# Beyond Decoupling: Geopolitical Shocks and the Reconfiguration of International Innovation Networks

**Abstract**

When bilateral relations deteriorate, does geopolitical hostility reach inventive activity controlled by multinational firms, or does common ownership shelter cross-border knowledge production? I study this question with a global panel of country pairs from 2000 to 2023 built from USPTO co-invention and patent references and linked to corporate ownership, and I measure bilateral hostility from the composition of political events on the Goldstein scale. Using abrupt, dyad-specific animosity shocks and stacked event studies, I find that a reconfiguration of innovation networks follows geopolitical hostility. Co-invention falls 24.6% in the shock year and 33.1% by year five, yet annual country-pair activity is largely unchanged and disclosed patent references rise 4.0% on average. Recorded ownership is highly concentrated: 90.1% of internationally co-invented patents have one assignee, and 98.0% of strictly parent-mapped patents are controlled by a common parent. Single-assignee activity falls 29.5%, close to the 33.2% point estimate for multiple-assignee activity. Because the dominant ownership category appears in 94.3% of active country-pair years, I do not infer a firm-boundary effect from category-level dissolution. At the parent-corridor level, co-invention controlled by the same multinational parent falls 28.1%, while the same parents expand co-invention in the same technology through third-country locations by 12.4%. The decline concentrates among rival and asymmetrically dependent country pairs. Weak corridors account for the extensive-margin response, and strategic and dual-use fields show a larger patent-reference response despite a similar co-invention decline. The findings show that geopolitical hostility thins international invention even under unified ownership while multinational geographic scope provides a margin for partial reconfiguration.

**Keywords.** Geopolitics, innovation networks, multinational firm, decoupling, co-invention, knowledge flows.

## 1. Introduction

International co-invention has grown rapidly, nearly tripling between 2000 and 2018, even as the bilateral animosity that surrounds it has risen sharply since 2016 (Figure 1). Geopolitical rivalry has become a first-order concern of corporate strategy (Caldara and Iacoviello 2022). Governments increasingly treat advanced technologies as strategic assets. Firms now manage cross-border research under the threat of retaliation, reputational backlash, and abrupt policy reversals (Fajgelbaum et al. 2020; Fajgelbaum and Khandelwal 2022). Much of this friction arrives before formal sanctions or export controls, as a deterioration in bilateral relations. Rhetoric, media tone, and public sentiment shift first, changing how partners read one another (Hassan et al. 2019). One possible outcome is decoupling, the severing of ties that bind national innovation systems (Fajgelbaum and Khandelwal 2022). Yet we know little about what happens to the microstructure of inventor networks when relations sour. Some patented inventions span recorded owners. Others are controlled within a single multinational group. The two need not respond alike (Berry and Donnelly 2025). Hostility might sever a corridor, erode its depth while it remains active, or reroute inventive work across locations. Whether geopolitical hostility reaches international invention under common corporate control, or whether organizational ownership insulates it, is the central question this study takes up.

Two contrasting perspectives bear on this question. One holds that the firm insulates it. Multinationals exist in part because internal organization governs knowledge exchange more reliably than markets (Kogut and Zander 1992). Intra-firm inventor ties are bound by common ownership, shared identity, managerial authority, and accumulated routines (Gupta and Govindarajan 2000; Alcácer and Zhao 2012; Almeida and Kogut 1999). These ties should outlast swings in the political climate. Deteriorating relations between two countries should do little to collaboration between inventors in different countries but the same firm. That is, organization buffers geopolitics.

The other perspective is relational in nature and holds that geopolitics may penetrate the firm. Once two countries turn hostile, cross-border inventive activity becomes a compliance, security, and reputational exposure (Hassan et al. 2019). Managers may restrict data access, travel, disclosure, and project assignment, while headquarters may move sensitive work elsewhere in the multinational network (Alcácer and Zhao 2012; Berry and Donnelly 2025). Macro animosity may then reshape invention under common corporate control because it changes the costs of maintaining particular cross-border relationships. Which force tends to dominate is unclear. Even if hostility penetrates an exposed corridor, however, multinational organization may matter in a different way. A geographically distributed firm can preserve an inventive capability without preserving every bilateral relationship through which it was previously produced. It can move related work across locations or place sensitive activity elsewhere in its network (Kogut and Kulatilaka 1994; Alcácer and Zhao 2012). The empirical tension is therefore not only whether organization insulates an exposed relationship, but whether the multinational responds by protecting that relationship or by reconfiguring the broader network around it.

I study this question on a global panel of country pairs observed annually from 2000 to 2023. It draws on 435,397 internationally co-invented patents, those with inventors in more than one country. I build the innovation network from USPTO patents, geolocating disambiguated inventors to countries and dating each patent by application year so the network reflects when invention occurred. At the country-pair level, I measure three margins of the network: the intensity of co-invention, whether the country corridor is active in a year, and disclosed patent references. I observe ownership of patented inventions rather than the employment relationship of every inventor. I therefore use assignee records to describe unified and multiple recorded ownership and historically dated ultimate-owner links to identify inventive activity controlled by the same multinational parent. This distinction lets me examine whether parent-owned activity on an exposed corridor contracts and whether the same parent expands related activity through other locations. I measure bilateral hostility from political events on the Goldstein conflict-cooperation scale (Goldstein 1992; Leetaru and Schrodt 2013), using the share of strong-conflict events in each dyad-period rather than average tone. The share is a composition-based measure that is less sensitive to news volume and recovers known episodes of bilateral conflict. The empirical strategy exploits abrupt, dyad-specific hostility shocks defined as large detrended surprises in bilateral hostility relative to a dyad-specific rolling baseline. I estimate dynamic effects with a stacked event study that uses only never- and not-yet-treated dyads as controls and absorbs country-pair-by-stack and endpoint-country-by-year-by-stack fixed effects. I use alternative shock definitions, worldwide patent-family validation, and a reduced-form global animosity-wave diagnostic as robustness.

Three findings emerge. First, the depth of joint inventive production falls after an animosity shock, while the country corridor usually remains active. Co-invention intensity falls by 24.6% in the shock year, by 30.0% on average over years 0 to 5, and by 33.1% by year five. Scaled against the pre-shock treated mean of 24.5 co-invented patents per year, these effects imply about 6.0 fewer co-inventions in the shock year and 8.1 fewer by year five. The probability that two countries co-invent at all changes by only 0.5 percentage points on average, a change too small to be economically meaningful. Disclosed patent references do not collapse. They rise by 4.0% on average after the shock and by 5.2% by year five. What contracts most clearly in the data is the volume of joint production, not annual activity on the country corridor.

Second, the contraction also appears in the dominant ownership form of international invention. Among assigned patents, 91.8% have one assignee of record and 8.2% have multiple assignees. Single-assignee co-invention falls by 29.5% on average after a shock and by 32.2% by year five. Multiple-assignee activity falls by 33.2% on average and by 34.7% by year five. The ownership categories have real common support in pre-shock intensity, but radically different coverage: single-assignee activity is present in 94.3% of active country-pair years, compared with 32.3% for multiple-assignee activity. I therefore do not interpret category-level dissolution as a causal firm-boundary effect. The organizational evidence instead comes from parent-specific corridors. Co-invention controlled by the same parent on an exposed corridor falls 28.1% over the next three years, while the same parents expand co-invention in the same technology through third-country locations by 12.4%. This increase is consistent with partial geographic reconfiguration rather than complete replacement of the lost activity.

Third, the decline concentrates where political exposure concentrates. The fall in co-invention is largest among rival dyads, where co-invention intensity falls by 34.3%, against a decline of 1.8% among non-rivals that is not economically meaningful. The decline is also largest under asymmetric dependence, reaching 47.2% in the top quartile against no detectable change among balanced pairs. The stable average on country-corridor activity masks turnover beneath it. Weak corridors are 13.1 percentage points more likely to become inactive after a shock, while strong corridors are unaffected. Strategic and dual-use fields show a more nuanced pattern. Their co-invention decline is similar to the rest of the network, but disclosed patent references rise more strongly in those fields. A generalized uncertainty shock would move all pairs more evenly. This pattern instead tracks the political structure of the relationship.

This paper makes three contributions. First, it speaks to the knowledge-based theory of the multinational firm, which holds that internal organization carries cross-border knowledge more reliably than markets (Kogut and Zander 1992; Gupta and Govindarajan 2000; Alcácer and Zhao 2012). The findings qualify a simple insulation interpretation. Inventive activity contracts sharply on politically exposed corridors even under unified ownership. The organizational response appears instead in the broader geographic network, where the same parents expand related activity through third-country locations. This distinction between relationship-level insulation and network-level adaptation identifies another way in which multinational organization can matter under geopolitical stress (Kogut and Kulatilaka 1994). Second, it reframes the geopolitics of innovation. Where prior work documents contraction in cross-border collaboration after conflict and sanctions (Makkonen and Mitze 2023; Vanino, Douch, and Çakmak 2026), I separate annual country-corridor activity, the depth of joint production, and the disclosed-reference channel, and ask which margin gives. Third, it connects firm-level political risk to the geographic organization of knowledge production (Hassan et al. 2019). Because I do not observe internal project assignments or managerial deliberations, I interpret the parent-level patterns as consistent with organizational reconfiguration rather than as direct evidence of a particular managerial decision.

The paper proceeds as follows. Section 2 describes the data and measures. Section 3 sets out the identification strategy and treatment diagnostics. Section 4 presents the results. These cover the three network margins, recorded ownership, parent-specific adjustment, heterogeneity, mechanisms, and robustness checks. Section 5 discusses the findings and concludes.

## 2. Data and Measures

I construct the data to distinguish three ways in which geopolitical decoupling can occur. A country pair can stop producing jointly in a year, remain active but produce less, or change its reliance on disclosed prior art. I therefore build the panel around three margins of the bilateral innovation relationship. I then decompose co-invention by recorded patent ownership and use historically dated parent links to examine whether geopolitical hostility reaches inventive activity controlled by multinational firms and whether related activity expands elsewhere in their geographic networks.

### 2.1. Co-invention network and sample

I build the innovation network from utility patents in PatentsView. This database links raw USPTO patent records to disambiguated inventors, assignees, organizations, locations, and patent documents. I restrict the analysis to application years 2000-2023 and map every disambiguated inventor location to a country. I date each patent by application year.

A patent enters the international co-invention network when its inventors reside in at least two countries. Each patent contributes total weight one to the set of country pairs that its inventors span. For each unordered country pair on a patent, I assign a weight proportional to the product of the number of inventors located on the two sides and then normalize those pair weights so that they sum to one within the patent. I also construct a whole-counting version in which every co-inventing country pair on a patent receives weight one, and I use it as a robustness check.

The main panel contains 435,397 internationally co-invented patents. The unit of observation is the country pair by year. Co-invention is undirected, so I use unordered country pairs for co-invention outcomes. Disclosed patent references are directed, so I keep both directions of each pair for that margin. The balanced panel spans a 196-country universe, yielding 19,110 unordered country pairs and 458,640 pair-years from 2000 to 2023. Of these countries, 180 appear with positive USPTO-visible inventive output and 173 have at least one international co-invention partner.

I build the network from USPTO patents for two reasons. First, it supplies a common patent venue with comparable inventor, assignee, publication, and patent-reference information across countries (Jaffe, Trajtenberg, and Henderson 1993; Singh 2005). PatentsView provides a consistent, high-quality record of a single patent system, with disambiguated inventors, assignees, and locations linked to grants, pre-grant publications, and citations, which lets me measure all three margins on a common footing across country pairs. Second, USPTO patents capture inventions that firms are likely to protect in the world's largest technology market (Alcácer and Zhao 2012; Berry and Donnelly 2025). The resulting network should be read as the U.S.-observed layer of international invention rather than its universe. It is likely that this scope is conservative for the present question. If geopolitical hostility thins collaboration even among the commercially significant inventions that firms protect in the United States, the effect reaches the part of the innovation network most likely to matter for multinational strategy (Kogut and Zander 1992; Gupta and Govindarajan 2000). A worldwide patent-family network broadens coverage, at the cost of harder cross-office comparability in family counting, applicant-to-owner matching, and citation practices. I thus keep USPTO as the main design and use PATSTAT as external validation, which I report in Section 4.5 and Appendix Table A5.

### 2.2. Margins in the inventor network

The first margin is co-invention intensity, the fractional count of co-invented patents in country pair *ij* and year *t*. This measure captures the depth of joint production. It is sparse and right-skewed, so the main count specifications use Poisson pseudo-maximum likelihood (Santos Silva and Tenreyro 2006). The second margin is annual country-corridor activity. I set *Activeijt* equal to one when pair *ij* co-invents at least one patent in year *t*, and zero otherwise. This measure records whether the bilateral corridor remains active and does not depend on the fractional allocation rule. The third margin is disclosed patent references. I count references from patents invented in country *i* to patents invented in country *j*, date each reference by the application year of the citing patent, and allocate references fractionally across inventor-country pairs on the citing and cited patents. Patent references are directed, so I observe *i* to *j* and *j* to *i* separately. This margin records references to disclosed prior art, not the full set of knowledge flows or a direct behavioral substitute for co-invention. In auxiliary specifications, I separate references under the same ultimate owner from those that cross ownership boundaries.

### 2.3. Patent ownership and corporate-parent mapping

I first classify the recorded ownership of each international patent. PatentsView assignees are legal owners of patent rights; they do not identify the employment relationship of every inventor. Among the 427,421 international patents with an observed assignee, 91.8% have one assignee of record and 8.2% have multiple assignees. Equivalently, single-assignee patents account for 90.1% of the full sample, multiple-assignee patents for 8.1%, and unassigned patents for 1.8%. I use these categories to describe unified and multiple recorded ownership, not collaboration within and across employer boundaries.

I then link assignees to historically dated ultimate corporate parents in Orbis. A patent is strictly mapped only when every relevant assignee links to an ultimate parent in the application year. Strict all-assignee mapping covers 54.3% of international co-invented patents; a loose rule requiring only one mapped assignee covers 54.5%. Coverage varies by assignee country, including 66.0% for the United States and 23.9% for China. In the strict mapped sample, approximately 98.0% of patents are controlled by a common parent and 2.0% span multiple parents; the latter category contains 4,790 patents. I aggregate these ownership-classified patent counts to the country-pair-year level using the same fractional weights as in the total co-invention measure.

The ownership decomposition remains informative for the intensive margin because pre-shock thickness separates single- from multiple-assignee activity only moderately (AUC = 0.706) and the categories retain common support. It cannot, however, identify a firm-boundary effect on country-pair dissolution. Single-assignee activity is present in 94.3% of years in which a country pair is active, compared with 32.3% for multiple-assignee activity. A category disappears while the aggregate corridor remains active in 2.7% of single-assignee observations and 29.9% of multiple-assignee observations. Those differences largely reflect coverage of the corridor itself. I therefore use ownership-classified intensity as a decomposition and move the organizational analysis to the parent-country-corridor level, where one parent's activity can contract independently of the aggregate country pair. Appendix Tables A8 and A9 report construct, coverage, and common-support diagnostics.

### 2.4. Bilateral hostility

I measure bilateral hostility with political events coded in GDELT on the Goldstein conflict-cooperation scale (Goldstein 1992, Leetaru and Schrodt 2013). The Goldstein scale assigns event types to a conflict-cooperation continuum. I collapse directed events to the country-pair level and compute a strong-conflict share: the number of bilateral events with Goldstein scores at or below -5 divided by the total number of bilateral political events in the same pair-period. The measure captures the composition of bilateral interaction rather than the amount of news coverage.

The dataset keeps the total event count as a separate diagnostic because a share can be noisy when the denominator is small. The shock construction in Section 3 therefore requires at least eight quarters of prior event history, and the validation exercises report results under alternative minimum-event thresholds. The conflict-share measure covers 81.8% of pair-years. I use the share rather than average tone because average tone combines coverage volume with valence and can move when attention rises even if the mix of bilateral interaction does not become more hostile. I validate the series by showing that it rises around known episodes of bilateral conflict, including U.S.-China relations in 2018, U.S.-Russia relations in 2014, and U.S.-Iran relations in 2020. I also report specifications that control for bilateral event volume, placebo shocks built from total event volume, pseudo-shock dates, raw conflict-count shocks, and mean Goldstein tone shocks.

### 2.5. Heterogeneity moderators

The analyses use four pre-specified moderators to locate where animosity should matter most. First, I identify rival dyads from pre-period bilateral hostility and distinguish them from allied dyads using formal defense-pact data. Rivalry captures whether a pair already sits in a politically exposed relationship before the shock.

Second, I measure asymmetric dependence in inventive collaboration. For each country in a pair, I calculate the share of its international co-invention that runs through the partner. I then summarize the two directional shares with an index that equals zero when dependence is balanced and approaches one when one side depends on the other far more than the reverse. This measure captures exposure to one-sided leverage in the innovation relationship.

Third, I measure strategic-field exposure. I define strategic and dual-use technologies using IPC subclasses in semiconductors, electronic circuits, telecommunications, computing, navigation and sensing, aerospace, and nuclear technology. I compute the strategic-field share as the share of a pair's co-invention in these fields before the shock whenever possible. These fields account for 40% of active co-invention.

Fourth, I sort active ties by pre-shock strength. I assign country pairs to terciles based on their cumulative co-invention before the shock. This split separates weak, peripheral ties from established corridors. A decoupling process that prunes fragile relationships should appear first on this margin.

### 2.6. Descriptive statistics

Table 1 reports descriptive statistics for patents and firms, countries, and country-pair-years. I report percentiles as well as means because the network is sparse and highly concentrated, and I also report conditional-on-positive statistics for the active part of the network. Only 13.2% of country pairs ever co-invent, and only 4.2% of pair-years contain any co-invention. Conditional on positive co-invention, the median pair-year contains 1.8 fractional patents, the 90th percentile contains 28, and the 99th percentile contains 418. The top 10 corridors account for 49.1% of all co-invention, and the top 50 account for 77.7%. Eleven of the twelve largest corridors involve the United States, which reflects the use of USPTO data. The five largest producing countries are the United States, Japan, South Korea, Germany, and China. The top 10 countries account for 90.7% of world inventive output in the USPTO-visible network.

*[ Table 1 about here ]*

The grant-based network grows through 2018 and then becomes mechanically right-censored. International co-invention rises from roughly 9,900 fractional patents in 2000 to about 28,100 in 2018, and codified patent references grow faster. Both series fall in recent application years because many applications had not yet issued as grants when I built the dataset. Figure 1 therefore shades the recent right-censored period. I restrict the main event-study sample to mature cohorts whose post-shock window ends by 2019 and use pre-grant publications to de-censor recent application years in Section 3. USPTO practice generally publishes utility and plant applications after eighteen months, subject to standard exceptions, so pre-grant publications substantially shorten the observation lag relative to grants (U.S. Patent and Trademark Office 2024).

*[ Figure 1 about here ]*

## 3. Identification Strategy

I estimate the dynamic effect of abrupt, dyad-specific hostility shocks on bilateral inventor collaboration. I compare shocked country pairs with clean never- and not-yet-shocked pairs inside stacked event windows, absorb country-pair-by-stack and endpoint-country-by-year-by-stack fixed effects, and validate that the treatment captures hostile composition rather than bilateral media salience.

### 3.1. Animosity shocks

In this study, I define an animosity shock as an abrupt rise in bilateral hostility relative to a country pair's own recent history. For each pair and quarter I take the quarter-to-quarter change in the strong-conflict share, subtract the mean of that change over the preceding eight quarters, and divide by its standard deviation over the same window:

$$Z_{ij,q} = ( \Delta C_{ij,q} - \mu_{ij,q} ) / \sigma_{ij,q} \tag{1}$$

A shock occurs when this standardized surprise is at least 2.5, so the pair experiences a sharp deterioration relative to its own path rather than a slow drift. I require a minimum number of bilateral events in the rolling window and in the current quarter, and I impose non-overlapping windows. Shocks dated to the first or second quarter enter the current year, and shocks dated to the third or fourth quarter enter the next year. The baseline uses each pair's first shock and treats repeated non-overlapping episodes as a robustness check. Detrending removes slow-moving, dyad-specific hostility and isolates the abrupt component that is less likely to be driven by gradual collaboration dynamics. I do not treat this residual as exogenous by construction, and I validate it in Section 3.4.

### 3.2. Estimating changes around animosity shocks

The estimand is the path of proportional changes in each network margin in the years around an animosity shock, relative to pairs not yet shocked. I estimate it with a stacked event study that avoids already-treated controls (Cengiz, Dube, Lindner, and Zipperer 2019). For each cohort of pairs first shocked in a given year, I form an event window from three years before to five years after the shock and attach as controls the pairs that are never shocked or not yet shocked through the end of that window. For count outcomes, I estimate by Poisson pseudo-maximum likelihood:

$$\mathrm{E}[Y^{s}_{ij,t}] = \exp\left( \Sigma_k \, \beta_k \, D^{s,k}_{ij,t} + \alpha_{ij}^{s} + \eta_{it}^{s} + \eta_{jt}^{s} \right) \tag{2}$$

Here $s$ indexes the cohort stack, the indicator $D$ equals one when a treated pair is $k$ years from its shock, the sum runs over event time from three years before to five years after and omits the reference period at $k$ equal to minus one, the term $\alpha$ is a stack-by-pair fixed effect, and the two $\eta$ terms are stack-by-country-by-year fixed effects applied to both endpoints of the undirected pair. The quantity $\exp(\beta) - 1$ is the proportional change in the margin at each event time. For directed citations I replace the common endpoint fixed effect with separate origin-country-year and destination-country-year fixed effects. I estimate tie existence by a linear probability model, so its coefficients are changes in the probability that the pair co-invents.

A pooled specification replaces the event-time indicators with a single post-shock indicator and reports the average effect over the five post-shock years. I report the average pre-shock coefficient separately as a placebo. I estimate equation (2) separately for unified- and multiple-owner patent activity and test the difference in proportional contraction. Because the ownership categories cover country corridors at very different rates, I do not use their country-pair dissolution gap as a firm-boundary estimand. The parent-level analysis instead follows a multinational parent's own activity on an exposed country corridor and related activity through third countries. Since the stacked Poisson event study is not itself a heterogeneity-robust estimator, I benchmark its dynamic pattern against linear heterogeneity-robust difference-in-differences estimators that use the same never- and not-yet-treated comparison groups, and I avoid contaminated comparisons from conventional two-way fixed effects (Goodman-Bacon 2021; de Chaisemartin and D'Haultfoeuille 2020). I cluster standard errors two ways on the endpoint countries, which captures correlation across pairs that share a country and serial correlation within a pair (Cameron, Gelbach, and Miller 2011; Aronow, Samii, and Assenova 2015).

The empirical specifications absorb time-invariant pair characteristics through country-pair fixed effects. These fixed effects account for distance, contiguity, language, colonial history, persistent migration links, baseline technological similarity, and any stable bilateral propensity to co-invent. The preferred event-study specifications also absorb country-year fixed effects for both endpoint countries. These fixed effects remove shocks to each country's overall innovation system, patenting propensity, sanctions exposure, macroeconomic conditions, and geopolitical environment in a given year.

In specifications that do not include endpoint-country-year fixed effects, I add time-varying pair and country characteristics that could move with both hostility and innovation. These include the log of each country's GDP, bilateral trade value, and an indicator for a preferential trade agreement, drawn from the World Bank, the IMF Direction of Trade Statistics, and CEPII. The estimates do not depend on adding these controls. I report the exact fixed-effect structure, control set, clustering level, and sample restrictions under each table.

Identification rests on conditional parallel trends and no anticipation. Absent the shock, treated and not-yet-treated dyads would have followed the same path after pair and country-year absorption, and collaboration does not adjust before the shock. The never- and not-yet-treated controls, together with the limited spillover of a dyad-specific shock to other dyads, make interference across pairs unlikely. Section 3.4 reports the tests of these conditions and of the requirement that treatment timing reflect hostile composition rather than news attention, and Section 3.6 addresses differential censoring of the patent outcomes.

### 3.4. Treatment support and diagnostics for the identifying variation

Table 2 shows that animosity shocks do not fall randomly across the network and reports the support for the mature-cohort event-study sample. The full panel contains 19,110 country pairs, of which 6,342 are ever shocked and 12,768 are never shocked. The mature-cohort sample contains 2,853 shock events across 12 treatment cohorts, with 1,269 treated co-inventing dyads and 12,768 unique clean control dyads. Pairs that ever receive a shock average 24.5 co-inventions per active dyad-year, compared with 2.0 among never-shocked pairs. Shocks land on large, active corridors rather than on the periphery. I therefore do not interpret treated-control balance as evidence of quasi-random assignment. The event-study design absorbs country-pair-by-stack fixed effects, uses never- and not-yet-treated pairs as controls, reports pre-trends, and in the preferred specifications absorbs endpoint-country-by-year-by-stack shocks for both endpoint countries.

*[ Table 2 about here ]*

The tables report event-study leads, so the coefficients before the shock test the parallel-trends assumption directly. I add pseudo-shock dates and placebo shocks built from total event volume, which separate hostile composition from media salience. The baseline conflict-share shock reduces co-invention intensity by 24.6% in the shock year. Controlling for bilateral event volume leaves the estimate essentially unchanged at 24.1%, while the event-volume placebo is close to zero (-1.2%) and pseudo-shock dates are also close to zero (-0.6%). I compare specifications with and without country-year fixed effects, specifications that drop the most connected countries, and alternative hostility measures. Because shocks land on large, active corridors, I also report pre-shock outcome levels and slopes by cohort. I bound the results with honest difference-in-differences, which evaluates how large a post-shock deviation from parallel trends, relative to the largest pre-shock deviation, would be needed to overturn a finding (Rambachan and Roth 2023; Roth 2022).

### 3.5. Global animosity-wave diagnostic

I complement the dyad-level event study with a global animosity-wave diagnostic, not as the main source of identification. The diagnostic interacts each pair's predetermined exposure to global animosity with the leave-both-countries-out worldwide animosity wave:

$$Z_{ij,t} = e_{ij} \times W^{-ij}{}_{t} \tag{3}$$

Exposure $e$ is fixed in the pre-sample period and the wave $W$ excludes both endpoint countries. I estimate the reduced form in a linear stacked event study for comparability and report it in the robustness table. The exclusion restriction for an instrumental-variable interpretation would be demanding, because a global animosity wave can affect collaboration through sanctions, export controls, visas, and managerial attention, not only bilateral hostility (Goldsmith-Pinkham, Sorkin, and Swift 2020; Borusyak, Hull, and Jaravel 2022). I therefore do not rely on the shift-share diagnostic for identification. Its reduced-form estimate is close to zero on the productive margin, which helps rule out a global animosity wave as the driver of the main dyad-level results.

### 3.6. Patent timing and censoring

The network is dated by application year but observed through granted patents, so recent application years are right-censored by the grant lag, as Figure 1 shows. I restrict the main sample to mature cohorts whose post-shock window ends by 2019. I also rebuild the outcomes from pre-grant publications, which the patent office releases about eighteen months after filing. These are an earlier-observed measurement system. I report an event-time support table giving treated and control pairs and observed patents by event time, and I note that grant lags vary by technology and country. The pre-grant-publication robustness produces an impact effect of -25.2% and a five-year effect of -32.9%, close to the main estimates.

## 4. Results

Table 3 and Figure 2 report the dynamic effect of an animosity shock on the three network margins. Table 4 develops the recorded-ownership decomposition and its measurement diagnostics. Table 5 reports the pre-specified heterogeneity tests. Table 6 reports organizational adjustment in the parent-mapped sample. Appendix Table A6 reports supplementary firm-response patterns, Appendix Table A7 reports robustness checks, and Appendix Tables A8 and A9 report ownership-classification and match-quality diagnostics. Throughout, the baseline specification absorbs country-pair-by-stack and endpoint-country-by-year-by-stack fixed effects, the reference period is the year before the shock, and standard errors are clustered two ways on the endpoint countries. I report effects at impact, over the short run of years one and two, the medium run of years three and four, the long run of year five, and the average post-shock period, and I scale each effect against the pre-shock mean of the same margin for treated dyads.

*[ Table 3 and Figure 2 about here]*

### 4.1. Where the adjustment occurs

Co-invention intensity falls immediately, and the decline deepens across the event window (Table 3, column 1, and Figure 2). The effect is -0.246 at impact, a 24.6% reduction. It grows to -0.287 over years 1 to 2, -0.318 over years 3 to 4, and -0.331 by year five, a 33.1% reduction. The average post-shock effect over years 0 to 5 is -0.300 ($p < 0.001$). Scaled against the pre-shock treated mean of 24.5 co-invented patents per year, these effects amount to about 6.0 fewer co-inventions on impact and 8.1 fewer by year five. The pre-shock coefficients are flat, and the post-shock coefficients emerge only after the shock.

The extensive margin does not move after a shock. The probability that two countries co-invent at all changes by 0.4 percentage points at impact and by 0.5 percentage points on average over the post-shock window ($p = 0.617$), against a pre-shock co-invention rate of 13.2% among treated dyads. The pre-shock coefficients are flat, and the post-shock coefficients are flat. So, the tie remains active while the volume of joint work it carries contracts.

Codified patent references do not contract. The effect is 0.013 at impact ($p = 0.516$), rises to 0.039 over years 1 to 2, 0.049 over years 3 to 4, and reaches 0.052 by year five. The average post-shock effect is 0.040 ($p = 0.022$), or about 1.3 additional references relative to the pre-shock treated mean of 31.8. The pre-shock reference coefficients are flat. I interpret this margin as disclosed patent references rather than as a complete measure of knowledge flow. The honest difference-in-differences bound implies that overturning the intensity result would require a post-shock departure from parallel trends about 3.8 times the largest pre-shock deviation.

### 4.2. Ownership structure and organizational adjustment

Table 4 decomposes co-invention by recorded ownership form. Single-assignee activity falls by 23.8% at impact, 29.5% on average after the shock, and 32.2% by year five. Multiple-assignee activity falls by 27.1% at impact, 33.2% on average, and 34.7% by year five. The point estimates are close, although multiple-assignee activity is much less prevalent. I interpret this comparison as evidence that the aggregate contraction is not confined to patents with multiple recorded owners, not as a causal estimate of the advantage of common ownership.

*[ Table 4 about here]*

Panel B reports the distributional facts that constrain the extensive-margin interpretation. Single-assignee activity appears in 94.3% of active country-pair years, whereas multiple-assignee activity appears in 32.3%. Conditional on the aggregate country corridor remaining active, the ownership category disappears in 2.7% of single-assignee observations and 29.9% of multiple-assignee observations. Among treated dyads under the strict presence rule, the corresponding post-shock dissolution rates are 15.5% and 22.9%. These are descriptive differences in category coverage and persistence. They do not identify whether common ownership causally shelters a relationship from geopolitical hostility.

The codified patent-reference margin provides a complementary channel result (Table 3, column 3, and Table 5). Codified patent references do not collapse after the shock. The average post-shock effect is 0.040 ($p = 0.022$), and the year-five effect is 0.052 ($p = 0.030$). In strategic and dual-use fields, the reference response is larger, rising 10.1% ($p = 0.001$), 4.9 percentage points above other fields. This pattern is consistent with substitution toward disclosed references, although I interpret patent references more cautiously than co-invention.

### 4.3. Heterogeneity. Where the adjustment is strongest

If hostility operates through political and structural exposure rather than through ambient uncertainty, the contraction should concentrate in the relationships most exposed to it. I examine four moderators fixed in the pre-shock period, each capturing a distinct source of exposure (Table 5).

*[ Table 5 about here ]*

4.3.1. Rivalry and pre-shock hostility. In theory, bilateral tensions should affect inventors hardest where relations are already adversarial, because geopolitical risk concentrates in conflictual dyads and rivals carry the sharpest compliance and reputational exposure (Caldara and Iacoviello 2022; Fajgelbaum and Khandelwal 2022). I test this by interacting the shock with pre-shock bilateral hostility. The intensity decline is concentrated among rivals. For dyads in the top tercile of pre-shock hostility, co-invention falls by 34.3% ($p < 0.001$). For non-rivals, the effect is -1.8% and near zero ($p = 0.561$). The rival-minus-non-rival difference is -32.5 percentage points ($p < 0.001$).

4.3.2. Asymmetric dependence. When one country depends on the other far more than the reverse, the dependent side is open to one-sided leverage, the feature that turns interdependence into a tool of coercion (Farrell and Newman 2019). The decline is largest where dependence is one-sided. In the top quartile of asymmetric dependence, co-invention falls by 47.2% ($p < 0.001$), against 0.6% for balanced pairs ($p = 0.836$). The top-quartile-minus-balanced difference is -47.8 percentage points ($p < 0.001$). The change is monotone across dependence quartiles.

4.3.3. Tie strength and the extensive margin. It is possible that weak ties are more fragile than strong ones and are the first to be pruned when a relationship comes under strain (Granovetter 1973), so the stable average on tie existence may hide turnover concentrated among them. Ties in the bottom tercile of pre-shock intensity are 13.1 percentage points more likely to dissolve after a shock ($p < 0.001$), ties in the middle tercile 4.1 points ($p = 0.040$), and ties in the top tercile are unaffected ($p = 0.617$). The bottom-minus-top difference is 12.7 percentage points ($p < 0.001$).

4.3.4. Strategic and dual-use fields. Past work suggests that technologies with security and dual-use significance draw sharp political scrutiny and are the focus of export controls (Fajgelbaum and Khandelwal 2022), so the response may differ by technology. In strategic and dual-use fields, co-invention falls by 26.2% ($p < 0.001$), similar to the 24.8% decline in other fields (difference = -0.014, 0.061). The patent-reference response differs. Codified patent references in strategic and dual-use fields rise by 10.1% ($p = 0.001$), compared with 5.2% in other fields ($p = 0.030$), and the difference of 4.9 percentage points is economically meaningful ($p = 0.041$).

**4.4. How firms adjust**

The results suggest that firms respond to hostility by reorganizing cross-border work rather than simply abandoning it. Thus, I explore three channels that are central to this adjustment as per extant research.

*[ Table 6 about here ]*

4.4.1. Geographic reconfiguration. Multinationals hold an option to move activity across their network when a location becomes costly (Kogut and Kulatilaka 1994), and they manage valuable knowledge across locations as they expand abroad (Alcácer and Zhao 2012; Berry and Donnelly 2025). I follow treated multinational parents and compare co-invention on the shocked corridor with co-invention through third-country inventor locations in the same technology (Table 6). On the shocked corridor, common-parent-controlled co-invention falls 28.1% over the three years after the shock ($p < 0.001$), while related third-country activity rises 12.4% ($p = 0.017$). Both movements are larger in strategic and dual-use fields (-34.7% on the shocked corridor and +18.6% in third-country locations). Parent-by-year and technology-by-year fixed effects hold the parent and technology fixed across locations. I interpret the increase as evidence consistent with partial geographic reconfiguration, not one-for-one replacement of the activity lost on the exposed corridor.

4.4.2. Team substitution and affiliate continuity. Firm-level political risk changes whom firms hire and where they invest (Hassan et al. 2019), and multinationals compartmentalize knowledge across locations to limit leakage (Alcácer and Zhao 2012). Managers may therefore narrow the cross-border composition of teams while keeping the affiliate active. In the mapped-parent sample, the foreign co-inventor share on surviving common-parent corridor patents falls 4.1 percentage points ($p = 0.003$) against a pre-shock share of 43%, while the hostile-country affiliate's domestic output ($p = 0.535$) and its co-invention on non-hostile corridors ($p = 0.519$) do not change (Appendix Table A6, Panel A). The contraction is specific to the politically exposed cross-border corridor.

4.4.3. Disclosure and compartmentalization. If hostility raises the cost of revealing knowledge to a rival jurisdiction, firms may disclose less on exposed corridors, consistent with the internal compartmentalization multinationals use to protect knowledge (Alcácer and Zhao 2012). In Appendix Table A6, Panel B, patents controlled by a common mapped parent on shocked corridors have 8.7% smaller international families ($p = 0.004$), reach first publication 10.6% later ($p = 0.007$), and cite 15.8% less partner-country prior art ($p = 0.001$). These patterns are consistent with filing in fewer jurisdictions, disclosing later, and sourcing less from the partner.

4.4.4. Parent home-country alignment. If the firm's own governance under home-country political pressure drives the response, rather than local friction alone, the effect should concentrate in parents whose home country is a party to the dispute, because these firms bear the compliance and national-security exposure that firm political risk captures (Hassan et al. 2019; Caldara and Iacoviello 2022). Table 6 splits the common-parent corridor estimate by the parent's home country. For parents headquartered in one of the two endpoints, common-parent-controlled co-invention falls 35.1% ($p < 0.001$). For neutral third-country parents, the effect is 4.9% and near zero ($p = 0.221$). The endpoint-parent-minus-neutral-parent difference is -30.2 percentage points ($p < 0.001$). Because historical-parent coverage differs across assignee countries, I treat this split as an exploratory alignment test rather than a population comparison of endpoint and neutral parents.

4.4.5. Timing relative to formal measures. Geopolitical risk is priced from perceptions well before events materialize (Caldara and Iacoviello 2022), so if firms respond to animosity itself, the contraction should not require formal policy. In Table 6, common-parent-controlled co-invention still falls 21.4% among shocks not followed by any sanction, export-control action, or Entity-List designation for the dyad ($p < 0.001$). Among dyads that later receive a formal measure, the common-parent corridor decline appears a median of three quarters earlier. The response does not require legal compulsion.

4.4.6. R&D reallocation. It is possible that firms simply cut research where animosity is high, since political and policy uncertainty depress corporate investment (Gulen and Ion 2016; Hassan et al. 2019). Appendix Table A6, Panel C reports the effect of a corridor shock on R&D at the parent and its affiliates. Parent consolidated R&D does not change ($p = 0.568$). The budget moves across locations instead. R&D in the hostile-partner affiliate falls 14.3% ($p = 0.006$), while third-country affiliate R&D rises 7.9% ($p = 0.038$). The affiliate decline reaches 21.1% in strategic and dual-use fields ($p = 0.002$).

4.4.7. Ownership and divestiture. At the extreme, firms exit locations that become politically exposed, and conflict raises the hazard of foreign subsidiary divestiture (Dai, Eden, and Beamish 2013). Appendix Table A6, Panel C shows that ownership restructuring occurs but remains the exception. The probability of divestiture or stake reduction in the hostile-partner affiliate rises 1.8 percentage points ($p = 0.010$) against a baseline rate of 5.1%, and the probability that the ultimate owner changes rises 1.1 percentage points ($p = 0.067$). Divestiture rises 3.4 percentage points in strategic and dual-use fields ($p = 0.009$).

### 4.5. Robustness and validity checks

I run a set of analyses to test whether the main results depend on the way I measure hostility, define the sample, count patents, or classify patent ownership. The results are summarized in Appendix Tables A7 through A9.

4.5.1. Pre-trends and placebo shocks. The pre-shock event-time coefficients are flat on all three margins, and placebo shocks assigned from total bilateral news volume produce dynamic coefficients clustered around zero. The event-volume placebo has an impact effect of -0.012 ($p = 0.699$) and a five-year effect of -0.019 ($p = 0.602$). Pseudo-shock dates also produce a precise zero ($p = 0.830$).

4.5.2. Salience versus hostility. Controlling for bilateral event volume leaves the impact effect on intensity at -0.241 ($p < 0.001$), close to the baseline estimate of -0.246. Replacing the conflict-share shock with mean Goldstein tone yields -0.230 ($p < 0.001$), and replacing it with a raw conflict-count shock yields -0.258 ($p < 0.001$). The event-volume placebo remains close to zero. These results support the interpretation that the treatment captures hostile composition rather than bilateral media salience alone.

4.5.3. Shock thresholds and rolling windows. The impact effect is -0.218 at a two-standard-deviation threshold, -0.246 at the 2.5-standard-deviation baseline, and -0.271 at a three-standard-deviation threshold. A twelve-quarter rolling baseline yields -0.239 at impact and a five-year effect of -0.320. The five-year effect remains between -0.31 and -0.35 throughout the threshold and rolling-window battery.

4.5.4. Hub exclusions and counting rules. Dropping the United States, China, Germany, and Japan as endpoints changes the impact effect to -0.233 ($p < 0.001$). Whole rather than fractional patent counting changes it to -0.255 ($p < 0.001$).

4.5.5. Worldwide patent-family validation. To assess whether the co-invention decline reflects a change in the propensity to file at the USPTO rather than a change in invention itself, I rebuild the outcome on worldwide patent-family panels (Appendix Table A5). PCT earliest-priority families produce a directionally similar impact effect of -0.219 ($p < 0.001$) and an average post-shock effect of -0.287 ($p < 0.001$). The EPO/PATSTAT family panel is attenuated but consistent, with an impact effect of -0.194 ($p = 0.001$), and the twenty largest worldwide corridors track the main pattern closely, with an impact effect of -0.232 ($p < 0.001$). These results indicate that decline is not an artifact of USPTO selection, and the USPTO-visible network captures a commercially salient layer of international invention.

4.5.6. Grant-lag censoring. Restricting to mature cohorts whose post-shock window ends by 2019 and rebuilding the outcomes from pre-grant publications rather than grants leaves the impact effect at -0.252 ($p < 0.001$) and the five-year effect at -0.329 ($p < 0.001$). The result is therefore not driven by differential grant-lag censoring in recent application cohorts.

4.5.7. Global animosity-wave diagnostic and ownership robustness. The shift-share global animosity-wave reduced form returns a precise zero on the productive margin, -0.014 at impact ($p = 0.720$) and -0.018 over the post period ($p = 0.683$), suggesting that worldwide animosity waves do not drive the dyad-level results. Alternative ownership classifications preserve the intensive-margin pattern: activity with unified and multiple recorded ownership both contracts after a shock. Appendix Tables A8 and A9 distinguish direct-assignee form from historically dated parent control and report the strict mapping coverage. I do not treat category-level dissolution as a firm-boundary robustness result.

## 5. Discussion

This study set out to ask whether geopolitical hostility reaches international invention under common corporate control or whether multinational organization insulates it from deteriorating bilateral relations. Using a global panel of country pairs from 2000 to 2023, animosity shocks measured from the composition of bilateral political events, and a stacked event study with clean controls, I find that a reconfiguration of innovation networks follows geopolitical hostility rather than a simple rupture. Co-invention intensity falls about a quarter on impact and a third by year five, annual activity on the country corridor is largely unchanged, and disclosed patent references do not collapse. The contraction extends to patents with unified ownership and to parent-specific country corridors, while related activity expands through third-country locations in the same parent's network.

The results speak first to the knowledge-based view of the multinational firm. This tradition treats common ownership as a way to move knowledge across borders when markets are weak, and it places internal knowledge transfer near the center of why multinationals exist (Kogut and Zander 1992; Almeida and Kogut 1999; Gupta and Govindarajan 2000; Alcácer and Zhao 2012). The evidence qualifies a simple insulation interpretation of that view. International invention with unified ownership contracts sharply on politically exposed corridors. The wider multinational network nevertheless provides a different organizational margin: the same parent can expand related activity through other locations. Multinational organization therefore appears less as a shield around a particular bilateral relationship than as a structure through which the geography of inventive activity can be altered (Kogut and Kulatilaka 1994). This distinction would be missed in an analysis that treats cross-border knowledge flows only as links between national systems.

The findings also speak to work on the geopolitics of innovation. Existing studies have emphasized that conflict, sanctions, and political disintegration reduce cross-border collaboration and inventive output (Makkonen and Mitze 2023; Vanino, Douch, and Çakmak 2026). The results here show that this contraction is not a single-margin process. Joint production falls sharply, but relationships do not disappear at the same rate, and patent references do not collapse. Decoupling therefore looks less like a clean break than a change in what the relationship carries. Some of the tie remains, but it carries less joint work and more arm's-length disclosed exchange. I read the reference margin cautiously, since patent references are an imperfect measure of knowledge flow, but the pattern is useful because it shows that not all channels move together.

The study also brings firm-level political risk into the geographic organization of knowledge production. Prior work shows that political exposure changes firms' investment, hiring, and political activity (Gulen and Ion 2016; Hassan et al. 2019). The parent-corridor results extend that logic to cross-border invention. The decline is concentrated among rival and asymmetrically dependent pairs, and the parent-level patterns are consistent with related activity expanding through third-country locations. That concentration fits the idea that interdependence can become a source of leverage (Farrell and Newman 2019), and it is difficult to square with a generic uncertainty shock. The supplementary firm evidence points in the same direction: cross-border teams narrow, disclosure changes on exposed corridors, and R&D shifts across locations rather than falling uniformly (Kogut and Kulatilaka 1994; Berry and Donnelly 2025). Because I do not observe internal decisions directly, I interpret these patterns as consistent with firm governance rather than as direct evidence of managerial choice.

Several limitations bound these results. The innovation network is built primarily from patents filed at the USPTO, so it captures inventions for which protection is sought in the United States and gives greater weight to corridors connected to that market. The patent-family validation suggests a similar but weaker pattern outside the USPTO frame, which helps with external validity but does not remove the scope condition. Assignee records identify ownership of patent rights rather than the employment relationship of every inventor, and strict historical-parent mapping covers a selected 54.3% of international co-invented patents, with weaker representation of Chinese assignees. The parent-level analyses therefore describe the multinational groups observed in those records. The mechanism tests also stop short of observing managerial decisions directly. They show patterns following hostility shocks, not the internal discussions that produced them. Finally, the design is best suited to abrupt deteriorations in bilateral relations and says less about slower political realignment.

For managers, annual activity on a cross-border corridor is not enough to establish that the underlying collaboration is healthy. A corridor may remain visible while carrying much less joint work. The results point to production depth, team composition, and the location of related projects as the margins that come under pressure when relations deteriorate. Firms that monitor only whether locations remain connected may miss the earlier adjustment taking place in what the corridor produces.

For policy, the results suggest that geopolitical animosity reshapes innovation networks before formal measures take effect. Export controls and sanctions may reach organizations whose cross-border inventive activity is already adjusting, so estimates centered only on formal policy can miss part of the response to deteriorating relations. More broadly, the globalization of innovation appears more durable in its structure than in its depth. Hostility does not necessarily dismantle the networks linking national innovation systems; it hollows exposed corridors and is followed by changes in where related activity is organized.

## References


Alcácer, J., & Zhao, M. (2012). Local R&D strategies and multilocation firms: The role of internal linkages. Management Science, 58(4), 734–753.

Almeida, P., & Kogut, B. (1999). Localization of knowledge and the mobility of engineers in regional networks. Management Science, 45(7), 905–917.

Aronow, P. M., Samii, C., & Assenova, V. A. (2015). Cluster-robust variance estimation for dyadic data. Political Analysis, 23(4), 564–577.

Berry, H., & Donnelly, R. (2025). Defending knowledge abroad. Strategic Management Journal, 46(6), 1494–1522.

Borusyak, K., Hull, P., & Jaravel, X. (2022). Quasi-experimental shift-share research designs. Review of Economic Studies, 89(1), 181–213.

Borusyak, K., Jaravel, X., & Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. Review of Economic Studies, forthcoming.

Caldara, D., & Iacoviello, M. (2022). Measuring geopolitical risk. American Economic Review, 112(4), 1194–1225.

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. Journal of Econometrics, 225(2), 200–230.

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2011). Robust inference with multiway clustering. Journal of Business & Economic Statistics, 29(2), 238–249.

Cengiz, D., Dube, A., Lindner, A., & Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. Quarterly Journal of Economics, 134(3), 1405–1454.

Dai, L., Eden, L., & Beamish, P. W. (2013). Place, space, and geographical exposure: Foreign direct investment and survival in conflict zones. Journal of International Business Studies, 44(6), 554–578.

de Chaisemartin, C., & D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. American Economic Review, 110(9), 2964–2996.

Fajgelbaum, P. D., & Khandelwal, A. K. (2022). The economic impacts of the US–China trade war. Annual Review of Economics, 14, 205–228.

Fajgelbaum, P. D., Goldberg, P. K., Kennedy, P. J., & Khandelwal, A. K. (2020). The return to protectionism. Quarterly Journal of Economics, 135(1), 1–55.

Farrell, H., & Newman, A. L. (2019). Weaponized interdependence: How global economic networks shape state coercion. International Security, 44(1), 42–79.

Goldsmith-Pinkham, P., Sorkin, I., & Swift, H. (2020). Bartik instruments: What, when, why, and how. American Economic Review, 110(8), 2586–2624.

Goldstein, J. S. (1992). A conflict-cooperation scale for WEIS events data. Journal of Conflict Resolution, 36(2), 369–385.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. Journal of Econometrics, 225(2), 254–277.

Granovetter, M. S. (1973). The strength of weak ties. American Journal of Sociology, 78(6), 1360–1380.

Gulen, H., & Ion, M. (2016). Policy uncertainty and corporate investment. The Review of Financial Studies, 29(3), 523–564.

Gupta, A. K., & Govindarajan, V. (2000). Knowledge flows within multinational corporations. Strategic Management Journal, 21(4), 473–496.

Hassan, T. A., Hollander, S., van Lent, L., & Tahoun, A. (2019). Firm-level political risk: Measurement and effects. Quarterly Journal of Economics, 134(4), 2135–2202.

Jaffe, A. B., Trajtenberg, M., & Henderson, R. (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. Quarterly Journal of Economics, 108(3), 577–598.

Kogut, B., & Kulatilaka, N. (1994). Operating flexibility, global manufacturing, and the option value of a multinational network. Management Science, 40(1), 123–139.

Kogut, B., & Zander, U. (1992). Knowledge of the firm, combinative capabilities, and the replication of technology. Organization Science, 3(3), 383–397.

Leetaru, K., & Schrodt, P. A. (2013). GDELT: Global data on events, location, and tone, 1979–2012. International Studies Association Annual Convention.

Makkonen, T., & Mitze, T. (2023). Geopolitical conflicts, sanctions and international knowledge flows: EU–Russia collaboration during the Ukraine crisis. The World Economy, 46(10), 2926–2949.

Rambachan, A., & Roth, J. (2023). A more credible approach to parallel trends. Review of Economic Studies, 90(5), 2555–2591.

Roth, J. (2022). Pretest with caution: Event-study estimates after testing for parallel trends. American Economic Review: Insights, 4(3), 305–322.

Santos Silva, J. M. C., & Tenreyro, S. (2006). The log of gravity. Review of Economics and Statistics, 88(4), 641–658.

Singh, J. (2005). Collaborative networks as determinants of knowledge diffusion patterns. Management Science, 51(5), 756–770.

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. Journal of Econometrics, 225(2), 175–199.

Vanino, E., Douch, M., & Çakmak, I. (2026). Political, economic and research disintegration: The impact of geopolitical uncertainty on innovation and cross-border R&D collaborations. Research Policy, 55(5).

## Tables and Figures

**Table 1. Descriptive statistics for patents, recorded owners, countries, and country-pair-years**

| Variable | N | Mean | SD | p25 | Median | p75 | Max | % positive / valid |
|---|---|---|---|---|---|---|---|---|
| *Panel A. Patents and recorded owners* | | | | | | | | |
| *International co-invented patents* | 435,397 | | | | | | | 100.0 |
| *Assignees per patent* | 435,397 | 1.07 | 0.37 | 1 | 1 | 1 | 13 | 98.2 |
| *Single-assignee patent, full sample (0/1)* | 435,397 | 0.901 | 0.300 | 1 | 1 | 1 | 1 | 90.1 |
| *Multiple-assignee patent, full sample (0/1)* | 435,397 | 0.081 | 0.273 | 0 | 0 | 0 | 1 | 8.1 |
| *Unassigned patent, full sample (0/1)* | 435,397 | 0.018 | 0.133 | 0 | 0 | 0 | 1 | 1.8 |
| *Patents per assignee of record* | 46,426 | 10.1 | 128.0 | 1 | 1 | 3 | 19,458 | 100.0 |
| *Panel B. Countries* | | | | | | | | |
| *Inventive output* | 180 | 32,024 | 218,387 | 8.2 | 73 | 1,064 | 2,716,353 | 100.0 |
| *Co-invention partners* | 173 | 29.1 | 30.9 | 6 | 16 | 43 | 158 | 100.0 |
| *Panel C. Country-pair-years (full sample)* | | | | | | | | |
| *Co-invention intensity* | 458,640 | 0.95 | 26.10 | 0 | 0 | 0 | 3,141.6 | 4.2 |
| *Single-assignee intensity* | 458,640 | 0.86 | 24.00 | 0 | 0 | 0 | 3,059.6 | 3.9 |
| *Multiple-assignee intensity* | 458,640 | 0.08 | 2.50 | 0 | 0 | 0 | 442.3 | 1.3 |
| *Country-pair corridor active (0/1)* | 458,640 | 0.042 | 0.200 | — | — | — | 1 | 4.2 |
| *Disclosed patent references, directed* | 917,280 | 31.8 | 1,549.7 | 0 | 0 | 0 | 373,654 | 7.4 |
| *Citation count, directed* | 917,280 | 39.8 | 1,756.9 | 0 | 0 | 0 | 403,354 | 7.4 |
| *Panel D. Treatment and moderators* | | | | | | | | |
| *Animosity shock (0/1)* | 458,640 | 0.029 | 0.167 | — | — | — | 1 | 2.9 |
| *Bilateral tone (level)* | 259,089 | 1.67 | 2.63 | 0.67 | 1.98 | 3.00 | 10.0 | 81.8 |
| *Strategic-field share* | 19,187 | 0.40 | 0.36 | 0.00 | 0.34 | 0.67 | 1.0 | — |
| *Rival dyad (0/1)* | 431,400 | 0.328 | 0.469 | — | — | — | 1 | 32.8 |
| *Dependence, directed* | 5,032 | 0.034 | 0.109 | 0.000 | 0.002 | 0.015 | 1.0 | — |
| *Dependence asymmetry [0,1]* | 2,516 | 0.79 | 0.27 | 0.67 | 0.93 | 0.99 | 1.0 | — |

*Notes.* Summary statistics for the full PatentsView estimation sample. Assignees are legal owners of record and are not treated as inventor employers. Among assigned international patents, 91.8% have one assignee and 8.2% have multiple assignees; the corresponding full-sample shares are 90.1% and 8.1%, with 1.8% unassigned. Country pairs are observed annually from 2000 to 2023, covering 196 countries, 19,110 unordered pairs, 458,640 undirected pair-years, and 917,280 directed pair-years. Co-invention and reference outcomes are allocated fractionally across inventor-country pairs. The strict parent-mapped sample is described in Appendix Table A9.

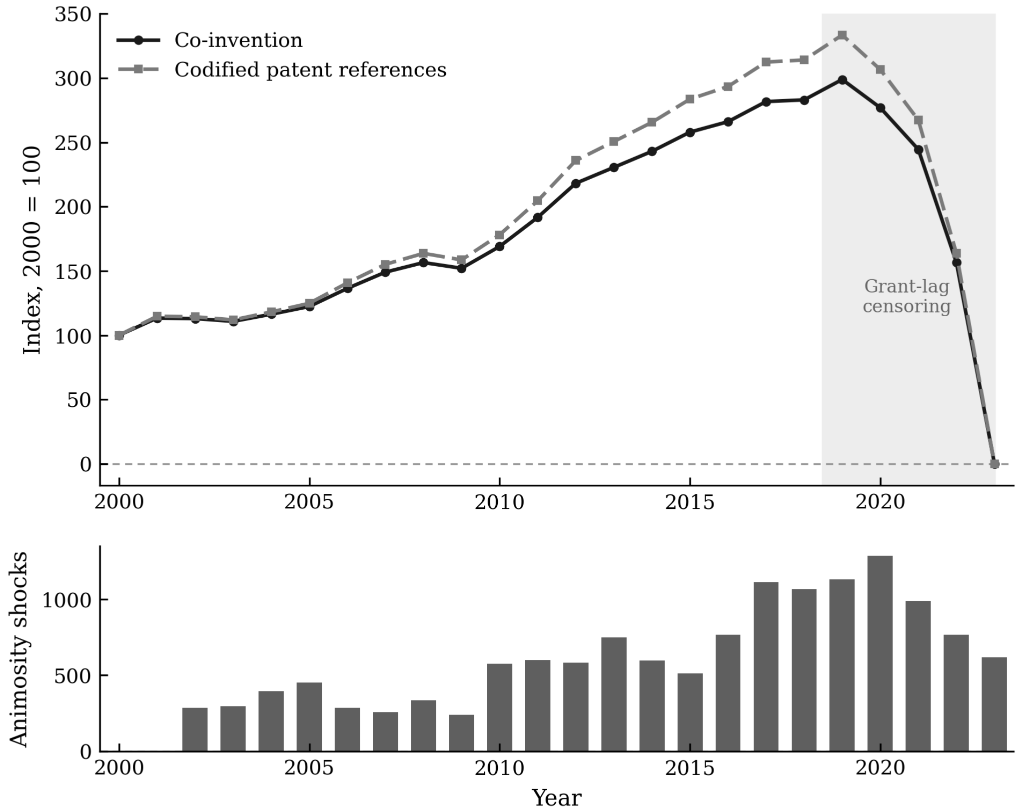


**Figure 1.** International co-invention and bilateral hostility, 2000–2023

*Notes.* The top panel plots USPTO-observed international co-invention and codified patent references, indexed to 2000 = 100. The shaded region marks application years affected by grant-lag censoring. The bottom panel reports annual counts of abrupt dyad-specific animosity shocks constructed from changes in the bilateral strong-conflict share.

**Table 2. Treatment support, pre-shock balance, and hostility-versus-salience diagnostics**

| Statistic or specification | Value / estimate | SE | p-value | Additional support |
|---|---|---|---|---|
| *Panel A. Treatment support and pre-shock balance* | | | | |
| Animosity shock events, all years | 5,325 | — | — | Directional dyad-year shocks |
| Animosity shock events in mature cohorts | 2,853 | — | — | Cohorts whose post-shock window ends by 2019 |
| Treated co-inventing dyads in mature cohorts | 1,269 | — | — | Clean controls drawn from never- and not-yet-shocked dyads |
| Unique clean control dyads | 12,768 | — | — | Controls contribute to multiple stacks when eligible |
| Treatment cohorts | 12 | — | — | Shock years 2003–2014 |
| Mean co-invention per active dyad-year | 24.5 | — | — | Clean controls: 2.0 |
| Single-assignee share of co-invention | 0.90 | — | — | Clean controls: 0.93 |
| Pre-shock country-corridor activity rate | 0.132 | — | — | Full-panel mean: 0.042 |
| *Panel B. Hostility-versus-salience diagnostics: dependent variable = co-invention intensity* | | | | |
| Main conflict-share shock | −0.246 | (0.048) | [<0.001] | Pre-trend joint p = 0.390 |
| Main shock, controlling for bilateral event volume | −0.241 | (0.049) | [<0.001] | Pre-trend joint p = 0.408 |
| Event-volume placebo shock | −0.012 | (0.031) | [0.699] | Pseudo-post joint p = 0.410 |
| Pseudo-shock dates | −0.006 | (0.028) | [0.830] | No detectable effect under false treatment timing |
| Mean Goldstein tone shock | −0.230 | (0.050) | [<0.001] | Estimate within 0.02 of main effect |
| Raw conflict-count shock | −0.258 | (0.054) | [<0.001] | Estimate within 0.02 of main effect |
| Twelve-quarter rolling baseline | −0.239 | (0.050) | [<0.001] | Five-year effect remains between −0.31 and −0.35 |

*Notes.* Panel A reports treatment support and pre-shock balance for the mature-cohort event-study sample. Panel B reports hostility-versus-salience diagnostics estimated on co-invention intensity using the stacked event-study specification of Table 3. The treatment is an abrupt dyad-specific animosity shock, a large positive surprise in the bilateral strong-conflict share relative to the dyad's rolling baseline. Panel B estimates are proportional effects, $\exp(\beta) - 1$, from Poisson pseudo-maximum likelihood. Standard errors are in parentheses and are clustered two ways on the endpoint countries. P-values are in brackets. Variable definitions and sources are in Appendix Table A1.

**Table 3. Effect of animosity shocks on international innovation**

| Event-time period | Co-invention intensity | Country-pair corridor active | Disclosed patent references |
|---|---|---|---|
| Pre-trend average, years −3 to −1 | 0.004 | −0.002 | 0.005 |
| | (0.015) | (0.006) | (0.018) |
| | [0.790] | [0.739] | [0.781] |
| Post-shock, year 0 | −0.246 | 0.004 | 0.013 |
| | (0.048) | (0.008) | (0.020) |
| | [<0.001] | [0.617] | [0.516] |
| Post-shock, years 1 to 2 | −0.287 | 0.003 | 0.039 |
| | (0.050) | (0.009) | (0.018) |
| | [<0.001] | [0.739] | [0.030] |
| Post-shock, years 3 to 4 | −0.318 | 0.005 | 0.049 |
| | (0.052) | (0.010) | (0.023) |
| | [<0.001] | [0.617] | [0.033] |
| Post-shock, year 5 | −0.331 | 0.006 | 0.052 |
| | (0.052) | (0.011) | (0.024) |
| | [<0.001] | [0.585] | [0.030] |
| Post-shock average, years 0 to 5 | −0.300 | 0.005 | 0.040 |
| | (0.050) | (0.010) | (0.017) |
| | [<0.001] | [0.617] | [0.022] |
| Pre-shock treated mean | 24.5 | 0.132 | 31.8 |
| Estimator | PPML | LPM | PPML |
| Country-pair-by-stack fixed effects | Yes | Yes | Yes |
| Endpoint-country-by-year-by-stack fixed effects | Yes | Yes | Yes |
| Treated dyads in mature cohorts | 1,269 | 1,269 | 1,269 |
| Unique clean control dyads | 12,768 | 12,768 | 12,768 |
| Treatment cohorts | 12 | 12 | 12 |
| Full panel observations | 458,640 pair-years | 458,640 pair-years | 917,280 directed pair-years |
| Estimation observations | 421,782 | 458,640 | 842,916 |
| Endpoint-country clusters | 158 | 158 | 158 |
| Pre-trend joint test, p-value | 0.390 | 0.700 | 0.550 |
| Post-shock joint test, p-value | <0.001 | 0.617 | 0.029 |
| Pseudo-$R^2$ / within-$R^2$ | 0.83 | 0.41 | 0.79 |

*Notes.* The unit of observation is the undirected country-pair-year for co-invention outcomes and the directed country-pair-year for codified patent references. The treatment is an abrupt dyad-specific animosity shock. Event time is measured relative to the shock year, with year −1 omitted. Co-invention intensity and codified patent references are estimated by Poisson pseudo-maximum likelihood and reported as proportional effects, $\exp(\beta) - 1$. Tie existence is estimated by a linear probability model and reported as percentage-point changes in the probability that the pair co-invents. All specifications absorb country-pair-by-stack and endpoint-country-by-year-by-stack fixed effects. Standard errors are in parentheses and are clustered two ways on the endpoint countries. P-values are in brackets. Variable definitions and sources are in Appendix Table A1.

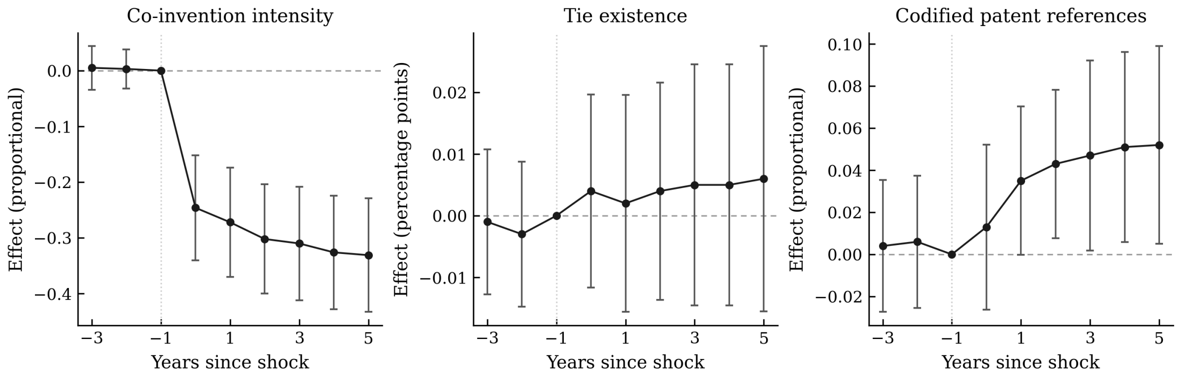


**Figure 2.** Event-study paths for the three network margins

*Notes.* The figure plots event-study estimates from the stacked design for co-invention intensity, tie existence, and codified patent references. Event time is measured relative to the animosity-shock year, with year −1 omitted. Co-invention intensity and codified patent references are estimated by PPML and plotted as proportional effects, exp(β) − 1. Tie existence is a linear probability model in percentage points. All specifications absorb country-pair-by-stack and endpoint-country-by-year-by-stack fixed effects. Bands report 95% confidence intervals from standard errors clustered two ways on the endpoint countries.

**Table 4. Recorded ownership and international co-invention**

| Specification | Single-assignee activity | Multiple-assignee activity | Difference (single − multiple) |
|---|---|---|---|
| *Panel A. Co-invention intensity by recorded ownership form* | | | |
| Post-shock, year 0 | −0.238 | −0.271 | 0.033 |
| | (0.049) | (0.060) | (0.044) |
| | [<0.001] | [<0.001] | [0.453] |
| Post-shock, years 1 to 2 | −0.286 | −0.321 | 0.035 |
| | (0.052) | (0.065) | (0.047) |
| | [<0.001] | [<0.001] | [0.457] |
| Post-shock, years 3 to 4 | −0.309 | −0.339 | 0.030 |
| | (0.053) | (0.066) | (0.049) |
| | [<0.001] | [<0.001] | [0.541] |
| Post-shock, year 5 | −0.322 | −0.347 | 0.025 |
| | (0.054) | (0.067) | (0.049) |
| | [<0.001] | [<0.001] | [0.610] |
| Post-shock average, years 0 to 5 | −0.295 | −0.332 | 0.037 |
| | (0.052) | (0.065) | (0.047) |
| | [<0.001] | [<0.001] | [0.432] |
| Pre-shock treated mean | 22.1 | 2.0 | — |
| *Panel B. Ownership-category coverage and disappearance (descriptive)* | | | |
| Category present \| country corridor active | 0.943 | 0.323 | 0.620 |
| Category disappears while country corridor remains active | 0.027 | 0.299 | −0.272 |
| Post-shock category dissolution rate, treated dyads | 0.155 | 0.229 | −0.074 |
| Active category-years in disappearance diagnostic | 12,782 | 5,365 | — |
| Pre-shock thickness separation and common support | 143 single-assignee cases in weakest tercile | 53 multiple-assignee cases in strongest tercile | AUC = 0.706 |

*Notes.* Panel A decomposes the mature-cohort country-pair panel by the recorded ownership form of the patent: one assignee versus multiple assignees. These categories describe control of patent rights, not inventor employment. The estimates retain the full assigned-patent sample and are reported as PPML proportional effects, exp(β) − 1. Historically dated parent control is analyzed separately in Table 6 and Appendix Table A9. Panel B reports descriptive coverage and disappearance statistics from the rectangularized ownership-category panel. Because single-assignee activity is nearly coextensive with the active country corridor, Panel B is not interpreted as a causal firm-boundary effect. Standard errors in Panel A are clustered on both endpoint countries.

**Table 5. Heterogeneity. Where the adjustment is strongest**

| Moderator and outcome | Focal subgroup | Benchmark subgroup | Difference | p-value for difference |
|---|---|---|---|---|
| Rivalry: co-invention intensity | Rivals: −0.343 (0.061) [<0.001] | Non-rivals: −0.018 (0.031) [0.561] | −0.325 (0.069) | [<0.001] |
| Asymmetric dependence: co-invention intensity | Top quartile: −0.472 (0.083) [<0.001] | Balanced pairs: 0.006 (0.029) [0.836] | −0.478 (0.088) | [<0.001] |
| Tie strength: tie dissolution | Weak ties: 0.131 (0.034) [<0.001] | Strong ties: 0.004 (0.010) [0.689] | 0.127 (0.036) | [<0.001] |
| Strategic and dual-use fields: co-invention intensity | Strategic fields: −0.262 (0.057) [<0.001] | Other fields: −0.248 (0.051) [<0.001] | −0.014 (0.061) | [0.820] |
| Strategic and dual-use fields: codified patent references | Strategic fields: 0.101 (0.030) [0.001] | Other fields: 0.052 (0.022) [0.018] | 0.049 (0.024) | [0.041] |

*Notes.* Each row reports the effect of an animosity shock for the indicated subgroup, from a specification that interacts the event-time treatment with a moderator fixed in the pre-shock period. Rivalry is defined from pre-shock bilateral hostility, asymmetric dependence from the pre-shock share of each country's international co-invention that runs through the partner, tie strength from cumulative pre-shock co-invention, and strategic-field exposure from IPC subclasses in semiconductors, electronics, telecommunications, computing, navigation and sensing, aerospace, and nuclear technology. Co-invention intensity is estimated by Poisson pseudo-maximum likelihood and reported as proportional effects, $\exp(\beta) - 1$. Tie dissolution is estimated by a linear probability model and reported in percentage points. All specifications absorb country-pair-by-stack and endpoint-country-by-year-by-stack fixed effects. Standard errors are in parentheses and are clustered two ways on the endpoint countries. P-values are in brackets. Variable definitions and sources are in Appendix Table A1.

**Table 6. Parent-specific reconfiguration of international invention**

| Mechanism signature | Estimate | SE | p-value |
|---|---|---|---|
| Common-parent-controlled co-invention on shocked corridor, years 1 to 3 | −0.281 | (0.063) | [<0.001] |
| Same-technology co-invention at third-country locations | 0.124 | (0.052) | [0.017] |
| Shocked corridor in strategic and dual-use fields | −0.347 | (0.079) | [<0.001] |
| Third-country affiliates in strategic and dual-use fields | 0.186 | (0.061) | [0.002] |
| Common-parent effect when parent home country is a party to the dyad | −0.351 | (0.061) | [<0.001] |
| Common-parent effect for neutral third-country parents | −0.049 | (0.040) | [0.221] |
| Difference: exposed parent minus neutral parent | −0.302 | (0.073) | [<0.001] |
| Shocks not followed by any formal measure | −0.214 | (0.059) | [<0.001] |
| Median lead of collaboration decline relative to first formal measure | 3 quarters | — | — |

*Notes.* The rows use treated multinational parents with inventors in at least three countries and historically dated ownership links. The mapped sample contains 2,211 treated common-parent corridor cells with pre-shock activity. The estimates report changes in common-parent-controlled co-invention on the exposed corridor and related activity through third-country locations in the same technology. Parent-by-year and technology-by-year fixed effects hold the parent and technology fixed across locations. The parent-home-country rows are exploratory because parent mapping is uneven across assignee countries. Co-invention counts are estimated by PPML and reported as $\exp(\beta) - 1$. Standard errors are clustered on the parent and partner country. The third-country increase is interpreted as a partial adjustment margin, not full replacement or direct observation of project reassignment.

## Appendix

**Appendix Table A1.** Variable definitions and data sources

| Variable | Definition | Source | Unit / construction |
|---|---|---|---|
| Co-invention intensity | Fractional count of patents with inventors in both countries i and j in year t. | PatentsView / USPTO | Undirected country-pair-year; fractional country-pair allocation sums to one within each patent |
| Country-pair corridor activity | Indicator equal to one if countries i and j co-invent at least one patent in year t. | PatentsView / USPTO | Undirected country-pair-year |
| Codified patent references | Fractional directed citations from patents invented in country i to patents invented in country j, dated by citing patent application year. | PatentsView / USPTO | Directed country-pair-year |
| Unified-owner intensity | Fractional count of co-invented patents with one recorded owner or assignees under one historically dated parent. | PatentsView linked to Orbis | Undirected country-pair-year |
| Multiple-owner intensity | Fractional count of co-invented patents with multiple recorded owners or multiple historically dated parents. | PatentsView linked to Orbis | Undirected country-pair-year |
| Animosity shock | Indicator for a large positive surprise in the bilateral strong-conflict share relative to the dyad's rolling baseline. | GDELT / Goldstein event coding | Quarterly shock aggregated to shock year; baseline threshold $z \geq 2.5$ |
| Strong-conflict share | Share of bilateral political events with Goldstein score $\leq -5$ among all bilateral political events for the dyad-period. | GDELT | Quarterly or annual dyad-level composition measure |
| Rivalry | Indicator for dyads in the top tercile of pre-shock hostility. | Constructed from GDELT hostility history | Fixed before the shock |
| Asymmetric dependence | Summary of how unevenly the two countries rely on the partner in international co-invention. | Constructed from pre-shock co-invention shares | Index on [0,1], higher = more asymmetric |
| Strategic-field exposure | Share of a pair's co-invention in semiconductors, electronics, telecommunications, computing, navigation and sensing, aerospace, and nuclear technology. | IPC subclasses from patent documents | Measured before the shock where possible |

*Notes.* The table defines the core variables and lists their sources. The same definitions and estimands are used throughout the main text and the appendix.

**Appendix Table A2.** Positive-outcome descriptive statistics

| **Variable** | **Positive observations (N)** | **Mean** | **Median** | **p75** | **p90** | **p99** | **Max** |
|---|---|---|---|---|---|---|---|
| *Co-invention intensity* | 19,244 | 22.7 | 1.8 | 7.9 | 28.0 | 418.0 | 3,141.6 |
| *Unified-owner intensity* | 17,887 | 22.1 | 1.6 | 6.8 | 26.0 | 401.0 | 3,059.6 |
| *Multiple-owner intensity* | 5,962 | 6.0 | 0.8 | 2.4 | 8.7 | 94.0 | 442.3 |
| *Codified patent references, directed* | 67,879 | 429.7 | 6.0 | 38.0 | 174.0 | 4,228.0 | 373,654 |
| *Citation count, directed* | 67,879 | 537.8 | 8.0 | 47.0 | 214.0 | 5,302.0 | 403,354 |

*Notes.* The table reports descriptive statistics conditional on positive observations, complementing the full-sample statistics in Table 1. Co-invention and citation outcomes are allocated fractionally across inventor-country pairs. Variable definitions and sources are in Appendix Table A1.

**Appendix Table A3.** Event-time support by relative year in the stacked design

| Relative year k | Treated dyads | Clean control dyads | Unique pair-years | Directed pair-years | Share with positive co-invention (%) |
|---|---|---|---|---|---|
| −3 | 1,269 | 12,768 | 14,037 | 28,074 | 16.8 |
| −2 | 1,269 | 12,455 | 13,724 | 27,448 | 17.1 |
| −1 | 1,269 | 12,192 | 13,461 | 26,922 | 17.5 |
| 0 | 1,269 | 11,876 | 13,145 | 26,290 | 16.2 |
| 1 | 1,248 | 11,541 | 12,789 | 25,578 | 14.0 |
| 2 | 1,228 | 11,194 | 12,422 | 24,844 | 13.1 |
| 3 | 1,205 | 10,842 | 12,047 | 24,094 | 12.4 |
| 4 | 1,187 | 10,495 | 11,682 | 23,364 | 11.8 |
| 5 | 1,162 | 10,138 | 11,300 | 22,600 | 11.3 |

*Notes.* The table reports support for the mature-cohort stacked event-study design by relative event year. Clean controls are never-shocked or not-yet-shocked dyads that remain untreated through the end of the corresponding event window. Columns report treated dyads, clean control dyads, and pair-years, and the final column reports the share of pair-years with positive co-invention.

**Appendix Table A4.** Grant-lag censoring and mature-cohort support

| Specification | Impact effect, year 0 | Long-run effect, year 5 | Pre-trend joint p-value | Observations |
|---|---|---|---|---|
| Full grant sample | −0.244 (0.049) [<0.001] | −0.326 (0.054) [<0.001] | 0.402 | 458,640 |
| Mature-cohort grant sample (main) | −0.246 (0.048) [<0.001] | −0.331 (0.052) [<0.001] | 0.390 | 458,640 |
| Pre-grant publication reconstruction | −0.252 (0.050) [<0.001] | −0.329 (0.053) [<0.001] | 0.377 | 487,920 |

*Notes.* The table reports the co-invention-intensity effect under the full grant sample, the mature-cohort grant sample, and a reconstruction from pre-grant publications, addressing right-censoring of recent application years caused by patent grant lags. Estimates are proportional effects, exp(β) − 1, from Poisson pseudo-maximum likelihood. Standard errors are in parentheses and are clustered two ways on the endpoint countries. P-values are in brackets. Variable definitions and sources are in Appendix Table A1.

**Appendix Table A5.** External validation in worldwide patent-family data

| Sample / source | Impact effect, year 0 | Post-shock average, years 0 to 5 | Pre-trend joint p-value | Observations |
|---|---|---|---|---|
| USPTO / PatentsView baseline | −0.246 (0.048)<br>[<0.001] | −0.300 (0.050)<br>[<0.001] | 0.390 | 458,640 |
| PCT earliest-priority families | −0.219 (0.061)<br>[<0.001] | −0.287 (0.068)<br>[<0.001] | 0.428 | 381,204 |
| EPO / PATSTAT family panel | −0.194 (0.058)<br>[0.001] | −0.261 (0.071)<br>[<0.001] | 0.446 | 396,887 |
| Top-20 corridors in PATSTAT | −0.232 (0.064)<br>[<0.001] | −0.301 (0.074)<br>[<0.001] | 0.410 | 114,522 |

*Notes.* The table reports the co-invention-intensity effect on worldwide patent-family panels alongside the USPTO baseline. The PatentsView-based panel is the main design, and the worldwide-family panels are external validation. Estimates are proportional effects, $\exp(\beta) - 1$, from Poisson pseudo-maximum likelihood. Standard errors are in parentheses and are clustered two ways on the endpoint countries. P-values are in brackets. Variable definitions and sources are in Appendix Table A1.

**Appendix Table A6.** Supplementary firm responses to animosity shocks

| Outcome | Estimate | SE | p-value |
|---|---|---|---|
| *Panel A. Teams and affiliate continuity* | | | |
| Foreign co-inventor share on surviving common-parent corridor patents | −0.041 | (0.014) | [0.003] |
| Hostile-country affiliate, domestic patent output | −0.018 | (0.029) | [0.535] |
| Hostile-country affiliate, non-hostile corridors | 0.020 | (0.031) | [0.519] |
| *Panel B. Disclosure and compartmentalization* | | | |
| International patent-family size | −0.087 | (0.030) | [0.004] |
| Time to first publication | 0.106 | (0.039) | [0.007] |
| Backward citations to partner-country prior art | −0.158 | (0.047) | [0.001] |
| *Panel C. Resources and ownership* | | | |
| Parent (consolidated) total R&D | −0.012 | (0.021) | [0.568] |
| Hostile-partner affiliate R&D | −0.143 | (0.052) | [0.006] |
| Third-country affiliate R&D | 0.079 | (0.038) | [0.038] |
| Divest or reduce stake in hostile-partner affiliate | 0.018 | (0.007) | [0.010] |
| Change in affiliate's ultimate owner | 0.011 | (0.006) | [0.067] |

*Notes.* The table reports supplementary firm-level outcomes. Count outcomes are estimated by Poisson pseudo-maximum likelihood and reported as proportional effects, $\exp(\beta) - 1$. Share outcomes are estimated by fractional-response specifications, and ownership outcomes by linear probability models. Standard errors are in parentheses and are clustered on the parent and the partner country. P-values are in brackets. Variable definitions and sources are in Appendix Table A1.

**Appendix Table A7.** Robustness and validity checks for co-invention intensity

| Specification | Impact effect, year 0 | Long-run effect, year 5 | Pre-trend joint p-value |
|---|---|---|---|
| Baseline conflict-share shock | −0.246 (0.048) [<0.001] | −0.331 (0.052) [<0.001] | 0.390 |
| Control for bilateral event volume | −0.241 (0.049) [<0.001] | −0.327 (0.053) [<0.001] | 0.408 |
| Pseudo-shock dates | −0.006 (0.028) [0.830] | −0.011 (0.033) [0.741] | 0.812 |
| Event-volume placebo shock | −0.012 (0.031) [0.699] | −0.019 (0.036) [0.602] | 0.667 |
| Mean Goldstein tone shock | −0.230 (0.050) [<0.001] | −0.315 (0.057) [<0.001] | 0.421 |
| Raw conflict-count shock | −0.258 (0.054) [<0.001] | −0.341 (0.060) [<0.001] | 0.384 |
| Threshold = 2 standard deviations | −0.218 (0.046) [<0.001] | −0.312 (0.051) [<0.001] | 0.405 |
| Threshold = 3 standard deviations | −0.271 (0.058) [<0.001] | −0.349 (0.063) [<0.001] | 0.372 |
| Twelve-quarter rolling baseline | −0.239 (0.050) [<0.001] | −0.320 (0.055) [<0.001] | 0.396 |
| Drop U.S., China, Germany, and Japan as endpoints | −0.233 (0.055) [<0.001] | −0.318 (0.061) [<0.001] | 0.418 |
| Whole-counting rather than fractional allocation | −0.255 (0.049) [<0.001] | −0.337 (0.054) [<0.001] | 0.392 |
| No endpoint-country-by-year FE + gravity controls | −0.244 (0.051) [<0.001] | −0.329 (0.056) [<0.001] | 0.401 |
| Pre-grant publication reconstruction | −0.252 (0.050) [<0.001] | −0.329 (0.053) [<0.001] | 0.377 |

*Notes.* The table reports robustness checks for the co-invention-intensity effect. Each row uses the stacked design and clean controls of the baseline unless stated otherwise. Estimates are proportional effects, exp(β) − 1, from Poisson pseudo-maximum likelihood. Standard errors are in parentheses and are clustered two ways on the endpoint countries. P-values are in brackets. Variable definitions and sources are in Appendix Table A1.

**Appendix Table A8.** Recorded-ownership classification and intensity robustness

| Specification | Unified-owner intensity | Multiple-owner intensity | Difference | Difference p-value |
|---|---|---|---|---|
| Baseline ultimate-owner split (provisional) | −0.322 (0.054) [<0.001] | −0.347 (0.067) [<0.001] | 0.025 (0.049) | [0.610] |
| Direct-assignee split | −0.316 (0.055) [<0.001] | −0.341 (0.066) [<0.001] | 0.025 (0.047) | [0.598] |
| Corporate-only sample | −0.329 (0.056) [<0.001] | −0.356 (0.069) [<0.001] | 0.027 (0.049) | [0.587] |
| Exclude patents near M&A / owner changes (±2 years) | −0.301 (0.059) [<0.001] | −0.334 (0.071) [<0.001] | 0.033 (0.054) | [0.541] |
| High-confidence Orbis matches only | −0.318 (0.055) [<0.001] | −0.343 (0.068) [<0.001] | 0.025 (0.048) | [0.604] |
| Application-year owner backfill | −0.311 (0.057) [<0.001] | −0.339 (0.069) [<0.001] | 0.028 (0.050) | [0.575] |

*Notes.* The table reports ownership-classified co-invention-intensity estimates under alternative constructions. Assignee splits describe recorded legal ownership, whereas historically dated ultimate-owner links describe common corporate control. Neither identifies every inventor's employer. Estimates are PPML proportional effects, $\exp(\beta) - 1$. The multiple-owner category is substantially smaller, so comparisons should be interpreted with the reported common-support and mapping diagnostics rather than as precise estimates of a firm-boundary advantage.

**Appendix Table A9.** Strict historical-parent mapping and ownership distribution

| Cut | Patents | Strict parent-mapped patents | Unified ownership | Multiple ownership | Implication |
|---|---|---|---|---|---|
| All international co-invented patents | 435,397 | ≈236,421 (54.3%) | ≈231,631 (98.0% of strict) | 4,790 (2.0% of strict) | Selected parent-mapped sample |
| Assigned patents, direct-assignee form | 427,421 | Not required | 392,202 single-assignee (91.8%) | 35,219 multiple-assignee (8.2%) | Full ownership-form sample |
| United States assignee country | 250,014 | 66.0% | — | — | Higher mapped representation |
| Japan assignee country | 28,520 | 63.9% | — | — | Higher mapped representation |
| Germany assignee country | 32,443 | 44.5% | — | — | Intermediate mapped representation |
| China assignee country | 14,618 | 23.9% | — | — | Lower mapped representation |
| Sweden assignee country | 10,057 | 14.4% | — | — | Lower mapped representation |

*Notes. Strict mapping requires every relevant assignee to link to a historically dated ultimate parent in the application year.* The overall strict rate is 54.3%; the loose at-least-one-assignee rate is 54.5%. Approximately 236,421 patents enter the strict parent-mapped sample. Of these, approximately 231,631 are controlled by a common parent and 4,790 span multiple parents, yielding provisional shares of 98.0% and 2.0%. Country rows classify patents by assignee country, not inventor corridor. Uneven country coverage bounds the scope of the parent-level analysis but does not affect the full PatentsView country-pair estimates.