in trade costs between blocs (Baur et al., 2025; Bonadio et al., 2025; Góes and Bekkers, 2022; Javorcik et al., 2024; Plummer, 2023; Seong et al., 2025). Because the magnitude of geopolitical trade barriers is uncertain, these exercises depend on the assumed size of the counterfactual barriers. We supply a disciplined estimate of that parameter. Our estimated tariff equivalents are considerably smaller than the trade barriers typically assumed in these simulation exercises, which range from 7.5 to 60 percent or amount to a complete prohibition of trade in strategic sectors (see Table A.1 in the Appendix for a list of assumptions). We further quantify the trade and welfare consequences of this competition with important implications particularly for third countries drawn between hegemons.

Leadership Transitions. Finally, our analysis relates to work on the economic consequences of leadership transitions. Beginning with the seminal work of Jones and Olken (2005), this literature has studied the consequences of both violent (coups, revolutions) and peaceful (democratizations, elections) leadership transitions (Acemoglu et al., 2019; Bazzi and Blattman, 2014; Bennett et al., 2021; Derpanopoulos et al., 2016; Dorsch and Maarek, 2018; Marx et al., 2025; Meyersson, 2016; Papaioannou and Siourounis, 2008). Closest in spirit to our work, Jones and Olken (2009) use variation in assassin success to identify the effect of leader transitions on domestic institutions. Relative to this work, which focuses on the domestic consequences of leadership turnovers, we study their international effects.

The rest of the paper proceeds as follows. Section 2 describes the historical setting and introduces the data. Section 3 describes our empirical framework and its implementation. Section 4 presents the main results. Section 5 examines mechanisms. Section 6 quantifies welfare consequences and extends the analysis to the contemporary U.S.–China rivalry. Section 7 concludes.

2 Setting and Data

2.1 Setting: Blocs, Realignment, and Coups

The Cold War provides substantial variation in countries' political alignment, both across countries and over time. This section first describes the structure of the two major blocs and then shows how abrupt political realignments, those generated by coups, provide the variation used to identify the effect of geopolitical alignment on trade.

2.1.1 The Structure of Cold War Blocs

The post-World War II international economic order quickly fractured along ideological lines. By the late 1940s, the global economy was divided into three broad groups, the U.S.-led Western bloc, the Soviet-led Eastern bloc, and a heterogeneous set of non-aligned nations that attempted to maintain economic relationships with both sides (Westad, 2005). We use this three-way split throughout the paper for descriptive analysis and to test for heterogeneous treatment effects, fixing each country's bloc membership over the full sample period to avoid endogeneity concerns. When we turn to our identification strategy in Section 3, we rely on a continuous measure of political distance, which captures gradual realignment and within-bloc heterogeneity.

The **Western bloc**, centered on the United States and its NATO allies, was organized around market principles and multilateral institutions. The 1947 General Agreement on Tariffs and Trade (GATT) established a framework for reducing trade barriers among market economies, while the Marshall Plan provided massive economic assistance to rebuild Western Europe. We define the Western bloc as NATO members plus Japan and Australia, the founding members (Belgium, Canada, Denmark, France, Iceland, Italy, Luxembourg, the Netherlands, Norway, Portugal, the United Kingdom, and the United States) together with the countries that acceded during the Cold War (Greece, Turkey, West Germany, Spain), plus Japan and Australia, which, though never formal NATO members, were closely integrated into the Western economic sphere through their post-war military and economic alignment with the United States. This grouping coincides with the membership of CoCom, the Coordinating Committee for Multilateral Export Controls, the multilateral regime established in 1949 that restricted Western exports of strategic goods and technology to the Soviet bloc. By the 1950s, the Western bloc accounted for approximately 70% of global GDP and was characterized by relatively open trade within the bloc, despite maintaining various barriers to trade with communist countries. Trade was economically substantial. The median country exported roughly 15–20% of its income, with intra-Western flows accounting for around 65% of all bilateral trade (see Appendix Figure A.13).

The **Eastern bloc**, led by the Soviet Union, operated under different economic principles. Central planning replaced market mechanisms, state enterprises monopolized foreign trade, and bilateral clearing agreements governed most international transactions. The Council for Mutual Economic Assistance (COMECON), established in 1949, coordinated economic relations among communist countries, conducting trade through bilateral protocols that specified quantities and

prices, often divorced from world market prices. We define the Eastern bloc as the formal members of the Warsaw Pact, the military alliance formed in 1955 that reinforced this economic integration, the Soviet Union, Poland, East Germany, Hungary, Romania, Bulgaria, Czechoslovakia, and Albania. Other countries maintained communist governments or close Soviet ties without joining the Pact. We exclude them from the Eastern bloc to keep the classification clean, treating their alignment instead through the continuous political distance measure introduced in Section 2.2.

Between these two poles lay the **Non-Aligned** world, whose organizing movement emerged following the 1955 Bandung Conference. Countries such as India, Egypt, Yugoslavia, Ghana, and Indonesia sought to maintain economic relationships with both superpowers while preserving political independence. For these nations, playing superpowers off against each other could yield substantial economic benefits in the form of aid, favorable trade terms, and development assistance from both sides (Meyer and Wesseler, 2026). However, their room for maneuver was often constrained by geographic proximity to one bloc or the other, historical colonial relationships, and superpower competition in their regions. We classify every country outside the Western and Eastern blocs as Non-Aligned, a deliberately heterogeneous residual category that includes both the movement's founding members and countries that maintained varying degrees of independence from both sides, some leaning consistently toward one superpower, others shifting repeatedly.

2.1.2 Coups as Sources of Political Realignment

While the basic structure of Cold War blocs remained relatively stable, individual countries' positions within this structure shifted substantially over time. These shifts provide the variation we exploit for identification.

Some realignments were dramatic and sudden. Egypt shifted from being a Soviet ally in the 1950s and 1960s to a U.S. partner following President Sadat's 1972 expulsion of Soviet advisors and the 1979 Camp David Accords. Chile, one of the most prominent examples, experienced two sharp reversals. The 1970 election of socialist Salvador Allende moved the country toward the Soviet sphere, while the 1973 military coup that brought Pinochet to power reversed this alignment toward the West. While Chile is a particularly well-known example, there were about five successful coups per year during the Cold War, most of them fomented by domestic political considerations, not foreign powers as we show below. Other shifts were more gradual. Many newly independent African and Asian nations experimented with different alignments during their first decades of independence. Some countries, such as India, maintained consistent non-aligned positions while varying the balance of their relationships with each superpower. Portugal and Spain moved from authoritarian neutrality toward fuller integration with Western Europe following their democratizations in the 1970s.

These shifts were driven by various factors, domestic political changes (coups, elections, revolutions), regional conflicts, and superpower interventions, among others. Importantly for our identification strategy, many alignment shifts occurred abruptly following political shocks, creating discontinuous changes in a country's political distance from various trading partners.

Among these abrupt events, coups d'état, the attempt by elites or the military to seize political power, are particularly useful for our identification strategy. First, they are difficult to predict beforehand. An extensive literature has proposed dozens of country characteristics as coup determinants, but testing these systematically across millions of model specifications finds that only a handful, such as slow economic growth, a history of prior coups and other political violence, are robust predictors, while most proposed determinants are not (Gassebner et al., 2016; Lehoucq, 2021). Coups are accordingly closer to unanticipated shocks than to the visible culmination of a predictable political trend, which lowers the likelihood that trading partners adjust in anticipation before the coup itself occurs. Second, they are discrete and dateable. Unlike the gradual electoral or diplomatic shifts described above, a coup typically installs a materially different government within days rather than a slow drift that is difficult to date precisely. Third, coups do not require the country experiencing them to be a democracy. Designs that instead rely on close elections to isolate quasi-random shifts in leadership (Liu and Yang, 2025) are only available where competitive, contested elections take place. During the Cold War (and today), however, a large share of the world was governed by single-party states, military juntas, or monarchies with no meaningful electoral contest. Coups occurred across regime types, allowing us to include much of the Third World in our sample.

A concern with treating coups as a source of identifying variation is that some of them may not be genuine, spontaneous expressions of domestic political change but events fabricated or directly orchestrated by a foreign power specifically to install a government aligned with its own interests. If a superpower engineers a coup precisely because it wants the resulting political realignment and the closer economic relationship that comes with it, then the coup is not an exogenous shock. Instead, it would be an instrument of the superpower's own objectives, which could confound our estimates of the trade response to realignment. We address this concern directly in our analysis, both through verifying the robustness of our results to the exclusion of potentially manipulated coups, as well as through our two identification strategies.

2.2 Data

Three main ingredients go into our empirical design, bilateral trade flows, a measure of political alignment, and the coup events. We describe each in turn below. For two of these three, bilateral trade and the classification of coup events by superpower involvement, existing sources are insufficient for our purposes, so we construct new data: digitized product-level Soviet trade statistics and an LLM-based classification of superpower backing across the full universe of Cold War coup attempts. All other data from our analysis are described as they are used.

2.2.1 Bilateral Trade

We obtain bilateral trade flow data from two main sources. The main dataset is the CEPII TradHist database, which provides bilateral trade flows from 1827 to 2014 (Fouquin and Hugot, 2016). For the period covered by our analysis, TradHist draws primarily on the IMF's Direction of Trade Statistics

(DOTS), supplemented by additional data from government publications and primary sources.[2] The database also provides standard gravity variables such as bilateral distance, contiguity, common language, and colonial relationships as well as data on national GDP. We retain zero trade flows as reported by TradHist in all our specifications.

However, TradHist has important limitations for studying Cold War trade. Soviet bloc countries did not consistently report trade statistics to international organizations, and their use of non-market pricing mechanisms complicates comparisons with market economies. More importantly, DOTS relies on partner-reported trade flows, resulting in particularly sparse coverage of Soviet trade with developing countries. Missing trade data may also be systematically related to political alignment. If socialist countries systematically underreported trade with Western partners, for instance, this could bias our estimates.

We address these concerns by complementing TradHist with product-level trade data digitized directly from Soviet annual statistical yearbooks (*Vneshnyaya Torgovlya SSSR*), which provide Soviet-reported bilateral trade flows. This approach follows Campos et al. (2024), who similarly supplement TradHist with primary Soviet sources to reconstruct comprehensive trade flows between the USSR and both Eastern and Western trading partners. Beyond the aggregates digitized by Campos et al. (2024), we digitize the full product-level trade data.[3] This allows us to construct bilateral sector-level trade flows needed for our analysis on heterogeneity.

2.2.2 Political Alignment

Our primary measure of political alignment comes from United Nations General Assembly (UNGA) voting records. We use the ideal point estimates of Bailey et al. (2017), who assign every country a single scalar value each year, estimated jointly from its full voting record across all UNGA resolutions. The underlying model accounts for missing votes and for variation in how salient different resolutions are, so that resolutions with widespread agreement do not carry the same weight as more divisive ones, and covers all UNGA participants from 1946 to 2023. Countries that consistently vote together are inferred to have similar ideal points, while countries that frequently vote in opposition are assigned distant ideal points. Ideal points are scaled so that higher values indicate positions closer to the Western bloc and lower values indicate positions closer to the Soviet bloc. Using UNGA idealpoints follows a long literature in international relations and, more recently, economics (Becko et al., 2025; Camboni and Porcellacchia, 2024; Gopinath et al., 2025; Kleinman et al., 2024).

We use these ideal points to construct measures of bilateral political distance. For any country

[2]To examine the effect of political distance on trade after the Cold War, we extend some of our analysis through 2022 using data from the CEPII Gravity database (Conte et al., 2022).

[3]Standard sector-level sources such as the NBER-UN World Trade Flows by Feenstra et al. (2005) are built from partner-reported mirror statistics. Because the Soviet Union did not report to the United Nations trade statistics system, these data almost entirely omit Soviet trade with Warsaw Pact countries, which accounted for 44 to 71 percent of total Soviet trade depending on the year.

pair (i, j) at time t, we define:

$$\text{IPD}_{ijt} = |\text{IdealPoint}_{it} - \text{IdealPoint}_{jt}|$$

where IPD_{ijt} is the ideal point distance between countries i and j. This measure captures how far apart two countries are in the foreign policy space. Larger values indicate greater political distance, suggesting less alignment on international issues.

The UN voting-based measure has several advantages for our purposes. First, it varies continuously both across countries and over time, capturing gradual realignments as well as abrupt shifts. Second, it is based on revealed preferences through actual voting decisions rather than subjective assessments. Third, it aggregates information across many issues, reducing the influence of any single vote or temporary disagreement. Fourth, participation in UNGA voting was nearly universal, providing coverage for both aligned and non-aligned countries.

Figure 1 shows the evolution of average ideal points for different groups during our sample period. The United States maintains a consistently high ideal point (around 2.5), while the Soviet Union has a consistently low ideal point (around -2). NATO allies cluster near the U.S. position, Warsaw Pact members cluster near the Soviet position, and non-aligned countries occupy intermediate positions with substantial variation. Importantly, individual countries' positions shift over time in response to political changes, providing the within-country variation we exploit for identification.

In Meyer and Wesseler (2026) we validate this measure of alignment both narratively and quantitatively. We provide case studies of countries that are well-known to have shifted alignment during the Cold War, such as Chile, Egypt, Ethiopia, or Somalia and show that the UNGA alignment measure recovers the timeline of these alignment shifts as documented by historians (David, 1991). We further verify robustness to using an alternative measure of political alignment constructed from geopolitical event data in Section 4.3.

2.2.3 Coups

Our primary source is the coup dataset of Bjørnskov and Rode (2020) (BR), an update to the Democracy and Dictatorship dataset of Cheibub et al. (2010), which provides comprehensive coverage of coup attempts between 1950 and 2018.

According to the BR data, coups and coup attempts have to fulfil the following definition to count as such. First, the objective must be to overthrow the executive branch. Second, actors have to be previously linked to the state apparatus in some way. Third, a coup or coup attempt cannot last longer than a week at most.[4]

Coups were a common feature of the Cold War period. BR record 296 coup attempts, of

[4]For robustness, we cross-validate this data with several alternative sources. These are the Center for Systemic Peace's dataset on regime changes, the coup data from Berger et al. (2013b) covering CIA-supported interventions, and the Colpus dataset from Chin and Kirkpatrick (2023). While definitions and coverage vary slightly across sources, the timing and basic facts of major coups are consistent.

Figure 1: Ideal Points of Countries and Blocs, 1948–2022

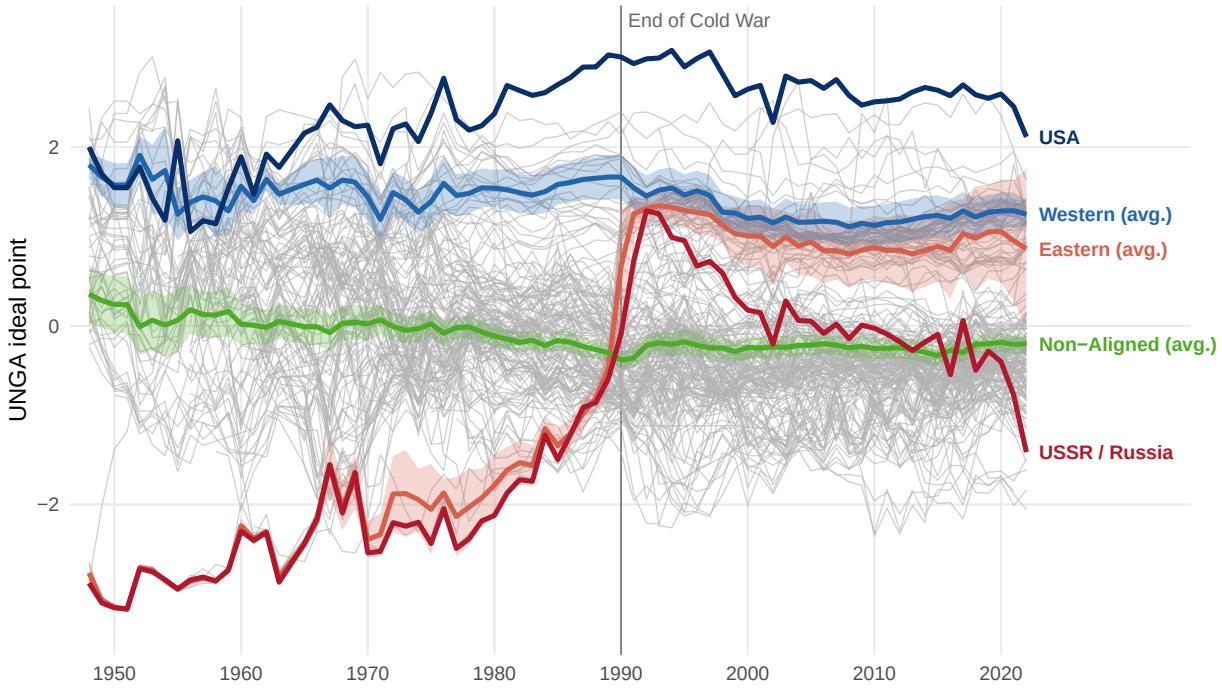

Note: This figure displays the evolution of UN General Assembly ideal points from Bailey et al. (2017) over time. Thick colored lines show bloc averages, Western (blue), Eastern (red-orange), and Non-Aligned (green), with shaded regions indicating 95% confidence intervals around each group mean. The United States and the USSR/Russia, both included in their respective bloc averages, are also highlighted individually. Thin gray lines trace individual country trajectories, illustrating within-group heterogeneity and countries shifting alignment over time. The vertical line marks the end of the Cold War (1990). Higher values indicate positions closer to the Western bloc, while lower values indicate positions closer to the Soviet bloc.

which 153, or 51.7%, succeeded in replacing the sitting government. Figure 2 maps the geographic distribution of these attempts. They are concentrated in Latin America, Sub-Saharan Africa, and the Middle East, with Bolivia (12), Iraq (10), and Sudan (10) being the most affected countries.

Figure 3 illustrates the resulting variation in political alignment directly. For each successful coup we plot the change in the country's UNGA idealpoint from one year before the coup to one year after, an arrow pointing up when the coup moves a country closer to the United States and down when it moves the country closer to the Soviet Union. Shifts occur in both directions in roughly equal numbers and vary substantially in size, from a decline of 2.7 to an increase of 2.1 on the idealpoint scale, with a mean absolute shift of 0.42. These shifts are not simply ordinary noise in the idealpoint series. Comparing every country-year in our sample, not only coup years, a successful coup is associated with an idealpoint swing about 50 percent larger than the typical year-to-year movement, and this gap survives country and year fixed effects (Appendix A.5).

An important question for our identification strategy is whether success itself, rather than the mere attempt, matters. For each year before and after an attempt, we measure the absolute change in a country's idealpoint relative to the year before the attempt and regress it on separate indicators for a successful and a failed attempt, with country and year fixed effects. Figure 4 plots the resulting coefficients at each event time. That is, the magnitude of the idealpoint swing

Figure 2: Coup attempts, 1955–1985

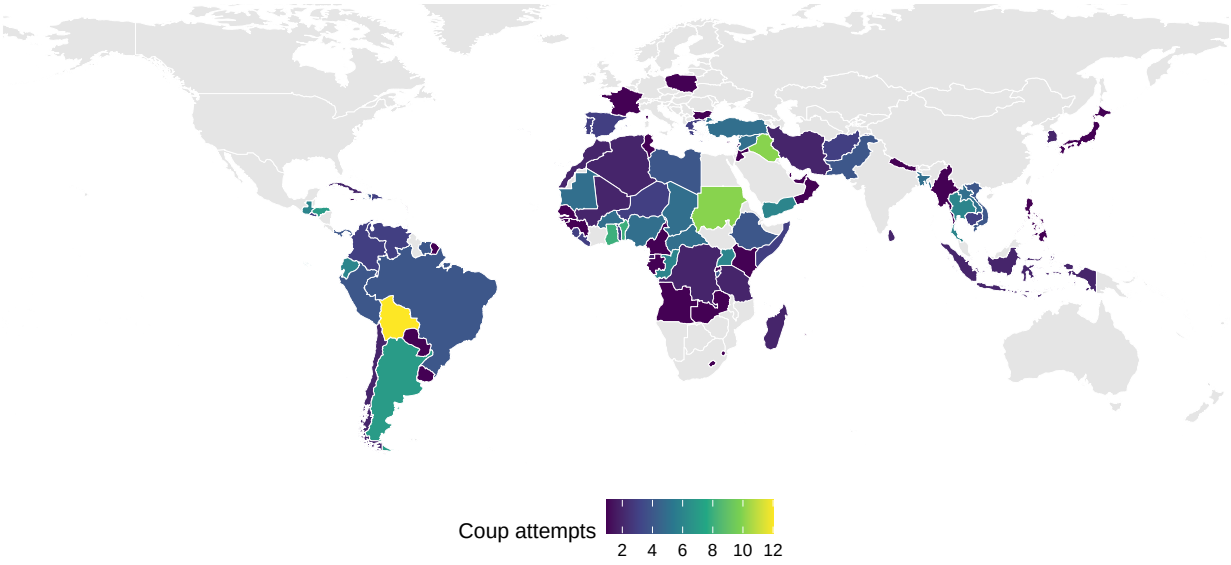

Note: This figure plots the total number of coup attempts recorded for each country between 1955 and 1985, the sample window used throughout the paper. Counts include both successful and failed attempts. Countries with no recorded attempt are shaded grey. Data from Bjørnskov and Rode (2020).

associated with success and with failure. The swing around successful coups is substantially larger than the swing around failed attempts at every post-attempt horizon. Appendix A.5 reports a table testing this difference directly, confirming that successful and failed coups are associated with significantly different idealpoint swings.

We also use these events to construct a new classification of superpower involvement. For every coup attempt in the BR data, we ask a large language model, acting as a Cold War historian, whether there is credible documented evidence, as opposed to an undocumented historical allegation, that the United States or the Soviet Union directly assisted, funded, or orchestrated that specific attempt. We then manually verify these cases. We use this classification in Section 4.3 to test whether our baseline results are driven by coups a superpower itself engineered, and Appendix A.3 reports the full classification together with its geographic distribution.

To provide the reader with a concrete sense of the variation in our data, Appendix A.4 presents four historical case studies, two successful coup episodes, and two failed, illustrating how a change in political leadership can produce sharp shifts in geopolitical alignment, and how such shifts need not reflect gradual changes in underlying political conditions but can instead result from the actions of a small number of individuals.

Figure 3: Change in Idealpoint Around Successful Coups, 1955–1985

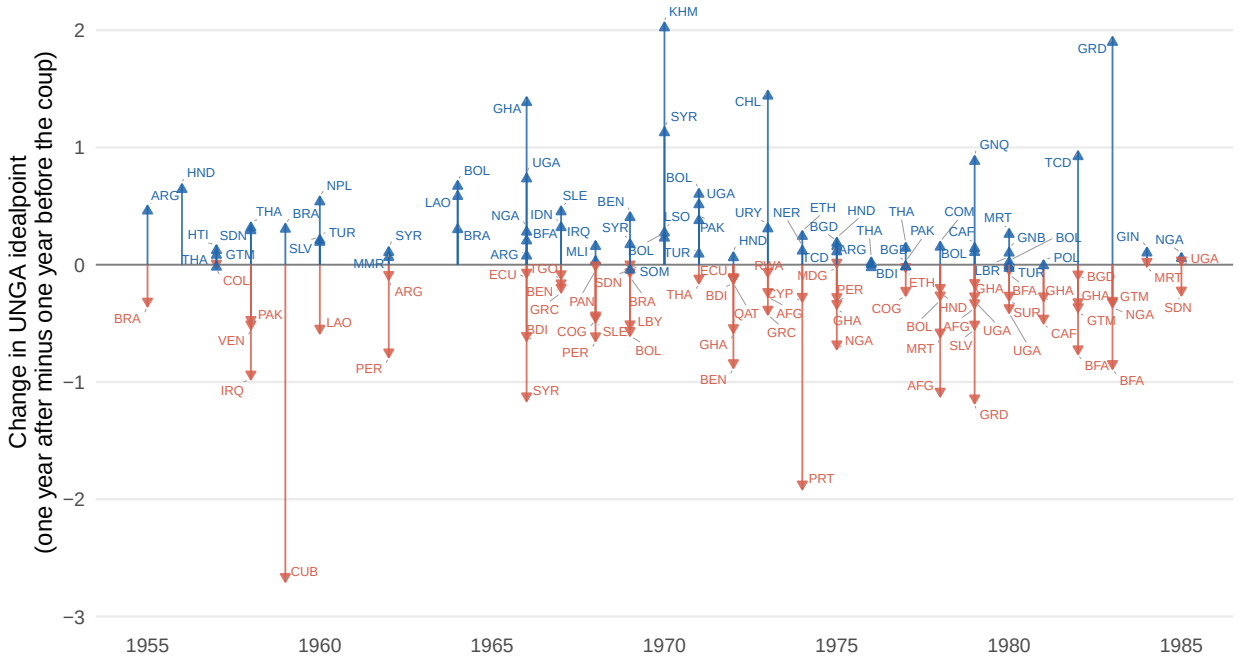

Note: This figure plots, for each successful coup between 1955 and 1985, the change in the country's UNGA idealpoint from one year before the coup to one year after. Countries with more than one successful coup in the same year are represented by the later coup only. Arrows point up (blue) when the idealpoint moves toward the United States and down (red) when it moves toward the Soviet Union. Idealpoints are from Bailey et al. (2017).

3 Empirical Framework and Implementation

3.1 Framework

This section derives the empirical specifications we take to the data. We start from a standard structural gravity model of trade, show how it delivers our baseline empirical specification and where political distance enters as a component of bilateral trade costs. The same structure then yields our event study specification, obtained by decomposing political distance around the timing of a coup.

3.1.1 Structural Gravity

We consider a standard Armington framework in the spirit of Anderson and van Wincoop (2003). The world consists of $i = 1, \ldots, N$ countries, and consumers in country n have CES preferences over the N national varieties, where q_{ni} denotes consumption in country n of the variety produced by country i, with elasticity of substitution $\sigma > 1$,

$$U_n = \left(\sum_{i=1}^{N} q_{ni}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}.$$

Country i produces the national variety at constant unit cost w_i, and delivering one unit to

Figure 4: Idealpoint Swing by Event Time

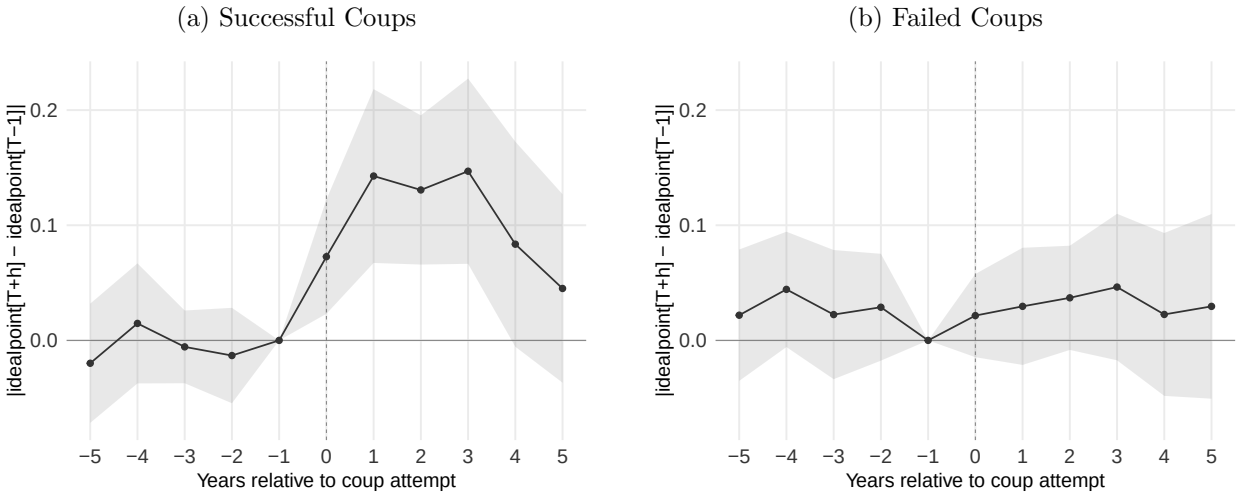

Note: Each panel plots the coefficient on an indicator for a coup attempt in a regression of $|\text{idealpoint}_{t_0-1+k} - \text{idealpoint}_{t_0-1}|$ on that indicator, with country and year fixed effects, estimated separately at each event time k, jointly for the successful-coup and failed-coup indicators. Shaded region is the 95% confidence interval from standard errors clustered by country and year.

country n requires shipping $\tau_{ni} \geq 1$ units, so the delivered price is $p_{ni} = w_i \tau_{ni}$. Combining CES expenditure shares with market clearing yields the structural gravity equation for bilateral trade flows X_{nit} from i to n at time t,

$$X_{nit} = \frac{Y_{it} X_{nt}}{\Omega_{it} \Phi_{nt}} \phi_{nit}, \qquad \phi_{nit} \equiv \tau_{nit}^{1-\sigma}, \tag{1}$$

where Y_{it} is country i's total output, X_{nt} is country n's total expenditure, and Ω_{it} and Φ_{nt} are outward and inward multilateral resistance terms, respectively. They summarize how bilateral trade depends not only on trade costs between i and n, but also on each country's trade costs with all alternative markets and suppliers.

Note that this structural gravity representation extends beyond the Armington model to a broad class of quantitative trade models, including Krugman (1980) and Eaton and Kortum (2002).

Taking logarithms of equation (1) gives

$$\log X_{nit} = \log Y_{it} - \log \Omega_{it} + \log X_{nt} - \log \Phi_{nt} + \log \phi_{nit}. \tag{2}$$

The multilateral resistance terms are unobserved and depend on the entire matrix of bilateral trade costs, but following Head and Mayer (2014) we can absorb them without solving for them explicitly. Grouping all terms that vary only by exporter and year into an exporter-year fixed effect γ_{it}, and all terms that vary only by importer and year into an importer-year fixed effect δ_{nt}, leaves the bilateral friction term $\log \phi_{nit}$ as the only object left to model. We decompose bilateral frictions into a time-invariant component, capturing geography, contiguity, common language, and colonial history, among others, and a time-varying component τ_{nit}^{TV}. Writing $\varepsilon \equiv \sigma - 1$ for the elasticity of

trade with respect to trade costs,

$$\log \phi_{nit} = -\varepsilon \log \tau_{nit} = -\varepsilon \, \bar{\tau}_{ni} - \varepsilon \, \tau_{nit}^{TV}. \tag{3}$$

The time-invariant piece is absorbed by a dyad fixed effect α_{ni}, yielding the empirical gravity specification

$$\log X_{nit} = \gamma_{it} + \delta_{nt} + \alpha_{ni} - \varepsilon \, \tau_{nit}^{TV} + \eta_{nit}. \tag{4}$$

Because trade flows are frequently zero, particularly for the Cold War dyads at the center of our analysis, and because OLS on the log-linearized gravity equation is inconsistent under heteroskedasticity, we estimate every specification derived from equation (4) by Poisson Pseudo-Maximum Likelihood (PPML), following Head and Mayer (2014).

3.1.2 Political Distance as a Trade Cost

Our object of interest is the time-varying trade cost τ_{nit}^{TV} isolated in equation (4). We let political distance account for this component

$$\tau_{nit}^{TV} = -\frac{\beta}{\varepsilon} \, \mathrm{IPD}_{nit}, \tag{5}$$

where IPD_{nit} is the ideal point distance between countries n and i defined in Section 2.2.2, and β is the parameter governing how political distance affects trade costs, to be estimated from the data below. Substituting equation (5) into equation (4) gives our baseline specification

$$\log X_{nit} = \gamma_{it} + \delta_{nt} + \alpha_{ni} + \beta \, \mathrm{IPD}_{ni,t} + \eta_{nit}, \tag{6}$$

The coefficient β is therefore a structural object, with $-\beta/\varepsilon$ giving the semi-elasticity of trade costs with respect to political distance. This mapping lets us translate any estimate of β into an ad valorem tariff equivalent, the uniform tariff that would raise trade costs by the same amount as a given level of political distance:

$$\mathrm{TE}_{nit} = 100 \times \left[\exp\left(-\frac{\beta \, \mathrm{IPD}_{nit}}{\varepsilon} \right) - 1 \right]. \tag{7}$$

We report tariff equivalents throughout the paper using this formula, together with the trade elasticity $\varepsilon = 5.03$ estimated by Head and Mayer (2014).

3.1.3 Coups as Exogenous Shocks

Equation (6) identifies β from all variation in political distance, including variation driven by reverse causality or omitted variables. To isolate a source of variation that is plausibly exogenous, we exploit coups, which generate sharp, discontinuous shifts in a country's political alignment. For a country pair (n, i) affected by a coup in country n at time t_0, we approximate the path

of political distance around the coup as a two-regime step process, equal to a pre-coup baseline $\text{IPD}_{ni,\text{pre}} \equiv \text{IPD}_{ni,t_0-1}$, the level of political distance in the year immediately preceding the coup, plus a coup-induced shift

$$\text{IPD}_{nit} = \text{IPD}_{ni,\text{pre}} + \Delta\text{IPD}_{ni} \cdot \mathbf{1}[t \geq t_0] + u_{nit}, \qquad \Delta\text{IPD}_{ni} \equiv \text{IPD}_{ni,t_0+1} - \text{IPD}_{ni,t_0-1}. \tag{8}$$

The step approximation captures the coup-induced jump in alignment, while u_{nit} absorbs any remaining, non-coup-driven movement in political distance – gradual drift, other geopolitical events, or measurement noise – that occurs before or after the coup. Substituting the decomposition into equation (6), and noting that $\text{IPD}_{ni,\text{pre}}$ is time-invariant and therefore absorbed by a dyad fixed effect $\tilde{\alpha}_{ni} \equiv \alpha_{ni} + \beta\,\text{IPD}_{ni,\text{pre}}$, gives a difference-in-differences specification

$$\log X_{nit} = \gamma_{it} + \delta_{nt} + \tilde{\alpha}_{ni} + \beta\,\Delta\text{IPD}_{ni} \cdot \mathbf{1}[t \geq t_0] + \beta\,u_{nit} + \eta_{nit}, \tag{9}$$

where $\beta\,u_{nit}$ is absorbed into the composite error term. This specification imposes that the effect of political realignment on trade is instantaneous and constant after the coup. To allow for gradual adjustment, and to test for pre-existing trends, we replace the single post-coup indicator with a full set of event-time dummies, arriving at our event study specification

$$\log X_{nit} = \sum_{k \neq -1} \beta_k \cdot \mathbf{1}[t - t_0 = k] \cdot \Delta\text{IPD}_{ni} + \gamma_{it} + \delta_{nt} + \tilde{\alpha}_{ni} + \tilde{\eta}_{nit}, \tag{10}$$

normalizing $\beta_{-1} = 0$ so that all coefficients are measured relative to the year before the coup. Note that each event-time dummy is interacted with the same coup-induced shift ΔIPD_{ni} rather than with the realized IPD_{nit} at each date t. The specification therefore traces out the dynamic trade response to a fixed initial dose of political realignment, rather than the contemporaneous elasticity of trade with respect to political distance. This design absorbs any approximation error u_{nit} into η_{nit}, and relies on u_{nit} being unrelated to event time around the coup. The coefficients β_k trace out how trade evolves in event time as a function of the political shift, and are directly comparable to the baseline coefficient β in equation (6). Under the gravity model, we would expect $\beta_k \approx 0$ for $k < -1$ if trade was not already trending differentially ahead of the coup and β_k converging to β if the baseline estimate reflects the same causal relationship rather than reverse causality or omitted variables. As with all gravity specifications in the paper, we estimate equation (10) by PPML. We therefore write the estimating equation in multiplicative form as

$$X_{nit} = \exp\left(\sum_{k \neq -1} \beta_k \cdot \mathbf{1}[t - t_0 = k] \cdot \Delta\text{IPD}_{ni} + \gamma_{it} + \delta_{nt} + \tilde{\alpha}_{ni}\right) \tilde{\eta}_{nit}, \tag{11}$$

where $\tilde{\eta}_{nit}$ is a multiplicative error term satisfying $E[\tilde{\eta}_{nit} \mid \mathbf{Z}] = 1$, with $\mathbf{Z}$ denoting the full set of regressors and fixed effects.

3.1.4 Two Complementary Identification Strategies

We take two distinct empirical strategies to the data, both built on the coup-induced shifts in political distance derived above, but differing in their control group, their identifying assumption, and the object they recover.

Strategy 1: Dose-Response Across Trading Partners. Our primary specification is equation (11), estimated on the trading partners of countries that experienced a successful coup. Identification comes from how much political distance a given partner gains or loses as a result of the coup. Dyads where neither side ever had a coup at all are excluded. Because the treatment is the continuous coup-induced shift ΔIPD_{ni}, β_k is a dose-response coefficient, the percentage change in trade per one-unit increase in political distance, k years after the coup. This is the object that enters the tariff equivalents in equation (7), and the welfare counterfactuals of Section 6, which require knowing how trade costs scale with political distance.

Strategy 1 requires that, absent any coup-induced shift in political alignment, a coup country's trade with its various partners would still have evolved at the same underlying rate after the coup, regardless of how much political distance each partner actually gained or lost. The exporter-year and importer-year fixed effects already absorb each country's overall trade trend, and the dyad fixed effect absorbs each dyad's starting level. What remains is the requirement that, among a coup country's partners, those who ended up more politically distant were not already on a different trade trajectory with that country for unrelated reasons. The assumption fails if whatever caused a partner to drift away or move closer is also what would have produced an unusual change in trade with that partner regardless of the coup.

Failed coup attempts provide a complementary placebo test for Strategy 1. A failed attempt shares many of the political and economic conditions surrounding a successful one, but produces no change in government. We estimate the identical dose-response specification around failed attempts. If the pattern we find around successful coups reflected political instability, anticipation of regime change, or some other shock common to all coup attempts, we would expect to see the same trade pattern around failed ones. Its absence supports the interpretation that the successful-coup results are driven by the realized realignment itself, not by the mere occurrence of an attempt. This placebo, however, cannot rule out shocks that arise only when a coup succeeds and are correlated with the resulting change in political distance, which motivates our second identification strategy.

Strategy 2: Success versus Failure. Strategy 1 never observes an empirical counterfactual for the coup countries it studies. A successful coup country's trade absent the realignment is inferred from fixed effects and from its other partners, not observed directly. Our second strategy supplies that counterfactual by holding fixed the underlying propensity for political upheaval and instead exploiting variation in its outcome, whether the coup succeeded or failed, as the source of identifying variation.

Let $D_n \in \{0, 1\}$ indicate whether country n experienced a coup attempt during our sample period, and let $S_n \in \{0, 1\}$, defined only for attempters, indicate whether that attempt succeeded. Our second specification restricts the sample to attempters, $D_n = 1$, and estimates the same

specification as equation (11), with the continuous coup-induced shift ΔIPD_{ni} replaced by the binary success indicator S_n. We denote the resulting coefficients β_k^S to distinguish them from Strategy 1's β_k. This compares countries where the attempt succeeded, $S_n = 1$, to those where it failed, $S_n = 0$. Because S_n is binary, β_k^S is an average treatment effect, the difference in trade trajectories between successful and failed coup countries, rather than a coefficient on the magnitude of political distance.

In practice, this comparison is carried out separately within direction-matched pools of partners, grouping attempts by whether the plotters were predicted to move the country closer to the United States or to the Soviet Union, as described in Section 3.2.2. A success is therefore only ever compared to a failure predicted to realign in the same direction, since pooling successes and failures across opposite predicted directions would mix dyads where the dose would have pushed trade in opposite directions.

Conditioning on $D_n = 1$ holds fixed any dyad-level trend that is driven by the mere fact that a coup was being planned or attempted. For instance, firms or governments might reposition trade with particular partners in anticipation of political change. This anticipatory response is absorbed as long as it is common to successful and failed attempts. The same logic speaks to the concern, raised in Section 2, that some coup attempts are not spontaneous domestic events but are themselves orchestrated by a foreign power seeking to engineer a particular realignment. Since that motive is already present when the attempt is made, it is common to successful and failed attempts alike and is absorbed by conditioning on $D_n = 1$, as long as foreign backing does not also strongly determine which attempts succeed.

This is an advantage of Strategy 2 over Strategy 1. Both the treated and control dyads come from countries where an attempt actually occurred, so any disruption common to attempting a coup, successful or not, cancels directly in the comparison between them. What remains is the requirement that, among attempters, which side wins is unrelated to dyad-specific trade trends. This is the sense in which the near coin-flip success rate, determined by tactical factors and luck, supports this assumption. These factors are unlikely to be correlated with which of a country's trading partners were already on a different trajectory. Section 3.2.1 tests this restriction directly by comparing pre-coup characteristics and trends across successful and failed attempts.

In short, Strategy 1 asks whether the magnitude and direction of political realignment is unrelated to differential counterfactual trade trends, while Strategy 2 asks whether success versus failure is unrelated to those trends among countries where an attempt actually occurred. The two strategies trade off assumption strength against the object they deliver. We treat Strategy 1 as our primary specification for magnitudes throughout the paper, together with the failed coup placebo described above, and use Strategy 2 as a validity check on the dyad-level parallel trends assumption described above.

3.2 Implementation

3.2.1 Is Coup Success Exogenous Conditional on Attempt?

The identifying assumption for Strategy 2 requires that, conditional on a coup attempt taking place, whether it succeeds is unrelated to the dyad-specific trade trends its trading partners would otherwise have followed. Two distinct concerns could violate this. First, success might not be the tactically determined, close to random outcome we have described. If, for instance, weaker states are both more likely to see a coup succeed and more dependent on particular trading partners, success would be correlated with exactly the dyad-specific trends the fixed effects cannot absorb. Second, a successful coup may bring more than a political realignment. If success is bundled with other country-level shocks, an economic crisis, or renewed conflict, and those shocks affect specific trading partners differently, β_k^S could pick up their effect rather than, or in addition to, the effect of political distance.

However, if successful and failed coup countries look similar on observable, pre-determined characteristics, that supports treating success as tactically determined rather than systematically predictable, which addresses the first concern. If country-level characteristics do not move differently after successful coups, that is evidence against the second concern, that successful coups bundle other consequential shocks alongside the realignment.

We compare successful and failed coup countries on population, trade openness, GDP per capita, the Polity2 democracy score, UNGA ideal point, inflation, civil war incidence, and superpower military intervention, using two-sided t-tests. Table 1 reports the pre-coup *level* of each variable, averaged over the two years before the coup (Panel A), which speaks to the first concern, and its *change* from before to after the coup (Panel B), which speaks to the second.

In Panel A, seven of the eight pre-coup levels are statistically indistinguishable between successful and failed coup countries. The exception is trade openness, which is higher in countries where the coup failed (0.282 versus 0.183, $p = 0.006$). This level difference does not threaten identification by itself, because our design relies on dyad-by-coup-event fixed effects that absorb any time-invariant difference in trade openness across successful and failed coup countries. What would be concerning is a differential change in trade openness around the coup itself, and Panel B shows no such pattern ($p = 0.633$). In Panel B, only the Polity2 democracy score changes significantly more after successful coups than after failed ones ($p = 0.001$), the expected mechanical consequence. None of the other variables change differentially, evidence against successful coups bringing a broader, confounding policy shift alongside the political realignment we study.

We further report the composition of coup types in Panel C. Military coups succeed at a markedly higher rate than civilian ones (81.5 percent versus 68.1 percent, $p = 0.005$), a pattern unsurprising given that militaries control the means of coercion. We view this imbalance as illustrative rather than a threat to identification, since coup type captures who seizes power rather than a channel with an independent direct effect on bilateral trade.

Table 1: Balance Test: Observables by Coup Outcome

Variable	Mean (Succ.)	Mean (Fail.)	Diff.	p-value
Panel A: Level (pre-coup)				
Log GDP per capita	1.714	1.868	-0.154	0.068
Log population	6.463	6.413	0.050	0.740
Trade openness [(X+M)/GDP]	0.183	0.282	-0.099	0.006*
Polity2 score	-2.363	-2.963	0.600	0.300
UNGA ideal point	0.022	-0.034	0.057	0.503
Log inflation rate	1.968	2.105	-0.137	0.358
Civil war (share)	0.118	0.148	-0.030	0.392
Superpower intervention (share)	0.253	0.204	0.048	0.265
Panel B: Change (post − pre)				
Log GDP per capita	0.033	0.029	0.004	0.635
Log population	0.073	0.079	-0.005	0.150
Trade openness [(X+M)/GDP]	0.025	0.020	0.005	0.633
Polity2 score	-1.790	0.260	-2.050	0.001*
UNGA ideal point	-0.050	-0.009	-0.041	0.485
Log inflation rate	0.289	0.025	0.263	0.054
Civil war (share)	0.047	-0.006	0.053	0.077
Superpower intervention (share)	0.054	0.003	0.051	0.156
Panel C: Coup type composition				
Military coup (share)	0.815	0.681	0.134	0.005*
Civil/civilian coup (share)	0.163	0.297	-0.134	0.004*
Royal coup (share)	0.022	0.022	0.001	1.000

Note: Sample includes all coup attempts, successful and failed, from Bjørnskov and Rode (2020). Panel A averages each variable over the two years before the coup ($t-2$ to $t-1$). Panel B is the change from that pre-coup average to the two-year post-coup average ($t+1$ to $t+2$). Panel C reports the share of attempts of each coup type. $^*p < 0.05$ from a two-sided t-test of equality of means.

3.2.2 Classifying Coup Direction

Strategy 2 requires signing the direction of a coup for every attempter, not only for successful ones, because we compare successful and failed coups that share the same intended direction. A successful coup reveals its direction through the realized post-coup shift in alignment, a failed coup produces no such shift.

We use a large language model (LLM), Anthropic's Claude Sonnet 5.5, to classify the Cold War ideological orientation of the coup leadership as pro-US, pro-Soviet, or unclear, based only on the identity of the plotters, their names, military rank, and the type of coup (military, civilian, or royal). Of 152 coup events, the model classifies 79 as pro-US and 35 as pro-Soviet, leaving 38 unclear, which we drop from the direction-dependent analysis. In many cases the political orientation is clear from party manifestos or leader identities. Appendix A.2 reproduces the exact prompt.[5]

[5]The value of this classification is greatest for failed attempts, where there is no governance record to reveal the plotters' orientation. For instance, the January 1986 uprising in South Yemen by Vice President Ali Ahmed Nasser Antar against President Abdul Fattah Ismail, both leaders of the ruling Marxist-Leninist Yemeni Socialist Party, is

This classification gives us the counterfactual orientation of the country's leadership,

$$\hat{o}_n \in \{\text{pro-US}, \text{pro-Soviet}\}, \tag{12}$$

but not yet the direction of the shift for a specific trading partner, since whether a given partner gains or loses distance also depends on where that partner already stood. We combine $\hat{o}_n$ with the pre-coup ideal points of the coup country n and of each trading partner i to sign the predicted change,

$$\widehat{\Delta\text{IPD}}_{ni} < 0 \iff \begin{cases} \hat{o}_n = \text{pro-US} & \text{and } \text{IdealPoint}_{i,\text{pre}} > \text{IdealPoint}_{n,\text{pre}}, \text{ or} \\ \hat{o}_n = \text{pro-Soviet} & \text{and } \text{IdealPoint}_{i,\text{pre}} < \text{IdealPoint}_{n,\text{pre}}, \end{cases}$$

using each country's average ideal point in the year before the coup. Intuitively, a partner that already sits on the side the leadership is expected to move toward sees distance shrink, while a partner on the other side sees it grow, regardless of whether the coup actually succeeds.

Applying this to every attempt, successful and failed alike, splits the Strategy 2 sample into two pools, dyads predicted to gain political distance and dyads predicted to lose it, within which we compare successful to failed attempts. Trade should fall more for successful than failed attempts in the distance-increasing pool and rise more in the distance-decreasing pool, giving a directional counterpart to Strategy 2's β_k^S for each pool.

3.2.3 Handling Multiple Coups per Country

Many countries in our sample experience more than one coup between 1955 and 1985, Bolivia, Iraq, and Sudan each see ten or more attempts. Estimating a single pooled two-way fixed effects regression across all of these events at once runs into a well-understood problem in the difference-in-differences literature, because with staggered treatment timing, observations from a country's post-coup periods can implicitly serve as controls for a different, later-timed coup elsewhere in the panel, and these comparisons can enter with negative weight, biasing the estimated event-time coefficients (Goodman-Bacon, 2021). A related, more direct problem also arises when the same country experiences two coups close together in time, the post-period we would use to trace out the dynamic effect of the first coup then overlaps with the run-up to, or aftermath of, the second.

We address both problems with a stacked event study design, following Cengiz et al. (2019). Rather than estimating one pooled regression, we build a separate sub-panel, or stack, around each individual coup event, and then estimate the event-time coefficients on the pooled collection of stacks. Each stack contains the treated country's dyads over a fixed event window from $k = -4$ to $k = 10$ years relative to that specific coup, with event time recentered to account for the coup's timing within the year. That is, a coup occurring in the second half of the calendar year has event time shifted back by one year, so that $k = 0$ consistently denotes the first full year of realignment

classified pro-Soviet even though it targets the Marxist-Leninists party. This is because the model correctly identifies the conflict as a factional fight within a Soviet-aligned party rather than a contest with any Western-leaning side.

rather than a partial one.

Before building the stacks, we drop any coup event contaminated by a later coup in the same country within the following ten years. A failed attempt is contaminated by any later coup, successful or failed, since even a failed later challenge introduces confounding instability into the earlier event's post-period trade window. A successful coup, by contrast, is contaminated only by a later success, a later failed challenge that never overturns the new government does not erase the earlier realignment from the sample. This filter leaves 251 of the original 296 coup events, 133 successful and 118 failed, which enter the stacked panel as treatment events. We verify robustness to alternative stacks and constructions below. We show that a plain dyad fixed effect in place of stacking, or using all successful coup regardless of contamination in place of the filtered sample, leaves sign and significance of the baseline result unchanged, though the coefficient magnitude varies somewhat across specifications (Appendix A.6).[6]

Within each stack, we use a dyad-by-event fixed effect rather than a single dyad fixed effect shared across the whole pooled sample. This matters because the same trading relationship can appear in more than one stack, if a country's partner separately experiences its own coup at a different time. Without event-specific dyad fixed effects that partner's already-treated observations in one stack could act as an invalid control in another, reintroducing the same contamination problem the stacking design is meant to solve. The exporter-year and importer-year fixed effects γ_{it} and δ_{nt}, by contrast, are defined on true calendar years rather than event time, since different stacks can span overlapping real years and only calendar-year fixed effects absorb the country-year shocks those terms are meant to capture. We cluster standard errors at the level of the true underlying dyad, accounting for the same trading relationship appearing in multiple stacks. We apply this procedure identically to Strategy 1 and Strategy 2.

4 Results

4.1 Strategy 1: Coups versus No Coups

We estimate equation (11) using the stacked design described in Section 3.2.3. Of the 251 coup events that survive the contamination filter, 133 are successful and enter the treatment stacks. As a placebo, we estimate the identical specification on the 118 failed coups that survive the same filter. Because a failed coup produces no realized shift in political alignment, ΔIPD_{ni} for these dyads reflects only the ordinary, non-coup-driven variation in ideal points, so a null result here is what Strategy 1's identifying assumption predicts.

For successful coups, the coefficients are close to zero and individually insignificant in the three years before the coup ($k = -4, -3, -2$: -0.055, 0.001, 0.046). The coefficient at $k = 0$ is negative but only marginally significant ($p = 0.087$), from $k = 1$ onward the coefficients are negative and

[6] In both cases the non-stacked specification yields a smaller coefficient than our baseline, in line with Nagengast and Yotov (2025), who similarly find that forgoing a staggered, stacking-based estimator biases standard gravity estimates of trade agreement effects downward.

Figure 5: Effect of Political Realignment on Trade Around Coups

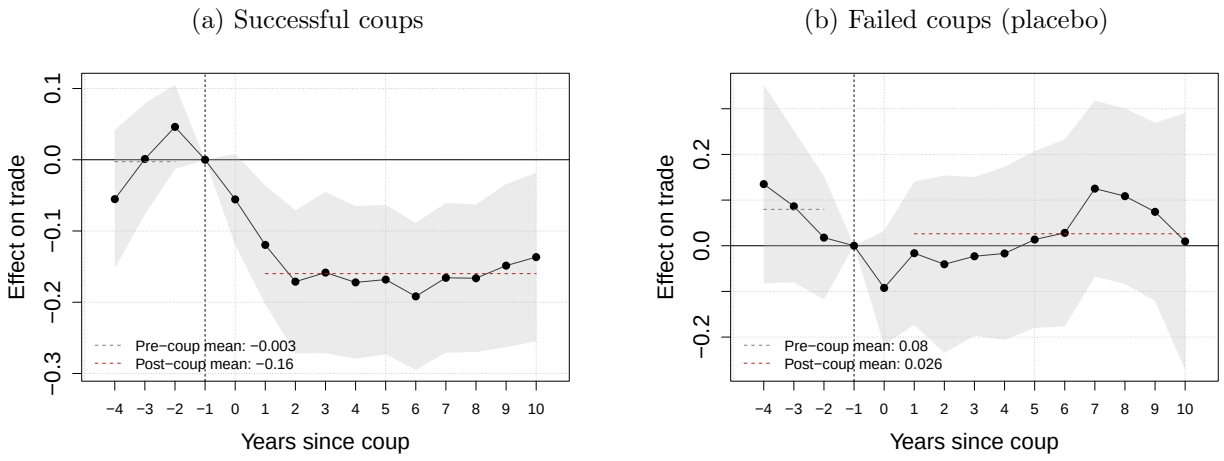

Note: Each panel plots $\hat{\beta}_k$ from equation (11), the effect on trade of a one-unit increase in political distance k years after the coup, relative to $k = -1$. Shaded regions are 95% confidence intervals from standard errors clustered by dyad. Dashed lines show the average pre-coup ($k < -1$) and post-coup ($k \geq 0$) coefficient.

individually significant, stabilizing between -0.14 and -0.19 log points from $k = 4$ through $k = 10$, and the joint test over $k = 0$ through $k = 5$ rejects the null of no effect ($p = 0.025$). A one-unit increase in political distance is therefore associated with roughly 13 to 17 percent less trade several years after a successful coup, in the direction implied by equation (5).

The failed-coup placebo shows no such pattern. Every coefficient, before and after the attempt, is individually insignificant, and the joint post-period test fails to reject the null of no effect ($p = 0.800$). This is consistent with Strategy 1's identifying assumption. Dyads that happen to gain or lose political distance around a coup that does not occur show no differential trend, which is what we would expect if the post-coup decline in trade among successful coups reflects the realignment itself rather than a feature of the estimation design.

Heterogeneity by Partner Type. We repeat the successful-coup event study separately for three partner groups, non-aligned partners that belong to neither bloc, partners that belong to a formal bloc (NATO or the Warsaw Pact), and partners that are themselves a superpower (the United States or the Soviet Union).[7]

Figure 6 shows that while geopolitical distance tends to impede trade between all types of countries, these impediments are larger for superpowers. For non-aligned partners the post-coup coefficients are small and the joint post-period test does not reject the null of no effect ($p = 0.316$). For bloc partners the joint post-period test rejects the null decisively ($p = 0.001$), with coefficients from $k = 2$ onward ranging from -0.19 to -0.32 log points. For superpower partners the joint post-period test also rejects the null decisively ($p < 0.001$), and the coefficients are substantially larger, jumping from -0.48 at $k = 1$ to as much as -0.87 at $k = 2$ and remaining between -0.69 and -0.85 for the rest of the post period. Political realignment therefore moves trade primarily through a country's relationship with blocs and the superpowers themselves.

[7]These groups can overlap in principle, though in practice a bloc member is rarely also coded as a superpower.

Figure 6: Effect of Political Realignment on Trade, by Partner Type

(a) Non-aligned partners

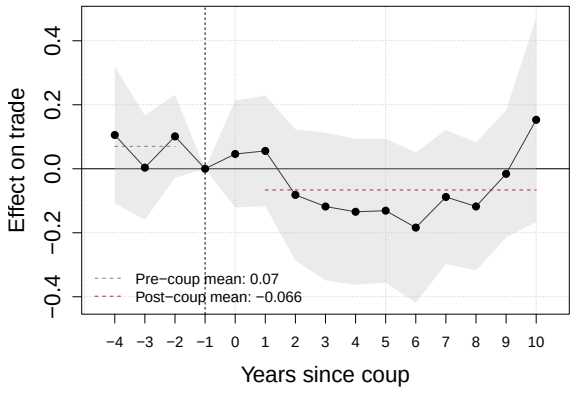

(b) Bloc partners

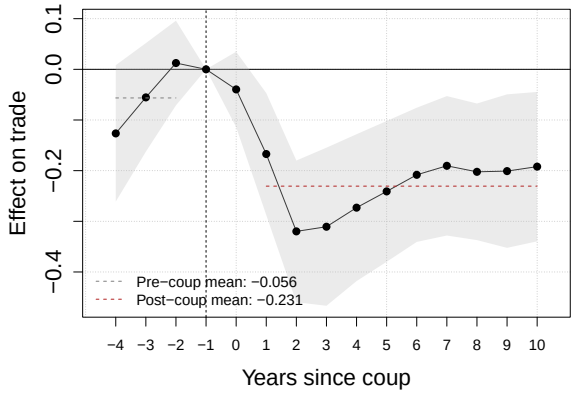

(c) Superpower partners

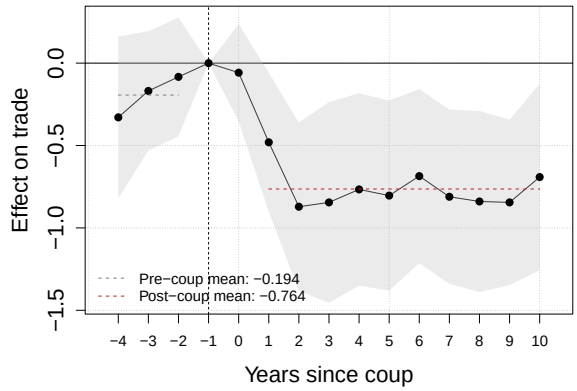

Note: Each panel plots $\hat{\beta}_k$ from equation (11), estimated separately on the subsample of dyads in each partner group, for successful coups only. Shaded regions are 95% confidence intervals from standard errors clustered by dyad.

Comparison to Naive Gravity Estimates. We compare the coup-identified estimates to a purely correlational gravity regression of trade on the one-year lagged change in political distance, estimated on the same trade panel and with the same exporter-year and importer-year fixed effects, but with a plain dyad fixed effect and no coup-based identification. To make the two specifications comparable, we collapse the coup-identified event study into a single post-period coefficient, an interaction of a post-coup indicator with the realized alignment shock ΔIPD_{ni}, rather than the full set of event-time dummies in equation (11). Table 2 reports both specifications, pooled and separately by partner type.

Pooled across all partners, the coup-identified estimate is smaller in magnitude than the naive gravity regression, -0.154 against -0.193. This is consistent with the identification concern that motivates our design. The gravity estimation cannot rule out reverse causality, whereby dyads that already trade heavily subsequently grow more politically aligned, or omitted shocks that move both trade and political distance together. Either channel would tend to inflate the naive gravity

Table 2: Naive Gravity versus Coup-Identified Difference-in-Differences

Dependent Variable:	Bilateral Trade Flows			
	All	Neutral	Bloc	Superpower
Model:	(1)	(2)	(3)	(4)
Panel A: Gravity Variables				
IPD_{t-1}	-0.193***	-0.240***	-0.142***	-0.447***
	(0.022)	(0.042)	(0.027)	(0.126)
Fixed-effects				
Dyad FE	Yes	Yes	Yes	Yes
Exporter-Year FE	Yes	Yes	Yes	Yes
Importer-Year FE	Yes	Yes	Yes	Yes
Fit statistics				
Observations	425,783	251,840	173,137	15,706
Pseudo R^2	0.98727	0.95040	0.98874	0.99558
Panel B: DiD Variables				
Post $\times$ ΔIPD	-0.154***	-0.122*	-0.186***	-0.546***
	(0.039)	(0.069)	(0.053)	(0.132)
Fixed-effects				
Dyad $\times$ Coup Event FE	Yes	Yes	Yes	Yes
Exporter-Year FE	Yes	Yes	Yes	Yes
Importer-Year FE	Yes	Yes	Yes	Yes
Fit statistics				
Observations	147,504	96,154	49,307	3,008
Pseudo R^2	0.97555	0.95665	0.98012	0.99277

Note: Panel A regresses trade on the one-year lagged change in political distance without any coup-based identification, using a dyad fixed effect together with exporter-year and importer-year fixed effects. Panel B collapses the coup-identified event study of equation (11) into a single post-period coefficient, using a dyad-by-event fixed effect together with exporter-year and importer-year fixed effects. Clustered (dyad) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

coefficient relative to the true causal effect. The partner-type comparison sharpens this picture. For neutral partners the naive gravity coefficient is sizable and significant (-0.240), while the coup-identified estimate is smaller in magnitude and only marginally significant (-0.122, $p < 0.10$), the same pattern as in the pooled comparison and consistent with reverse causality or omitted variables inflating the correlational estimate for these dyads. For bloc and superpower partners the opposite pattern holds, the coup-identified estimate is larger in magnitude than the naive correlation for both (bloc, -0.186 against -0.142, superpower, -0.546 against -0.447), so for these partners the naive correlation, if anything, understates the true effect of realignment on trade. The difference between superpowers and bloc partners is statistically significant, which we show in a version of the specification that interacts the post-coup dose with a partner type indicator in Appendix A.7.

Thus far we have comapred the coup-identified estimate to a gravity regression estimated on

the full trade panel because that is the comparison available to a researcher without a coup-based identification strategy. To check whether the gap between the two instead reflects the differing samples, Appendix A.8 replicates the gravity regression restricted to the same dyad-years used in the coup-identified sample. There we find that restricting the sample pulls the pooled, neutral, and bloc coefficients toward the coup-identified estimates, but it pushes the superpower coefficient the other way, to an insignificant -0.257, while the coup-identified estimate remains large and significant. Sample composition cannot account for the superpower result.

To gauge the economic magnitude of the effect, the superpower coefficient implies that a one-unit increase in political distance reduces trade with a superpower by about 42 percent, a magnitude similar to Glick and Taylor (2010)'s estimates of the trade effects of interstate war. Trade between newly warring countries falls by about 83 percent during the conflict itself and remains roughly 42 percent below its counterfactual level five years after the war ends. Political estrangement from a superpower, short of open conflict, is associated with a comparable order of magnitude of trade disruption.

Using equation (7) and the trade elasticity $\varepsilon = 5.03$ estimated by Head and Mayer (2014), we convert the coup-identified coefficients in Panel B of Table 2 into the ad valorem tariff that would raise trade costs by an equivalent amount. Pooled across all partners, a one-unit increase in political distance corresponds to a tariff equivalent of 3.1 percent. The partner-type estimates diverge sharply, 2.5 percent for neutral partners, a modest effect consistent with the underlying coefficient's only marginal significance, 3.8 percent for bloc partners, and 11.5 percent for superpower partners. A coup that pushes a superpower relationship toward political estrangement therefore raises trade costs by roughly three times as much as an equivalent shift with a bloc partner, and more than four and a half times as much as with a neutral one.

4.2 Strategy 2: Success versus Failure

We now turn to Strategy 2, which isolates the effect of political realignment by comparing coup attempts that succeeded to attempts that failed, rather than relying on cross-partner variation in the size of the realignment as Strategy 1 does. We estimate Strategy 2 by replacing the continuous coup-induced shift ΔIPD_{ni} in equation (11) with the binary success indicator S_n, comparing successful to failed coup attempts within the direction-matched pools described in Section 3.2.2. Across all coups, 114 have a usable LLM direction classification, 76 successful and 38 failed.

A coup predicted to move a country's alignment toward the United States decreases political distance with partners that are themselves pro-US and increases it with pro-Soviet partners, and conversely for a coup predicted to move toward the Soviet Union. We therefore group dyads into two pools on this basis, those predicted to see political distance decrease and those predicted to see it increase, each pooling both directions of the coup itself. Table 3 reports the resulting collapsed pre/post difference-in-differences for both pools and all partner-type columns, with the post period spanning the five years following the attempt.

In the pool where predicted political distance increases (Panel A), a successful coup lowers

Table 3: Success versus Failure: Collapsed Difference-in-Differences, by Partner Type

Dependent Variable:	Bilateral Trade Flows			
	All	Neutral	Bloc	Superpower
Model:	(1)	(2)	(3)	(4)
Panel A: Increasing distance				
Variables				
Post × Success	-0.025	-0.035	-0.284**	-0.635*
	(0.054)	(0.069)	(0.135)	(0.379)
Fixed-effects				
Dyad × Coup Event FE	Yes	Yes	Yes	Yes
Exporter-Year FE	Yes	Yes	Yes	Yes
Importer-Year FE	Yes	Yes	Yes	Yes
Fit statistics				
Observations	120,700	96,967	23,140	704
Pseudo R^2	0.97637	0.97613	0.97908	0.99778
Panel B: Decreasing distance				
Variables				
Post × Success	-0.086	0.201***	-0.284*	0.231
	(0.097)	(0.076)	(0.165)	(0.203)
Fixed-effects				
Dyad × Coup Event FE	Yes	Yes	Yes	Yes
Exporter-Year FE	Yes	Yes	Yes	Yes
Importer-Year FE	Yes	Yes	Yes	Yes
Fit statistics				
Observations	95,929	64,282	31,167	636
Pseudo R^2	0.98795	0.97295	0.99034	0.99914

Note: Panel A restricts to dyads predicted to gain political distance from the attempt, Panel B to dyads predicted to lose it, following the direction classification of Section 3.2.2. Post × Success is the coefficient on the interaction of a post-attempt indicator with a successful-coup dummy, estimated by PPML with dyad-by-event, exporter-year, and importer-year fixed effects. Standard errors, in parentheses, are clustered by dyad. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

trade relative to a failed one specifically among bloc partners (-0.284, $p < 0.05$), an effect more than twice as large when the sample is narrowed to dyads where a superpower is the direct partner (-0.635, $p < 0.10$). The pooled and neutral-partner coefficients are not distinguishable from zero.

In the pool where predicted political distance decreases (Panel B), a successful coup raises trade relative to a failed one among neutral partners (0.201, $p < 0.01$), while the pooled and superpower coefficients are not distinguishable from zero. The bloc coefficient here runs the opposite way, negative and marginally significant. Appendix A.9 traces this to a single event, a failed 1982 coup attempt in Spain whose trade with its Western European bloc partners surged in the following years for a reason unrelated to the coup, its accession to the European Community. Dropping this event resolves the anomaly and strengthens every increase-distance coefficient as well.

In short, a coup predicted to increase political distance from a partner lowers trade relative to

a failed attempt, an effect concentrated among bloc partners and strongest when the partner is itself a superpower. A coup predicted to decrease political distance shows a more mixed pattern. One possible reason is that trade relationships are easier to unwind than to build, existing supply chains and contracts can collapse quickly once alignment deteriorates, while new relationships after a favorable realignment may take longer to form.

4.3 Robustness

This section shows that the trade response to political realignment is robust to alternative measurement approaches and different sources of variation.

4.3.1 Excluding Superpower Interference

A natural concern with our baseline results is that some successful coups were not spontaneous domestic events but were themselves directly orchestrated by the United States or the Soviet Union. If a superpower installs a friendly government, the subsequent jump in bilateral trade with its patron may reflect direct sponsorship or reconstruction support rather than the arm's-length political-realignment channel we aim to isolate. We therefore test whether our baseline coefficient survives the exclusion of coups with credible evidence of direct foreign backing.

Existing sources, such as the data of Berger et al. (2013a) on superpower-backed interventions, are not built for this purpose. It records country-year episodes of ongoing superpower support for a regime, not backing of any specific coup event, so a country can be flagged as under superpower influence for decades after an episode's onset (indeed, many countries are coded as seeing a CIA intervention throughout the Cold War, as the CIA is continuously backing the government). Merging it directly onto our coup events therefore risks misattributing foreign backing.

To assess foreign backing at the level of the individual coup, we instead ask a large language model (Claude Sonnet 4.6), acting as a Cold War historian, whether there is credible documented evidence, declassified intelligence records, participant testimony, or a strong historian consensus, that the United States or the Soviet Union directly assisted, funded, or orchestrated each specific coup attempt in our sample. The prompt explicitly instructs the model to distinguish this from an undocumented allegation, since a claim that a superpower was behind a given coup is one of the most common tropes in twentieth-century political history and is frequently asserted with no real evidentiary basis. The model classifies each coup as "foreign backed", "no evidence", or "unclear", together with a confidence tier, "high", "medium", or "low", reflecting the strength of the underlying evidence. We review these classifications and apply it to the full universe of 546 coup events in the underlying data. Appendix A.3 reproduces the exact classification prompt, reports the complete classification (Table A.2), lists each individual US- and Soviet-backed event (Table A.3), and maps their geographic distribution by power (Figure A.1).

Of the 133 successful coups, the model flags 6 as foreign-backed at high confidence and 15 at medium confidence, 21 in total. This is consistent with recent historical work that revises the narrative that Cold War political events were driven by superpowers only and highlights the

Table 4: Coup-Identified DiD, Excluding Superpower-Backed Coups

Dependent Variable:	Bilateral Trade Flows		
	All	Excl. High	Excl. High + Medium
Model:	(1)	(2)	(3)
Variables			
Post × ΔIPD	-0.154***	-0.146***	-0.145***
	(0.039)	(0.042)	(0.044)
Fixed-effects			
Dyad × Coup Event FE	Yes	Yes	Yes
Exporter-Year FE	Yes	Yes	Yes
Importer-Year FE	Yes	Yes	Yes
Fit statistics			
Observations	147,504	142,683	120,396
Pseudo R^2	0.97555	0.97607	0.97209

Note: Column (1) reports our baseline coefficient using all 133 successful, uncontaminated coup events. Columns (2) and (3) progressively exclude coups the LLM classifies foreign backed at high confidence, and at high or medium confidence. No coup in our sample is classified foreign backed at low confidence only, so there is no additional column for that threshold. All columns use the collapsed pre/post specification of Panel B of Table 2, with dyad-by-event, exporter-year, and importer-year fixed effects. Standard errors, in parentheses, are clustered by dyad. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

local autonomy of the actors involved in coups around the world (Harmer, 2011; Westad, 2005). Table 4 shows that excluding these events leaves our estimate essentially unchanged: −0.154 at baseline, −0.146 excluding high-confidence foreign-backed coups, and −0.145 excluding both high- and medium-confidence cases, all significant at the 1 percent level. The trade response to political realignment we document is therefore not concentrated among coups with superpower involvement, but is quantitatively identical for coups without credible evidence of foreign interference.

4.3.2 Alternative Measures of Political Alignment

As an alternative measure of bilateral political alignment, we use the geopolitical score recently constructed by Fan et al. (2025a) and Fan et al. (2025b), built from the Global Geopolitical Events database. This database applies a large language model to code significant bilateral political events between countries from 1950 to 2024, drawing on government sources, international organizations, and news archives. Each event receives a Goldstein score between negative ten and positive ten, reflecting its degree of conflict or cooperation, and these scores are aggregated into a dynamic yearly measure for each country pair. Unlike our baseline alignment measure, which is based on similarity in United Nations General Assembly voting, this measure is constructed directly from coded political actions such as trade policy changes, diplomatic engagements, and security incidents, providing an independent check on our results. For comparability, we invert the sign of the geopolitical score so that higher values correspond to greater political distance. The geopolitical score is available for 75 percent of the dyad-years in our trade panel.

We reestimate equation (11) with this measure in place of ΔIPD_{ni}, holding fixed the same contamination filter, event window, and stacking procedure described in Section 3.2.3, and repeat the partner-type split of Section 4.1. Figure A.4 in the Appendix reports the results.

The alternative measure confirms our baseline finding. Bloc partners show no pre-trend ($p = 0.616$) and a strongly significant post-period effect ($p < 0.001$), with coefficients ranging from -0.46 to -0.77 log points from $k = 1$ through $k = 10$. Superpower partners show an equally strong post-period effect ($p < 0.001$), with much larger coefficients, from -1.75 at $k = 1$ to as much as -2.91 at $k = 5$. For neutral partners, where our baseline event study detected no effect, the coefficients turn negative from $k = 2$ onward and reach as much as -0.63 log points at $k = 5$, and the joint post-period test strongly rejects the null of no effect ($p < 0.001$), with most individual coefficients significant on their own from $k = 3$ onward. If anything, this alternative measure strengthens our main finding: political realignment is associated with less trade across all three partner groups, most strongly for bloc and superpower partners.

We further verify that our results are not sensitive to the choice of ideal-point measure used to construct the political alignment dose. In addition to our baseline measure, which uses ideal points estimated from all United Nations General Assembly votes, we reconstruct the dose using ideal points restricted to salient foreign-policy votes only and using an alternative legacy estimation of ideal points.[8] Re-estimating our stacked event study with each alternative measure, for both successful coups and the failed-coup placebo, yields coefficients that are quantitatively and statistically very similar to our baseline results. This holds regardless of which measure is used, indicating that our findings do not depend on a particular choice of ideal-point estimation. Figure A.5 in the Appendix reports the results.

4.3.3 Leader Deaths as an Alternative Source of Variation

We also study an alternative, plausibly exogenous source of variation in political alignment, the death of a national leader from natural or accidental causes. Jones and Olken (2005) assemble a cross-country dataset of leader transitions and identify episodes in which a sitting leader died in office of a natural cause or an accident. We apply the same collapsed pre/post difference-in-differences design used in our main analysis to these events, treating each death as the event date and re-estimate the same specification.

Table 5 reports the results. Column (1) reproduces our coup-based baseline coefficient of -0.154 for reference. Column (2) uses leader deaths alone and yields a coefficient roughly half as large, -0.074, and not statistically significant. Columns (3) and (4) show that this gap reflects sample composition rather than a weaker shock. Successful coups in our underlying catalog are overwhelmingly non-democratic, with 94.8 percent of them occurring in non-democracies (Polity2 score

[8] Both alternatives come from Bailey et al. (2017). The first restricts the estimation to votes the State Department's annual *Voting Practices* reports flag as important, which excludes routine and procedural resolutions that otherwise make up most of the UNGA record and could dilute the signal in the all-votes measure. The second uses an older, differently estimated vintage of UNGA ideal points that predates that update, included to confirm our results are not an artifact of a particular estimation procedure.

five or below). The Jones and Olken leader-death sample is more balanced, with 40 percent of deaths occurring in democracies. When we split the leader-death sample along this line, deaths in autocracies and anocracies yield a coefficient of -0.146, almost identical to the coup baseline and significant at the 1 percent level, while deaths in democracies yield a near-zero, statistically insignificant coefficient of -0.027. Succession in a stable democracy may be more orderly and preserve policy continuity, leaving bilateral trade undisturbed. If so, the realignment in ideal points following a leader's death should itself be larger in autocracies and anocracies than in democracies, a possibility we check directly. The mean absolute change in ideal point distance around a leader's death is 0.374 in autocracies and anocracies, compared with 0.266 in democracies, a gap that is highly significant ($p < 0.001$), while the signed mean change (whether deaths in democracies lead more often to realignments into a particular direction) is close to zero and statistically indistinguishable across the two groups. The larger trade response we estimate for autocracies is therefore accompanied by a correspondingly larger realignment.

Column (5) pools coups and leader deaths into a single sample. Pooling tightens the standard error relative to either source alone, 0.031 pooled versus 0.039 for coups only and 0.048 for deaths only, while leaving the point estimate, -0.132, within the range spanned by the two sources individually. The fact that an independent source of political-realignment shocks, once restricted to the same kind of regime where coups occur, reproduces the magnitude of our main results suggests that the effects we document are not specific to coups, but extend to other large political realignments more generally.

Table 5: Coups, Leader Deaths, and the Pooled Sample: Collapsed Pre/Post DiD

Dependent Variable:	Bilateral Trade Flows				
	Coups	Leader deaths			Pooled
		All	Autocracy	Democracy	
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Post × ΔIPD	-0.154***	-0.074	-0.146***	-0.027	-0.132***
	(0.039)	(0.048)	(0.047)	(0.074)	(0.031)
Fixed-effects					
Dyad × Event FE	Yes	Yes	Yes	Yes	Yes
Exporter-Year FE	Yes	Yes	Yes	Yes	Yes
Importer-Year FE	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	147,504	111,346	61,019	43,968	260,822
Pseudo R^2	0.97555	0.99274	0.98393	0.99662	0.98906

Note: Column (1) reproduces our coup-based baseline. Column (2) uses leader deaths from Jones and Olken (2005) as the event in place of a coup. Columns (3) and (4) split the leader-death sample by the leader's country's Polity2 score immediately before death, autocracy/anocracy (≤ 5) versus democracy (> 5). Column (5) pools coups and leader deaths into a single sample. All columns use the collapsed pre/post specification of Panel B of Table 2, with dyad-by-event, exporter-year, and importer-year fixed effects. Standard errors, in parentheses, are clustered by dyad. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

4.3.4 Additional Robustness Checks

We run three further checks, each detailed in full in the Appendix. First, we re-estimate our stacked event study using an alternative trade data source, the CEPII Gravity database, in place of the TRADHIST data used in our baseline specification. Despite a smaller matched sample, due to CEPII Gravity's sparser historical coverage, coefficients retain the same sign throughout the post-coup period, though smaller in magnitude and short of conventional significance, which we attribute to reduced statistical power rather than a weaker effect. The contrast with the failed-coup placebo still holds (Appendix A.12). Second, we narrow the event window from $k = -4, \ldots, 10$ to $k = -3, \ldots, 5$, which changes both which coups survive the contamination filter and how many years of post-coup trade enter the stacked panel. Our main results survive this alternative window. Post-period coefficients again turn negative and significant from $k = 2$ onward, while the failed-coup placebo shows no pattern (Appendix A.13). Third, to rule out a secular drift in ideal-point distance confounding our identification, we detrend the alignment measure by subtracting each year's cross-sectional mean ideal-point distance from every observation before constructing the pre- and post-coup dose. The detrended dose is highly correlated with the raw dose (correlation of 0.99) and yields coefficients nearly identical to our baseline results (Appendix A.14).

4.4 Heterogeneity

We next examine whether the trade effects documented above extend beyond the Cold War and how they vary across sectors. The analyses that follow use identification Strategy 1, which provides greater statistical power than the more restrictive specification of Strategy 2.

4.4.1 Post-Cold War Effects

We first ask whether coup-induced political realignments continue to affect trade after the end of the Cold War. We re-estimate equation (11) using coups and bilateral trade flows from the post-Cold War period (year ≥ 1990), while otherwise maintaining the same empirical design.

We find no evidence of a systematic trade response to coups after the Cold War. Using our baseline measure of political alignment, United Nations ideal point distance, the post-coup coefficients are small and jointly indistinguishable from zero (joint Wald $p = 0.777$). Appendix Figure A.10 reports the corresponding event-study estimates.

A potential concern is that UN ideal point distance may be less informative about geopolitical alignment after the collapse of the Cold War's bipolar structure. We therefore repeat the analysis using the geopolitical score developed by Fan et al. (2025a) and Fan et al. (2025b) which does not rely on a bipolar conception of geopolitical alignment. Here the results are not null. The post-coup coefficients are jointly significant (joint Wald $p = 0.020$), with the same gradually-emerging, negative pattern as our baseline estimates. However, the effect remains about 50-75 percent smaller than its Cold War-period counterpart under the same measure (for example, -0.55 against -0.14 at $k = 2$, and -0.82 against -0.37 at $k = 5$), as shown in Appendix Figures A.11 and A.12.

Taken together, these results present a mixed picture. Under UN ideal point distance, the effect vanishes entirely after 1990. Under the geopolitical score, a trade response to coup-induced realignment persists in the post-Cold War period, but at a fraction of the magnitude of the Cold War-period effect under the same measure. This is consistent with the Cold War's bipolar structure amplifying the economic consequences of political realignment. Competing superpowers may have had both the incentives and the instruments to reorganize economic relationships along geopolitical lines to a degree that only partially persists once that bloc structure dissolved.

4.4.2 Sector-Level Effects

We next ask whether the trade response is concentrated in particular sectors. No single source provides sector-level bilateral trade data with sufficiently broad Cold War coverage for both superpowers, so we draw on two complementary sources, one for each side. For the Soviet Union, we use the Soviet trade yearbooks digitized for this paper and described in Section 2.2, which report bilateral trade directly from Soviet records. For the United States, we use the NBER-UN World Trade Flows data (Feenstra et al., 2005), the standard source for historical disaggregated bilateral trade. Together, the two sources let us examine the sectoral composition of the trade response for both superpowers, each drawn from the best available source for the respective side.

We classify bilateral trade into four broad categories common to both exercises: primary goods, mineral fuels, chemicals and manufactures, and machinery and transport equipment, together with an aggregate across all four. For each category and for the aggregate, we pool the USSR-partner and USA-partner dyads into a single regression and estimate a collapsed pre/post version of our event-study specification, allowing the coefficient and the year fixed effect to differ by hegemon,

$$X^h_{nst} = \exp\left(\beta^h_s \cdot \text{Post}_{nt} \cdot \Delta\text{IPD}_{n,h} + \gamma_{ht} + \tilde{\alpha}_{n,\text{event}}\right) \tilde{\eta}_{nst}, \tag{13}$$

where X^h_{nst} denotes bilateral trade between coup country n and hegemon $h \in \{\text{USSR}, \text{USA}\}$ in category s in year t, β^h_s is allowed to differ by hegemon, and γ_{ht} is a hegemon-by-year fixed effect. The USSR-partner and USA-partner samples are pooled into a single regression per category, with $\tilde{\alpha}_{n,\text{event}}$ a dyad-by-coup-event fixed effect.

Table 6 shows that political distance reduces trade across almost all sectors, with the aggregate response significant for both hegemons. The "All Sectors" coefficient is negative and significant both for the Soviet Union (-0.697, $p < 0.001$) and for the United States (-0.158, $p = 0.047$).[9] At the individual-category level, roughly half of the categories reach significance on each side, though not the same ones. For the Soviet Union, the decline is concentrated in machinery and transport equipment (-0.841, $p < 0.001$) and primary goods (-0.547, $p < 0.001$), while mineral fuels (-0.374, $p = 0.361$) and chemicals and manufactures (-0.065, $p = 0.836$) show no significant response. For the United States, the decline is instead concentrated in chemicals and manufactures

[9]Note that these results differ from the main effect we document in Table 2 since the United States sector-level results rely on the NBER-UN data rather than TradHist. The former excludes zero trade flows, begins only in 1972, and was never supplemented with our digitized Soviet trade data the way TradHist was.

Table 6: Sector-Level Response of U.S. and Soviet Bilateral Trade to Coup-Induced Alignment Shifts

Dependent Variable:	Sector Trade (USD)				
	All Sectors	Primary	Mineral Fuels	Chem. & Manuf.	Machinery & Transp.
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Post × ΔIPD (USA)	-0.158**	-0.098	0.309	-0.225**	-0.309*
	(0.080)	(0.100)	(0.275)	(0.093)	(0.167)
Post × ΔIPD (USSR)	-0.696***	-0.547***	-0.374	-0.065	-0.841***
	(0.163)	(0.119)	(0.402)	(0.306)	(0.162)
Fixed-effects					
Dyad × Coup Event FE	Yes	Yes	Yes	Yes	Yes
Hegemon × Year FE	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	1,600	1,286	1,181	1,437	1,330
Pseudo R^2	0.98484	0.97841	0.93213	0.98290	0.98634

Note: Each column reports $\hat{\beta}_s$ from equation (13), estimated via PPML jointly on the pooled USA-partner and USSR-partner samples of dyad-events around successful coups, either for one category or for the aggregate across all four ("All Sectors"). The two rows come from a single regression per column, with hegemon-by-year fixed effects in place of a common year fixed effect so each side retains its own year trend; the USA and USSR coefficients are therefore directly comparable within each column. Clustered (dyad) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

(-0.225, $p = 0.015$) and, at the 10 percent level, machinery and transport equipment (-0.309, $p = 0.062$), while primary goods (-0.098, $p = 0.320$) and mineral fuels (0.309, $p = 0.259$) show no significant response.

Product-Level Classification Tests. The four-category split above is coarse. The NBER-UN data's underlying product-level detail lets us ask a sharper version of the same question. That is, within the USA-partner sample, which kinds of trade are most exposed to political realignment. We combine the trade panel with five external classifications of traded goods, which each target a different notion of exposure. Revealed comparative advantage and product complexity, from the Growth Lab's Atlas of Economic Complexity (Hausmann et al., 2014), test whether the response is stronger in goods the partner is naturally suited to produce or that require sophisticated productive capabilities. Rauch (1999)'s distinction between organized-exchange, reference-priced, and differentiated goods tests whether goods that require relationships, information, and matching are more exposed than standardized commodities. Nunn (2007)'s industry-level contract-intensity measure asks the same question from the input side, whether industries relying on relationship-specific intermediate inputs see a larger decline. The UN's Broad Economic Categories test whether the decline concentrates in capital goods, intermediate inputs, or consumer goods. Finally, UNCTAD's Lall (2000) technological classification tests whether the decline grows with a good's technological sophistication. For each classification we estimate a product-level analogue of equation (13),

$$X_{npt} = \exp\left(\beta \cdot \text{Post}_{nt} \cdot \Delta\text{IPD}_{n,\text{USA}} + \theta \cdot \text{Post}_{nt} \cdot \Delta\text{IPD}_{n,\text{USA}} \cdot Z_p + \gamma_t + \tilde{\alpha}_{n,\text{event}} + \mu_p\right) \tilde{\eta}_{npt}, \quad (14)$$

where X_{npt} is USA-partner bilateral trade in SITC2 product p, μ_p a product fixed effect, and Z_p

the classification measure. Where a classification is instead a genuine partition of the product space (Broad Economic Categories or Lall's technology tiers), we replace the single interaction with a fully saturated set of category-specific coefficients. Full construction detail and all underlying tables are reported in Appendix A.22.

Figure 7: Sector/Product Classification Tests: Summary of Coefficients

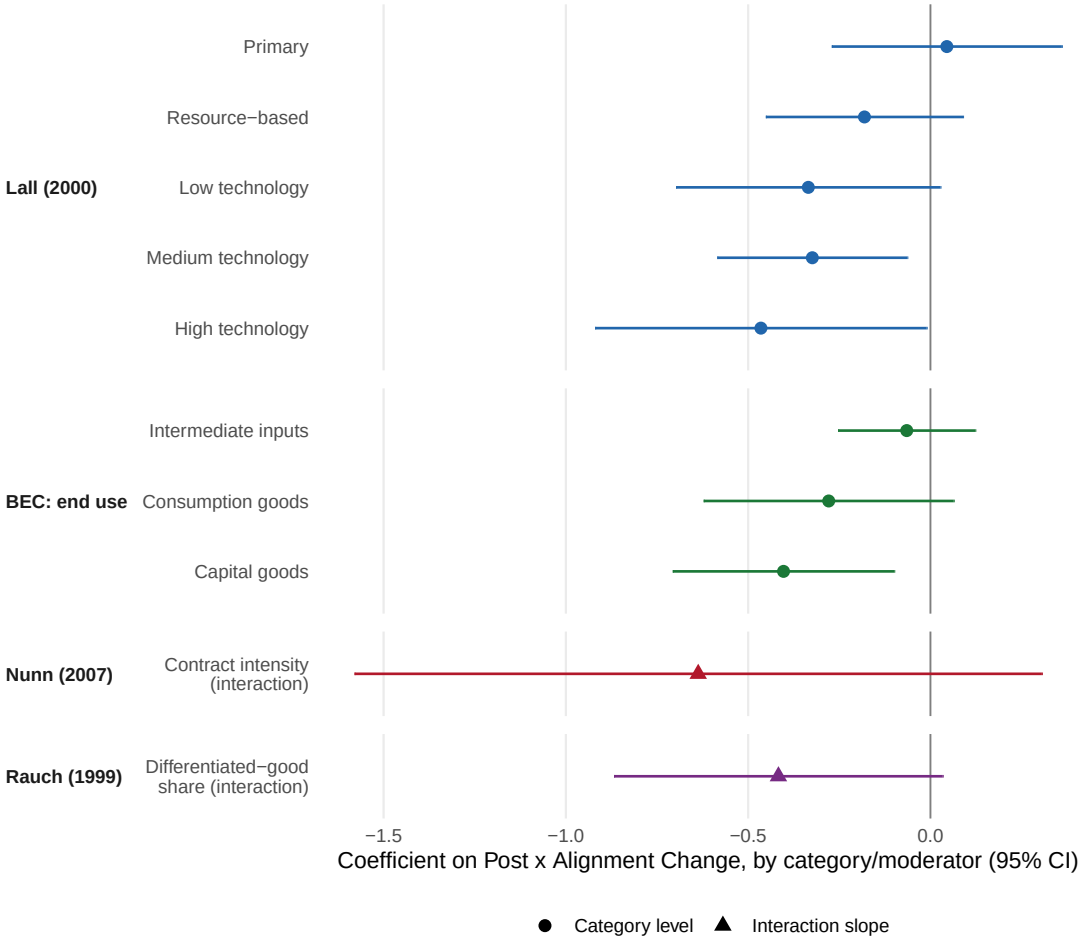

Note: Point estimates and 95% confidence intervals from the Rauch, Nunn, Broad Economic Category, and Lall classification tests described in the text (comparative advantage and complexity omitted, see Appendix A.22 for that result). Triangles are single interaction-slope coefficients ($\hat{\theta}$ in equation (14)), circles are category-specific dose coefficients from the fully saturated designs. All estimates are from the USA-partner sample.

Figure 7 shows that a consistent pattern emerges across all four independently constructed classifications. Homogeneous, commodity-like, or low-technology trade is largely unaffected by political distance. Lall's Primary category is statistically indistinguishable from zero, and BEC's Intermediate inputs category is the smallest and least significant of the three end-use groups. Trade in relationship-dependent or technologically sophisticated goods declines the most. Rauch's differentiated goods, Nunn's contract-intensive industries, BEC's Capital goods, and Lall's High-technology manufactures all show the largest point estimates, and Lall's five categories trace out a near-monotonic gradient from Primary to High technology. Separately, using comparative advantage,

we find that the trade response to realignment is amplified, in both directions, in sectors where the partner itself has a comparative advantage, consistent with a standard trade-cost/integration channel (see details in Appendix A.22). Taken together, these classifications point in the same direction. Political realignment does not reduce trade uniformly across the economy, it falls hardest on trade that depends on trust, relationships, and sophisticated productive capability, and least on trade in standardized, easily substitutable commodities.

5 Mechanisms

We have established that political realignment causes a large, persistent decline in bilateral trade, concentrated in the Cold War and strongest for relationships with the superpowers themselves. This reduced-form relationship is silent, however, on the specific channels through which political distance translates into lower trade. The appropriate policy response to political estrangement, and its welfare consequences, differ substantially depending on whether the trade decline reflects government policy tools such as aid and sanctions, or broader private-sector responses to political risk that are harder to target. We therefore ask through which concrete channels political realignment moved trade. We focus on foreign aid and sanctions, the two most commonly used instruments of Cold War economic statecraft (Baldwin, 1985). Other channels may also matter, including changes in shipping and insurance, firms' perceptions of political risk, or informal government pressure, and we return to this question after presenting the aid and sanctions results below.

5.1 Foreign Aid

Foreign aid was an important tool of superpower influence during the Cold War. Both the United States and the Soviet Union used aid strategically to reward allies and influence the political orientation of recipient countries (Meyer and Wesseler, 2026). Aid can also directly shape bilateral trade, particularly when assistance finances imports from the donor country or supports projects that rely on donor-country goods, equipment, and expertise (Levchenko et al., 2026; Pettersson and Johansson, 2013). Changes in aid following a political realignment may therefore provide one channel through which political alignment affects bilateral trade. We test this mechanism by asking whether coup-induced realignments led the superpowers to reallocate foreign assistance toward countries that became politically closer and away from those that became more distant.

To study this channel, we assemble bilateral aid data for both superpowers. For Soviet aid, we use data from declassified CIA files documenting roughly 4,000 Communist-bloc economic and military aid projects to developing countries between 1954 and 1990 from our prior work (CIA, 1955-86; Meyer and Wesseler, 2026).[10] Economic aid is recorded at the individual project level, with each entry listing the donor, recipient, date, project description, and the value extended and drawn. Military aid is recorded as an annual donor-recipient total, split by financing instrument.

[10]We focus on Soviet aid, which accounts for 75% of total Communist assistance in the CIA data. Appendix A.15 describes the digitization, valuation, and classification of both aid datasets in full.

For U.S. aid, we use the State Department's Greenbook (*U.S. Overseas Loans and Grants*), which records annual obligations of loans and grants by recipient country and aid account from 1945 onward, along with a broad distinction between economic and military assistance.

To make the two sources comparable, we construct a common classification and assign aid to four categories, military, food and humanitarian assistance, education and technical assistance, and general finance and other assistance.[11] We aggregate both series to the recipient-year level and convert all values to constant 2011 U.S. dollars. The resulting data cover aid from both superpowers to more than 120 recipient countries over the Cold War.

We first ask whether total aid responds to the same coup-induced alignment shocks that drive our trade results. We use a shorter event window of $k = -4, \ldots, 5$, focusing on the years in which aid is most likely to respond to political realignment, while retaining the same canonical coup-event sample as the rest of the paper. We combine each superpower's economic and military aid into a single outcome at the recipient-year level and match aid from superpower $i \in \{\text{USA}, \text{RUS}\}$ to the corresponding coup-induced change in political distance, ΔIPD_{ni}. Stacking the U.S. and Soviet aid panels, we estimate

$$\text{Aid}_{nit} = \exp\left(\sum_{k \neq -1} \beta_k \cdot \mathbf{1}[t - t_0 = k] \cdot \Delta\text{IPD}_{ni} + \gamma_t + \tilde{\alpha}_{ni}\right) \tilde{\eta}_{nit}, \tag{15}$$

where Aid_{nit} is total aid received by coup country n from superpower i in year t, $\tilde{\alpha}_{ni}$ is the dyad-by-coup-event fixed effect described in Section 3.2.3, and $\tilde{\eta}_{nit}$ is a multiplicative error term. Because the outcome is bilateral aid from a single superpower rather than bilateral trade, we include calendar-year fixed effects γ_t in place of the exporter-year and importer-year fixed effects in equation (11).

Panel (a) of Figure 8 shows that aid from the estranged superpower declines after a coup. The coefficients are flat and small before the coup, at $k = -4, -3, -2$ (0.018, 0.095, 0.112, all $p > 0.2$), consistent with no differential pre-trend, and turn uniformly negative from $k = 0$ onward, with $k = 1$ reaching marginal significance (-0.33, $p = 0.053$). Splitting combined aid into its economic and military components in Panel (b) shows that this decline is concentrated in military rather than economic aid. Military aid's coefficients at $k = 1$, $k = 3$, and $k = 4$ are each individually significant (-0.48, -0.47, and -0.63, all $p < 0.05$), while the economic-aid coefficients are not individually significant.

We test whether this result is specific to the two superpowers or reflects broader bloc-level engagement using bilateral aid from every individual Western donor reporting to the OECD's DAC2a database and every individual Warsaw Pact donor in the CIA data. Aid from DAC donors and from Warsaw Pact donors declines by a statistically indistinguishable amount after a coup-induced realignment, so the pattern is not confined to the two superpowers, but the bilateral

[11] For Soviet economic aid, we classify each of the roughly 2,550 individual project descriptions. For U.S. aid, the Greenbook records obligations by account (e.g. "Public Law 480 Title II Grants," "Peace Corps") rather than by individual project, so we classify the roughly 40 distinct account names and apply this mapping to all obligations recorded under each account.

Figure 8: Event Study of Foreign Aid Around Coups

(a) Combined (economic + military)

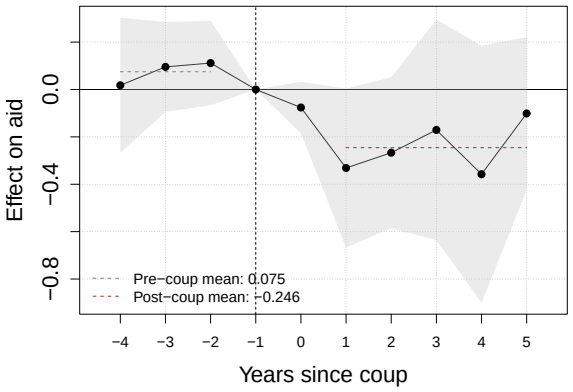

(b) Economic versus military

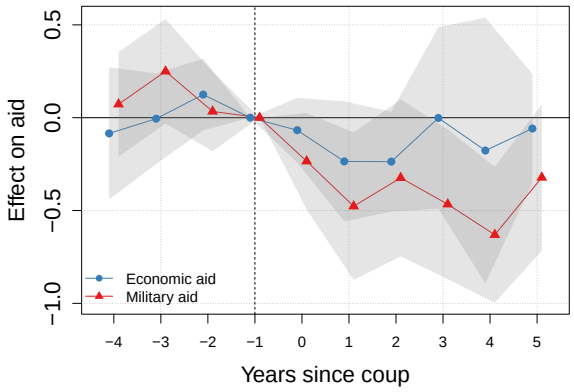

Note: Each panel plots $\hat{\beta}_k$ from equation (15). Panel (a) pools economic and military aid from both superpowers. Panel (b) reports economic and military aid separately, still pooled across both superpowers. Standard errors are clustered by dyad.

estimates are noisier. Appendix A.16 reports the full results.

As an additional check, we test whether formal aid sanctions respond to coup-induced alignment shifts using the aid-sanction data of Gibson et al. (2025) in Appendix A.17.[12] Applying the same collapsed pre/post specification as above, we find no evidence that political realignment affects the imposition of U.S. aid sanctions ($\beta = 0.005$, $SE = 0.043$, $p = 0.913$).[13] The aid-sanctions margin therefore provides no additional evidence beyond the changes in aid flows documented above.

To gauge the importance of the aid reduction for trade realignment, we combine our estimated aid response with an aid-to-trade elasticity from the literature. Averaging the post-coup coefficients in Panel (a) of Figure 8 ($k = 0, \ldots, 5$) implies that a one-unit increase in political distance reduces superpower aid by 19.5 percent. Pettersson and Johansson (2013) estimate that a 10 percent increase in bilateral aid is associated with roughly 1 percent higher recipient exports to the donor, an aid-to-trade elasticity of 0.1. Combining these two estimates implies a 2.0 percent mechanical decline in trade through the aid channel alone, compared with the 15.4 percent decline implied by our main coup-identified estimate. Based on this back-of-envelope calculation, changes in foreign aid can account for roughly 14 percent of the trade response we document. We view this figure as suggestive rather than precise. Pettersson and Johansson (2013)'s elasticity is estimated on a modern (1990–2005), globally representative donor-recipient panel rather than on the Cold War superpowers specifically. Taken at face value, it suggests that foreign aid is a plausible but modest channel through which political realignment affects trade.

[12] We thank Yoto Yotov for sharing the data.

[13] We cannot conduct the corresponding analysis for the Soviet Union since only five Soviet aid-sanction episodes appear in the data over our entire sample period, none of which falls within our estimation window.

5.2 Sanctions

Sanctions are another potential instrument of Cold War statecraft, used to punish adversaries or induce changes in their behavior. We ask whether the same coup-induced political realignments that affect trade and foreign aid also lead to the imposition of new sanctions. We use the Global Sanctions Data Base (Felbermayr et al., 2020), which records the sender, target, year, and type of sanction, including trade, financial, travel, arms, military, and other sanctions.

For each sanction type, we construct an onset indicator equal to one when a sanction becomes active in a directed sender-target pair. We retain the full stacked panel described in Section 3.2.3 and estimate a collapsed pre/post specification,

$$\text{Onset}_{nit} = \beta \cdot \text{Post}_{nt} \cdot \Delta\text{IPD}_{ni} + \gamma_{it} + \delta_{nt} + \tilde{\alpha}_{ni} + \tilde{\eta}_{nit}, \tag{16}$$

where $\text{Onset}_{nit} = 1$ if a sanction of the given type becomes newly active between n and i in year t, $\text{Post}_{nt} = 1$ in the years following the coup. We estimate equation (16) by OLS separately for each sanction type.

Table 7: Sanction Onset Response to Coup-Induced Alignment Shifts, by Type

Dependent Variable:	New Sanction					
	Any Type	Trade	Financial	Travel	Other	Arms & Military
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Post × ΔIPD	0.002**	0.001***	0.0009**	0.002***	0.002***	-0.0007*
	(0.0007)	(0.0005)	(0.0005)	(0.0003)	(0.0004)	(0.0004)
Fixed-effects						
Dyad × Coup Event FE	Yes	Yes	Yes	Yes	Yes	Yes
Exporter-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Importer-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	159,500	159,500	159,500	159,500	159,500	159,500
R^2	0.41344	0.34331	0.20280	0.29518	0.39009	0.62477

Note: Each column reports $\hat{\beta}$ from equation (16), estimated separately by sanction type. Sanction data comes from the Global Sanctions Data Base (Felbermayr et al., 2020). Clustered (dyad) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

Sanction onsets are rare, accounting for at most 0.3 percent of dyad-year observations for any individual sanction type, but they respond systematically to coup-induced changes in political distance as Table 7 shows. A one-unit increase in political distance raises the probability of a new sanction of any type by 0.19 percentage points ($p = 0.012$), relative to a baseline probability of 0.3 percent. The effect is broad-based, travel restrictions (0.20 percentage points, $p < 0.001$), other sanctions (0.16 percentage points, $p < 0.001$), trade sanctions (0.14 percentage points, $p = 0.007$), and financial sanctions (0.09 percentage points, $p = 0.047$) all rise significantly. Arms and military sanctions, which we combine because neither can be estimated precisely on its own, move in the

opposite direction (-0.07 percentage points, $p = 0.064$).

We next ask whether the increase in trade sanctions is sufficient to explain a meaningful share of the main trade effect. To do so, we combine two estimates. First, using the same onset specification as Table 7, a one-unit increase in political distance raises the probability of a new trade sanction by 0.14 percentage points ($p = 0.007$), relative to a baseline onset probability of 0.13 percent. Second, Felbermayr et al. (2020) estimate that an active trade sanction reduces bilateral trade by 14.5 percent in their preferred specification. Combining these estimates implies a trade decline of approximately 0.02 percent through the formal trade-sanctions channel, compared with the 14.3 percent decline implied by our main coup-identified estimate. On this calculation, formal trade sanctions account for less than one percent of the overall trade response.

Thus, political realignment does lead to greater use of economic sanctions, but the magnitude of this response is far too small to explain the decline in bilateral trade. This suggests that the trade response operates through other channels than through the imposition of trade sanctions.

5.3 A Case Study Without Formal Barriers

A substantial share of the trade response remains unexplained by the channels we study. To explore whether political realignment can move trade even in the complete absence of formal policy instruments, we turn to a historical episode in which political relations deteriorated sharply without any accompanying change in formal trade policy, the 1964 Johnson Letter and the subsequent deterioration in US-Turkey relations.

In June 1964, US President Lyndon B. Johnson sent a letter to Turkish Prime Minister İsmet İnönü warning that NATO might not defend Turkey against Soviet retaliation if Turkey proceeded with a planned military intervention in Cyprus. The letter was received negatively in Turkey and is widely regarded as an important turning point in US-Turkey relations, contributing to a more independent Turkish foreign policy in the years that followed. Importantly for our purposes, the deterioration in relations was not accompanied by a tariff, sanction, export control, or change in alliance commitment in either direction.[14] A declassified CIA assessment of Turkish reactions to the letter, reproduced in Appendix A.21, provides evidence that the episode represented a significant political shock. The Appendix also reports complementary descriptive evidence on trade and alignment. The bilateral trade friction implied by the baseline gravity specification increases in both directions of US-Turkey trade after 1964, alongside the increase in ideal point distance between the two countries. The episode provides an opportunity to examine whether bilateral trade can change in the absence of an observable formal trade-policy response.

We study the episode using a synthetic control analysis (Abadie et al., 2010). We construct a synthetic Turkey from a weighted combination of donor countries in Europe (excluding members of the communist bloc) and the MENA region, additionally matched on pre-1964 level of GDP and trade openness. We consider different variants of constructing the donor pool using weights

[14] Turkey eventually invaded Cyprus in 1974, after which the United States imposed an arms embargo that lasted until 1975. We restrict attention to the years immediately following the 1964 letter, well before this later episode.

selected with Lasso, Ridge Regression or an L_2 constrained synthetic control, which penalizes excessive weights on single countries.

Figure 9: US-Turkey Trade: Actual versus Synthetic Control

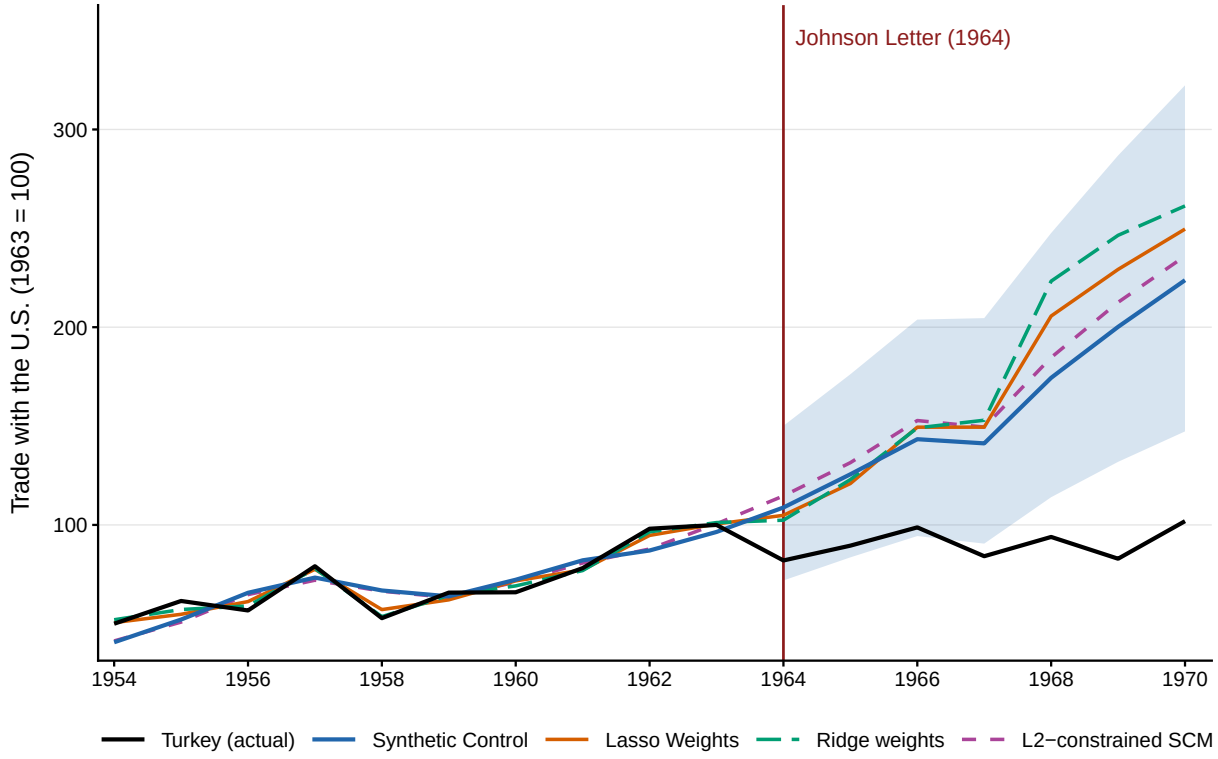

Note: The black line plots total U.S. trade with Turkey in black and its synthetic control in blue, estimated using the synthetic control method over 1953-1963. Shaded area indicate 90% confidence intervals bootstrapped following Cattaneo et al. (2021), dashed lines indicate alternative constructions of the donor weights. The vertical line marks the Johnson Letter (1964).

Figure 9 shows that U.S.-Turkey trade falls strongly relative to its synthetic counterfactual. By construction, synthetic trade tracks actual US-Turkey trade before the Johnson letter. After 1964, however, the two series diverge. Even during a time of global trade expansion and robust economic growth in both the U.S. and Turkey, U.S.-Turkey trade stalls at its 1963 level. In contrast, each synthetic control continues along the upward trajectory observed before the letter. By 1970, Turkey's Trade with the U.S. is about half the value of its synthetic counterfactual. Table A.12 shows the donor weights underlying the synthetic counterfactuals. The baseline assigns a high weight to Iraq, however, this is not the case for the other synthetic counterfactuals.

The episode complements the mechanism evidence above because it shows that geopolitical deterioration can lead to a decline in bilateral trade without an accompanying sanction, change in aid, or formal trade restriction.

6 Welfare Effects of Fragmentation

Our findings establish that political distance imposed substantial trade costs during the Cold War. We now embed our coup-identified estimate in the structural gravity model of Section 3 to quantify the resulting welfare consequences. We conduct two counterfactual exercises. The first measures the aggregate welfare losses from Cold War geopolitical fragmentation. The second asks whether countries' political alignments were consistent with their economic interests, both during the Cold War and in the contemporary US-China rivalry.

6.1 Extending the Model to Welfare

Section 3.1.1 derives bilateral trade flows X_{nit} from the structural gravity model in equation (1) and, in equation (5), the mapping from political distance to the trade cost component τ_{nit}^{TV}. Welfare counterfactuals require two further objects, bilateral trade shares and a solution for general equilibrium counterfactuals.

Trade Shares and Welfare. Dividing equation (1) by total expenditure X_{nt} gives country n's import share from i

$$\lambda_{nit} \equiv \frac{X_{nit}}{X_{nt}} = \frac{Y_{it}\,\phi_{nit}}{\Omega_{it}\,\Phi_{nt}}. \tag{17}$$

Real income in country n is $W_n = w_n/P_n$, where P_n is the CES price index. Under the Armington structure of Section 3.1.1, the domestic share satisfies $\lambda_{nn} = (w_n/P_n)^{1-\sigma}$, so that, using $\varepsilon \equiv \sigma - 1$,

$$W_n \;\propto\; \lambda_{nn}^{-1/\varepsilon}. \tag{18}$$

Arkolakis et al. (2012) show that the domestic expenditure share and the trade elasticity are sufficient statistics for welfare changes across a broad class of quantitative trade models. Welfare therefore depends on changes domestic trade share λ_{nn}, the fraction of a country's expenditure that stays at home. A fall in λ_{nn} reveals that a country has gained access to foreign varieties, and equation (18) is a sufficient statistic for the resulting welfare gain.

Counterfactual Trade Costs. For any variable x, let $\hat{x} \equiv x^{\text{CF}}/x^{\text{actual}}$ denote its counterfactual change. Each exercise below perturbs IPD_{nit} to some counterfactual configuration $\text{IPD}_{nit}^{\text{CF}}$. By equation (5), the implied change in bilateral trade costs is

$$\hat{\tau}_{nit} \;=\; \exp\!\left(\frac{\beta}{\varepsilon}\Big(\text{IPD}_{nit}^{\text{actual}} - \text{IPD}_{nit}^{\text{CF}}\Big)\right), \tag{19}$$

using the coup-identified estimate $\beta = -0.154$ from Table 2, Panel B. Domestic trade costs are held fixed throughout ($\hat{\tau}_{nnt} = 1$), since political distance does not apply to within-country transactions.

Solution Method. We solve for the counterfactual equilibrium using the exact hat algebra of Dekle et al. (2008), which recovers $\{\hat{w}_{it}, \lambda_{nit}^{\text{CF}}\}$ from observed trade shares and trade cost changes

alone. Counterfactual trade shares satisfy

$$\lambda_{nit}^{\mathrm{CF}} \;=\; \frac{\hat{w}_{it}^{-\varepsilon}\,\hat{\tau}_{nit}^{-\varepsilon}\,\lambda_{nit}}{\sum_{k}\hat{w}_{kt}^{-\varepsilon}\,\hat{\tau}_{nkt}^{-\varepsilon}\,\lambda_{nkt}}, \tag{20}$$

and, assuming balanced trade, goods market clearing in changes requires

$$\hat{w}_{it}\,Y_{it} \;=\; \sum_{n}\lambda_{nit}^{\mathrm{CF}}\,\hat{w}_{nt}\,Y_{nt}. \tag{21}$$

We solve equations (20) and (21) jointly for $\{\hat{w}_{it}, \lambda_{nit}^{\mathrm{CF}}\}$ by iterating to convergence, normalizing the GDP-weighted world wage index to one, separately for each year in our sample, using that year's observed bilateral trade flows and GDP. The resulting welfare change follows from equation (18),

$$\hat{W}_{nt} \;=\; \left(\hat{\lambda}_{nnt}\right)^{-1/\varepsilon}, \qquad \hat{\lambda}_{nnt} \equiv \frac{\lambda_{nnt}^{\mathrm{CF}}}{\lambda_{nnt}^{\mathrm{actual}}}, \tag{22}$$

which we implement using the trade elasticity $\varepsilon = 5.03$ from Head and Mayer (2014), the same value used for the tariff equivalents in Section 4.1. A reduction in the home trade share ($\hat{\lambda}_{nnt} < 1$) reflects expanded trade with newly accessible partners and implies a welfare gain ($\hat{W}_{nt} > 1$). We report all welfare changes as percentage changes in real income.

Data. The counterfactuals require a bilateral trade and GDP panel, including domestic (within-country) trade, which is not needed for the dyadic estimation in Section 4. We construct this panel from the CEPII Gravity dataset and Tradhist. Since neither database directly reports domestic trade X_{nn}, we impute it following Kleinman et al. (2024). Gross output is computed as 2.2 times GDP, and domestic expenditure is this figure net of total exports.[15] In modern data, this imputation yields figures in line with textbook levels of gains from trade (Costinot and Rodríguez-Clare, 2014).

6.2 The Economic Costs of the Cold War

Our first counterfactual asks what real incomes would have been during the Cold War if geopolitical alignment had not affected bilateral trade costs. We eliminate the trade-cost component associated with political distance, that is setting $\mathrm{IPD}_{nit}^{\mathrm{CF}} = 0$ for every dyad and year while leaving all other trade costs unchanged. We then compute the resulting welfare change.

Figure 10 shows that geopolitical barriers imposed non-trivial real income losses throughout the Cold War. Averaged across all countries and years, the GDP-weighted welfare cost amounts to 0.20% of real income. The burden is sharply unequal across blocs. Non-Aligned countries bear an average cost of 0.43%, roughly three times the 0.15% and 0.13% borne by Western- and Eastern-bloc members, respectively. The intuition follows from the structure of trade. Bloc members traded

[15] Datasets that directly report domestic expenditures either do not cover the Cold War era, such as Borchert et al. (2021), or have many missing observations, particularly for countries in the Global South that are central to our analysis, such as Head and Mayer (2021).

Figure 10: Welfare costs of Cold War geopolitical barriers

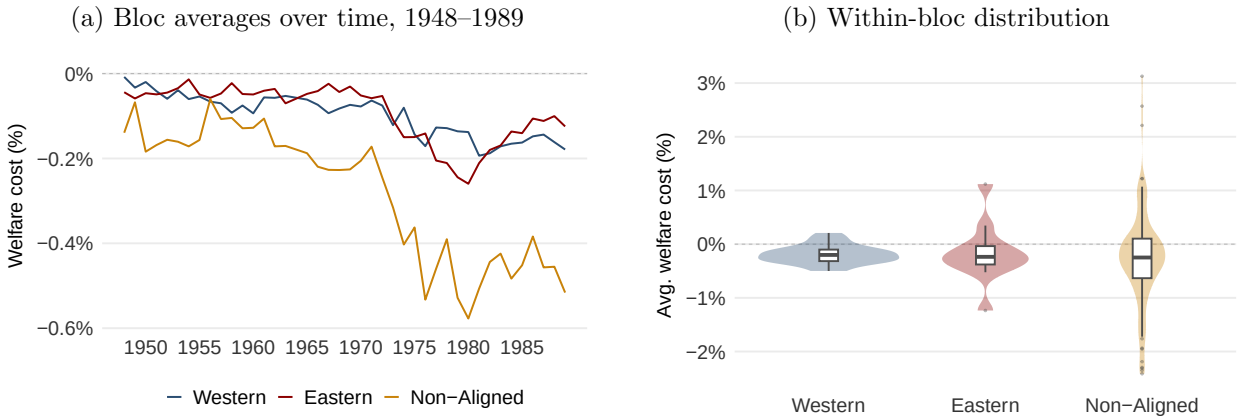

Notes: Both panels measure welfare costs as the percentage reduction in real income relative to a counterfactual in which political alignment does not affect bilateral trade costs ($\text{IPD}_{nit} = 0$ for all dyads), computed from equations (20)–(22) with $\varepsilon = 5.03$ (Head and Mayer, 2014) and $\beta = -0.154$ (Table 2, Panel B). Country groups: Western = NATO members plus Japan and Australia, Eastern = Warsaw Pact members, Non-Aligned = all others. *Panel (a):* Annual averages weighted by country GDP. *Panel (b):* Each observation is a country's GDP-weighted average cost across all sample years, the violin shows the kernel density and the embedded boxplot shows the median and interquartile range.

primarily within their own bloc (see Appendix Figure A.13), where political distance was relatively low, so geopolitical barriers with the opposing bloc and Non-Aligned countries affected a comparatively small share of their potential trade. Non-Aligned countries, by contrast, were politically distant from both blocs and therefore faced geopolitical trade costs across a much broader set of economically relevant partners. Removing these barriers consequently produces substantially larger welfare gains. The economic burden of Cold War fragmentation therefore fell disproportionately on countries outside the two formal blocs.

The welfare cost increased substantially over time. It averaged 0.09% of real income in 1948–1959 and tripled up to 0.27% in 1975–1989. The annual series shows a particularly pronounced increase around 1973–1974, after which welfare costs remain substantially above their earlier levels. This timing coincides with a period of intensified superpower competition in the developing world and growing geopolitical polarization during the 1970s. Costs peak around 1980, during the renewed deterioration in East-West relations surrounding the Soviet invasion of Afghanistan and the end of détente, before declining somewhat in the mid-1980s.

Panel (b) of Figure 10 also reveals substantial heterogeneity within blocs. For 40 of the 170 countries in our sample, removing geopolitical barriers would have resulted in welfare losses. This is a general equilibrium effect. Countries that were politically close to major trading partners effectively enjoyed preferential access relative to more politically distant competitors. Removing geopolitical barriers erodes this advantage by redirecting demand toward previously disadvantaged suppliers, generating a terms-of-trade loss that can outweigh gains from lower import costs. These 'winners' of geopolitical fragmentation tend to be among small, highly open economies with concentrated trade relationships, such as Djibouti, Cape Verde, Jordan, the Bahamas, and Hong Kong.

In Appendix Section A.20, we further examine the sources of these heterogeneous welfare costs by decomposing the full counterfactual into three partial counterfactuals, each removing geopolitical

barriers between a different pair of blocs. The decomposition reveals markedly different sources of welfare losses across groups. Eastern-bloc losses were driven primarily by barriers across the Iron Curtain, while Non-Aligned losses arose overwhelmingly from barriers to Western markets. Western-bloc losses, by contrast, were driven primarily by political frictions within the Western bloc rather than by barriers with the Eastern bloc.

6.3 The Economics of Bloc Loyalty

We can also study how the economic incentives implied by our estimates compare to geopolitical outcomes. Given the actual political and economic configuration of the Cold War, would a country have benefited more from adopting the geopolitical alignment of the United States or that of the Soviet Union? This exercise reveals whether countries' political and economic incentives reinforced or cut against each other.

Design. For each country n, we consider two counterfactuals, adopting the geopolitical alignment of the United States or that of the Soviet Union. Under alignment with superpower $a \in \{\text{US}, \text{USSR}\}$, we replace country n's political distance to every trading partner i with the corresponding distance between superpower a and country i, so that

$$\text{IPD}_{ni}^{\text{CF}} = \text{IPD}_{ai}.$$

We then use equation (19) to obtain the implied changes in bilateral trade costs. This exercise changes only trade costs involving country n, holding the political configuration of the rest of the world fixed. We report the net welfare advantage of Western over Eastern alignment, $\Delta W_n = \hat{W}_n^{\text{West}} - \hat{W}_n^{\text{East}}$, in percentage points of real income.

Results. Figure 11 studies whether economic incentives aligned with geopolitical bloc configurations during the Cold War. For each country, we show political alignment on the vertical axis, where countries above zero voted closer to the US in the UN General Assembly and those below zero closer to the USSR, and economic incentives on the horizontal axis, where $\Delta W_n > 0$ means Western alignment would have delivered a larger welfare gain. Countries in the upper-right and lower-left quadrants faced consistent incentives. Those in the lower-right faced an economic pull toward the West despite political alignment with the East, and those in the upper-left the reverse. Three findings stand out.

First, for the vast majority of countries, the welfare gain from alignment with the West exceed alignment with the Communist bloc. Non-Aligned countries gain an average of 0.52 percentage points more from full Western than Eastern alignment, while Western-bloc countries gain 0.39 percentage points more from maintaining a Western rather than adopting an Eastern orientation. This asymmetry reflects the greater economic weight of the Western bloc. Adopting the US alignment profile reduces geopolitical trade barriers toward a larger and economically more important set of trading partners.

Figure 11: Political Alignment and Economic Incentives during the Cold War

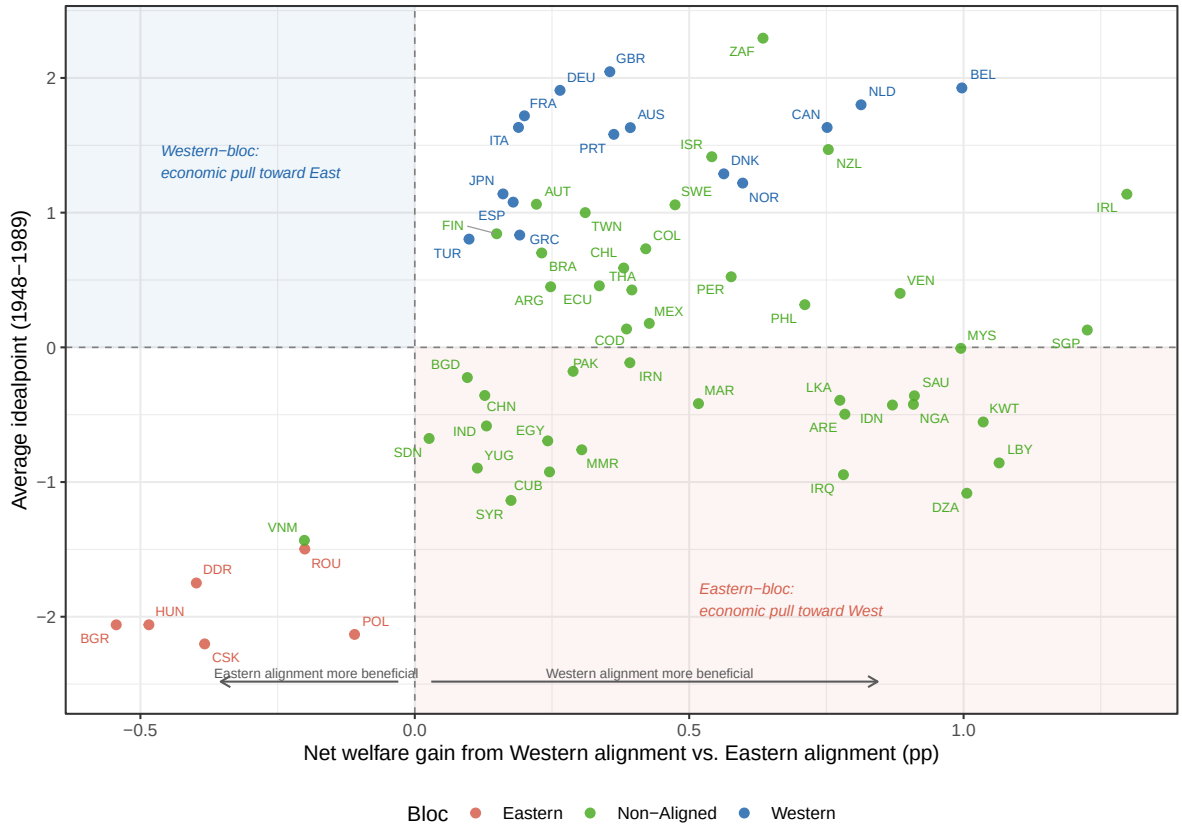

Note: Each point is one country's average over 1948–1989. The horizontal axis plots $\Delta W_n = \hat{W}_n^{\text{West}} - \hat{W}_n^{\text{East}}$, the difference in welfare gain (percentage points of real income) between full alignment with the US and full alignment with the USSR, computed from equations (20)–(22) with $\varepsilon = 5.03$ and $\beta = -0.154$. Bloc classification: Western = NATO members plus Japan and Australia, Eastern = Warsaw Pact members, Non-Aligned = all others. Idealpoints from Bailey et al. (2017).

Second, the pattern is sharply asymmetric. Not a single Western-bloc country would have gained more from adopting the Soviet rather than the US alignment profile. Political and economic incentives therefore pointed in the same direction for every Western-bloc member, making the bloc economically self-reinforcing in this sense.

Third, the pattern is more mixed among countries politically closer to the Soviet Union. For formal Eastern-bloc members, political and economic incentives coincided perfectly. The tension instead arises primarily outside the formal bloc. Of the 30 countries with below-zero ideal points, 22, most of them Non-Aligned countries that voted somewhat closer to the Soviet Union in the UN General Assembly, would nevertheless have gained more from Western alignment.

Taken together, the two exercises provide a consistent picture of the economic geography of the Cold War. The aggregate counterfactual shows that geopolitical barriers imposed their largest welfare costs on Non-Aligned countries. The alignment counterfactual provides a complementary perspective. Many of these countries faced a substantial economic pull toward Western markets despite remaining politically outside the Western bloc, and in many cases voting closer to the Soviet

Union. By contrast, political and economic incentives were mutually reinforcing for all Western-bloc members and for most Eastern-bloc members. The tension between geopolitical alignment and economic incentives was therefore concentrated primarily in the Non-Aligned world.

6.4 Contemporary Fragmentation: US-China Rivalry

In recent years, political tensions between the United States and China have risen sharply, prompting a growing literature on whether the world economy is drifting toward a new bipolarity (Aiyar et al., 2024; Gopinath et al., 2025). Recent studies also document declining trade between politically distant countries (Bosone and Stamato, 2024; de Soyres et al., 2025). We apply the same framework to the contemporary US-China rivalry to ask whether political and economic incentives remain aligned and how the resulting pattern compares with the Cold War.

Design. We use 2019 data as the economic baseline, the last year with complete bilateral data that is not confounded by the Covid Pandemic, and 2022 ideal points to capture the more recent geopolitical configuration following Russia's invasion of Ukraine. Following Aiyar et al. (2024), we classify countries by ideal point proximity, the 25% closest to the US form the US bloc, the 25% closest to China form the China bloc, and the remainder are Non-Aligned. The US anchor ideal point is 2.12 and the China anchor is -1.41. As in the Cold War exercise, under alignment with anchor country $a \in \{\text{US}, \text{China}\}$, we replace country n's political distance to every trading partner i with the corresponding distance between anchor a and country i, so that $\text{IPD}_{ni}^{\text{CF}} = \text{IPD}_{ai}$. All structural parameters are identical to the Cold War exercise, allowing for direct comparison.

Results. Figure 12 plots ΔW_n against each country's 2022 alignment. Three findings stand out.

First, aligning with one bloc rather than the other is substantially less consequential for welfare than during the Cold War. US-bloc countries gain an average of 0.16 percentage points more from US than from China alignment, compared with a 0.39 percentage point advantage of Western alignment for Western-bloc countries during the Cold War. Among Non-Aligned countries, the corresponding difference is only 0.08 percentage points today, compared with 0.52 percentage points during the Cold War. The economic divide between the two alignment choices is therefore considerably smaller today than under Cold War bipolarity.

Second, the contemporary US bloc is no longer uniformly economically self-reinforcing. During the Cold War, every Western-bloc country would have gained more from Western than from Eastern alignment. Today, 6 of 27 US-bloc countries would gain more from adopting China's alignment profile than that of the United States, including Australia and Japan who both have extensive trade ties to China. Political alignment and economic incentives therefore no longer point in the same direction for every member of the US bloc.

Third, economic incentives within the China bloc are evenly divided. China alignment is, on average, slightly more beneficial for China-bloc members, but this masks substantial heterogeneity. 7 of 14 members would gain more from US alignment, while the other 7 would gain more from

Figure 12: Political Alignment and Economic Incentives under Contemporary Fragmentation

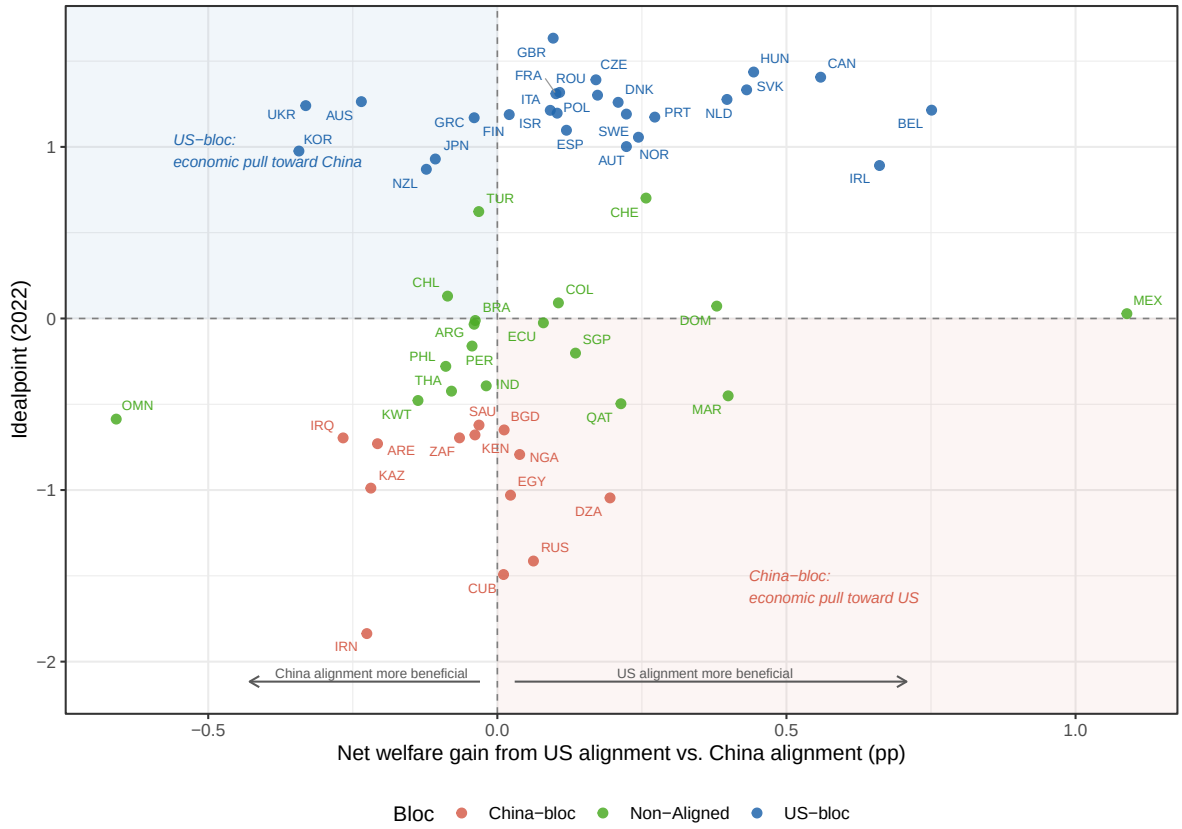

Note: Each point is one country. The horizontal axis plots $\Delta W_n = \hat{W}_n^{\text{US}} - \hat{W}_n^{\text{China}}$, the difference in welfare gain (percentage points of real income) between full alignment with the United States and full alignment with China. Baseline trade flows and GDP are from 2019, idealpoints and bloc classification from 2022, computed from equations (20)–(22) with $\varepsilon = 5.03$ and $\beta = -0.154$. Bloc classification follows Aiyar et al. (2024), top-25% idealpoint proximity to the US = US-bloc, top-25% to China = China-bloc, remainder = Non-Aligned. Idealpoints from Bailey et al. (2017).

China alignment. This contrasts sharply with the Cold War Eastern bloc, where all seven members were economically better off maintaining an Eastern orientation. Political and economic alignment is therefore considerably weaker within today's China bloc than it was within the Soviet bloc.

7 Conclusion

Geopolitical rivalry is increasingly shaping the organization of the world economy, yet the economic consequences of political alignment are difficult to identify because political and economic relationships are jointly determined. This paper uses the Cold War, and the abrupt political realignments produced by coups, to break that joint determination and estimate the causal effect of geopolitical alignment on international trade. Two complementary empirical strategies, one built on the size of a coup's realignment and the other on whether an attempted coup succeeded, point to the same conclusion. A one-standard-deviation increase in political distance reduces trade by approximately

15 percent in the medium run, equivalent to an ad valorem tariff of 3.1 percent on average and nearly four times as large for relationships involving the Cold War superpowers. These estimates help move the study of politics and trade from correlation to identification, and supply a trade-cost parameter that can discipline the growing theoretical literature on geopolitical fragmentation.

The effects of political distance extend beyond any single instrument of economic statecraft. Countries that become more politically distant receive less foreign assistance and are more likely to face new economic sanctions. The deterioration in US-Turkey relations following the 1964 Johnson Letter provides complementary evidence that trade can also respond to political estrangement in the absence of observable changes in formal economic policy. Taken together, these results suggest that geopolitical relationships shape economic integration through both government policy and broader responses to changes in political relations.

A structural gravity model translates these trade effects into welfare consequences. The burden of Cold War fragmentation fell unevenly. Non-Aligned countries experienced the largest welfare losses, reflecting in particular the economic cost of political distance from the large Western market. The same framework also reveals a distinction between the economic structure of Cold War rivalry and that of the emerging US-China rivalry. During the Cold War, political and economic incentives were largely mutually reinforcing within the two formal blocs, while tensions between political positioning and economic interests were concentrated in the Non-Aligned world. Today, choosing between geopolitical blocs is less consequential for welfare, but political and economic incentives are also less closely aligned. Both the US and China blocs contain countries whose existing trade relationships create an economic pull toward the opposing side.

References

Abadie, A., A. Diamond, and J. Hainmueller (2010): "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program," *Journal of the American Statistical Association*, 105, 493–505.

Acemoglu, D., S. Naidu, P. Restrepo, and J. A. Robinson (2019): "Democracy does cause growth," *Journal of political economy*, 127, 47–100.

Aiyar, S., D. Malacrino, and A. F. Presbitero (2024): "Investing in friends: The role of geopolitical alignment in FDI flows," *European Journal of Political Economy*, 83, 102508.

Anderson, J. E. and E. van Wincoop (2003): "Gravity with gravitas: A solution to the border puzzle," *American Economic Review*, 93, 170–192.

Are, J. (2020): "Obasanjo: How I Escaped Being Killed During 1976 Coup," *TheCable*.

Arkolakis, C., A. Costinot, and A. Rodríguez-Clare (2012): "New trade models, same old gains?" *American Economic Review*, 102, 94–130.

Ayyagari, M., J. Gao, and P. Ma (2025): "Partisan friendshoring," *Working Paper*.

Bailey, M. A., A. Strezhnev, and E. Voeten (2017): "Estimating dynamic state preferences from United Nations voting data," *Journal of Conflict Resolution*, 61, 430–456.

Baldwin, D. A. (1985): *Economic Statecraft: New Edition*, Princeton, NJ: Princeton University Press, new edition ed.

Baur, A., F. Dorn, L. Flach, and C. Fuest (2025): "Geoeconomic Fragmentation and the Role of Non-Aligned Countries," *Economics Letters*, 255, 112533.

Bazzi, S. and C. Blattman (2014): "Economic shocks and conflict: Evidence from commodity prices," *American Economic Journal: Macroeconomics*, 6, 1–38.

Becko, J. S., G. M. Grossman, and E. Helpman (2025): "Optimal Tariffs with Geopolitical Alignment," NBER Working Paper 34108, National Bureau of Economic Research.

Becko, J. S. and D. O'Connor (2024): "Strategic (dis) integration," *Princeton University*.

Bennett, D. L., C. Bjørnskov, and S. F. Gohmann (2021): "Coups, regime transitions, and institutional consequences," *Journal of Comparative Economics*, 49, 627–643.

Berger, D., A. Corvalan, W. Easterly, and S. Satyanath (2013a): "Do superpower interventions have short and long term consequences for democracy?" *Journal of Comparative Economics*, 41, 22–34.

Berger, D., W. Easterly, N. Nunn, and S. Satyanath (2013b): "Commercial imperialism? Political influence and trade during the Cold War," *American Economic Review*, 103, 863–896.

Bjørnskov, C. and M. Rode (2020): "Regime types and regime change: A new dataset on democracy, coups, and political institutions," *The Review of International Organizations*, 15, 531–551.

Bolhuis, M. A., J. Chen, and B. R. Kett (2023): "Fragmentation in Global Trade: Accounting for Commodities," IMF Working Paper 2023/073, International Monetary Fund.

Bonadio, B., Z. Huo, E. Kang, A. A. Levchenko, N. Pandalai-Nayar, H. Toma, and P. Topalova (2025): "Playing with blocs: Quantifying decoupling," *Journal of International Economics*, 104204.

Borchert, I., M. Larch, S. Shikher, and Y. V. Yotov (2021): "The International Trade and Production Database for Estimation (ITPD-E)," *International Economics*, 166, 140–166.

Bosone, C. and G. Stamato (2024): "Beyond Borders: How Geopolitics Is Reshaping Trade," Tech. Rep. 2960, European Central Bank.

Broner, F., A. Martin, J. Meyer, and C. Trebesch (2025): "Hegemonic globalization," Tech. rep., Kiel Working Paper.

Camboni, M. and M. Porcellacchia (2024): "Spheres of Influence," Tech. rep., available at SSRN.

Campos, R. G., B. Heid, and J. Timini (2024): "The economic consequences of geopolitical fragmentation: Evidence from the Cold War," .

Cattaneo, M. D., Y. Feng, and R. Titiunik (2021): "Prediction Intervals for Synthetic Control Methods," *Journal of the American Statistical Association*, 116, 1865–1880.

Cemlyn-Jones, B. (1981): "King Orders Army to Crush Coup," *The Guardian*, accessed July 7, 2025.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019): "The effect of minimum wages on low-wage jobs," *The Quarterly Journal of Economics*, 134, 1405–1454.

Cerdeiro, D. A., J. Eugster, R. C. Mano, D. Muir, and S. J. Peiris (2021): "Sizing Up the Effects of Technological Decoupling," IMF Working Paper 2021/069, International Monetary Fund.

Cheibub, J. A., J. Gandhi, and J. R. Vreeland (2010): "Democracy and dictatorship revisited," *Public Choice*, 143, 67–101.

Chin, J. J. and J. Kirkpatrick (2023): "African coups in the COVID-19 era: A current history," *Frontiers in Political Science*, 5, 1077945.

CIA (1955-86): "Communist Aid to Non-Communist Less Developed Countries," Tech. rep., Central Intelligence Agency, Langley, VA, declassified report series.

Clayton, C., M. Maggiori, and J. Schreger (2025): "A Theory of Economic Coercion and Fragmentation," .

——— (2026a): "A framework for geoeconomics," *Econometrica*, 94, 105–136.

——— (2026b): "The Great Game: A Model of Geoeconomic Competition," Tech. rep., Center for Open Science.

Conte, M., P. Cotterlaz, T. Mayer, et al. (2022): *The CEPII gravity database*, vol. 5, CEPII Paris.

Costinot, A. and A. Rodríguez-Clare (2014): "Trade theory with numbers: Quantifying the consequences of globalization," in *Handbook of international economics*, Elsevier, vol. 4, 197–261.

David, S. R. (1991): *Choosing Sides: Alignment and Realignment in the Third World*, Baltimore: Johns Hopkins University Press.

de Soyres, F., K. Richards, A. M. Santacreu, and F. Airaudo (2025): "Measuring geopolitical fragmentation: Implications for trade, financial flows, and economic policy," *International Finance Discussion Paper*.

Dekle, R., J. Eaton, and S. Kortum (2008): "Global rebalancing with gravity: Measuring the burden of adjustment," Tech. rep., National Bureau of Economic Research.

Derpanopoulos, G., E. Frantz, B. Geddes, and J. Wright (2016): "Are Coups Good for Democracy?" *Research & Politics*, 3.

Dixon, W. J. and B. E. Moon (1993): "Political similarity and American foreign trade patterns," *Political Research Quarterly*, 46, 5–25.

Dorsch, M. T. and P. Maarek (2018): "Rent extraction, revolutionary threat, and coups in non-democracies," *Journal of Comparative Economics*, 46, 1082–1103.

Eaton, J. and S. Kortum (2002): "Technology, geography, and trade," *Econometrica*, 70, 1741–1779.

Fan, T., M. Wo, and W. Xiang (2025a): "Geopolitical Barriers to Globalization," *arXiv preprint arXiv:2509.12084*.

——— (2025b): "Geopolitical Barriers to Globalization," ArXiv:2509.12084.

Feenstra, R. C., R. E. Lipsey, H. Deng, A. C. Ma, and H. Mo (2005): "World Trade Flows: 1962–2000," Tech. Rep. 11040, National Bureau of Economic Research.

Felbermayr, G., A. Kirilakha, C. Syropoulos, E. Yalcin, and Y. V. Yotov (2020): "The global sanctions data base," *European Economic Review*, 129, 103561.

Feng, L., Q. Huang, Z. Li, and C. M. Meissner (2026): "The "Peace Dividend" of International Trade: A New Empirical Approach," Tech. rep.

Fernández-Villaverde, J., T. Mineyama, and D. Song (2024): "Are we fragmented yet? Measuring geopolitical fragmentation and its causal effect," Tech. rep., National Bureau of Economic Research.

Fouquin, M. and J. Hugot (2016): "Two Centuries of Bilateral Trade and Gravity Data: 1827–2014," CEPII Working Paper 2016-14, CEPII.

Fuchs, A. and N.-H. Klann (2013): "Paying a visit: The Dalai Lama effect on international trade," *Journal of International Economics*, 91, 164–177.

Gassebner, M., J. Gutmann, and S. Voigt (2016): "When to expect a coup d'état? An extreme bounds analysis of coup determinants," *Public Choice*, 169, 293–313.

Gibson, R. M., P. H. Wise, J. L. Dieleman, Y. V. Yotov, A. Kirilakha, G. L. Darmstadt, E. Bendavid, C. Syropoulos, M. Barry, and S. Bradley (2025): "The impact of aid sanctions on maternal and child mortality, 1990–2019: a panel analysis," *The Lancet Global Health*, 13, e820–e830.

Glick, R. and A. M. Taylor (2010): "Collateral damage: Trade disruption and the economic impact of war," *The Review of Economics and Statistics*, 92, 102–127.

Góes, C. and E. Bekkers (2022): "The Impact of Geopolitical Conflicts on Trade, Growth, and Innovation," WTO Staff Working Paper ERSD-2022-09, World Trade Organization.

Goodman-Bacon, A. (2021): "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 225, 254–277.

Gopinath, G. (2024): "How Policymakers Should Handle a Fragmenting World," *Foreign Policy*.

Gopinath, G., P.-O. Gourinchas, A. F. Presbitero, and P. Topalova (2025): "Changing global linkages: A new Cold War?" *Journal of International Economics*, 153, 104042.

Harmer, T. (2011): *Allende's Chile and the inter-American cold war*, Univ of North Carolina Press.

Hausmann, R., C. A. Hidalgo, S. Bustos, M. Coscia, and A. Simoes (2014): *The atlas of economic complexity: Mapping paths to prosperity*, Mit Press.

Head, K. and T. Mayer (2014): "Gravity equations: Workhorse, toolkit, and cookbook," in *Handbook of International Economics*, Elsevier, vol. 4, 131–195.

——— (2021): "The United States of Europe: A gravity model evaluation of the four freedoms," *Journal of Economic Perspectives*, 35, 23–48.

Hinz, J. (2023): "The ties that bind: geopolitical motivations for economic integration," *Review of World Economics*, 159, 51–100.

Hirschman, A. O. (1945): *National power and the structure of foreign trade*, vol. 105, Univ of California Press.

Horn, S., C. M. Reinhart, and C. Trebesch (2021): "China's overseas lending," *Journal of International Economics*, 133, 103539.

Javorcik, B., L. Kitzmueller, H. Schweiger, and M. A. Yildirim (2024): "Economic costs of friendshoring," *The World Economy*, 47, 2871–2908.

Jones, B. F. and B. A. Olken (2005): "Do leaders matter? National leadership and growth since World War II," *The Quarterly Journal of Economics*, 120, 835–864.

——— (2009): "Hit or miss? The effect of assassinations on institutions and war," *American Economic Journal: Macroeconomics*, 1, 55–87.

Kandeh, J. D. (2004): *Coups from Below: Armed Subalterns and State Power in West Africa*, New York: Palgrave Macmillan.

Kempf, E., M. Luo, and M. Tsoutsoura (2025): "The Political Economy of Firm Networks: CEO Ideology and Global Trade," Tech. rep., National Bureau of Economic Research.

Khalil, M., D. Osten, and F. Strobel (2025): "Trade dynamics under geopolitical risk," .

Kleinman, B., E. Liu, and S. J. Redding (2024): "International friends and enemies," *American Economic Journal: Macroeconomics*, 16, 350–385.

Kooi, O. (2024): "Power and resilience: An economic approach to national security policy," Tech. rep., working paper.

Krugman, P. (1980): "Scale economies, product differentiation, and the pattern of trade," *American Economic Review*, 70, 950–959.

Lagarde, C. (2023): "Central Banks in a Fragmenting World," European Central Bank, speech at the Council on Foreign Relations, New York, April 17.

Lall, S. (2000): "The Technological structure and performance of developing country manufactured exports, 1985-98," *Oxford development studies*, 28, 337–369.

Lehoucq, F. (2021): "Military Coups d'État and Their Causes," in *Oxford Research Encyclopedia of Politics*, ed. by E. Hannah, Oxford University Press.

Levchenko, A. A., N. Pandalai-Nayar, and H. Young (2026): "Geopolitical Payoffs," .

Liu, E., S. J. Redding, B. Smith-Worthington, and D. Y. Yang (2026): "A Network Approach to Geopolitics," .

Liu, E. and D. Y. Yang (2025): "International power," Tech. rep., National Bureau of Economic Research.

Madsen, E. and A. Prat (2026): "Competing Powers," .

Martin, P., T. Mayer, and M. Thoenig (2008): "Make Trade Not War?" *The Review of Economic Studies*, 75, 865–900.

Marx, B., V. Pons, and V. Rollet (2025): "Electoral turnovers," *Review of Economic Studies*, 92, 3306–3339.

Meyer, T. and N. Wesseler (2026): "Hegemonic Competition with Carrots and Sticks," Working Paper.

Meyersson, E. (2016): "Political man on horseback: coups and development," *Stockholm Institute for Transition Economics (SITE)*, 5, 1–27.

Nagengast, A. J. and Y. V. Yotov (2025): "Staggered difference-in-differences in gravity settings: Revisiting the effects of trade agreements," *American Economic Journal: Applied Economics*, 17, 271–296.

Neri-Laine, M. (2026): "Sovereign Gravity: The Military Alliance Effect on Trade," Tech. rep.

Niemann, K. (2025): "Political distance and international trade," *Available at SSRN 5698582*.

Nunn, N. (2007): "Relationship-specificity, incomplete contracts, and the pattern of trade," *Quarterly Journal of Economics*, 122, 569–600.

Ofosu-Appiah, L. H. (1972): *The Life of Lt.-General E. K. Kotoka, the Hero of Ghana's 24th February Revolution*, Accra: Waterville Pub. House.

Papaioannou, E. and G. Siourounis (2008): "Democratisation and growth," *The Economic Journal*, 118, 1520–1551.

Pettersson, J. and L. Johansson (2013): "Aid, aid for trade, and bilateral trade: An empirical study," *The Journal of International Trade & Economic Development*, 22, 866–894.

Plummer, M. G. (2023): "Estimating the Costs of Geoeconomic Fragmentation to the Global Economy," in *Geoeconomic Fragmentation: The Economic Risks from a Fractured World Economy*, ed. by S. Aiyar, A. F. Presbitero, and M. Ruta, Paris and London: CEPR Press.

Pollins, B. M. (1989): "Does trade still follow the flag?" *American Political Science Review*, 83, 465–480.

Rauch, J. E. (1999): "Networks versus markets in international trade," *Journal of International Economics*, 48, 7–35.

Seong, J., O. White, S. Gupta, A. Padhi, B. Chhibber, J. Condon, and T. Devesa (2025): "A new trade paradigm: How shifts in trade corridors could affect business," Research report, McKinsey & Company, accessed July 7, 2025.

Tabellini, M. and G. Magistretti (2024): "Economic integration and the transmission of democracy," *Review of Economic Studies*.

Thoenig, M. (2024): "Trade in the shadow of war: A quantitative toolkit for geoeconomics," in *Handbook of the Economics of Conflict*, Elsevier, vol. 1, 325–380.

Westad, O. A. (2005): *The Global Cold War: Third World Interventions and the Making of Our Times*, Cambridge University Press.

A Appendix

A.1 Summary of Assumptions in Recent Trade Fragmentation Models

A rapidly growing literature uses structural trade models to quantify the potential economic costs of geopolitical fragmentation. The assumed bloc configurations vary considerably across papers. Some studies adopt a bipolar structure with clearly defined China-aligned and US-aligned blocs, while others allow for more complex groupings or focus specifically on Russia-Ukraine fragmentation. The treatment of non-aligned or swing countries also differs, some papers assign every country to one bloc or another, while others model an intermediate group that maintains relationships with both sides.

Even more variation exists in the assumed magnitude of trade barriers. Assumptions range from complete prohibitions on trade in strategic sectors ($\tau \to \infty$) to relatively modest tariff increases of 7.5 to 15 percent. Some studies model economy-wide barriers, while others focus specifically on high-technology goods or critical inputs. The most extreme scenarios consider trade cost increases of 160 percentage points or tariffs of 60 percent on critical goods, while more conservative approaches assume 20 to 25 percent tariff equivalents.

Table A.1 summarizes the key assumptions from these studies. The diversity of approaches reflects uncertainty about future geopolitical developments, and underscores the value of historical evidence, of the kind this paper provides, in grounding these exercises.

Table A.1: Assumed Bloc Configurations and Trade Barriers in Recent Trade Fragmentation Models

	Assumed Blocs		**Assumed Policy Shock**
	Bloc A	**Bloc B**	
Cerdeiro et al. (2021)	China	OECD	$\tau \to \infty$ in high-tech sectors
Góes and Bekkers (2022)	China-bloc	US-bloc	$\tau \uparrow$ 160pp or $t \uparrow$ 32pp
Bolhuis et al. (2023)	China + Russia	US + EU	$\tau \to \infty$ in high-tech sectors
Plummer (2023)	China + allies	US + allies	$t = 7.5\%$ between blocs $t = 15\%$ on sensitive imports
Bonadio et al. (2025)	China-bloc	US-bloc	General decoupling, no single tariff figure
Javorcik et al. (2024)	Pro-Russia	Pro-Ukraine	$\tau \uparrow$ 20pp
Baur et al. (2025)	China	US + ROW	$t = 25\%$ between blocs doubling of NTBs
Seong et al. (2025)	China + Russia	Adv. Economies	$t = 60\%$ on critical goods $t = 20\%$ other goods

Note: t = tariffs, τ = trade costs, pp = percentage points, NTBs = non-tariff barriers.

A.2 LLM Classification of Coup Direction

This appendix reproduces the exact prompt used for the coup direction classification described in Section 3.2.2. For every coup event, the model (Claude Sonnet 5.5) is given the country, year, coup leader(s), and coup type, and is prompted with web search access as follows, with bracketed fields filled in per event:

```
You are a Cold War historian with access to web search. I will give you
a coup attempt and its leader(s)/plotters.

Country: [country]
Year: [year]
Coup leader(s)/plotters: [leader]
Coup type: [category]

[If successful:] This coup SUCCEEDED, so you can also use what the new
government actually did in office (who it aligned with diplomatically
and militarily) as evidence of the plotters' own ideology.

[If failed:] This coup FAILED, so there is no governance record. Rely
on the plotters' declared affiliations, backing, and documented
platform instead.

Use web search to find out about the PLOTTERS THEMSELVES: their
declared political ideology or party affiliation (e.g. socialist,
communist, Marxist, nationalist, conservative, monarchist), any
foreign backing, training, or patronage, and their Cold War alignment.

Classify their collective Cold War orientation, based only on who they
are, as exactly one of:
- pro_us: Plotters were anti-communist, right-wing, Western-aligned,
  or otherwise identifiably part of the US/Western bloc orientation
- pro_soviet: Plotters were communist, Marxist, left-wing, or
  otherwise identifiably part of the Soviet/Eastern bloc orientation
- unclear: Despite searching, not enough found about the plotters to
  classify them, OR the coup was driven by dynamics with no clear Cold
  War ideological content at all (e.g. purely ethnic, personal,
  palace, or succession disputes with no left-right valence)

IMPORTANT ON EVIDENCE: you do NOT need an explicit statement like
```

```
'backed by the USSR' or 'a CIA asset' to make a call. Reasonable
ideological inference is valid and expected. A plotter who is a
declared or avowed socialist, communist, or Marxist, or a documented
member of a communist/socialist party, should normally be classified
pro_soviet absent specific contrary evidence. A plotter who is a
declared anti-communist, closely tied to a conservative/business
elite, trained by or allied with Western militaries, or a
monarchist/traditionalist aligned with the West should normally be
classified pro_us. Weigh this kind of ideological and biographical
evidence the same way you would weigh direct evidence of foreign
sponsorship. Do not default to 'unclear' just because you lack a
document explicitly describing foreign backing, only use 'unclear'
when you genuinely cannot find enough, or the coup truly has no
ideological content either way.

Also provide:
- confidence: high (clear direct evidence, e.g. explicit party
  membership/declared ideology or documented foreign ties), medium
  (reasonable ideological inference from partial evidence), low
  (weak/circumstantial, mostly guesswork)
- sources_found: a brief note on what, if anything, your search
  actually turned up
- rationale: one to two sentences explaining the classification,
  citing what you found

Do NOT write a summary, findings list, or any other commentary after
you finish searching. Skip straight from your last search to the
final answer. Your very last message must contain ONLY a valid JSON
object in this exact format, no other text before or after it:
{"classification": "pro_us", "confidence": "high", "sources_found":
"...", "rationale": "..."}
```

The model is allowed up to three web searches per event. Of 152 coup events in the clean, non-contaminated sample, the model classifies 79 as pro-US, 35 as pro-Soviet, and 38 as unclear, which are dropped from the direction-dependent analysis.

A.3 LLM Classification of Foreign-Backed Coups

This appendix reports full detail on the large language model classification described in Section 4.3, used to identify coups with credible documented evidence of direct superpower backing. We classify all 546 coup events in the underlying Bjørnskov and Rode (2020) data, both successful and failed,

including the second coup attempt some country-years record in addition to the main one. For each event, the model (Claude Sonnet 4.6) is given the country, year, leader(s), coup type, and success or failure, and prompted as follows, with bracketed fields filled in per event.

```
You are a Cold War historian. I will give you a coup attempt. I need
you to assess whether there is CREDIBLE HISTORICAL EVIDENCE that a
foreign superpower (the United States or the Soviet Union) directly
assisted, orchestrated, funded, or otherwise materially backed THIS
SPECIFIC coup attempt, as distinct from merely having a general
relationship with the country or having supported earlier or later
governments there.

Country: [country]
Year: [year]
Coup leader(s): [leader]
Coup type: [category_txt]
[This attempt SUCCEEDED. / This attempt FAILED.]

CRITICAL: 'a foreign power backed this coup' is one of the most
common tropes in 20th-century political history and is frequently
alleged with no real evidentiary basis, as a domestic political
accusation, or as after-the-fact rumor. You must distinguish between:
(a) DOCUMENTED evidence: declassified intelligence records, memoirs
or testimony from participants/officials, or a strong consensus among
serious historians citing primary sources, specifically tying that
power's intelligence services, military, or diplomats to the
planning, funding, or execution of THIS coup;
(b) UNDOCUMENTED allegation: claims that circulate in journalism,
partisan political rhetoric, or popular belief, without the kind of
sourcing in (a).
Only classify as foreign_backed for case (a). If the strongest
evidence you know of is case (b), classify as no_evidence and say so
in the rationale, do not round a popular allegation up to a confident
classification.

Classify as exactly one of:
- foreign_backed: Credible documented evidence (per (a) above)
  directly ties a foreign power to this specific coup's planning,
  funding, or execution
- no_evidence: No credible documented evidence of direct foreign
```

```
  involvement in this specific coup, even if the country had a
  broader relationship with a superpower at other times, or even if
  popular allegations exist without solid sourcing
- unclear: Insufficient information about this coup to assess either
  way

If foreign_backed, also identify which power:
- power: "us", "soviet", "other" (a different foreign power
  entirely), or "none" (if classification is not foreign_backed)

Also provide:
- confidence: high (declassified/primary-source documentation or
  strong historian consensus), medium (credible secondary-source
  evidence but less definitive), or low (plausible but thin or
  contested evidence)
- rationale: one sentence citing the TYPE of evidence (e.g.
  declassified CIA/KGB records, participant testimony, historian
  consensus, or 'no credible sourcing beyond popular allegation')

Respond ONLY with a valid JSON object in this exact format, no other
text:
{"classification": "foreign_backed", "power": "us", "confidence":
"high", "rationale": "..."}
```

Table A.2 reports the full classification. Of the 546 events, 66 are classified foreign backed at medium or high confidence, 30 attributed to the United States, 4 to the Soviet Union, and 32 to another foreign power, outside the scope of the superpower rivalry this paper studies. No coup is classified foreign backed at low confidence only. Table A.3 lists each of the 34 US- and Soviet-backed events individually, and Figure A.1 maps their geographic distribution.

Table A.2: LLM Classification of All 546 Coup Events, by Power and Confidence

	High	Medium	Low	Total
Foreign-Backed: US	6	24	0	30
Foreign-Backed: Soviet	0	4	0	4
Foreign-Backed: Other Power	11	21	0	32
Foreign-Backed: Total	**17**	**49**	**0**	**66**
No Evidence	380	35	11	426
Unclear	0	0	54	54
Total	**397**	**84**	**65**	**546**

Note: $n = 546$ coup events, both successful and failed, including second-coup-slot events. No coup is classified foreign backed at low confidence, and no "no evidence"/"unclear" classification carries a power attribution.

Table A.3: Coups Classified Foreign Backed by a Superpower

Country	Year	Successful	Power	Confidence
Egypt	1952	Yes	US	Medium
Iran	1953	Yes	US	High
Guatemala	1954	Yes	US	High
Iraq	1958	No	US	Medium
Congo, Dem. Rep.	1960	Yes	US	High
France	1961	No	US	Medium
Iraq	1963	Yes	US	Medium
Vietnam	1963	No	US	High
Bolivia	1964	Yes	US	Medium
Brazil	1964	Yes	US	High
Laos	1964	Yes	US	Medium
Vietnam	1964	Yes	US	Medium
Congo, Dem. Rep.	1965	Yes	US	Medium
Ghana	1966	Yes	US	Medium
Indonesia	1967	Yes	US	Medium
Cambodia	1970	Yes	US	Medium
Bolivia	1971	Yes	US	Medium
Uganda	1971	Yes	US	Medium
Chile	1973	Yes	US	High
Uruguay	1973	Yes	US	Medium
Portugal	1975	No	US	Medium
Argentina	1976	Yes	US	Medium
Turkey	1980	Yes	US	Medium
Chad	1982	Yes	US	Medium
Libya	1984	No	US	Medium
Libya	1985	No	US	Medium
Panama	1989	No	US	Medium
Afghanistan	1992	Yes	US	Medium
Iraq	1995	No	US	Medium
Gambia, The	2014	No	US	Medium
Venezuela	1962	No	Soviet	Medium
Angola	1977	No	Soviet	Medium
Yemen	1978	Yes	Soviet	Medium
Poland	1981	Yes	Soviet	Medium

Note: The 34 events (of 546 total) classified foreign backed at medium or high confidence with the United States or the Soviet Union as the responsible power. Events classified foreign backed but attributed to another power are omitted, as is every event classified no evidence or unclear. Successful indicates whether the coup attempt itself succeeded, not whether the backing power's preferred outcome was achieved.

Figure A.1: Geographic Distribution of LLM-Classified Foreign-Backed Coups

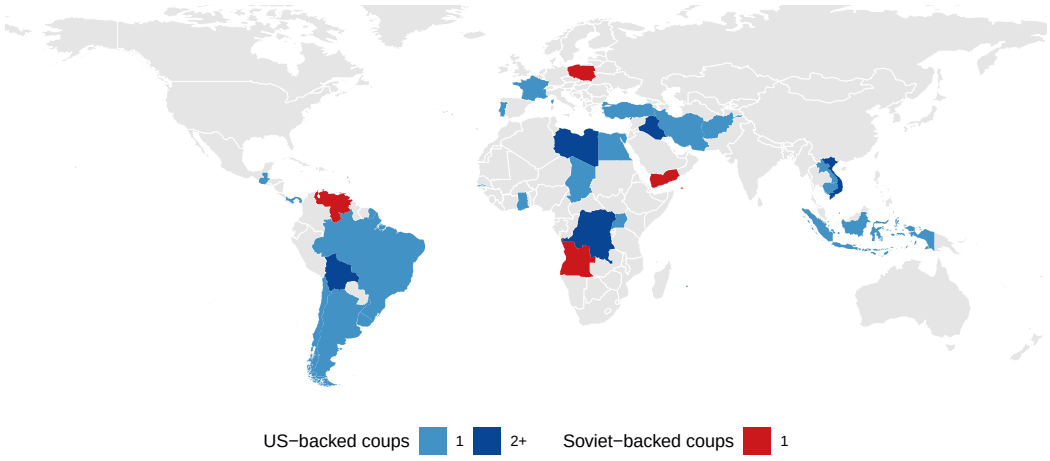

Note: Count of coups per country classified foreign backed at medium or high confidence, split by power. US-backed counts are bucketed to 1 versus 2+ (the darker shade), since Iraq is the only country with more than two US-backed coups (three). No country has more than one Soviet-backed coup, so the Soviet legend shows only a single shade. US-backed and Soviet-backed countries never overlap in this data, so both counts are shown on a single map (blue = US, red = Soviet).

A.4 Coup Case Studies

This appendix presents four historical case studies for coups described in Section 2.2, two successful and two failed. These examples illustrate how a change in political leadership can produce sharp shifts in geopolitical alignment, and how such shifts need not reflect gradual changes in underlying political conditions but can instead result from the actions of a small number of individuals.

A.4.1 Successful Coups

Mali 1968. On 19 November 1968, Lieutenant Moussa Traoré led a bloodless coup against President Modibo Keïta in Mali, deposing the country's founding father who had ruled a socialist one-party state since independence in 1960. The immediate causes were severe economic difficulties, compounded by a poor harvest in 1968 and growing popular dissatisfaction with Keïta's increasingly authoritarian rule. Traoré's new military government partially reversed Keïta's socialist policies and moved Mali closer to the Western sphere of influence, a shift clearly visible in Mali's UN voting record in Figure A.2a.

Crucially, this coup received no foreign support. Jacques Foccart, de Gaulle's closest adviser for African affairs, was caught completely off guard by the putsch, as Foccart's own memoirs reveal:[16]

[16] Jacques Foccart, *Mémoires*, vol. II, p. 446 (19 November 1968). Original French: "Mon Général, ... ce qui, à mes yeux, signifie que le coup est fait par des modérés en réaction contre la position de plus en plus prochinoise et

19 November 1968.

[De Gaulle:] *"But how did this happen? I have seen no telegram, I have nothing!"*

[Foccart:] *"General, they have just arrived, here they are."*

[De Gaulle:] *"They fought?"*

[Foccart:] *"Yes, they even brought in armored vehicles. [...] In my view, this means the coup was carried out by moderates, in reaction against the increasingly pro-Chinese and pro-Soviet position of Madeira Keïta. The affair is being led by lieutenants and captains. From the first indications I have, these are people trained at the military school of Fréjus."*

20 November 1968.

"I confirm that Modibo has been arrested, as well as Ousmane Bâ and Madeira Keïta. I make him read the proclamations of the military, which all go in a favorable direction to the Western world. [...] Among all these subordinate officers who somehow make up the new government, none have been formed in the East."

Burkina Faso 1983. On 4 August 1983, Captain Thomas Sankara came to power in Upper Volta (today Burkina Faso) through a coup led by his close ally Captain Blaise Compaoré, who commanded 250 paratroopers in a march on the capital and deposed President Jean-Baptiste Ouédraogo (Kandeh, 2004). The trigger was Ouédraogo's arrest of Sankara in May 1983, after tensions within the ruling military council over ideological direction had grown irreconcilable. Sankara established a National Council of the Revolution and immediately pursued a radical pan-Africanist and anti-imperialist agenda, renaming the country Burkina Faso, nationalizing land, rejecting foreign debt, and distancing the country from its former colonial power France. The shift in geopolitical alignment was sharp and persistent, as reflected in Burkina Faso's UN voting record in Figure A.2b.

Figure A.2: Successful Coups

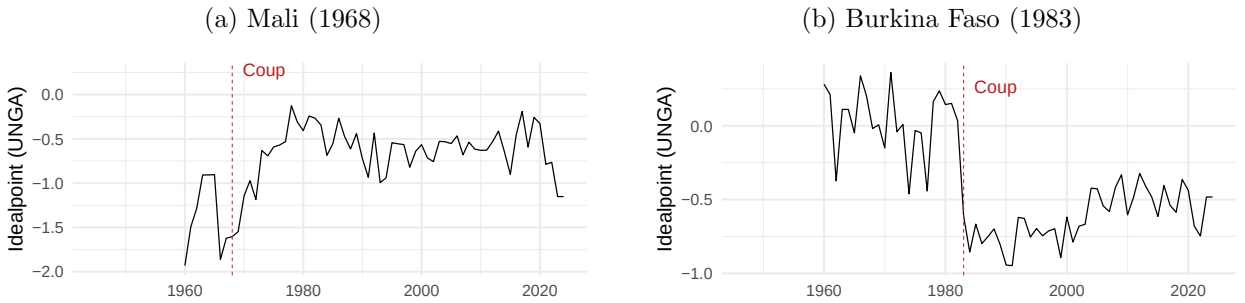

Notes: The figure plots the UN ideal point of Mali (panel a) and Burkina Faso (panel b) over time. The dashed vertical line marks the year of the coup, 1968 for Mali and 1983 for Burkina Faso. A higher ideal point indicates closer alignment with the United States.

prosoviétique ... de Madeira Keïta. L'affaire est menée par des lieutenants et des capitaines. D'après les premières indications que j'ai, ce sont des gens formés à l'école militaire de Fréjus."

A.4.2 Failed Coups

Spain 1981. On 23 February 1981, Lieutenant-Colonel Antonio Tejero led 200 armed Civil Guard officers into the Congress of Deputies in Madrid during the parliamentary vote to confirm Leopoldo Calvo-Sotelo as Prime Minister, taking all deputies and cabinet ministers hostage. Simultaneously, Captain General Jaime Miláns del Bosch deployed tanks onto the streets of Valencia and declared a state of emergency, hoping to prompt other military regions to join the uprising. The plotters had structured their plan around the assumption that King Juan Carlos I would either endorse the coup or remain silent. On February 24, the King appeared on national television in military uniform, explicitly denounced the coup, and ordered the armed forces to uphold the constitutional order (Cemlyn-Jones, 1981). Isolated and without broader military support, the conspirators surrendered that morning. Spain's UN voting record shows no discernible shift in geopolitical alignment around the coup, as seen in Figure A.3a.

Ghana 1967. On 17 April 1967, three junior officers of the Ghana Armed Forces (Lieutenant Samuel Arthur, Lieutenant Moses Yeboah, and Second Lieutenant Osei-Poku) led 119 soldiers in an attempted coup against the National Liberation Council (NLC). The operation, code-named "Guitar Boy", seized the national broadcasting house in Accra and broadcast an announcement of the takeover. However, the coup unravelled almost immediately. Lieutenant Arthur killed a captain who refused to hand over the keys to the ammunition depot. Lieutenant Yeboah shot and killed General Emmanuel Kotoka, Ghana's Chief of the Defence Staff, at the airport. Loyal units quickly moved to suppress the rebellion, and the coup collapsed within hours. The failure stemmed not from foreign intervention but from a thin operational base. Arthur later admitted his primary motivation was personal ambition. He wanted to be the first subaltern in Africa to stage a successful coup (Ofosu-Appiah, 1972). Ghana's geopolitical alignment remained largely unchanged, as seen in Figure A.3b.

Figure A.3: Failed Coups

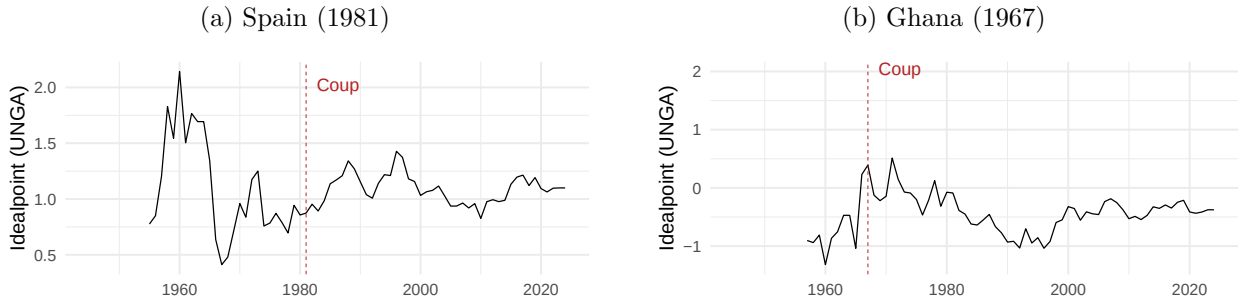

Notes: The figure plots the UN ideal point of Spain (panel a) and Ghana (panel b) over time. The dashed vertical line marks the year of the coup, 1981 for Spain and 1967 for Ghana. A higher ideal point indicates closer alignment with the United States.

A.5 Idealpoint Swings Around Coups Are Larger Than Ordinary Year-to-Year Movement

Figure 3 shows that successful coups coincide with sizable shifts in a country's UNGA idealpoint, but a mean absolute shift could in principle still be within the range of ordinary year-to-year noise in the series. Table A.4 checks this directly, regressing the same two-year-window idealpoint change used in the figure, idealpoint one year after minus one year before, in absolute value, on an indicator for a successful coup, for every country-year in the 1955–1985 window where both neighboring years are observed, not only coup years.

Column 1 shows that the average country-year sees an absolute idealpoint swing of 0.275, while a successful-coup country-year sees one 0.144 larger, about 52 percent bigger, and highly significant. Column 2 adds country and year fixed effects, so the comparison is no longer across different countries or eras but within the same country relative to its own typical volatility, net of any common shock in that particular year. The coefficient barely moves, to 0.130, and remains significant at the 1 percent level. The larger swings around coups are therefore not an artifact of a handful of chronically volatile countries or turbulent periods driving the comparison, they reflect something systematically different about the coup years themselves.

Table A.4: Idealpoint Swings Are Larger in Successful-Coup Years

Dependent Variable:	|Δ Idealpoint|	
	All Country-Years	
Model:	(1)	(2)
Variables		
Constant	0.275***	
	(0.006)	
Successful Coup	0.144***	0.130***
	(0.039)	(0.038)
Fixed-effects		
Country FE		Yes
Year FE		Yes
Fit statistics		
Observations	3,517	3,516
R^2	0.00996	0.21459

Note: Dependent variable is the absolute value of the idealpoint change from one year before to one year after, for every country-year in 1955–1985 with both neighboring years observed. Successful Coup is an indicator for a country-year with a successful coup, matching the sample underlying Figure 3. Standard errors, in parentheses, are clustered by country. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

Table A.5 reports the same comparison underlying Figure 4, at event time $k = 1$, splitting

out failed attempts as their own regressor rather than leaving them in the comparison group, and testing directly whether the successful-coup and failed-coup coefficients differ from each other. The successful-coup coefficient is significant at the 1 percent level. The failed-coup coefficient is not distinguishable from zero. The difference between the two, reported in the bottom row, is itself statistically significant at the 1 percent level.

Table A.5: Idealpoint Swings by Coup Outcome

	Estimate	Std. Error	p-value
Successful coup	0.143***	0.037	<0.001
Failed coup	0.030	0.025	0.247
Successful − Failed	0.113***	0.042	0.007

Note: Dependent variable is the absolute value of the idealpoint change from one year before to one year after, for every country-year in 1946–1989, with country and year fixed effects throughout. Successful Coup and Failed Coup are indicators for a country-year with a successful or failed attempt, respectively. Successful − Failed tests equality of the two coefficients. Standard errors, in parentheses, are clustered by country and year. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

A.6 Stacking and the Contamination Filter

Section 3.2.3 addresses two related problems that arise from pooling every coup into a single regression. A dyad-by-event fixed effect prevents one coup's post-period from serving as an invalid control for another, and a contamination filter drops any coup that is itself contaminated by a second coup in the same country within ten years. Table A.6 checks whether either choice drives the result, by removing one safeguard at a time from the baseline specification. Replacing the dyad-by-event fixed effect with a single dyad fixed effect, holding the contamination filter fixed, lowers the coefficient from -0.154 to -0.125, still negative and significant. Dropping the contamination filter instead, keeping the stacked design, lowers it further to -0.091, again significant. Stacking and the contamination filter are therefore not required individually to find a qualitatively similar effect.

Table A.6: Baseline Result Without Stacking or the Contamination Filter

Dependent Variable:	Bilateral Trade Flows		
	Stacked		Pooled
	Clean	Unclean	
Model:	(1)	(2)	(3)
Variables			
Post × ΔIPD	-0.154***	-0.091***	-0.125***
	(0.039)	(0.021)	(0.037)
Fixed-effects			
Dyad × Coup Event FE	Yes	Yes	
Exporter-Year FE	Yes	Yes	Yes
Importer-Year FE	Yes	Yes	Yes
Dyad FE			Yes
Fit statistics			
Observations	147,504	315,576	137,059
Pseudo R^2	0.97555	0.97430	0.97531

Note: Column 1 reproduces the pooled Strategy 1 coefficient from Table 2. Column 2 keeps the stacked, dyad-by-event design but uses every successful coup from the unpacked, unfiltered event list in place of the contamination-filtered sample. Column 3 keeps the contamination filter but replaces the dyad-by-event fixed effect with a single dyad fixed effect, collapsing each dyad's overlapping event windows to one observation per calendar year at the nearest coup. All columns use PPML with exporter-year and importer-year fixed effects. Standard errors, in parentheses, are clustered by dyad. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

A.7 Pooled Test of the Superpower-Bloc Difference

Table 2 estimates the coup-identified DiD separately for neutral, bloc, and superpower partners. This lets us compare point estimates across groups, but it does not test whether they differ. Table A.7 instead pools all partners into one regression and interacts the post-coup dose with group indicators, nesting the superpower group inside bloc since every superpower dyad is itself a bloc dyad. The interaction between the dose and the superpower indicator is negative and significant, $p = 0.001$, confirming that the superpower effect is statistically distinguishable from the rest of the bloc sample rather than only larger in point estimate.

Table A.7: Pooled Difference-in-Differences with Group Interactions

Dependent Variable:	Bilateral Trade Flows
	Pooled, with Interactions
Model:	(1)
Variables	
Post × ΔIPD	-0.099
	(0.070)
Post × ΔIPD × Bloc	0.038
	(0.077)
Post × ΔIPD × Superpower	-0.218***
	(0.067)
Fixed-effects	
Dyad × Coup Event FE	Yes
Exporter-Year FE	Yes
Importer-Year FE	Yes
Fit statistics	
Observations	147,504
Pseudo R^2	0.97559

Note: Single pooled regression on the stacked event panel of successful coups, with dyad-by-event, exporter-year, and importer-year fixed effects. Bloc nests the superpower indicator, since every superpower dyad is also a bloc dyad, so the coefficient on Post × ΔIPD × Superpower tests whether the superpower effect differs from the rest of the bloc sample, not from the neutral baseline directly. Standard errors, in parentheses, are clustered by dyad. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

A.8 Restricting the Gravity Comparison to the Coup Sample

Table 2 compares the coup-identified DiD to a naive gravity regression estimated on the full 1948–1989 trade panel, so the two panels differ in both identification strategy and sample. To check how much of the gap between them is sample composition rather than identification, Table A.8 re-estimates the gravity regression restricted to the same dyad-years the DiD stacks actually use.

Restricting the sample pulls the gravity coefficient toward the DiD estimate for the pooled, neutral, and bloc columns, consistent with part of the gap in the main text reflecting which dyad-years each specification looks at rather than only the identification strategy. For superpower partners, however, restricting the sample pushes the gravity coefficient the other way, from -0.447 to an insignificant -0.257, while the coup-identified estimate remains large and significant at -0.546. Sample composition therefore cannot account for the superpower result, if anything it sharpens the contrast between the two approaches for exactly the partner type where the paper's heterogeneity result is strongest.

Table A.8: Gravity Regression Restricted to the Coup-Identified Sample

Dependent Variable:	Bilateral Trade Flows			
	All	Neutral	Bloc	Superpower
Model:	(1)	(2)	(3)	(4)
Variables				
IPD_{t-1}	-0.125***	-0.164***	-0.138***	-0.257
	(0.022)	(0.046)	(0.030)	(0.158)
Fixed-effects				
Dyad FE	Yes	Yes	Yes	Yes
Exporter-Year FE	Yes	Yes	Yes	Yes
Importer-Year FE	Yes	Yes	Yes	Yes
Fit statistics				
Observations	157,697	103,512	51,906	3,486
Pseudo R^2	0.97517	0.95289	0.97926	0.99235

Note: Same specification as Panel A of Table 2, trade regressed on the one-year lagged change in political distance with a dyad fixed effect and exporter-year and importer-year fixed effects, but restricted to the dyad-years inside the stacked coup event windows used in Panel B, deduplicated across overlapping stacks. Standard errors, in parentheses, are clustered by dyad. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

A.9 Robustness: An Outlier Driving the Bloc Coefficient in the Decrease-Distance Pool

The bloc coefficient in the decrease-distance pool of Table 3 is negative and marginally significant, the opposite of what the other columns of that table suggest a predicted decrease in political distance should do to trade. We traced this to a single coup event, Spain's failed coup attempt of October 1982.

The event enters the decrease-distance pool because the plotters are classified pro-US while Spain's own pre-coup ideal point, 0.874 in 1981, sat below that of most of its NATO partners, 2.69 for the United States, 2.04 for the United Kingdom, 1.90 for Germany, and 1.73 for France. Had the coup succeeded and realized this classified direction, Spain would have moved closer to these partners, not further away. The coup failed, however, so this event sits on the failure side of the within-cell comparison. Spain's trade with exactly these Western European partners grew sharply in the following years because Spain was in the middle of its accession to the European Community, which it joined in 1986. This unrelated trade boom, landing on the failure side of the comparison, is what pulls the pooled bloc coefficient negative.

Table A.9 reproduces Table 3 after dropping this one event. The anomalous bloc coefficient in the decrease-distance pool falls from -0.284 to an insignificant -0.131. Every coefficient in the increase-distance pool also grows larger and more significant, consistent with Spain's idiosyncratic trade boom diluting the effect there as well.

Table A.9: Success versus Failure, Excluding Spain's 1982 Coup Attempt

Dependent Variable:	Bilateral Trade Flows			
	All	Neutral	Bloc	Superpower
Model:	(1)	(2)	(3)	(4)
Panel A: Increasing distance				
Variables				
Post × Success	-0.070	-0.080	-0.398**	-0.635*
	(0.056)	(0.071)	(0.176)	(0.379)
Fixed-effects				
Dyad × Coup Event FE	Yes	Yes	Yes	Yes
Exporter-Year FE	Yes	Yes	Yes	Yes
Importer-Year FE	Yes	Yes	Yes	Yes
Fit statistics				
Observations	117,765	94,243	22,920	704
Pseudo R^2	0.97682	0.97635	0.97948	0.99778
Panel B: Decreasing distance				
Variables				
Post × Success	0.010	0.195***	-0.131	0.231
	(0.101)	(0.075)	(0.127)	(0.203)
Fixed-effects				
Dyad × Coup Event FE	Yes	Yes	Yes	Yes
Exporter-Year FE	Yes	Yes	Yes	Yes
Importer-Year FE	Yes	Yes	Yes	Yes
Fit statistics				
Observations	95,445	64,106	30,859	636
Pseudo R^2	0.98607	0.97198	0.98880	0.99914

Note: Reproduces Table 3 after excluding Spain's failed October 1982 coup attempt (see discussion above). Post × Success is the coefficient on the interaction of a post-attempt indicator with a successful-coup dummy, estimated by PPML with dyad-by-event, exporter-year, and importer-year fixed effects. Standard errors, in parentheses, are clustered by dyad. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

A.10 Alternative Measure of Political Alignment

Figure A.4: Effect of Political Realignment on Trade, by Partner Type, Geopolitical Score Measure

(a) Non-aligned partners

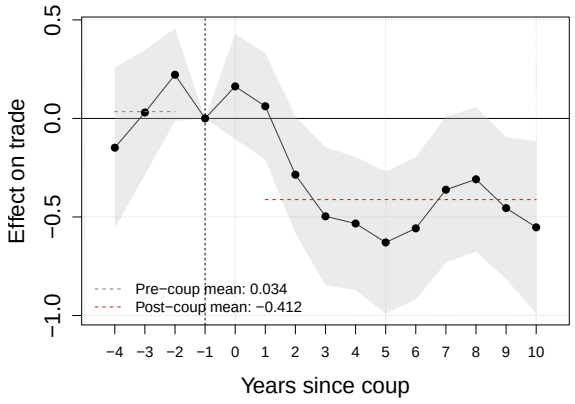

(b) Bloc partners

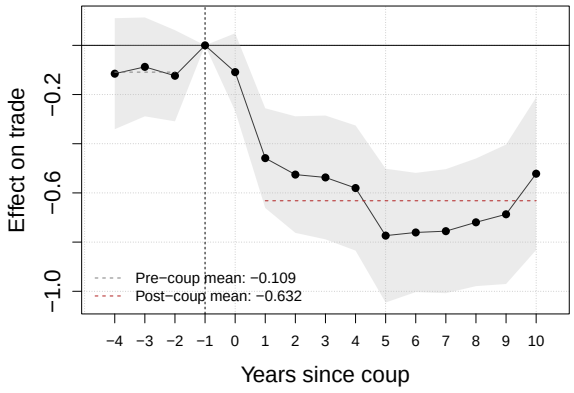

(c) Superpower partners

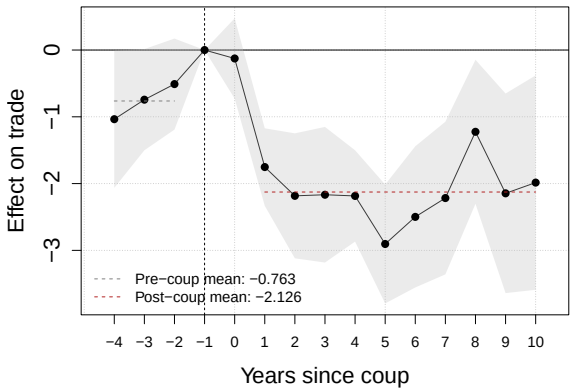

Note: Each panel plots $\hat{\beta}_k$ from equation (11), using the geopolitical score of Fan et al. (2025a) and Fan et al. (2025b) in place of ΔIPD_{ni}, estimated separately on the subsample of dyads in each partner group, for successful coups only. Shaded regions are 95% confidence intervals from standard errors clustered by dyad.

A.11 Alternative UNGA Idealpoints

Figure A.5: Effect of Political Realignment on Trade Around Coups - Alternative Idealpoints

(a) Successful coups

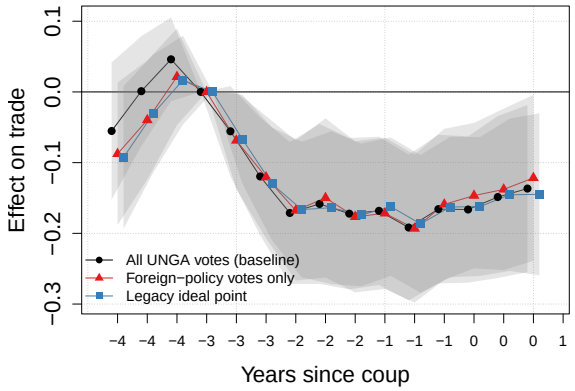

(b) Failed coups (placebo)

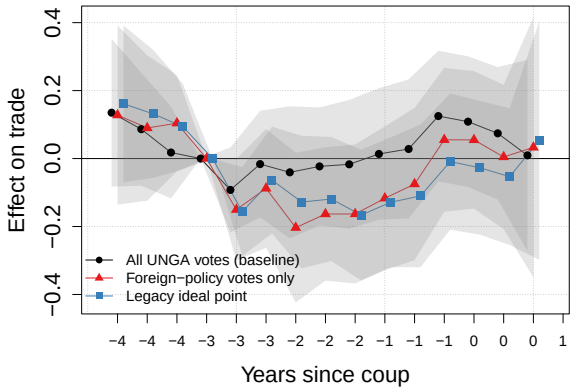

Note: Each panel plots $\hat{\beta}_k$ from equation (11), the effect on trade of a one-unit increase in political distance k years after the coup, relative to $k = -1$. Shaded regions are 95% confidence intervals from standard errors clustered by dyad. Dashed lines show the average pre-coup ($k < -1$) and post-coup ($k \geq 0$) coefficient.

A.12 Alternative Trade Data

We further validate our results using an alternative trade data source, the CEPII Gravity database, in place of the TRADHIST data used in our baseline specification. Bilateral trade flows are constructed by combining COMTRADE and IMF DOTS reports, prioritizing values reported by the importing country over the exporting country, since import statistics are generally considered more reliable due to customs duty enforcement. Because CEPII Gravity offers substantially sparser historical coverage than our baseline trade data, particularly before 1962, the estimation sample shrinks considerably once we require a matched Gravity observation for each dyad-year. Re-estimating our stacked event study on this reduced sample, we recover coefficients of the same sign throughout the post-coup period, though smaller in magnitude, and the post-coup effect no longer reaches conventional levels of statistical significance. The contrast with the failed-coup placebo still holds. Successful coups show a consistent negative coefficient at every post-period horizon, while the placebo coefficients have no consistent sign or trend. We interpret the loss of significance for successful coups as a consequence of reduced statistical power from the smaller sample rather than evidence against our baseline finding.

Figure A.6: Effect of Political Realignment on Trade Around Coups - Gravity Data

(a) Successful coups

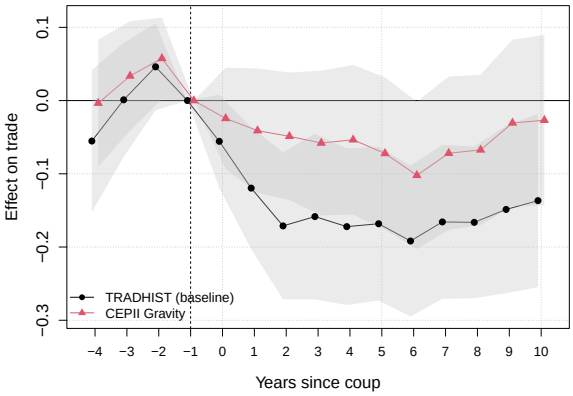

(b) Failed coups (placebo)

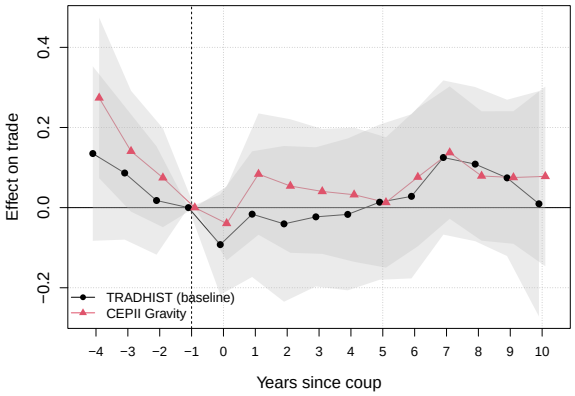

Note: Each panel plots $\hat{\beta}_k$ from equation (11), the effect on trade of a one-unit increase in political distance k years after the coup, relative to $k = -1$. Shaded regions are 95% confidence intervals from standard errors clustered by dyad. Dashed lines show the average pre-coup ($k < -1$) and post-coup ($k \geq 0$) coefficient. Trade data from CEPII Gravity database.

A.13 Alternative Time Window

Because the contamination filter looks forward over the same horizon as the event window, our estimation sample is sensitive to this choice. A shorter window is less restrictive about what counts as a nearby contaminating coup, so it changes both which coups survive the filter and how many years of post-coup trade enter the stacked panel. As a further check, we therefore narrow the event window from $k = -4, \ldots, 10$ to $k = -3, \ldots, 5$. Of the 546 coup events, 334 survive the shorter contamination filter, 169 successful and 165 failed.

Our main results survive this alternative window. For successful coups, the pre-trend coefficients at $k = -3$ and $k = -2$ remain jointly insignificant ($p = 0.172$), and the post-period coefficients turn negative and significant from $k = 2$ onward, ranging from -0.10 to -0.13 log points, with the joint post-period test strongly rejecting the null of no effect ($p = 0.001$). The failed-coup placebo again shows no pattern, with neither the pre-trend nor the post-period coefficients jointly significant (pre-trend $p = 0.215$, post-period $p = 0.391$).

Figure A.7: Effect of Political Realignment on Trade Around Coups

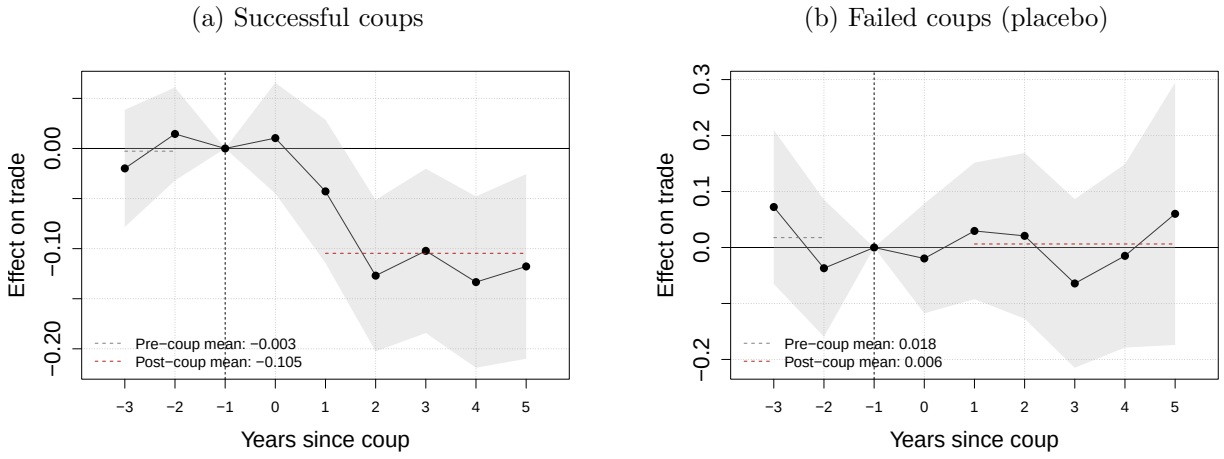

Note: Each panel plots $\hat{\beta}_k$ from equation (11), the effect on trade of a one-unit increase in political distance k years after the coup, relative to $k = -1$. Shaded regions are 95% confidence intervals from standard errors clustered by dyad. Dashed lines show the average pre-coup ($k < -1$) and post-coup ($k \geq 0$) coefficient.

A.14 Detrending the Alignment Data

We further address the concern that the observed post-coup decline in trade could reflect a secular drift in ideal-point distance from the United States rather than coup-induced realignment. Because ideal-point distance trends gradually over the sample period (see Figure 1) independent of any single coup, we detrend the alignment measure by subtracting each year's cross-sectional mean ideal-point distance from every observation before constructing the pre- and post-coup dose. This year-demeaning removes the average secular drift common to the whole cross-section, isolating each dyad's deviation from that year's typical change. The detrended dose is highly correlated with the raw dose (correlation of 0.99) and re-estimating our stacked event study with this alternative dose yields coefficients nearly identical to our baseline results. This indicates that the secular trend in political distance shifts all coup events by a similar amount and therefore does not differentially confound our identification, which relies on cross-sectional variation in the realized alignment shock across coup events.

Figure A.8: Effect of Political Realignment on Trade Around Coups - Detrended Alignment Data

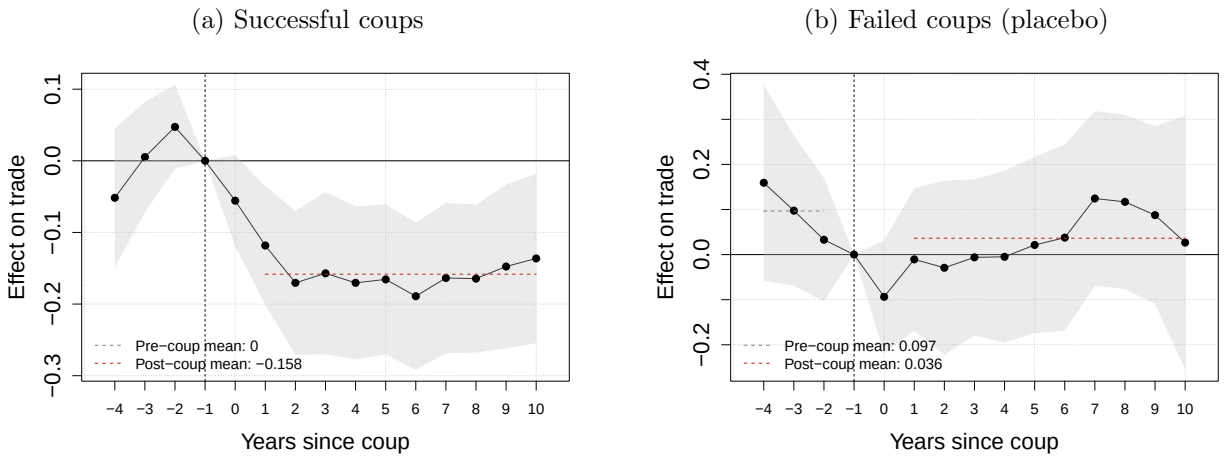

Note: Each panel plots $\hat{\beta}_k$ from equation (11), the effect on trade of a one-unit increase in political distance k years after the coup, relative to $k = -1$. Shaded regions are 95% confidence intervals from standard errors clustered by dyad. Dashed lines show the average pre-coup ($k < -1$) and post-coup ($k \geq 0$) coefficient.

A.15 Foreign Aid Data

This appendix describes the construction of the foreign aid data used in Section 5.1 in full.

Figure A.9: Soviet Foreign Aid

COMMUNIST COUNTRY	DATE EXTENDED	PROJECTS	MILLION CURRENT US $ EXTENDED	DRAWN	STATUS*
ALGERIA					
ROMANIA	1972 MAR				
		NINE STATE FARMS			UC
		DAM, EL FAKIA			UC
		TOTAL	100.0	45.0	
USSR	1962 AUG	GRANT -- WHEAT (10,000 TONS)	0.6	0.6	C
	1963 OCT	LINE OF CREDIT	100.0	100.0	
		TRUCKS			C
		TWO IL-18 AIRCRAFT ($3.4 MILLION)			
		AGRICULTURAL EQUIPMENT			C
		TECHNICAL ASSISTANCE			
		EQUIPMENT FOR TEXTILE INSTITUTE			C
		INDUSTRIAL TRAINING CENTER			C

Note: This figure displays an excerpt of the CIA reports on Soviet aid flows to non-aligned countries during the Cold War. The data includes information on the type, amount, date, origin and recipients of aid.

U.S. Aid. For American foreign assistance, we use the "U.S. Overseas Loans and Grants" database, commonly known as the Greenbook. This official source records annual commitments of foreign loans and grants from 1945 onward at the instrument level, providing comprehensive coverage of U.S. official capital flows across multiple agencies. The database distinguishes between economic assistance (development aid, food aid, technical cooperation) and military assistance, defined as "foreign aid primarily for the benefit of armed forces of recipient governments, or aid which subsidizes or substantially enhances military capability."

Comparison and Measurement. Both datasets record commitments rather than actual disbursements, making them comparable in timing. However, they differ in temporal structure. U.S. aid commitments are recorded annually, while Soviet commitments typically refer to multi-year agreements. Following Horn et al. (2021)'s approach for modern Chinese aid data, we assume Soviet aid is evenly distributed over five years, the most common project horizon when such information is available in the CIA reports. We aggregate both U.S. and Soviet aid flows to the recipient-year level and convert all values to constant 2011 U.S. dollars using the Consumer Price Index. Soviet aid is originally reported in current dollars by the CIA, facilitating direct comparison.

A.16 Bilateral Foreign Aid by Individual Donor

Section 5.1 asks whether the decline in superpower aid after a coup-induced realignment is specific to the two superpowers or reflects broader bloc-level engagement. We extend the analysis to bilateral aid from every individual Western donor reporting to the OECD's DAC2a database (20 countries, covering the NATO allies, Japan, Australia, and the remaining DAC members with usable pre-1990 data) and every individual Warsaw Pact donor in the CIA data (the USSR together with East Germany, Czechoslovakia, Poland, Hungary, Bulgaria, and Romania). We have one observation per coup event, individual donor, and year, so that each donor's own bilateral change in alignment with the coup country is matched to that donor's own bilateral aid. At this shorter window the non-zero observations available to identify the flexible event-time coefficients become sparse enough that the specification is unstable for some donors, so we collapse the event study into a pre/post difference-in-differences specification with a single coefficient,

$$\text{Aid}_{nit} = \exp\left(\beta \cdot \text{Post}_{nt} \cdot \Delta\text{IPD}_{ni} + \gamma_t + \tilde{\alpha}_{ni}\right) \eta_{nit}, \tag{A.1}$$

where $\text{Post}_{nt} = 1$ in the years following the coup, n indexes the coup country, and i indexes the individual donor. We estimate this specification five times, pooling every DAC and Warsaw Pact donor together under one shared coefficient (Combined), restricting to DAC donors only, restricting to Warsaw Pact donors only, and isolating each superpower individually, the US and the Soviet Union. For the first three columns, Western aid is measured using the OECD's DAC2a bilateral ODA data, which by definition excludes military assistance. For the Warsaw Pact column we correspondingly use CIA *economic* aid only, omitting the military-aid file, so that the two sides of the DAC/Warsaw Pact comparison remain scope-consistent. The US and Soviet columns instead use this project's standard total (economic and military) aid sources for each superpower, the Greenbook and the CIA data respectively. The Soviet column does not converge at this bilateral granularity. With only 40 dyad-events after fixed-effects absorption we omit it from Table A.10 rather than report the resulting degenerate estimate.

Table A.10 shows that aid declines by a similar amount across the three scope-consistent columns, DAC alone (-0.28), Warsaw Pact alone (-0.27), and the two pooled together (Combined, -0.28) are all close in magnitude, and a formal test of the DAC/Warsaw Pact difference

Table A.10: Bilateral Foreign Aid Response to Coup-Induced Alignment Shifts

	Combined	DAC	Warsaw Pact	US	Soviet
Post × ΔIPD	-0.276	-0.276	-0.275	-0.558*	–
	(0.250)	(0.250)	(0.443)	(0.322)	
Dyad × Coup Event FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	7,172	6,261	773	615	–

Note: Each column reports $\hat{\beta}$ from equation (A.1), estimated on a panel with one observation per coup event, individual donor country, and year, pooling all DAC and Warsaw Pact donors under one shared coefficient (Combined), DAC donors only, Warsaw Pact donors only, the US alone, and the Soviet Union alone. Combined/DAC/Warsaw Pact use DAC2a (economic aid only) and CIA economic-aid data, while US uses the Greenbook and Soviet uses CIA economic plus military aid, neither on the same aid-scope basis as the other three columns. A dash indicates a specification that does not converge. Standard errors clustered by the (coup country, donor) dyad in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

confirms they are statistically indistinguishable ($p = 0.998$). None of the three, however, reaches conventional significance ($p = 0.27$ for Combined and DAC, $p = 0.54$ for Warsaw Pact alone), reflecting that this bilateral specification asks a broader and noisier question, whether aid from *any* individual Western or Communist-bloc donor responds to a coup-induced change in alignment with that donor specifically. The US column, using the Greenbook's broader economic-and-military aid measure, is both larger in magnitude (-0.56) and marginally significant ($p = 0.083$), consistent with military aid being the component of the superpower-specific relationship that carries the strongest signal. Restricting the Western donors to the DAC2a economic-aid-only concept, as the first three columns do for comparability with the Warsaw Pact, likely understates the true response for exactly this reason. The corresponding Soviet-specific column is not estimable at this bilateral granularity.

A.17 Aid Sanctions

Table A.11: Aid Sanctions Response to Coup-Induced Alignment Shifts

Dependent Variable:	Aid Sanctioned
	US
Model:	(1)
Variables	
Post × ΔIPD	0.005
	(0.043)
Fixed-effects	
Dyad × Coup Event FE	Yes
Year FE	Yes
Fit statistics	
Observations	1,192
R^2	0.42582

Clustered (dyad) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: Collapsed pre/post specification analogous to equation (A.1), with a binary aid-sanction indicator from Gibson et al. (2025) as the outcome, estimated by OLS with dyad-by-coup-event and year fixed effects. Russia is omitted, only 5 aid-sanction episodes appear in the underlying data across the full sample period, none within our estimation window.

A.18 Post-Cold War Effects

Figure A.10: Effect of Political Realignment on Trade Around Coups, Post-Cold War Sample, UN Idealpoints

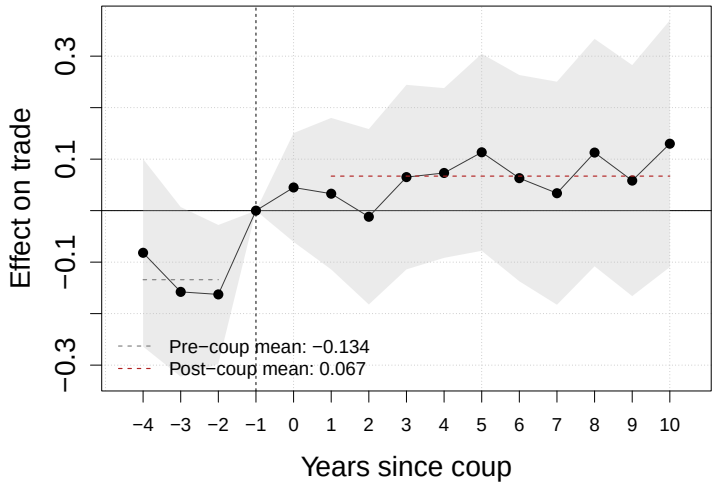

Note: Plots $\hat{\beta}_k$ from equation (11), restricted to coups and trade with year $\geq$ 1990. Shaded region is the 95% confidence interval from standard errors clustered by dyad.

Figure A.11: Effect of Political Realignment on Trade Around Coups, Post-Cold War Sample, Geopolitical Score

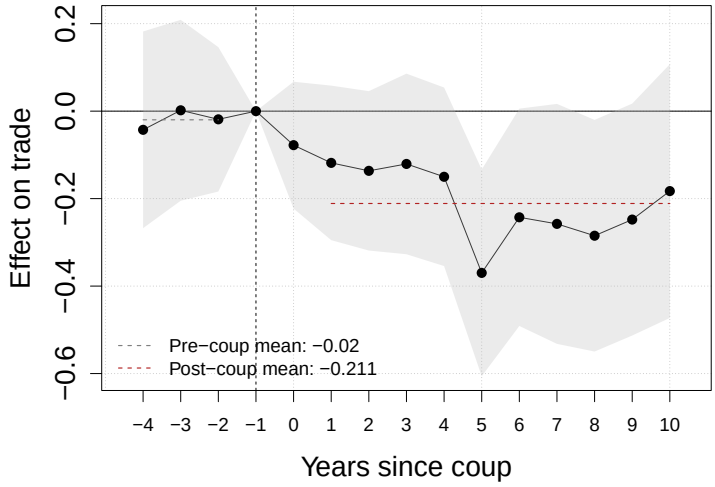

Note: Plots $\hat{\beta}_k$ from equation (11), using the geopolitical score of Fan et al. (2025a) and Fan et al. (2025b) in place of ΔIPD_{ni}, restricted to coups and trade with year $\geq$ 1990. Shaded region is the 95% confidence interval from standard errors clustered by dyad.

Figure A.12: Effect of Political Realignment on Trade Around Coups, Cold War Sample, Geopolitical Score

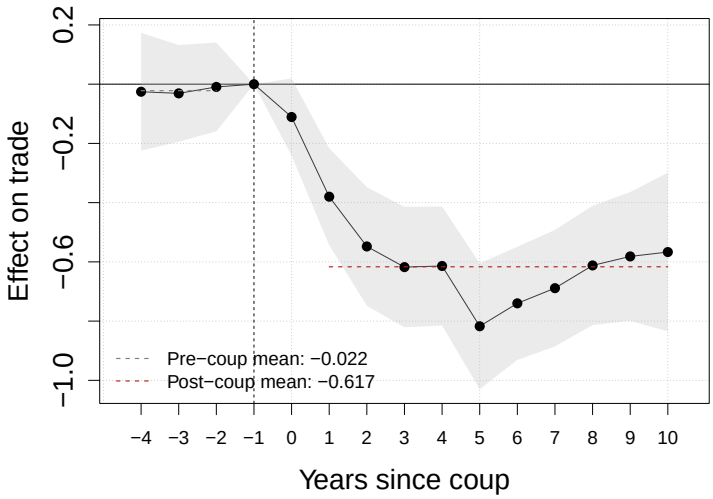

Note: Cold War-period counterpart to Figure A.11, same specification and geopolitical score measure but restricted to coups and trade with year $\leq$ 1989. Shaded region is the 95% confidence interval from standard errors clustered by dyad.

A.19 GDP and Trade Shares

Figure A.13: Cold War Bloc Structure, 1948–1989

(a) Share of World GDP

(b) Share of Bilateral Exports

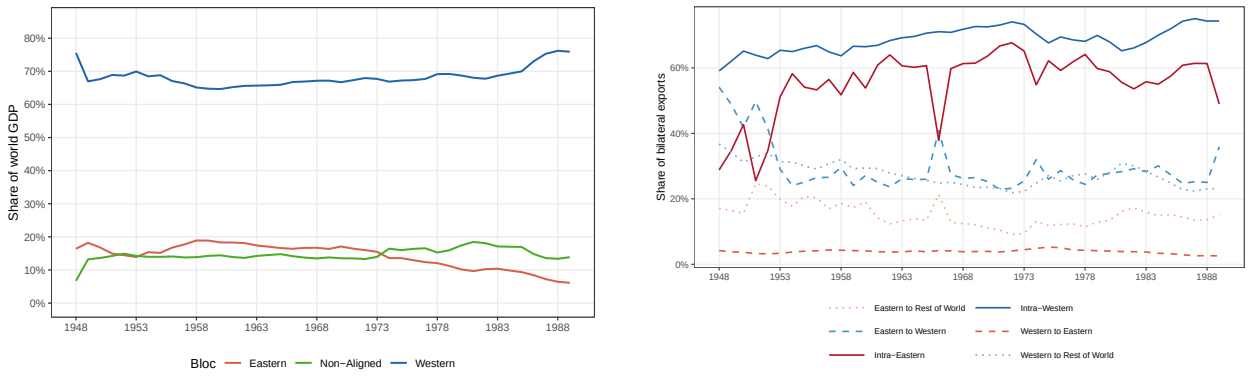

(c) (Exports + Imports) / GDP

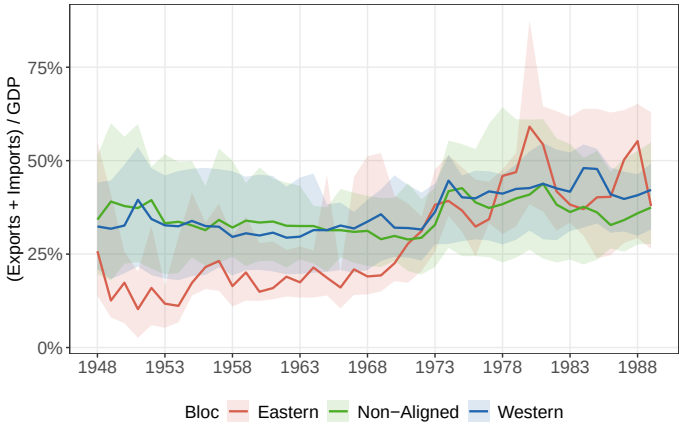

Note: Western bloc = NATO members plus Japan, Eastern bloc = Warsaw Pact members, Non-Aligned = all others. Top-left: each bloc's share of total world GDP. Top-right: share of bilateral (non-domestic) exports by origin bloc and destination. Bottom: median exports-to-GDP ratio by bloc, shaded band = interquartile range across countries. GDP from the Maddison Project / Penn World Tables, bilateral trade from CEPII.

A.20 Decomposing the Sources of Welfare Costs

To isolate the mechanisms behind the heterogeneous burdens from Section 6.2, we run three additional partial counterfactuals, each removing the barriers of a single cross-bloc dyad pair, East-West (the Iron Curtain), West-Non-Aligned, and East-Non-Aligned. Within-bloc barriers are held at their observed values throughout, so the partials do not sum to the full counterfactual. The gap reflects both the omitted within-bloc component and general-equilibrium interactions. Figure A.14 presents the results.

For Eastern bloc countries, the Iron Curtain is the dominant channel. Removing E-W barriers alone recovers 76% of the Eastern welfare cost (0.10% of real income against a full cost of 0.14%). Removing W-NA or E-NA barriers in isolation produces a welfare *loss* for Eastern countries, a

general-equilibrium terms-of-trade effect in which Non-Aligned and Western trade reorients away from Eastern suppliers, depressing Eastern export prices without opening Eastern goods to new markets.

The most striking result concerns Non-Aligned countries. Of their 0.43% welfare cost, 93% is attributable to political barriers to the *Western* bloc alone (W-NA, 0.40%), while political distance to the Eastern bloc accounts for only 14% (E-NA, 0.06%). This asymmetry reflects the West's dominant economic weight. With the large majority of global trade flowing through Western-affiliated markets during this period, exclusion from them was far more costly than exclusion from the smaller Eastern trading system.

The Western bloc presents a different picture. Summing all three cross-bloc partial counterfactuals, their net contribution to Western welfare costs is approximately −0.01%, essentially zero, against a full counterfactual cost of 0.15%. The individual partials are small and partly offsetting, with two of the three producing a terms-of-trade loss rather than a gain, E-W at −0.02%, W-NA at +0.04%, and E-NA at −0.04%. The implication is that the bulk of Western welfare costs arose from political frictions *within* the Western bloc itself. Because intra-Western trade dominated Western expenditure, political distance among Western countries, distortions that the full counterfactual removes but no partial counterfactual touches, was the primary channel through which the Cold War imposed real income losses on the West.

Figure A.14: Decomposition of welfare costs by barrier type

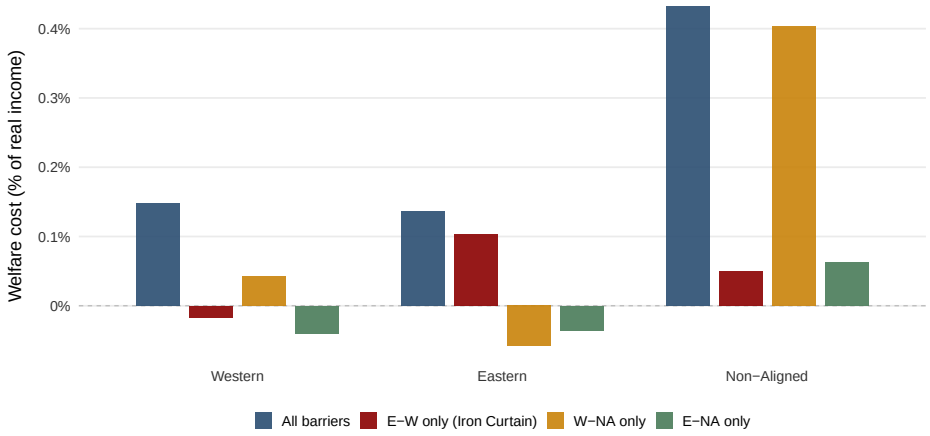

Notes: Each bar shows the GDP-weighted average welfare cost (1948–1989) from a partial counterfactual that removes the indicated subset of geopolitical barriers, computed from equations (20)–(22) with $\varepsilon = 5.03$ and $\beta = -0.154$. "All barriers" replicates the baseline counterfactual (IPD = 0 for all dyads). The three partial scenarios each zero out IPD for a single cross-bloc dyad pair, E-W (the Iron Curtain), W-NA, and E-NA, while holding within-bloc barriers at their observed values. Because within-bloc barriers are never removed, the three partials do not sum to the "All barriers" bar. The gap reflects the omitted within-bloc component and general-equilibrium interaction terms. Bars below zero indicate a welfare *loss* from removing that subset of barriers, reflecting general-equilibrium terms-of-trade effects on the third bloc.

A.21 The Johnson Letter

Figure A.15 reproduces an excerpt from a declassified CIA intelligence report, classified *Secret* at the time of writing, assessing Turkish reactions to President Johnson's June 1964 letter to Prime Minister İnönü. Portions of the document remain redacted under Executive Order 12958. The visible text records that reactions to the letter ranged from shock and disbelief to complete disappointment, with a general consensus that the letter made it almost mandatory for Turkey to pursue greater independence from the United States in the field of international relations, corroborating the interpretation in Section 5.3 that the letter was a discrete political shock with no accompanying formal trade restriction.

Figure A.15: Turkish Reaction to the Johnson Letter

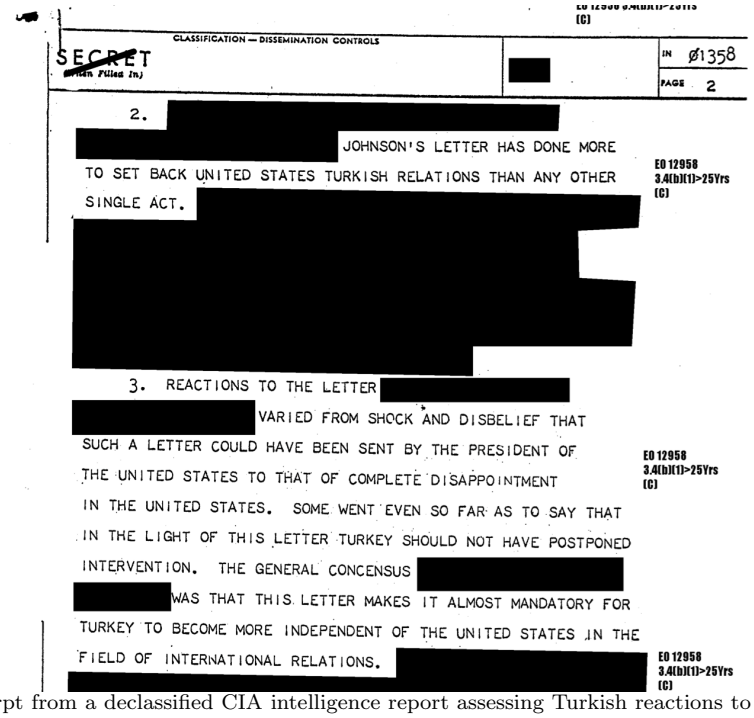
(C)

CLASSIFICATION — DISSEMINATION CONTROLS

SECRET
(When Fills In)

IN Ø1358

PAGE 2

2. [redacted] JOHNSON'S LETTER HAS DONE MORE TO SET BACK UNITED STATES TURKISH RELATIONS THAN ANY OTHER SINGLE ACT. [redacted]

E0 12958 3.4(b)(1)>25Yrs (C)

3. REACTIONS TO THE LETTER [redacted] VARIED FROM SHOCK AND DISBELIEF THAT SUCH A LETTER COULD HAVE BEEN SENT BY THE PRESIDENT OF THE UNITED STATES TO THAT OF COMPLETE DISAPPOINTMENT IN THE UNITED STATES. SOME WENT EVEN SO FAR AS TO SAY THAT IN THE LIGHT OF THIS LETTER TURKEY SHOULD NOT HAVE POSTPONED INTERVENTION. THE GENERAL CONCENSUS [redacted] WAS THAT THIS LETTER MAKES IT ALMOST MANDATORY FOR TURKEY TO BECOME MORE INDEPENDENT OF THE UNITED STATES IN THE FIELD OF INTERNATIONAL RELATIONS. [redacted]

E0 12958 3.4(b)(1)>25Yrs (C)

E0 12958 3.4(b)(1)>25Yrs (C)

Note: Excerpt from a declassified CIA intelligence report assessing Turkish reactions to the Johnson Letter, classified *Secret* at the time of writing. Portions of the document remain redacted under Executive Order 12958.

Figure A.16 plots a second, independent piece of evidence. The dyadic trade friction $\hat{\tau}_{nit}$ implied by a baseline gravity regression along the lines of equation (6), estimated separately in each direction between the United States and Turkey, alongside the ideal point distance between the two countries. Before 1964, both directions of trade friction sit well below their sample average and are roughly flat, consistent with the close Cold War alliance between the two NATO members. From 1964 onward, both directions turn upward and continue rising through the late 1960s and 1970s, tracking the same period over which UN ideal point distance between the two countries also climbs.

Figure A.16: Bilateral Trade Frictions and Political Distance: USA-Turkey

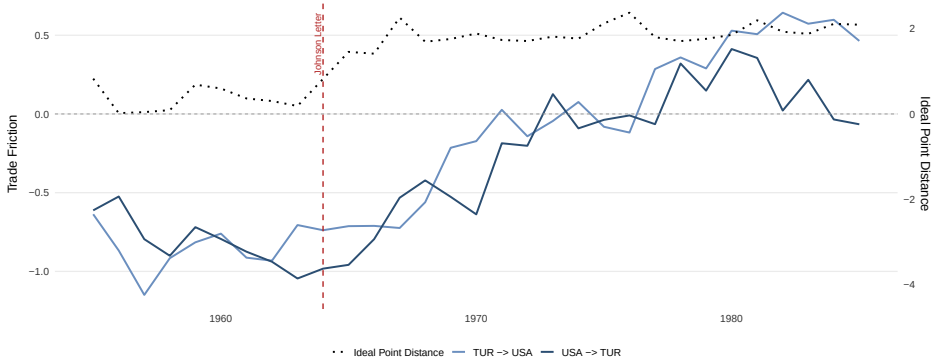

Note: Solid lines (left axis) plot the dyadic trade friction $\hat{\tau}_{nit}$ implied by a baseline gravity regression, shown separately for Turkish exports to the United States (TUR $\rightarrow$ USA) and US exports to Turkey (USA $\rightarrow$ TUR), restricted to 1955-1985. The dotted line (right axis) plots UN ideal point distance between the two countries. The vertical dashed red line marks the year of the Johnson Letter (1964).

Table A.12 reports the five largest donor country weights for each of the four synthetic control specifications underlying Figure 9.

Table A.12: Five largest donor weights

	Standard SCM		Lasso		Ridge		L2-constrained SCM	
Rank	Country	(%)	Country	(%)	Country	(%)	Country	(%)
1	Iraq	66.65	Switzerland	22.80	Saudi Arabia	18.06	Iraq	36.86
2	Spain	23.31	Saudi Arabia	20.43	Switzerland	13.73	Spain	28.34
3	Italy	7.74	Spain	13.06	Italy	11.50	Greece	14.37
4	Egypt	1.63	Italy	6.53	Spain	11.33	Lebanon	9.19
5	United Kingdom	0.67	United Kingdom	3.72	Belgium	10.56	Italy	5.15

Notes: Each column reports the five positive donor weights, expressed in percent. Displayed weights are not re-normalized to sum to 100.

A.22 Sector/Product Classification Tests

This appendix reports full detail on the five product-classification tests summarized in Section 4.4.2 and Figure 7. All test use 2-digit SITC level granularity, we aggregate revealed comparative advantage (RCA) and product complexity (PCI) from the 4-digit level by taking averages. We demean the continuous variables (RCA, PCI, Rauch's differentiated-good share, contract intensity) on the estimation sample before interacting with the post-coup dose, so that the uninteracted dose coefficient $\hat{\beta}$ is evaluated at the sample average rather. Standard errors are clustered by dyad throughout.

Revealed Comparative Advantage and Product Complexity. These measures come from the Growth Lab's Atlas of Economic Complexity. RCA follows Balassa's standard formula. Its distribution is extremely right-skewed, with a median of 0.13 and a maximum in excess of 4,000 for hyper-specialized micro-economies, so we use $\log(1 + \text{RCA})$ rather than the raw measure. PCI is the Hidalgo-Hausmann product complexity index, under which high-PCI products are exported

mainly by diversified economies that few other countries also export. Both measures are fixed at their pre-coup value. Table A.13 reports the result. The post-coup dose interacted with the partner's own (importer) RCA is negative and significant, trade falls more, and grows more when relations improve, in sectors where the partner itself has a comparative advantage, consistent with a standard trade-cost/integration channel.

Table A.13: Sector-Level Response by Comparative Advantage and Product Complexity

Dependent Variable:	SITC2 Product Trade (USD)
	USA-partner dyads
Model:	(1)
Variables	
Post × ΔIPD	-0.203*
	(0.104)
Post × ΔIPD × log(1+Exporter RCA)	-0.024
	(0.122)
Post × ΔIPD × log(1+Importer RCA)	-0.362*
	(0.185)
Post × ΔIPD × Product Complexity	-0.099
	(0.082)
log(1+Exporter RCA)	1.30***
	(0.057)
log(1+Importer RCA)	-0.687***
	(0.154)
Product Complexity	0.279
	(0.252)
Fixed-effects	
Dyad × Coup Event FE	Yes
Year FE	Yes
SITC2 Product FE	Yes
Fit statistics	
Observations	39,259
Pseudo R^2	0.70982

Note: Product-level analogue of equation (13) with Z_p equal to log(1 + RCA) (exporter and importer separately) and PCI, estimated jointly on the USA-partner sample. Clustered (dyad) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

Rauch (1999) Differentiated-Goods Classification. We use Rauch (1999)'s classification of SITC4 goods as organized-exchange (for example, wheat or copper), reference-priced (for example, certain chemicals and steel products), or differentiated (for example, clothing or machinery), aggregated to a trade-weighted differentiated-good share at the 2-digit level. Table A.14 shows that trade in more differentiated, relationship-dependent sectors falls significantly more as align-

ment worsens.

Table A.14: Sector-Level Response by Rauch (1999) Differentiated-Goods Share

Dependent Variable:	SITC2 Product Trade (USD) USA-partner dyads
Model:	(1)
Variables	
Post × ΔIPD	-0.228***
	(0.084)
Post × ΔIPD × Differentiated-Good Share	-0.417*
	(0.230)
Fixed-effects	
Dyad × Coup Event FE	Yes
Year FE	Yes
SITC2 Product FE	Yes
Fit statistics	
Observations	41,868
Pseudo R^2	0.51077

Note: Product-level analogue of equation (13) with Z_p equal to the trade-weighted differentiated-good share, estimated on the USA-partner sample. Clustered (dyad) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

Nunn (2007) Contract Intensity. We use Nunn (2007)'s industry-level contract-intensity measure, the share of an industry's intermediate inputs that are themselves differentiated, relationship-specific goods. The measure is available for 29 manufacturing ISIC Rev. 3 industries, which we bridge to SITC via the standard UN correspondence tables and then aggregate to the 2-digit level by trade-weighted average. Because the measure is defined only for manufacturing, SITC codes for agriculture, mining, and raw materials have no match and are dropped, so the estimation sample is smaller than Rauch's, which covers all goods. Table A.15 shows the same sign as the Rauch result but a larger standard error, plausibly reflecting this reduced sample and the coarser industry-level (rather than product-level) source of the underlying measure.

Table A.15: Sector-Level Response by Nunn (2007) Contract Intensity

Dependent Variable:	SITC2 Product Trade (USD)
	USA-partner dyads
Model:	(1)
Variables	
Post × ΔIPD	-0.184**
	(0.079)
Post × ΔIPD × Contract Intensity (Nunn)	-0.637
	(0.481)
Fixed-effects	
Dyad × Coup Event FE	Yes
Year FE	Yes
SITC2 Product FE	Yes
Fit statistics	
Observations	39,530
Pseudo R^2	0.50575

Note: Product-level analogue of equation (13) with Z_p equal to trade-weighted contract intensity, estimated on the USA-partner sample, manufacturing only. Clustered (dyad) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

Broad Economic Categories (BEC): End Use. We use the UN Comtrade SITC Revision 2 — BEC correspondence table, mapped to three end-use categories, Capital goods, Intermediate inputs, and Consumption goods, following the standard BEC basic-class definitions. Each SITC2 group is assigned to its single trade-weighted dominant end-use category, with a mean dominant-category share of 0.89. Because this is a genuine partition rather than a matter of degree, we fully saturate, estimating one post-coup dose coefficient per category with no omitted reference group. Table A.16 shows that the decline is concentrated in Capital goods.

Table A.16: Sector-Level Response by Broad Economic Category End Use

Dependent Variable:	SITC2 Product Trade (USD) USA-partner dyads
Model:	(1)
Variables	
Post × ΔIPD × Capital Goods	-0.403***
	(0.155)
Post × ΔIPD × Intermediate Inputs	-0.065
	(0.096)
Post × ΔIPD × Consumption Goods	-0.279
	(0.175)
Fixed-effects	
Dyad × Coup Event FE	Yes
Year FE	Yes
SITC2 Product FE	Yes
Fit statistics	
Observations	40,243
Pseudo R^2	0.50782

Note: Fully saturated product-level analogue of equation (13), one post-coup dose coefficient per end-use category, no omitted reference group, estimated on the USA-partner sample. Clustered (dyad) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

Lall (2000) Technological Classification. We use UNCTAD's Lall (2000) classification of SITC Revision 3 goods into Primary, Resource-based, Low-, Medium-, and High-technology manufactures, bridged to SITC2 via the standard UN correspondence table and, as with BEC, assigned by trade-weighted dominant category, with a mean dominant-category share of 0.86. As with BEC, this is a genuine partition, so we again fully saturate. Table A.17 shows a near-monotonic gradient. The post-coup dose coefficient grows more negative moving from Primary goods, statistically indistinguishable from zero, up through High-technology manufactures, the single largest point estimate of any category across all five classifications.

Table A.17: Sector-Level Response by Lall (2000) Technological Classification

Dependent Variable:	SITC2 Product Trade (USD)
	USA-partner dyads
Model:	(1)
Variables	
Post × ΔIPD × Primary	0.045
	(0.161)
Post × ΔIPD × Resource-Based	-0.181
	(0.138)
Post × ΔIPD × Low Technology	-0.335*
	(0.185)
Post × ΔIPD × Medium Technology	-0.324**
	(0.133)
Post × ΔIPD × High Technology	-0.465**
	(0.232)
Fixed-effects	
Dyad × Coup Event FE	Yes
Year FE	Yes
SITC2 Product FE	Yes
Fit statistics	
Observations	40,594
Pseudo R^2	0.51183

Note: Fully saturated product-level analogue of equation (13), one post-coup dose coefficient per technology tier, no omitted reference group, estimated on the USA-partner sample. Clustered (dyad) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.