Geopolitical Tensions and the Recomposition of International Invention

Luis Ballesteros

Boston University

Abstract

Rising geopolitical tensions are widely expected to fragment international innovation, yet it remains unclear whether firms end cross-border invention relationships or reorganize what and where they invent. Using USPTO patent records for 6,645 firms between 2000 and 2023, I compare invention relationships affected by abrupt rises in bilateral tensions with the same firms' other preexisting international relationships in a stacked event study. Within exposed relationships, firms reduce geopolitically sensitive co-invention by 12.2% while increasing co-invention in other technologies by 6.0%. Contraction reaches 14% for technologies prominent in the firm's prior patenting, against 8% for peripheral ones. Firms expand sensitive co-invention by 10.8% in their existing relationships with third countries, particularly those covering these technologies and one of the affected countries. Yet, the bilateral relationships under tension rarely end and their total co-invention barely changes. These patterns are consistent with firms preserving established inventive trajectories through technological recomposition across their international portfolios, while the network's visible structure remains stable. Conventional measures of relationship survival and aggregate output can conceal substantial technological and geographic reconfiguration that follows geopolitical tensions.

Keywords: geopolitical fragmentation, international innovation, technological recomposition, multinational flexibility, resource redeployment

1. Introduction

Innovation has become increasingly international, with patents listing inventors in more than one country nearly tripling between 2000 and 2018, even as relations among many of the countries involved have deteriorated (Kerr and Kerr 2018, Gopinath et al. 2025). When tensions rise between two countries, the invention that links them is exposed even when no new regulation or sanction follows. Firms respond to anticipated security review and to the prospect of abrupt policy reversal (Hassan et al. 2019, Jia et al. 2024). With no binding restriction yet in force, firms retain discretion over whether and how to adjust. Exposure is uneven within a single cross-border invention relationship, since technologies differ in how directly they bear on national security and industry priorities (Han et al. 2024, Chatterji and Murray 2026). Firms also differ in how much of their own invention a given technology accounts for. An invention relationship can therefore remain valuable in the aggregate while part of what it produces becomes exposed, which makes the firm's problem finer than whether to keep or abandon it. Do firms reconfigure the technological content and geography of international invention when an abrupt increase in bilateral tensions exposes geopolitically sensitive invention?

Prior research documents contractions in cross-border science and invention following sanctions, political divergence, travel restrictions, and rivalry between technological blocs (Makkonen and Mitze 2023, Charpin et al. 2024, Milani and Pham 2025, Koski 2026). Related work shows changes in who remains involved, including weaker new international research ties and the substitution of domestic for foreign inventors on continuing patents (Aghion et al. 2023, Vanino et al. 2026). Firms may also respond to bilateral political tension by increasing local patenting to signal their usefulness to the host-country government (Zhou et al. 2024). This evidence leaves unresolved whether firms respond to rising bilateral tensions through technological recomposition of their cross-border invention relationships.

A deterioration in bilateral relations can raise the expected cost of some invention conducted through a cross-border relationship without necessarily eroding the value accumulated in the relationship. Abandoning the relationship can remove the exposure but may sacrifice specialized knowledge, inventor familiarity, coordination routines, and value accumulated across locations (Singh 2005, Kellogg 2011, Alcácer and Zhao 2012, Berry 2014). An abrupt rise in tensions can also increase uncertainty because the firm cannot know whether tensions will harden into controls or recede (Caldara and Iacoviello 2022), which raises the value of deferring adjustments that are costly to reverse (Gulen and Ion 2016). Ending the relationship is typically among the most difficult responses to reverse. Reducing invention across the relationship uniformly may preserve those assets, but it would also curtail research that remains productive and comparatively unexposed. Thus, when geopolitical sensitivity varies across technologies and no formal restriction dictates the response, selective recomposition may allow the firm to reduce exposure in certain domains while maintaining the international relationship. This possibility aligns with research suggesting that firms restrict the scope of collaboration in international R&D alliances when hazards make exchange costly (Oxley and Sampson 2004), but comparative adjustment costs favor the preservation of valuable organizational arrangements (Argyres et al. 2019).

The stakes of this recomposition should depend on how prominent the technology carried by the bilateral relationship is in the firm's established inventive trajectory (Helfat 1994, 1997, Aghion et al. 2016). When an affected relationship carries a prominent technology, the resulting exposure reaches more deeply into the firm's prior inventive activity, raising the consequences of adjustment. Geographic reconfiguration can alleviate this problem when the firm already has alternative country relationships (Kogut and Kulatilaka 1994, Glennon 2023). However, their usefulness may also depend on whether they already carry the affected technologies (Sakhartov and Folta 2014, Sohl and Folta 2024).

To explore whether firms adjust the content and geography of cross-border invention when bilateral tensions rise, I construct a firm-country-pair-year panel from 435,397 internationally co-invented USPTO patents with application years 2000 through 2023. Using disambiguated legal-assignee records, I identify 6,645 firms with at least one invention relationship that experienced a bilateral tension shock. I define geopolitical sensitivity ex ante by mapping IPC subclasses to the technological domains covered by the Wassenaar Arrangement's control lists (Wassenaar Arrangement 2024). Separately, I measure each technology's pre-shock prominence as its share of the firm's patenting during the preceding five years, following patent-share measures of technological focus (Song et al. 2003, Ziedonis 2007, Galasso and Schankerman 2018).

I measure bilateral tensions from machine-coded news records of international interactions. For each country pair, I calculate the share of interactions that the system classifies as strongly conflictual, so the measure captures composition rather than volume. I identify a shock when the conflictual share rises abruptly relative to the pair's recent history. A stacked event study compares each affected relationship with the same firm's preexisting, unshocked relationships. I include firm-relationship-event and firm-event-year fixed effects to absorb relationship differences and shocks common to the firm's observed relationships. Equivalence tests assess aggregate stability, and an independent country-pair analysis provides a benchmark that does not use the firm's other relationships as controls. To examine geographic reconfiguration, I compare exposed and unexposed firms already active on the same alternative relationships and distinguish relationships retaining one affected country from those retaining neither.

Three empirical patterns follow an abrupt increase in bilateral tensions. First, affected invention relationships persist while what they produce changes. Durable dissolution remains rare through year 4, its predicted probability rising by less than one percentage point. Total co-invention changes little relative to the same firm's other eligible relationships, and the country-pair comparison yields the same aggregate conclusion, with an estimated decline of 2.6%. However, this stability conceals substantial reconfiguration. Relative to the same firm's eligible comparison relationships, geopolitically sensitive co-invention falls by 12.2% on average. Other classified co-invention rises by approximately 6% on average after the shock, while sensitive invention's share of classified co-invention on affected relationships falls by approximately 4.1 percentage points. The country-pair analysis yields a similar result, with sensitive co-invention declining by 9.8% while other classified co-invention changes little. Separately, across all classified technologies, co-invention falls by 14% when a technology's share of the firm's pre-shock patenting is at the 75th percentile, compared with 8% at the 25th percentile. Affected relationships therefore carry a materially different mix of invention even as total co-invention remains economically stable.

Second, affected firms expand geopolitically sensitive co-invention on their preexisting relationships with third countries. Relative to unexposed firms that were active on the same alternative country relationship before the shock, affected firms increase such co-invention by 10.8% on average. Geopolitically sensitive co-invention grows by approximately 15% on relationships that retain one affected country and by 8% on relationships that retain neither. A separate within-firm comparison estimates 6.3% greater growth on relationships retaining one affected country. Restricting the outcome to the same IPC subclasses yields corresponding estimates of 19% and 9%. Other classified co-invention expands almost evenly across destinations, growing only about 1.5% faster on relationships retaining an affected country, compared with 6.3% for sensitive co-invention. This geographic pattern suggests expansion outside the bilateral relationship under tension while preserving involvement with one of the affected countries.

Third, not every firm affected by bilateral tensions can make this adjustment. Firms with pre-shock relationships already active in the affected technologies are predicted to reduce sensitive co-invention on the affected relationship by approximately 15% and expand it elsewhere by approximately 17%. Firms with many alternatives but none active in those technologies reduce it by 12.5%, while firms with few alternatives show little change. Even as geopolitically sensitive co-invention contracts, citations from sensitive applications to earlier patents associated with the other affected country rise by approximately 8%, compared with 3% in other fields. This evidence suggests that prior art associated with that country remains relevant to subsequent invention.

The combined pattern is difficult to reconcile with several plausible alternative explanations. First, contemporaneous formal restrictions do not account for the estimates, which remain similar in direction and magnitude when affected-pair shocks that coincide with sanctions or export controls are excluded (Farrell and Newman 2019, Crosignani et al. 2026). Second, if the estimates reflected the general contraction in cross-border collaboration documented when bilateral relations deteriorate (Makkonen and

Mitze 2023, Milani and Pham 2025), total and other classified co-invention should also decline, yet both remain stable or rise. Third, political scrutiny broad enough to induce withdrawal from both countries in the affected pair (Jia et al. 2024, Furnas et al. 2026) should shift expansion toward relationships retaining neither country, yet expansion concentrates on relationships retaining one. Fourth, redirection of invention toward other technologies (Aghion et al. 2016, Moscona and Sastry 2023) would place the expansion outside the technologies previously carried by the affected relationship, yet the results show that expansion in third-country relationships concentrates exactly there. Finally, if firms were simply shedding marginal invention activity (Vanino et al. 2026), declines should be largest in technologies peripheral to their prior patenting, yet they are larger in the most prominent technologies. The results are robust to an external country-pair comparison, an earlier eligibility window, fixed pre-shock comparison portfolios, alternative measures of bilateral tension and geopolitical sensitivity, and analyses of application timing and bilateral news volume.

These findings contribute to two literatures. First, the study advances research on geopolitical fragmentation and international innovation by identifying technological recomposition within a continuing cross-border relationship as a distinct margin of organizational adjustment (Makkonen and Mitze 2023, Aghion et al. 2023, Koski 2026, Vanino et al. 2026). Prior research generally locates hazards in the collaborator or in exchange with that collaborator, leading firms to narrow scope, adopt protective governance, or alter inventor participation (Oxley and Sampson 2004, Ryu et al. 2018, Palomeras and Wehrheim 2021). Bilateral tensions create a different type of hazard because exposure attaches selectively to geopolitically sensitive technologies passing through an affected country relationship. Firms can therefore preserve the relationship and its accumulated relationship-specific assets while changing the technologies assigned to it. The larger contraction among technologies prominent in the firm's prior patenting shows that this adjustment reaches established inventive trajectories. Persistent ties and stable aggregate output can therefore conceal selective fragmentation, making technological composition essential for assessing the depth of international integration.

Second, the findings contribute to research on multinational operating flexibility and resource redeployment. Existing work shows that geographic portfolios create adaptive options (Kogut and Kulatilaka 1994, Glennon 2023), that options lose value when exposures are correlated (Belderbos et al. 2014), and that international breadth can enable reallocation once policy restrictions bind (Ma and Clougherty 2026). This study distinguishes geographic breadth from technology-matched relational capacity, the degree to which other relationships already carry the exposed technologies when tensions rise before any formal restriction is implemented. A usable alternative is likely to be separated from the affected relationship in geopolitical exposure and close to it in technological experience (Helfat and Eisenhardt

2004, Sakhartov and Folta 2014, Sohl and Folta 2024, Branstetter et al. 2026). In identifying these conditions, the study extends multinational flexibility from the distribution of assets and operating units to the technology configuration of established relationships through which firms combine knowledge across countries.

2. Data and Measures

2.1. International Co-Invention, Legal Assignees, and Country Relationships

I build the invention panel from PatentsView, which links USPTO records to disambiguated inventors, assignees, locations, technology classifications, and citations. The data begin with 435,397 patent records with application years 2000 through 2023 that list inventors in at least two countries. After excluding 946 reissues and 329 plant patents, the patent sample contains 434,122 unique patents. Online Appendix Table A1 reports the complete patent and assignee reconciliation, including the relationship between the corporate patent sample and the broader corporate extraction. I aggregate inventive activity by application year because this date is closer to when the underlying work occurred. The data measure USPTO-visible patented co-invention in a common patent system used across countries (Jaffe et al. 1993, Singh 2005, Alcácer and Zhao 2012). Because the source consists of granted patents dated by application year, grant lags make the most recent application cohorts incomplete. With application years through 2023, 2019 is the latest treatment cohort in my sample that provides the full years 0–4 calendar window needed for the analyses.

Given that co-invention is undirected, I represent an international relationship as an unordered pair of inventor countries. I define a firm-country-pair relationship as co-invention on patents assigned to a firm that involve inventors located in both countries. The measure captures the bilateral geography of the firm's patented inventive activity, irrespective of whether the participating inventors share a formal organizational affiliation. A patent with inventors in more than two countries can connect several pairs. For each pair on a patent, I calculate the product of the number of inventors at its two endpoints and normalize these products to sum to one within the patent. This procedure assigns greater weight to the country pairs that account for more of the inventive team while preserving one patent equivalent across all pairs. Whole counting provides a sensitivity analysis.

The organizational actor is the normalized matched legal assignee. I retain PatentsView corporate-assignee types 2 and 3, standardize disambiguated organizations to stable legal-assignee identifiers, and preserve all distinct matched assignees recorded on a patent. When a patent has multiple matched assignees, each receives an equal share. The weight of a patent-assignee-pair observation is therefore the country-pair weight divided by the number of distinct matched legal assignees. This allocation avoids a first-assignee rule and preserves the patent totals when the firm-level panel is aggregated. Patents without a corporate assignee remain in the external country-pair benchmark but do not enter the firm-level panel. I refer to the matched legal assignee as the firm in the remainder of the paper. The main unit is the firm-country-pair-year, which places each affected relationship inside the same firm's international portfolio. The legal assignee remains the organizational unit because assigning patents to later corporate parents would attribute inventions predating an acquisition to organizations that did not own them at the time. Analyses separating single-assignee and multiple-assignee patents provide measurement decompositions. A single assignee does not establish that all participating inventors worked within one organizational boundary.

2.2. Bilateral Tension Shocks

I measure bilateral tensions from political interactions recorded by the Global Database of Events, Language, and Tone (GDELT). GDELT machine-codes reported interactions between countries and assigns each event type a score on the Goldstein conflict-cooperation scale (Goldstein 1992, Leetaru and Schrodt 2013). I pool interactions in both directions of a country pair and define strongly conflictual events as those with Goldstein scores at or below -5. For each pair-quarter, the strong-conflict share is the number of strongly conflictual events divided by all recorded bilateral political events. The share isolates the composition of bilateral interaction from the amount of news attention. I examine event volume separately as a measure of media salience and data coverage.

A shock in bilateral tensions is a sharp increase in the strong-conflict share relative to the pair's own recent path. I first remove a pair-specific linear time trend from the quarterly conflict-share series. I then calculate the quarterly change and standardize the detrended change against the mean and standard deviation of the preceding eight quarterly changes. The quarter qualifies when the standardized increase is at least 2.5 and the focal quarter contains at least 25 bilateral events. The main treatment is the first qualifying shock for each unordered country pair. I record first- and second-quarter shocks in the current application year and third- and fourth-quarter shocks in the following application year. This convention prevents applications filed before a late-year deterioration from being coded as post-shock invention. Section 3.7 defines the exact-date and first-complete-post-year specifications used to address that timing issue. The retained treatment cohorts span 2003 through 2019. The analysis window extends from three years before through four years after the treatment year, with the year immediately before the shock as the reference period. Comparison relationships must remain unshocked through event year 0 and are right-censored at the first subsequent qualifying shock.

This construction identifies abrupt, pair-specific deteriorations relative to persistent tensions. A prespecified audit captures 21 of 24 recognizable bilateral escalations within one quarter of their documented dates. Online Appendix Table A4 and Figure A1 report this audit, the shock magnitude, the subsequent conflict path, the cohort distribution, and an event-volume placebo. I assess its sensitivity to alternative conflict thresholds, minimum-event requirements, and tone measures. The shock is observational. Section 3 states the identifying assumptions and reports the pre-event paths and treatment diagnostics.

I construct a directional country-pair record of formal sanctions, export controls, and investment restrictions using official records from the U.S. Office of Foreign Assets Control, Bureau of Industry and Security, and Committee on Foreign Investment in the United States, together with records from the European Union, United Nations, and United Kingdom. For each measure, I record the implementing and target countries, restriction type, announcement and effective dates, and termination where available. I preserve direction during coding and then match restrictions in either direction to the unordered inventor-country pair using official country fields. The contemporaneous indicator covers restrictions announced or effective from 90 days before through 90 days after the tension-shock quarter. Entity names are not matched to patent assignees. Online Appendix Table A14 reports the sources, coding rules, pair matches, and exclusion results.

2.3. Geopolitically Sensitive Invention and Relationship Outcomes

I identify geopolitically sensitive invention by constructing a crosswalk between normalized International Patent Classification (IPC) subclasses and technological domains represented in the Wassenaar Arrangement's control lists. The Arrangement coordinates export controls among participating states for conventional arms and dual-use goods and technologies (Wassenaar Arrangement 2024). The mapping covers documented subclasses in advanced electronics and semiconductors, telecommunications and information security, computing, navigation and sensing, aerospace, and nuclear technologies. It captures domains with security and dual-use relevance. The legal control status and commercial importance of an individual patent remain separate attributes. Online Appendix Table A2 reports the complete crosswalk, documentary source and rationale for each mapping, IPC normalization rules, and any-code, exclusive-code, mixed-code, leave-one-domain-out, and alternative critical-technology classifications.

I normalize IPC codes before applying the mapping, including restoration of leading zeros in subclasses such as G06F and H04L. When a patent is assigned to more than one distinct normalized subclass, I divide its fractional weight equally across those subclasses. Each patent-assignee-pair-subclass observation therefore combines the country-pair, assignee, and subclass weights. I classify the resulting fractional patent counts as geopolitically sensitive or other classified invention. Together, the two categories sum exactly to classified patenting at every level of aggregation. Patents without a usable normalized IPC code remain in all-patent totals but do not enter the primary sensitive-versus-other comparison. In the classified corporate sample, 34.68% of fractional patenting falls in geopolitically sensitive domains.

The outcome system separates relationship scale, continuity, and technological composition. Total co-invention is the relationship's annual fractional patent count across all patents. Sensitive co-invention, other classified co-invention, and their sum measure technological content on the same classified-patent sample. The sensitive share divides sensitive co-invention by total classified co-invention. Annual activity equals one when the relationship records positive co-invention in a year. Among relationships active before the shock and observed through year 4, durable dissolution equals one when the relationship records at least two consecutive years without co-invention, beginning no later than year 3, and does not resume activity through year 4.

2.4. Technological Prominence

Pre-shock technological prominence measures the share of the firm's prior observed invention accounted for by a focal classified IPC subclass. The analysis covers all classified subclasses that meet the sample criteria and is estimated separately from the comparison between geopolitically sensitive and other classified invention. I construct the measure from the broader universe of the firm's classified USPTO utility patents, including domestic and international invention. I build the extraction from 1995 to provide a complete prior window for the earliest eligible events. For firm f, focal classified IPC subclass k, and event e, prominence is the firm's fractional patent count in subclass k divided by its total fractional classified patent count during the twenty complete grant quarters before the actual shock quarter. The denominator requires at least ten fractional patent equivalents. A focal subclass absent from an otherwise valid portfolio receives a prominence value of zero. Firm-events with a missing classified-patent denominator or fewer than ten patent equivalents remain missing rather than being coded as zero. Thresholds of five and twenty assess sensitivity to thin portfolios.

The numerator divides the fractional patent count across distinct matched legal assignees and normalized IPC subclasses using the same rules as the outcome data. The denominator spans all subclasses and inventor geographies, which prevents the affected country relationship from defining the firm's technological profile. The primary measure uses grant timing, which attributes only patents whose legal assignee is observed before the shock. An application-timed measure and a twenty-quarter window ending before the first reported event-study lead assess timing and temporal overlap.

Prominence ranges from zero to one and remains continuous in the analyses. A larger value indicates that the focal subclass occupied a larger share of the firm's prior USPTO-observed patenting. Its scope leaves product-market importance, managerial priority, and the technological core of a consolidated corporate group outside the construct. The measure follows prior work that uses the share of a firm's patents in a focal field to characterize technological focus or core technological areas (Song et al. 2003, Ziedonis 2007, Galasso and Schankerman 2018).

2.5. Pre-Shock Relationship Portfolios, Alternative Relationships, and Technology-Matched Relational Capacity

For each firm-event, I define two fixed pre-shock relationship portfolios. The recent-window sample admits every international country relationship with positive co-invention during event years −3 through −1 and fixes in that window all activity-dependent criteria for relationship membership, mixed-content eligibility, alternative-relationship eligibility, and focal-subclass eligibility. Sample relationships remain in the panel after later zero outcomes. Eligible alternatives must remain free of a competing shock through event year 0 and are right-censored at the first subsequent competing shock. Because the qualification window overlaps the displayed leads, these lead estimates are conditional on sample construction. The earlier-window sample applies the same rules during years −6 through −4 and follows the resulting relationships from event year −3 through year 4. It provides the primary pre-event diagnostic for relationships active before the displayed window and defines the across-firm geographic analysis. In that analysis, affected and unexposed firms must both have been active on the same alternative country relationship during years −6 through −4. Eligibility is fixed in that window, and later zeros are retained. The recent-window sample supplies the main within-firm post-shock estimates. Online Appendix Table A5 and Figure A2 compare complete post-shock paths across the two samples, including among firm-events eligible under both rules. Relationships formed after treatment enter descriptive patent-equivalent accounting but never define the main comparison set.

I distinguish alternative relationships by whether they retain either country in the affected pair. For an affected pair A-B, endpoint-retaining alternatives connect A or B with a third country (C), while endpoint-free alternatives connect two countries outside the affected pair. A patent listing inventors in A, B, and a third country still involves the affected pair. I therefore exclude such patents from alternative-relationship outcomes even though fractional country-pair allocation also creates A-C and B-C dyads. Technology-matched outcomes further restrict alternative activity to the normalized IPC subclasses carried by the affected relationship during the relevant pre-shock window.

Technology-matched relational capacity measures the share of the affected sensitive technology bundle that the firm already pursued on another eligible relationship. For each geopolitically sensitive subclass carried by the affected relationship during the relevant pre-shock window, I record whether another fixed, shock-free relationship carried that subclass during the same window. I weight this indicator by the subclass's share of sensitive invention on the affected relationship and sum across subclasses. The resulting measure ranges from zero to one. Relational capacity differs from prominence. Prominence measures how much the technology accounts for in the firm's complete prior patent portfolio. Capacity measures whether the firm's other preexisting country relationships already carried the technologies exposed on a particular bilateral relationship. I use raw relationship breadth and breadth in unrelated technologies as benchmarks.

2.6. Citations to Prior Art Associated with the Other Affected Country

I construct a separate directed firm-country-pair-year panel for citations. An observation records citations from patents assigned to firm f and associated with inventor country i to earlier patents associated with inventor country j. I date a citation by the application year of the citing patent, allocate it across the citing and cited inventor countries and the citing patent's legal assignees, and classify technological sensitivity from the citing patent. The panel fixes directed country links with positive citations before the shock and retains later zero-citation cells. The primary specification excludes same-family citations, limits cited prior art to patents first publicly disclosed before the shock, and removes citing patents that also contain inventors from country j. These restrictions separate citations to preexisting disclosed prior art from citations mechanically embedded in continuing joint invention on the affected relationship.

The primary citation analysis retains only citations that USPTO grant records classify as "cited by applicant," focusing the measure on prior art disclosed through the applicant side of prosecution. Citations with missing source information remain missing. The analysis excludes examiner-added citations, which partly reflect patent-office search and examination practices and vary across applicants, countries, and technologies (Alcácer and Gittelman 2006, Alcácer et al. 2009). The primary outcome is an opportunity-adjusted citation rate, implemented by including the logarithm of the number of eligible citing patents as the PPML exposure. The rate model asks whether each eligible patent becomes more likely to cite earlier patents associated with the other affected country. A companion count model captures the aggregate number of qualifying citations, which can also change when the number of eligible citing patents changes. Citations provide a bounded measure of connection to disclosed prior art. Their interpretation stops short of interpersonal knowledge transmission or the full geography of knowledge flows. Online Appendix Table A8 documents citation provenance, disclosure timing, same-family and other-country-inventor exclusions, common samples, and the count and opportunity-adjusted specifications.

2.7. Sample Composition and Descriptive Statistics

The broad relationship sample includes firm-events with a matched legal assignee, an affected relationship, at least one clean preexisting comparison relationship, and calendar coverage from event year −3 through event year 4. Later competing shocks can right-censor otherwise eligible relationships. I use this sample for total co-invention and relationship continuity. The mixed-content sample additionally requires pre-shock activity in both geopolitically sensitive and other classified invention and uses identical observations for both outcomes. This requirement allows me to conduct a direct comparison without selecting relationships on post-shock activity. The prominence, same-IPC geographic, relational-capacity, and citation analyses impose the corresponding pre-shock eligibility criteria described above and retain eligible zeros after treatment. Online Appendix Table A3 reports sample-construction counts and PPML separation losses for each analysis.

Within the 24-year analytic window, the patent sample represents 173 inventor countries, 14,878 possible unordered country pairs, and 357,072 pair-years. Among the 434,122 utility patents, 391,093 have one assignee record, 35,103 have multiple assignee records, and 7,926 have none. A further 2,844 assigned patents have no matched corporate assignee. The corporate patent sample contains 423,352 patents and 44,302 normalized legal assignees. Of these patents, 423,083 contain a usable normalized IPC subclass. After the tension-history and event-coverage restrictions, the external benchmark includes 593 affected first-shock country pairs from treatment cohorts spanning 2003 through 2019, together with eligible never-treated and not-yet-treated comparison pairs. The firm-level construction identifies 6,645 legal assignees active on an affected relationship and 25,841 legal-assignee-by-affected-pair shock events. Of these events, 16,894 have positive classified patenting in the twenty-quarter prominence window and 12,758 satisfy the ten-patent-equivalent requirement. Mature technology-resolved prominence data include 2,020 legal assignees and 7,829 distinct legal-assignee-by-IPC-subclass pairs. These are pre-separation coverage counts rather than final regression observations; Online Appendix Tables A3 and A9 report the corresponding annual sample. In the recent-window raw-eligibility construction, the broad geographic sample contains 2,725 affected firm-events involving 2,106 firms and 249 affected country pairs, generating 10,313 eligible firm-event-relationship cells and 48,511 affected annual rows. In the recent-window raw-eligibility construction, the same-IPC geographic sample contains 1,787 affected firm-events involving 1,513 firms and 201 affected pairs, generating 15,844 eligible firm-event-relationship-subclass cells and 72,766 affected annual rows. Online Appendix Tables A1–A3 and A7 report the detailed construction, comparison counts, strata restrictions, annual observations, and PPML separation losses.

International co-invention observed in USPTO patents is concentrated in a relatively small set of established country relationships. Of the 14,878 possible country pairs, 2,514 appear at least once, representing 16.90% of possible pairs. These relationships produce 19,169 positive pair-years, or 5.37% of the 357,072 possible pair-years. The ten largest country pairs account for 49.2% of fractional co-invention, and nearly all of the largest involve the United States, consistent with the USPTO sampling frame. Table 1 reports network density and distributional statistics for the full patent sample and each estimation sample, including pre-shock sensitive shares, technological prominence, relational capacity, and activity across exposed and alternative relationships. For each design, the table compares exposed observations with the eligible comparison observations that identify the estimates.

[Insert Table 1 about here]

3. Empirical Strategy

3.1. Main Within-Assignee Stacked Design

The baseline design places each affected country relationship inside the same firm’s fixed pre-shock portfolio. The main estimates define relationships and other activity-based eligibility criteria using years −3 through −1. Because this window overlaps the displayed leads, coefficients for years −3 and −2 are conditional on sample construction. An earlier-window analysis instead fixes eligibility in years −6 through −4 and estimates the same path from year −3 through year 4. Its leads provide the primary pre-event diagnostic for relationships already active before the displayed window. I report both paths and a comparison among firm-events eligible under both rules.

For firm *f*, relationship *r*, shock event *e*, and year *t*, I estimate

$$\mathbb{E}\left[Y_{fret} \mid \mathcal{F}\right] = \exp\left\{\sum_{\substack{k=-3 \\ k\neq-1}}^{4} \beta_k \, D_{fret}^{k} + \alpha_{fre} + \lambda_{fet}\right\} \tag{1}$$

The indicator *D* marks the affected relationship at event time *k*. Firm-relationship-event fixed effects absorb stable relationship differences. Firm-event-year fixed effects absorb annual changes shared across the firm's relationships, including patenting scale, financial condition, and research priorities. Firm-year covariates add no identifying variation and would restrict coverage, so I omit them. I estimate fractional patent outcomes by Poisson pseudo-maximum likelihood, which retains zeros and remains consistent under general conditional heteroskedasticity (Santos Silva and Tenreyro 2006). I report exponentiated coefficients minus one as proportional effects. A pooled model uses one post-shock indicator for years 0 through 4.

Equation (1) measures how activity on the affected relationship changes relative to the same firm’s other eligible relationships in the same year. If those other relationships expand in response to the shock, that expansion becomes part of the estimated contrast. The coefficient therefore cannot recover the absolute change that the affected firm-relationship would have experienced without the shock. I use the external country-pair design to provide an aggregate bilateral benchmark relative to eligible comparison pairs and the geographic models to identify where activity expands within and beyond the firm's portfolio.

3.2. External Country-Pair Benchmark

The external benchmark estimates total co-invention and classified technological content at the country-pair level without using affected firms’ other relationships as controls. The model aggregates fractional patent counts across legal assignees, other organizations, and patents without a matched corporate assignee. The content model estimates sensitive and other classified co-invention jointly on the same pair-years and comparison observations. For each treatment cohort, I combine pairs first shocked in that year with country pairs that are never shocked or remain unshocked through the end of the cohort's event window. Following the stacked-event approach in Cengiz et al. (2019), I estimate

$$\mathbb{E}\left[Y_{ijt}^{s} \mid \mathcal{F}\right] = \exp\left\{\sum_{\substack{k=-3 \\ k\neq -1}}^{4} \delta_k\, D_{ijt}^{s,k} + \mu_{ij}^{s} + \eta_{it}^{s} + \eta_{jt}^{s}\right\} \tag{2}$$

Here, s denotes the cohort stack, and i-j is an unordered country pair. Pair-by-stack fixed effects absorb persistent bilateral attributes. Endpoint-country-by-year-by-stack fixed effects absorb annual changes in each country's innovation system, patenting, and geopolitical environment. I impose these effects symmetrically so that estimates do not depend on endpoint ordering. The model uses PPML with dyadic-robust inference that allows dependence between observations whose country pairs share an endpoint (Aronow et al. 2015). I estimate total co-invention separately. For technological content, I interact exposure and event time with an indicator for sensitive invention while allowing the fixed effects to vary by content category. The interaction directly estimates the sensitive-versus-other divergence on identical observations. The coefficients describe aggregate change on shocked country pairs relative to eligible comparison pairs. They do not identify absolute change for a treated firm or convert the within-firm relative estimate into an absolute firm-level effect. The baseline risk set allows a comparison pair to share one country with the shocked pair. Because adjustment involving an affected country could then enter the comparison group, Table 7 reports specifications that remove affected firms' own contributions and exclude every comparison pair sharing either affected country. The endpoint-disjoint model is also estimated without endpoint-country-by-year fixed effects.

3.3. Technological Composition, Continuity, and Equivalence

The primary composition test uses the fixed mixed-content sample. Eligibility requires pre-shock activity in both categories on the affected relationship and its comparison portfolio. Admitted relationships remain after later zeros. I expand each firm-relationship-event-year into geopolitically sensitive and other classified invention and estimate both outcomes together. Content-specific firm-relationship-event fixed effects allow different baseline levels, while content-specific firm-event-year fixed effects absorb annual changes in the firm's allocation to each category across relationships. The coefficient of interest is the interaction between the affected relationship, event time, and sensitive content. Affected-relationship-by-event-time terms capture changes common to both content categories. The interaction therefore compares the sensitive-versus-other change on the affected relationship with the corresponding contrast in the firm's eligible comparison portfolio. This direct event-time contrast is the primary test of selective recomposition. Its exponentiated value gives the change in the sensitive-to-other ratio on the affected relationship relative to the corresponding ratio elsewhere in the firm's portfolio. I report the common observations and the number of rows removed through PPML separation.

I use the broader relationship sample for all-patent total co-invention and continuity. For durable dissolution, I collapse the completed post-shock window to one observation for each relationship active before treatment and observed through year 4. I estimate the outcome defined in Section 2.3 using a linear probability model with firm-event fixed effects and pretreatment measures of the relationship's activity history. I report standardized counterfactual predictions for the estimation sample after setting exposure to zero and one, together with their percentage-point difference.

To estimate the change in the sensitive share, I use the joint model to predict geopolitically sensitive and other classified co-invention for each observation with the exposure indicator set first to one and then to zero. I calculate the sensitive share under each condition by dividing predicted sensitive co-invention by total predicted classified co-invention and aggregate these estimates using the model's weights. I use these measures to estimate how much of the model-implied decline in sensitive co-invention is offset by growth in other classified co-invention within the joint classified sample. This procedure retains pairs with zero post-shock co-invention and estimates both quantities on the same sample.

I evaluate near-zero coefficients using an economically meaningful multiplicative band from 0.90 to 1.10, corresponding to coefficient bounds ln(0.90) and ln(1.10). Equivalence at the 5% level requires the 90% confidence interval to lie inside both bounds and both one-sided tests to reject the relevant boundary (Schuirmann 1987). I report the conventional 95% confidence interval separately and assess sensitivity to 5% and 7.5% boundaries. Passing the within-firm test means that the affected relationship changes by less than 10% relative to the same firm's eligible comparison relationships. Passing the country-pair test means that aggregate co-invention on shocked pairs changes by less than 10% relative to eligible comparison pairs. These tests establish economic stability only within their respective comparisons, but do not establish absolute stability for a treated firm or equivalence in every event year.

3.4. Geographic Reconfiguration and Endpoint-Retaining Relationships

The geographic analysis combines an across-firm comparison between affected firms and unexposed firms on the same alternative country relationship with a within-firm comparison between relationships retaining one affected country and those retaining neither. The across-firm model uses the earlier-window sample. Within each affected-pair-by-cohort stack, affected and comparison firms must satisfy the same positive-activity rule on the same alternative country relationship during years −6 through −4. Comparison firms must remain unexposed to the focal shocked pair under the same exposure rule. I fix all activity-based eligibility and balancing variables during this window, leaving years −3 and −2 as holdout leads. Admitted relationships remain after later zero outcomes. Alternative relationships must remain free of a competing shock through event year 0 and are right-censored at the first subsequent shock. Among otherwise eligible comparison cells, 17.3% experience a later shock before year 4. These cells remain through the prior year and are censored when that shock occurs. Firm-relationship-stack fixed effects absorb stable firm-relationship differences, while relationship-year-stack fixed effects absorb annual changes shared by firms on the same country relationship. The stack is defined by the affected pair and treatment cohort. This across-firm design holds the alternative country relationship fixed and estimates relative growth separately for relationships retaining one affected country and those retaining neither.

The weights focus on the average treatment effect for affected firm-relationship-stack cells. I estimate the probability of exposure using only qualification-window relationship output and slope, firm patenting scale, portfolio breadth, technology mix, and relationship-by-stack fixed effects. Affected cells receive weight one. Comparison cells receive inverse-odds weights e/(1-e), where e is the estimated exposure probability. I impose common support within each relationship-by-stack stratum and drop affected cells without a supported comparison. Within each retained stratum, I normalize comparison weights to sum to the affected-cell weight. The Online Appendix reports affected and comparison cells, dropped strata, balance, weight distributions, effective sample sizes, and sensitivity to trimming, capping, alternative balancing specifications, and unweighted estimation on the same observations. Overlap weighting provides a separate sensitivity analysis that targets the overlap population.

The second design compares relationship types inside affected firms whose fixed pre-shock portfolios contain at least one relationship retaining an affected country and one retaining neither. Every admitted relationship remains after later zeros. Let *r* index an alternative relationship and let *R* identify relationships that retain one country from affected pair *e*. I estimate

$$\mathbb{E}\left[Y_{fret} \mid \mathcal{F}\right] = \exp\left\{\sum_{\substack{k=-3 \\ k\neq -1}}^{4} \theta_k \, R_{re} \mathbf{1}[t - T_e = k] + \kappa_{fre} + \lambda_{fet}\right\} \tag{3}$$

Firm-relationship-event fixed effects absorb persistent differences across alternative relationships. Firm-event-year fixed effects absorb annual changes common to the firm's alternatives. Each event-time coefficient estimates how the change on relationships retaining one affected country differs from the change on relationships retaining neither within the same affected firm, relative to year −1. It does not estimate either relationship type's separate growth rate because the firm-event-year fixed effects absorb their common component. A stacked-content model tests whether this difference varies between geopolitically sensitive and other classified invention. The comparison across all eligible alternative relationships is primary. Estimates limited to the same IPC subclasses use identical technological eligibility across relationship types and provide corroborating evidence.

I translate source contraction and destination growth into patent equivalents within their respective designs. I do not divide one by the other because the within-firm source model and the across-firm destination model use different comparisons. The quantities describe changes across relationship sets and do not trace individual patents, projects, inventors, or capabilities between them.

3.5. Technology-Matched Relational Capacity and Technological Prominence

I estimate whether preexisting portfolio structure conditions the source and destination responses. Technology-matched relational capacity and raw relationship breadth enter continuously and jointly through interactions with pooled and dynamic exposure terms. I report model-based percentage changes for three combinations: capacity zero and breadth two, capacity zero and breadth five, and capacity 0.60 and breadth five. Capacity of 0.60 is the 75th percentile, while breadth of two and five relationships corresponds to the 25th and 75th percentiles. The direct breadth contrast compares two with five relationships at capacity zero. The direct capacity contrast compares capacity zero with 0.60 at breadth five. All predictions and contrasts use identical observations within the relevant model and its joint covariance matrix. I estimate source and destination responses separately because the models use different observations and comparison groups. Supplementary analyses compare capacity with breadth in unrelated technologies, pre-shock firm patenting, technological diversity, source-relationship intensity, and the pre-shock sensitive share.

I estimate technological prominence separately across all classified IPC subclasses, regardless of geopolitical sensitivity. The panel is indexed by firm, relationship, event, year, and focal IPC subclass. Firm-relationship-event-subclass fixed effects absorb stable differences across relationships within a technology. Firm-event-subclass-year fixed effects compare the affected relationship with the fixed pre-shock comparison relationships in the same subclass and year, retaining later zeros. I interact event time with the continuous pre-shock prominence measure and report percentage changes at the 25th and 75th percentiles among unique eligible firm-event-subclass observations. These are presentation values, not groups. Online Appendix Table A9 applies the same interaction to focal-subclass activity on alternative relationships. Because the main qualification and prominence windows overlap the displayed leads, those lead estimates are conditional on sample construction. An earlier-window specification fixes eligibility in years −6 through −4, and a lagged prominence measure ending before year −3 provides a separate pre-event check. The primary prominence analysis remains the separate all-classified test. Prominence and capacity are pre-shock characteristics, not randomly assigned causes.

3.6. Citations

The directed citation design follows the same event timing but changes the comparison to reflect the production of citing patents and the availability of cited-country prior art. For citations from firm f and inventor country i to prior art associated with country j, I estimate PPML models with firm-directed-pair-event fixed effects, firm-origin-country-event-year fixed effects, and cited-country-year-by-stack fixed effects. The first set absorbs stable citation intensity on the directed link. The second absorbs changes in the firm's eligible citing patenting from the origin country. The third absorbs annual changes in the visibility and relevance of prior art associated with the cited country. Directed links are selected from pre-shock citation activity and remain in the panel after zero-citation years.

I estimate sensitive and other classified citation rates on the same observations and test their event-time difference within one model. The primary outcome counts citations that USPTO grant records classify as "cited by applicant" from eligible citing patents to pre-shock patents associated with the other affected country. I include the logarithm of the number of eligible citing patents as the PPML exposure, so the coefficients estimate changes in qualifying citations per eligible patent. A companion count model uses the same eligibility rules and measures changes in the aggregate number of qualifying citations. Both specifications exclude examiner-added citations, same-family citations, and citing patents with inventors from the other affected country. Directed links are selected using the same pre-shock windows as the other analyses and remain after later zero-citation years. The recent-window sample supplies the main post-shock estimates, while the earlier-window sample supplies the pre-event diagnostic. The coefficients describe changes in citation intensity to disclosed prior art associated with the other affected country. They do not measure complete knowledge flows, continued access to tacit knowledge, or substitution for joint invention.

3.7. Identifying Assumptions and Inference

Identification requires comparison-specific conditional parallel trends and no anticipation. Equation (1) requires parallel relative paths for a firm's affected and eligible comparison relationships absent the focal shock. Because comparison relationships may respond after treatment, the coefficients estimate changes relative to those relationships rather than absolute changes. When years −3 through −1 determine sample eligibility, the displayed leads are conditional because those same outcomes helped select the sample. In the earlier-window sample, eligibility is fixed by year −4, so years −3 and −2 provide pre-event comparisons for relationships already active in that period. I report the full observation period, a comparison among firm-events eligible under both rules, and joint lead tests. The external benchmark requires parallel aggregate paths after pair and endpoint-country-year absorption. The across-firm geographic design requires parallel paths between affected and symmetrically eligible comparison firms, with all activity rules and balancing variables fixed during years −6 through −4.

Firm-event-year fixed effects cannot absorb a relationship-specific change that coincides with bilateral tensions and independently changes invention. I assess treatment validity by reporting the magnitude and subsequent path of the tension measure, auditing recognizable geopolitical episodes, and comparing the shocks with an independently constructed conflict-composition measure. Alternative conflict thresholds and composition definitions assess sensitivity to the treatment rule. Bilateral event counts enter as a news-coverage and data-coverage diagnostic, not as an alternative measure of tension. I compare persistent and transitory shocks as heterogeneity while retaining the same initial treatment definition, which avoids selecting the main treatment on future tension.

My design fixes the comparison set before the event and separates the focal tension shock from subsequent shocks. Comparison relationships remain unshocked through event year 0 and are right-censored when a later qualifying tension shock occurs. The main specification uses the first qualifying tension shock for each country pair, and a sensitivity analysis censors affected pairs at any later qualifying tension shock. In the contemporaneous-restriction exclusion analysis, I omit affected-pair shocks when a qualifying formal restriction described in Section 2.2 was already in force or became effective during the shock quarter. In a separate pre-policy-horizon analysis, I censor the affected-pair event path when a qualifying restriction takes effect later. I do not enter later restrictions as controls because they may follow from the initial deterioration.

I address partial exposure in event year 0 by excluding applications filed before the shock date and, separately, by defining year 0 as the first complete post-shock application year. Mature grant cohorts and publication-based comparisons address grant truncation and patent-observation timing.

Firm-level models use two-way clustering by firm and affected country pair, while country-pair models use dyadic-robust inference. Wild-cluster bootstrap inference at the affected-pair level and firm-event block-bootstrap inference accompany the principal effects and direct comparisons. Constrained pseudo-event assignments preserve country pair, calendar cohort, treatment timing, eligibility, and event-time coverage. I use 1,999 successful assignments for each reported pseudo-event test and report empirical p-values and exceedance counts. The Online Appendix reports stack reuse, event-time coverage, common-support diagnostics, and PPML metrics.

4. Results

4.1. Stable Relationships, Reconfigured Technological Content

The results indicate that international invention relationships usually continue after bilateral tensions rise, although durable dissolution increases slightly. Among relationships active before the shock and observed through year 4, the standardized predicted probability of durable dissolution rises from 11.6% without exposure to 12.3% with exposure. Thus, the estimate indicates modest deterioration at the margin ($p = .020$ for the difference), while nearly 88% of affected relationships continue without durable dissolution.

After the geopolitical shock, total co-invention also changes little relative to the firm's eligible comparison relationships. The pooled estimate for years 0 through 4 corresponds to a 1.8% decline ($p = .391$). The estimate meets the 7.5% and 10% equivalence criteria, although it narrowly misses the tighter 5% criterion. In the external country-pair model, aggregate total co-invention declines by 2.6% relative to eligible comparison pairs ($p = .298$). This estimate also meets the 7.5% ($p = .020$) and 10% ($p = .0008$) criteria, but not the 5% criterion ($p = .156$). The first comparison is relative to the firm's fixed portfolio. The second includes invention outside the affected firms and does not use their other relationships as controls. Both designs place the aggregate estimate inside the 7.5% equivalence band under their respective comparisons, while neither identifies an absolute counterfactual for an affected firm.

This apparent stability conceals substantial technological recomposition in the relationship affected by the geopolitical shock. Relative to the same firm's eligible comparison relationships, geopolitically sensitive co-invention falls by 12.2% across years 0 through 4 ($p < .001$), while other classified co-invention rises by 6.0% ($p = .008$). The estimated sensitive decline is 10.0% in the shock year ($p = .001$) and reaches 13.0% in year 4 ($p < .001$).

The joint content model estimates the change on the same observations. Its pooled interaction implies a 17.2% decline in sensitive co-invention relative to other classified co-invention on affected relationships, compared with the corresponding difference across the firm's eligible comparison relationships ($p < .001$). Because other classified co-invention has a larger predicted baseline, its 6.0% increase offsets most of the model-implied decline in sensitive co-invention within the mixed-content sample. This offsetting pattern is consistent with the limited change in total co-invention, although the total and content models use different samples and do not provide an exact decomposition. Online Appendix Table A20 reports corresponding model-based patent-equivalent accounting on the separately identified earlier-window diagnostic sample and linear sensitive-share checks.

The external content model shows the same technological change. Sensitive co-invention falls by 9.8% on affected country pairs relative to eligible comparison pairs ($p = .013$), while other classified co-invention changes by 1.2% ($p = .566$). The direct sensitive-versus-other contrast is −10.9% ($p = .009$), and the sensitive share declines by 2.6 percentage points ($p = .004$). Table 2 reports the continuity, total co-invention, and technological-content estimates. Figure 1 presents the within-firm and external event paths through year 4. Under both comparisons, affected relationships carry a materially different mix of invention even as total co-invention remains economically stable.

[Insert Table 2 and Figure 1 about here]

4.2. Greater Contraction in the Firm's Prominent Technologies

Contraction on an affected relationship is larger when the focal technology accounted for a greater share of the firm's prior patenting. This analysis covers all classified IPC subclasses that satisfy the eligibility criteria, regardless of geopolitical sensitivity. I enter technological prominence continuously as the subclass's share of the firm's classified patenting during the twenty complete quarters before the shock.

The pooled interaction between exposure and technological prominence is −0.281 ($p = .003$). At the 25th percentile of prominence, where the focal subclass accounts for 3% of the firm's prior classified patenting, the model predicts an 8.0% decline in co-invention on the affected relationship ($p = .003$). At the 75th percentile, where the corresponding share is 27%, the predicted decline reaches 14.0% ($p < .001$). The direct difference is economically important ($p = .003$).

The percentiles identify values of a continuous pre-shock measure, not separate groups of technologies. The estimates therefore describe a steeper contraction as technological prominence rises. The prominence and category models use different units and samples. On observations eligible for both analyses, weighting the category estimates by each subclass's share implies an 11.4% decline, while the prominence model predicts an 11.3% decline at mean prominence. The difference is statistically indistinguishable from zero ($p = .886$). Technological prominence is observed rather than randomly assigned, so the evidence describes moderation in the firm's response and does not establish an effect attributable to prominence itself. Table 3 reports the continuous interaction, the common-sample reconciliation, and the percentile predictions. Figure 2 presents predicted changes across the observed prominence distribution.

[Insert Table 3 and Figure 2 about here]

4.3. Expansion Occurs through Other Country Relationships

After a sharp rise in bilateral tensions, geopolitically sensitive co-invention expands on affected firms' other preexisting country relationships. Relative to unexposed firms previously active on the same country relationship, affected firms increase sensitive co-invention by 10.8% across years 0 through 4 ($p = .001$). The increase reaches 13.9% in year 4 ($p = .005$), while a joint test finds no pre-shock path difference across the two displayed leads ($p = .897$). Table 4 reports the pooled destination estimate and the direct geographic comparisons. Figure 3 presents the complete event path through year 4. Online Appendix Table A6 and Figure A3 report the complete endpoint and same-IPC paths.

[Insert Table 4 and Figure 3 about here]

The estimated expansion is stronger on relationships that retain one country from the affected pair. Relative to unexposed firms active on the same relationships, sensitive co-invention increases by 15.0% when the relationship retains one affected country ($p < .001$) and by 8.0% when it retains neither ($p = .028$).

Naturally, these across-firm estimates do not themselves identify the difference between the two types of relationships. Among affected firms possessing both types of alternatives, the direct within-firm comparison shows 6.3% greater expansion on relationships retaining one affected country ($p = .011$).

The same technologies display a clearer geographic pattern. Restricting the outcome to IPC subclasses previously carried by the affected relationship yields increases of 19.0% on relationships retaining one affected country ($p = .001$) and 9.0% on relationships retaining neither ($p = .045$). A joint model estimates an 8.8% direct difference between the two relationship types within affected firms ($p = .013$). Expansion therefore occurs within technologies that the firm previously pursued on the affected relationship.

Other classified co-invention also expands on alternative relationships by 7.8%, on average ($p = .004$). Its difference across the two relationship types is 1.5% ($p = .430$), about one quarter of the difference for sensitive co-invention. The direct difference between the two content contrasts is sizeable ($p = .028$). Expansion extends across technological fields, but its concentration on relationships retaining one affected country is substantially stronger for geopolitically sensitive invention.

The 15% and 8% estimates come from across-firm comparisons between affected firms and unexposed firms active on the same relationships within each relationship type. The 6.3% estimate comes from a within-firm comparison of the two relationship types among affected firms. These portfolio-level patterns show where sensitive co-invention expands after bilateral tensions rise. They do not track individual inventions across relationships or quantify how much of the decline on the affected relationship is matched by expansion elsewhere.

4.4. Technology-Matched Relational Capacity

The models distinguish the number of alternative country relationships from their technological fit. Raw relationship breadth counts the firm's other eligible preexisting country relationships. Technology-matched relational capacity measures the share of the geopolitically sensitive technologies on the affected relationship that those alternatives had previously carried. Both measures enter continuously and jointly. I report predictions at two and five alternative relationships, the 25th and 75th percentiles of the breadth distribution, and at capacity values of zero and 0.60. Zero indicates that the eligible alternatives had carried none of the affected technologies. For example, a capacity of 0.60, the 75th percentile of the separate capacity distribution, indicates that the alternatives had previously carried 60% of the affected sensitive technology bundle.

At zero technological overlap and two alternative relationships, geopolitically sensitive co-invention on the affected relationship declines by 2.1% ($p = .432$). At zero overlap and five alternative relationships, the predicted decline reaches 12.5% ($p = .002$). Holding breadth at five and increasing capacity to 0.60

yields a predicted decline of 15.0% ($p < .001$). On identical observations, the direct breadth contrast at zero capacity is −0.112 log points ($p = .004$), while the direct capacity contrast at breadth five is −0.029 log points ($p = .038$). This indicates that the source-side estimates vary with breadth even at zero overlap, while the additional capacity contrast is smaller.

The destination estimates vary more steeply with technological overlap. At zero capacity and two alternative relationships, the destination prediction is −0.7% ($p = .831$). Increasing breadth to five while holding capacity at zero yields a 2.8% increase ($p = .410$). At breadth five and capacity 0.60, sensitive co-invention on other country relationships increases by 17.0% ($p < .001$). The direct capacity contrast at breadth five is 13.8% ($p = .001$), while the direct breadth contrast at zero capacity is 3.6% ($p = .279$). The difference between these two contrasts is 0.095 log points ($p = .028$). On the same observations and in the same destination model, expansion varies more with technological overlap than with an interquartile change in the number of alternative relationships. Table 5 reports the three prespecified combinations and their direct comparisons. Figure 4 plots the source and destination predictions across the observed capacity distribution. Two falsification tests distinguish technological fit from generic destination growth. Capacity predicts little growth in technologies outside the affected bundle (2.6%, $p = .444$) or in a size-matched placebo bundle (1.4%, $p = .697$). When I recalculate capacity without using the focal destination, the capacity contrast for the affected bundle is 12.5% ($p = .004$) and exceeds its matched placebo by 11.0% ($p = .021$). The association is therefore concentrated in the technologies used to construct the pre-shock capacity measure.

[Insert Table 5 and Figure 4 about here]

The models use distinct comparisons. The within-firm source model compares the affected relationship with the same firm's eligible relationships in the same year. The across-firm destination model compares affected firms with unexposed firms previously active on the same other country relationship. The estimates therefore show complementary source and destination patterns rather than a one-for-one transfer. Breadth is associated with source contraction even at zero capacity, whereas destination expansion varies much more with capacity than with breadth. Counts of international relationships can therefore overstate usable flexibility when those relationships lack prior activity in the affected technologies.

4.5. Citations and Continuing Technological Relevance

Relative to eligible comparison directed links, citation rates from geopolitically sensitive patent applications to earlier patents associated with the other affected country rise by 8.0% across years 0 through 4 ($p < .001$). The corresponding rate for applications in other classified fields rises by 3.0% ($p = .065$). The joint model estimates a 4.9% larger increase for sensitive applications ($p = .019$). In the earlier-window sample, the estimated sensitive-versus-other differences are small before the shock and jointly indistinguishable from zero across the two displayed leads ($\chi^2(2) = 0.69$, $p = .707$).

Measuring citations as total counts produces the same pattern. Citations from sensitive applications rise by 9.2% ($p = .001$), compared with 3.8% for applications in other classified fields ($p = .078$). The direct sensitive-versus-other difference is 5.2% ($p = .041$). I treat citation rates as the primary outcome because they measure how frequently each eligible application cites prior art associated with the other country while accounting for changes in the number of eligible applications. Table 6 reports both sets of estimates.

Therefore, following a geopolitical shock, sensitive co-invention contracts even as sensitive applications cite earlier patents associated with the other affected country more frequently. This divergence indicates that inventions associated with that country remain technologically relevant to the firm's subsequent patenting despite the contraction in sensitive co-invention. This evidence documents continuing relevance in disclosed prior art while leaving post-shock knowledge transmission, access to tacit knowledge, and substitution between citations and joint invention unresolved.

[Insert Table 6 about here]

4.6. Alternative Explanations and Robustness

I first consider whether compliance with binding policy accounts for the central pattern. This concern is especially relevant because the sensitivity measure identifies technological domains represented in export-control lists. I repeat the affected-relationship and other-country-relationship analyses after excluding tension shocks with a formal restriction announced or effective from 90 days before through 90 days after the shock quarter (Farrell and Newman 2019, Cao et al. 2024, Crosignani et al. 2026, Fajgelbaum et al. 2024). The estimates remain similar in direction and magnitude. Sensitive co-invention declines by 10.8% on affected relationships and expands by 10.2% on other country relationships. The pattern also remains when I censor event paths at the onset of a later restriction and retain only the preceding period. Later restrictions may follow from the initial deterioration, so entering them as ordinary post-shock controls would condition on a potential consequence of the event. The central pattern is therefore not confined to shocks accompanied by an already binding formal restriction.

Second, it may be possible that bilateral news volume, independent of how conflictual the reported events are, accounts for the observed response. Greater media attention can make an issue more salient to managers and external stakeholders and prompt changes in resource allocation (Bednar et al. 2013), while greater political-risk salience is associated with changes in firm investment and hiring (Hassan et al. 2019, Caldara and Iacoviello 2022). If bilateral salience alone drives the estimates, the contraction on affected relationships and expansion elsewhere should attenuate after controlling for the total number of recorded bilateral events. With bilateral event volume included, sensitive co-invention declines by 11.3% on affected relationships and expands by 10.5% on other country relationships, close to the main estimates. This stability reduces concern that bilateral news salience alone accounts for the pattern.

Third, political scrutiny and enforcement risk may discourage collaboration involving either affected country (Jia et al. 2024, Furnas et al. 2026, Koski 2026). If firms respond by withdrawing from both countries, expansion should favor relationships retaining neither. Instead, sensitive co-invention expands by 15% on relationships retaining one affected country and by 8% on those retaining neither, with a 6.3% direct within-firm difference. This evidence weighs against withdrawal from both affected countries. Because the reported estimates do not compare relationships retaining one affected country with relationships retaining the other, they cannot distinguish unilateral withdrawal from one country. Hence, the most plausible inference is that expansion avoids the affected pair while preserving invention involving at least one of its countries.

Fourth, changes in technological incentives or production location could redirect invention toward other technologies (Aghion et al. 2016, Moscona and Sastry 2023, Branstetter et al. 2026). Such redirection does not predict stronger expansion within IPC subclasses previously carried by the affected relationship. In these subclasses, co-invention increases by 19% on relationships retaining one affected country and by 9% on those retaining neither, with an 8.8% direct difference. Destination expansion also reaches 17% when other relationships already carry 60% of the affected technology bundle, whereas additional relationship breadth at zero overlap is associated with a small and imprecise response.

Fifth, firms may respond to geopolitical uncertainty by concentrating invention in their established technological strengths and reducing more peripheral activity. For example, UK organizations reduced patenting following the Brexit referendum in fields where UK applicants lacked revealed comparative advantage, a pattern consistent with technological refocusing (Vanino et al. 2026). An analogous firm-level response in my setting would produce greater contraction in technologies occupying smaller shares of the firm's prior patenting. However, co-invention declines by 8% at the 25th percentile of prominence and by 14% at the 75th percentile, placing the larger contraction in technologies more prominent in the firm's established inventive trajectory. The same pattern appears when I measure prominence earlier, leave a gap between that measurement and the shock, account for whether patenting in each technology was already rising or falling, and repeat the analysis using placebo shock dates. Firms therefore do not appear to be simply cutting technologies that had accounted for little of their prior invention or were already declining. Because my design observes prominence rather than randomly assign it, I interpret the evidence as moderation rather than an independent effect of prominence.

Sixth, technology-specific opportunity shocks could lower sensitive co-invention independently of bilateral tensions (Aghion et al. 2016, Moscona and Sastry 2023, Branstetter et al. 2026). Such shocks should produce corresponding declines in the same technologies elsewhere. Sensitive co-invention expands in the same IPC subclasses on firms' other preexisting relationships, including relative to unexposed firms active on those relationships. The specification reported in Table 7 that absorbs normalized IPC-subclass-by-year fixed effects retains a 15.0% sensitive-versus-other decline on the affected relationship and a 9.9% expansion elsewhere. This joint source and destination pattern weighs against a common technology-specific opportunity shock.

Seventh, sanctions, travel restrictions, and political disintegration have been associated with lower bilateral publication or patent collaboration (Makkonen and Mitze 2023, Milani and Pham 2025, Vanino et al. 2026). A general contraction account would therefore predict lower continuity and total co-invention, with declines extending beyond sensitive technologies. Instead, in my data, durable dissolution rises by only 0.7 percentage points, both total co-invention models establish equivalence within a 10% boundary, and other classified co-invention increases as sensitive co-invention declines. This selective contrast is difficult to reconcile with a uniform loss of relationship productivity.

I next assess the robustness of the findings to alternative measures of bilateral tension, event timing, comparison construction, weighting, patent measurement, and inference. I first re-estimate the models using different Goldstein thresholds, minimum event-count requirements, alternative cutoffs for the standardized shock measure, and an independently constructed ICEWS measure. Across these specifications, the decline in sensitive co-invention ranges from 9.1% to 14.8%, the relative sensitive-versus-other decline ranges from 13.8% to 19.5%, and destination expansion ranges from 8.9% to 12.1%.

I separately examine whether the response varies with the subsequent persistence of conflict. Shocks followed by sustained conflict are associated with greater technological divergence than transitory spikes. The sensitive-versus-other difference is equivalent to an additional 8.1% relative decline ($p = .038$). Because persistence is observed after the shock, this comparison describes heterogeneity in realized conflict severity and provides no independent validation of the original shock classification.

I check whether the findings are stable to alternative event timing and sample construction. I exclude applications filed before the shock quarter from event year 0. In a separate specification, event year 0 begins with the first complete post-shock application year. The earlier-window analysis determines all activity-based eligibility using years −6 through −4, before the displayed leads. Under this construction, sensitive co-invention declines by 11.2%, other classified co-invention rises by 5.4%, and the direct relative decline is 15.7%. The estimates for years −3 and −2 are jointly indistinguishable from zero ($p = .670$). Comparison portfolios remain fixed before the event, and later zero outcomes remain in the panel. I censor a comparison cell when it subsequently experiences its own shock, which affects 17.3% of otherwise eligible cells.

I also test whether the geographic expansion estimates depend on how affected and comparison firms are balanced or on the repeated use of some comparisons across event stacks. For the across-firm destination model, I cap extreme inverse-odds weights, estimate an unweighted model on the same observations, vary the characteristics used to balance affected and comparison firms, apply overlap weights, and reduce the influence of comparison cells that appear in many stacks. These specifications place destination expansion between 9.5% and 11.3%. Overlap weights preserve the same ordering among affected and comparison firms with the most similar pre-shock characteristics. This specification describes a narrower population than the main inverse-odds model, which targets firms exposed to a bilateral shock. The endpoint contrast is estimated within affected firms and therefore does not use inverse-odds weights. Reweighting observations to limit repeated use across stacks places this contrast between 5.0% and 7.0%.

I test whether the technology classification or construction of the patent data drives the findings. Five classification exercises change how patents with multiple IPC codes are assigned, treat mixed-content patents separately, remove one sensitive domain at a time, or replace the Wassenaar crosswalk with an independent critical-technology mapping. They imply relative declines in sensitive co-invention ranging from 14.5% to 18.5%. Restricting the sample to grant cohorts with sufficient time to mature, assigning exposure using exact application dates, and reconstructing the analysis from publication records preserve the decline on affected relationships and expansion elsewhere. These checks address grant truncation, filing delays, and dependence on a single sensitivity mapping. The empirical scope remains limited to patented co-invention because firms may protect some inventions through secrecy or confidentiality (Gross 2022, Cunningham and Kapacinskaite 2025).

I also examine whether inventor relocation, assignee changes, or ownership events mechanically produce the geographic pattern. Inventor mobility can change the countries attached to inventive relationships and the routes through which knowledge moves (Almeida and Kogut 1999, Song et al. 2003). Using inventor-country assignments established before the shock, restricting the sample to stable legal-assignee histories, and ending follow-up before major ownership changes place destination expansion between 8.5% and 12.5% and the endpoint contrast between 4.5% and 7.5%. Separate estimates for patents with one and multiple assignees retain the same ordering. This split provides a measurement check because assignee counts do not reliably distinguish collaboration within an organization from collaboration across organizations.

I also tighten the comparison sets directly. Removing the affected firms' own contributions from the country-pair benchmark yields a 9.8% sensitive-versus-other decline ($p = .028$). Excluding every comparison pair that shares either affected country yields a 9.1% decline (p = .073). Removing endpoint-country-by-year fixed effects from that endpoint-disjoint sample yields an 8.5% decline (p = .112). In the within-firm design, restricting comparison relationships to those that share neither affected country nor any affected subclass yields a 13.2% decline (p = .020). The estimates retain their direction and approximate magnitude, although the narrower external comparisons are less precise.

Table 7 summarizes selected specifications spanning the principal outcomes. The Online Appendix reports the corresponding sample construction, comparison reuse, weighting, balance, retained-zero, PPML-separation, and resampling diagnostics. The main firm-level standard errors allow errors to correlate within firms and affected country pairs, while the country-pair models use dyadic-robust inference. Alternative inference preserves the sign and magnitude of the direct recomposition and destination estimates. The endpoint, capacity, and citation contrasts retain their direction but have alternative two-sided p-values between .058 and .074 (Online Appendix Table A18). Constrained pseudo-event assignments place direct recomposition, destination expansion, and prominence in the tails of their respective distributions. Endpoint and capacity are less unusual under reassignment, with empirical p-values of .061 and .071 (Online Appendix Table A19).

The within-firm models estimate changes relative to the firm's other preexisting relationships, which may also respond after bilateral tensions rise. The country-pair benchmarks use separate comparison pairs and provide an aggregate bilateral comparison, but they do not identify the absolute effect for a particular firm. Evidence across the outcomes remains difficult to reconcile with the seven alternative explanations considered above. An unobserved change specific to the affected relationship and coincident with rising bilateral tensions could still contribute to the estimates.

[Insert Table 7 about here]

5. Discussion

This study shows that abrupt rises in bilateral tensions are followed by technological recomposition within international invention relationships that remain active, even as total co-invention changes little relative to relevant comparisons. The technological content of affected relationships shifts away from geopolitically sensitive work, while sensitive co-invention expands on other preexisting relationships. The results show that expansion is strongest where those relationships already carried the affected technologies and, in a separate comparison, on relationships joining one affected country to a third country.

The coexistence of continuity and recomposition is consistent with a selective, comparatively reversible response to geopolitical uncertainty. Abrupt deterioration can raise uncertainty about future restrictions and their persistence, while established invention relationships may embody coordination routines and specialized knowledge that are costly to rebuild (Singh 2005, Kellogg 2011, Berry 2014, Gulen and Ion 2016). Under those conditions, changing the technological content and geographic configuration of co-invention may be easier to reverse than dismantling a relationship. The study does not observe managerial beliefs or intent, so this is an interpretation of the empirical pattern. This interpretation can explain why the small changes in relationship survival and total output constitute a secondary result. In the face of geopolitical tensions, the more consequential adjustment occurs inside continuing relationships and across their portfolios. These patterns support two contributions.

5.1. Technological Recomposition as a Margin of Geopolitical Fragmentation

This study contributes to research on geopolitical fragmentation and international innovation by identifying technological recomposition within continuing relationships as a distinct margin of cross-border adjustment. Prior studies examine fragmentation through fewer research ties, lower cross-border co-invention, or changes in inventor networks as political relations deteriorate (Makkonen and Mitze 2023, Aghion et al. 2023, Charpin et al. 2024, Milani and Pham 2025, Koski 2026, Vanino et al. 2026). Those outcomes describe whether cross-border activity continues, how much remains, or what entities participate, but they do not isolate technological change within a continuing firm-country relationship. The findings show that a relationship may persist and remain economically stable in total co-invention relative to the relevant comparisons while shifting away from geopolitically sensitive technologies. Research on fragmentation should therefore incorporate technological recomposition alongside relationship dissolution and aggregate contraction. Persistent ties and stable totals may otherwise be read as continued integration even when the technological content of cross-border invention has changed.

Research on R&D collaboration scope explains how appropriation concerns and the demands of cooperation shape the activities included in an alliance (Oxley and Sampson 2004, Argyres et al. 2020), and related work traces partner- and location-specific hazards into alliance design and the allocation of inventors to collaborations (Ryu et al. 2018, Palomeras and Wehrheim 2021). In my setting the hazard originates outside the collaboration and differs across technologies pursued through the same country relationship. The findings therefore extend scope research from the design of a collaboration at its formation to the technological content of a relationship that continues, showing that firms adjust what a relationship carries when the source of exposure cannot be governed away. The data observe patented co-invention linked to a legal assignee rather than alliance governance or formal organizational boundaries, and relationship-specific coordination may help account for continuity, but the patents do not reveal whether that mechanism explains the pattern.

When a technology was prominent in a firm's earlier patenting, the larger reduction in co-invention concerns work in which the firm had substantial prior experience. Research on technological change shows that established inventive activity can shape subsequent technological direction (Helfat 1994, 1997, Aghion et al. 2016). The observed recomposition therefore reaches technologies that accounted for substantial shares of the firm's prior patenting. It does not identify the organizational changes behind the pattern or the firm's total activity in the technology.

Bilateral citation flows distinguish current co-invention from the continuing relevance of earlier patented inventions. Following the rise in tensions, citations to earlier patents associated with the other affected country increase more for geopolitically sensitive applications than for applications in other fields, even as new geopolitically sensitive co-invention contracts. Earlier patented inventions may therefore remain relevant even after new firm invention contracts. Fewer co-invented patents therefore provide an incomplete measure of technological separation. Applicant-disclosed patent citations record references to prior patents and remain an imperfect measure of knowledge flows (Jaffe et al. 1993, Alcácer and Gittelman 2006, Alcácer et al. 2009). The evidence does not establish continuing interpersonal exchange, access to tacit knowledge, or ongoing collaboration.

5.2. Technology-Matched Capacity and Multinational Flexibility

The second contribution concerns the literature on multinational flexibility and resource redeployment. Research on multinational networks treats geographic dispersion as a source of options when conditions change in one location (Kogut and Kulatilaka 1994, Belderbos et al. 2014, Glennon 2023), while research on R&D adaptation and redeployment emphasizes related knowledge and resources elsewhere in the firm (Helfat 1997, Helfat and Eisenhardt 2004, Sakhartov and Folta 2014, Sohl and Folta 2024). The findings connect these ideas at the level of international invention relationships. International breadth measures the number of potential alternatives. Technology-matched relational capacity captures how much of the affected sensitive technology bundle those other relationships already carried. A broad international portfolio may therefore contain many potential destinations, but only some have prior experience in the relevant technologies.

The results inform why the distinction between breadth and capacity matters. At zero capacity, increasing breadth from two to five relationships is associated with source contraction rising from 2.1% to 12.5%, but destination expansion changes only from −0.7% to 2.8%. At five relationships and capacity of 0.60, destination expansion reaches 17.0%. Breadth is therefore associated with the option set for reducing exposure on the affected relationship, while capacity is more strongly associated with where expansion occurs. The pre-shock portfolio characteristics are not randomly assigned, so the evidence does not identify a causal organizational capability. It shows that the technology-country configuration of established relationships is more strongly associated with destination expansion than relationship counts alone.

This distinction also separates my study from evidence on MNE adjustment after policy restrictions. Ma and Clougherty (2026) examine the scaling down of subsidiary operations in China and scaling up in surrogate countries after U.S.–China tariff escalation. In my setting, the response follows tension shocks even without contemporaneous formal restrictions and occurs in the technological content and geography of continuing invention relationships. Expansion is larger on relationships joining one affected country to a third country than on relationships involving neither, especially within the same IPC subclasses. The pattern is consistent with a comparatively reversible reconfiguration that occurs outside the affected bilateral interface while continuing to involve one of its countries. The evidence remains at the portfolio level and does not identify the retained country or trace projects and inventor teams.

5.3. Boundaries and Future Research

The patent data define the study's empirical reach. They capture co-invention disclosed in USPTO applications that later received grants and exclude invention kept secret or never patented (Gross 2022, Cunningham and Kapacinskaite 2025). The firm-level unit is the normalized corporate legal assignee matched to each patent, and country relationships describe the geography of patented co-invention associated with that assignee. Patent records leave consolidated corporate boundaries and the governance of the underlying collaboration unobserved. The findings apply most directly to legal assignees with established multinational patenting portfolios and to invention visible in the U.S. patent system. International co-invention is concentrated in major country corridors, so settings with sparse or institutionally different relationships may show other patterns.

Because a shock is defined against each pair's own recent history, pairs with persistently conflictual relations rarely qualify, and the treated sample tilts toward deteriorations between countries with previously cooperative relations. The findings therefore speak to abrupt deteriorations of workable relationships rather than to entrenched rivalries. The sensitivity classification draws on technological domains represented in the Wassenaar Arrangement's control lists. Alternative mappings preserve the main ordering, although no classification can fully represent how firms and governments assess technological sensitivity as those assessments evolve (Wassenaar Arrangement 2024).

The design is observational. The within-firm estimates compare an affected relationship with eligible relationships that may also change after tensions rise, while the external analysis uses separate country-pair comparisons and describes aggregate bilateral patterns. Neither design identifies how a particular firm's affected relationship would have changed without the rise in tensions, and an unobserved relationship-specific change may coincide with the event. Excluding contemporaneous formal restrictions and censoring observations when later restrictions begin reduces concern about recorded policy changes. Anticipated or unrecorded policy responses remain possible.

The characteristics used to compare firms are measured before tensions rise, but firms are not randomly assigned to them. The estimates show how later changes vary with a technology's share of prior patenting, the number of other country relationships, and prior activity in the same technologies on those relationships. They do not identify the causal effect of any characteristic or establish that firms distributed related technological work across countries in anticipation of geopolitical disruption. The affected-relationship and other-relationship analyses use separate comparisons, so contraction and expansion cannot be interpreted as a one-for-one transfer. The citation results concern disclosed prior art and do not establish later knowledge transmission, access to tacit knowledge, or substitution between citations and joint invention. Managerial intent and the internal processes behind the observed changes remain outside the design.

Future research could examine whether firms deliberately distribute related technological work across country relationships before disruption and how centralized versus decentralized decision making shapes redeployment, extending work on internal markets for redeployment (Sohl and Folta 2024, Sakhartov and Helfat 2026). Project records and inventor histories could trace how projects and teams move across technologies and country relationships, while internal records and interviews could show how firms weigh established relationships against exposure in particular technologies. Such evidence could also assess whether changing the geography of joint invention affects its speed, cost, or quality. Studies spanning multiple patent offices and firms' choices between patenting and secrecy could test whether the patterns extend beyond USPTO-visible invention. Longer follow-up could examine whether technological recomposition reverses when tensions recede or accumulates into durable separation. Comparative work on sparser country relationships could show how the patterns differ across institutional and geographic conditions.

6. Conclusion

Geopolitical tensions can change what firms invent together and where much more than whether they continue inventing together. Across two distinct comparisons, affected relationships usually survive and total co-invention changes little relative to the relevant comparison group. Yet geopolitically sensitive co-invention contracts on affected relationships, other content grows, and sensitive co-invention expands across other preexisting country relationships. Apparent stability therefore can mask technological recomposition and geographic reconfiguration. International breadth alone also does not reveal where expansion can occur. Destination growth is strongest among firms whose other relationships already carry the affected technologies, the portfolio condition captured by technology-matched relational capacity. Relationship survival and aggregate patent counts are therefore incomplete measures of the firm-specific consequences of geopolitical tensions. These strategic consequences can reside inside continuing relationships and in the technology-country configuration of the firm’s portfolio.

References

Aghion P, Antonin C, Paluskiewicz L, Strömberg D, Wargon R, Westin K (2023) Does Chinese research hinge on US co-authors? Evidence from the China Initiative. CEP Discussion Paper 1936, Centre for Economic Performance, London School of Economics and Political Science.

Aghion P, Dechezleprêtre A, Hémous D, Martin R, Van Reenen J (2016) Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. J. Political Econom. 124(1):1–51.

Alcácer J, Gittelman M (2006) Patent citations as a measure of knowledge flows: The influence of examiner citations. Rev. Econom. Statist. 88(4):774–779.

Alcácer J, Gittelman M, Sampat BN (2009) Applicant and examiner citations in U.S. patents: An overview and analysis. Res. Policy 38(2):415–427.

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

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

Argyres N, Bercovitz J, Zanarone G (2020) The role of relationship scope in sustaining relational contracts in interfirm networks. Strategic Management J. 41(2):222–245.

Argyres N, Mahoney JT, Nickerson J (2019) Strategic responses to shocks: Comparative adjustment costs, transaction costs, and opportunity costs. Strategic Management J. 40(3):357–376.

Aronow PM, Samii C, Assenova VA (2015) Cluster-robust variance estimation for dyadic data. Political Anal. 23(4):564–577.

Bednar MK, Boivie S, Prince NR (2013) Burr under the saddle: How media coverage influences strategic change. Organization Sci. 24(3):910–925.

Belderbos R, Tong TW, Wu S (2014) Multinationality and downside risk: The roles of option portfolio and organization. Strategic Management J. 35(1):88–106.

Berry H (2014) Global integration and innovation: Multicountry knowledge generation within MNCs. Strategic Management J. 35(6):869–890.

Branstetter L, Chen J-R, Glennon B, Zolas N (2026) Does offshoring production reduce innovation? Firm-level evidence from Taiwan. Management Sci., articles in advance. https://doi.org/10.1287/mnsc.2024.04944.

Caldara D, Iacoviello M (2022) Measuring geopolitical risk. Amer. Econom. Rev. 112(4):1194–1225.

Cao Y, de Nicola F, Mattoo A, Timmis J (2024) Technological decoupling? The impact on innovation of U.S. restrictions on Chinese firms. World Bank Policy Research Working Paper 10950.

Cengiz D, Dube A, Lindner A, Zipperer B (2019) The effect of minimum wages on low-wage jobs. Quart. J. Econom. 134(3):1405–1454.

Charpin R, London J, Vincent N (2024) The effect of geopolitical tensions on international research collaborations and its implications for global operations management. Internat. J. Production Econom. 268:109120.

Chatterji AK, Murray F (2026) How geopolitics is changing the economics of innovation. Entrepreneurship Innovation Policy Econom. 5:211–244.

Crosignani M, Han L, Macchiavelli M, Silva AF (2026) Securing technological leadership? The cost of export controls on firms. J. Financial Econom. 175:104192. https://doi.org/10.1016/j.jfineco.2025.104192.

Cunningham C, Kapacinskaite A (2025) Keeping invention confidential. Management Sci. 72(2):1658–1679.

Fajgelbaum P, Goldberg P, Kennedy P, Khandelwal A, Taglioni D (2024) The US-China trade war and global reallocations. Amer. Econom. Rev. Insights 6(2):295–312.

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

Furnas AC, Jia R, Roberts ME, Wang D (2026) Geopolitics in the evaluation of international scientific collaboration. NBER Working Paper 34789.

Galasso A, Schankerman M (2018) Patent rights, innovation, and firm exit. RAND J. Econom. 49(1):64–86.

Glennon B (2023) How do restrictions on high-skilled immigration affect offshoring? Evidence from the H-1B program. Management Sci. 70(2):907–930.

Goldstein JS (1992) A conflict-cooperation scale for WEIS events data. J. Conflict Resolution 36(2):369–385.

Gopinath G, Gourinchas P-O, Presbitero AF, Topalova P (2025) Changing global linkages: A new Cold War? J. Internat. Econom. 153:104042.

Gross DP (2022) The hidden costs of securing innovation: The manifold impacts of compulsory invention secrecy. Management Sci. 69(4):2318–2338.

Gulen H, Ion M (2016) Policy uncertainty and corporate investment. Rev. Financial Stud. 29(3):523–564. https://doi.org/10.1093/rfs/hhv050.

Han P, Jiang W, Mei D (2024) Mapping U.S.-China technology decoupling: Policies, innovation, and firm performance. Management Sci. 70(12):8386–8413.

Hassan TA, Hollander S, van Lent L, Tahoun A (2019) Firm-level political risk: Measurement and effects. Quart. J. Econom. 134(4):2135–2202.

Helfat CE (1994) Evolutionary trajectories in petroleum firm R&D. Management Sci. 40(12):1720–1747.

Helfat CE (1997) Know-how and asset complementarity and dynamic capability accumulation: The case of R&D. Strategic Management J. 18(5):339–360.

Helfat CE, Eisenhardt KM (2004) Inter-temporal economies of scope, organizational modularity, and the dynamics of diversification. Strategic Management J. 25(13):1217–1232.

Jaffe AB, Trajtenberg M, Henderson R (1993) Geographic localization of knowledge spillovers as evidenced by patent citations. Quart. J. Econom. 108(3):577–598.

Jia R, Roberts ME, Wang Y, Yang E (2024) The impact of US-China tensions on US science: Evidence from the NIH investigations. Proc. Natl. Acad. Sci. USA 121(19):e2301436121.

Kellogg R (2011) Learning by drilling: Interfirm learning and relationship persistence in the Texas oilpatch. Quart. J. Econom. 126(4):1961–2004.

Kerr SP, Kerr WR (2018) Global collaborative patents. Econom. J. 128(612):F235–F272.

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

Koski H (2026) Geopolitical fragmentation and the geography of 5G innovation. ETLA Working Paper 140.

Leetaru K, Schrodt PA (2013) GDELT: Global data on events, location and tone, 1979–2012. Paper presented at the International Studies Association Annual Convention, San Francisco, April 3–6.

Ma H, Clougherty JA (2026) MNE retrenchment and redeployment in the face of the U.S.–China trade war: Real options and strategic flexibility. Global Strategy J., articles in advance. https://doi.org/10.1002/gsj.70027.

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

Milani S, Pham L (2025) Sanctions and global collaboration in innovative activities. World Econom. 48(4):690–724.

Moscona J, Sastry K (2023) Does directed innovation mitigate climate damage? Evidence from U.S. agriculture. Quart. J. Econom. 138(2):637–701.

Oxley JE, Sampson RC (2004) The scope and governance of international R&D alliances. Strategic Management J. 25(8–9):723–749.

Palomeras N, Wehrheim D (2021) The strategic allocation of inventors to R&D collaborations. Strategic Management J. 42(1):144–169.

Ryu W, McCann BT, Reuer JJ (2018) Geographic co-location of partners and rivals: Implications for the design of R&D alliances. Acad. Management J. 61(3):945–965.

Sakhartov AV, Folta TB (2014) Resource relatedness, redeployability, and firm value. Strategic Management J. 35(12):1781–1797.

Sakhartov AV, Helfat CE (2026) Resource redeployment in multi-business firms: Centralized or decentralized? Strategic Management J. 47(7):1878–1906.

Santos Silva JMC, Tenreyro S (2006) The log of gravity. Rev. Econom. Statist. 88(4):641–658.

Schuirmann DJ (1987) A comparison of the two one-sided tests procedure and the power approach for assessing the equivalence of average bioavailability. J. Pharmacokinetics Biopharmaceutics 15(6):657–680.

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

Sohl T, Folta TB (2024) Do firms proactively build internal markets for redeployment? Evidence from a natural experiment. Management Sci. 71(3):2487–2507.

Song J, Almeida P, Wu G (2003) Learning-by-hiring: When is mobility more likely to facilitate interfirm knowledge transfer? Management Sci. 49(4):351–365.

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. Res. Policy 55(5):105471.

Wassenaar Arrangement (2024) List of Dual-Use Goods and Technologies and Munitions List. Wassenaar Arrangement on Export Controls for Conventional Arms and Dual-Use Goods and Technologies, Vienna.

Ziedonis AA (2007) Real options in technology licensing. Management Sci. 53(10):1618–1633.

Table 1. Sample Construction, Design Coverage, and Descriptive Statistics

Panel A. Sample and network coverage

Statistic	Value	Statistic	Value
Internationally co-invented patent records	435,397	Corporate legal assignees	44,302
Utility patents	434,122	Firms active on an affected relationship	6,645
Corporate-assignee patents	423,352	Firm-by-affected-pair events	25,841
Classified corporate patents	423,083	Affected country pairs, external benchmark	593
Sensitive share of classified patenting	34.68%	Inventor countries	173
Observed / possible country pairs	2,514 / 14,878 16.90%	Positive / possible pair-years	19,169 / 357,072 5.37%
Fractional co-invention on ten largest pairs	49.2%	Event window / treatment cohorts	Years −3 to 4 / 2003–2019

Panel B. Analysis-specific coverage

Estimation sample	Affected N	Comparison N	Total N	Firms / clusters	Events or pairs	PPML separation
Durable dissolution	16,502	62,916	79,418	5,984 firms / 541 pairs	16,502 firm-events	n.a.
Within-firm total co-invention	190,256	1,256,664	1,446,920	6,501 firms / 590 pairs	23,782 events	65,884
Country-pair total co-invention	4,128	84,636	88,764	166 countries; dyadic	529 affected / 13,402 eligible comparison pairs	3,422
Within-firm technological content (expanded rows)	253,376	1,293,352	1,546,728	6,408 firms / 588 pairs	15,836 events	74,708
Country-pair technological content (expanded rows)	8,256	169,272	177,528	162 countries; dyadic	529 affected / 13,402 eligible comparison pairs	9,296
Technological prominence	62,632	349,936	412,568	2,020 firms / 331 affected-pair clusters	7,829 firm–IPC-subclass pairs	18,932
Broad geographic expansion, earlier-window baseline	16,008	637,604	653,612	825 treated / 8,022 total firms; 142 affected-pair clusters	1,006 firm-events; 3,085 treated / 116,835 comparison relationship cells	5,350
Exact-IPC geographic split samples (combined support)	71,682	218,694	290,376	1,513 affected firms / 201 pairs	1,787 events / 15,844 subclass cells	7,298
Relational capacity, source / destination	65,928 / 71,598	220,816 / 330,918	286,744 / 402,516	Source: 3,918 firms / 512 clusters; destination: 6,812 total firms / 134 clusters	Source: 41,806 relationship-event cells; destination: 77,408 comparison cells after support	12,604 / 19,340
Citation rate / count	107,584 / 104,128	486,532 / 474,774	594,116 / 578,902	5,218 firms / 487 affected-pair clusters	26,114 directed-link events	27,332 / 28,382

Panel C. Pre-shock construct distributions

Construct	Affected mean (SD)	Comparison mean (SD)	25th percentile	75th percentile	Analysis N
Sensitive share	0.412 (0.218)	0.356 (0.202)	0.22	0.57	773,364
Technological prominence (annual estimation rows)	0.154 (0.201)	0.148 (0.198)	0.03	0.27	2,020 firms; 7,829 firm-subclass pairs
Technology-matched relational capacity, common-outcome falsification sample	0.226 (0.263)	0.164 (0.218)	0.00	0.60	259,372
Pre-shock relationship activity	1.92 (4.06)	1.34 (3.12)	0.28	1.96	773,364

Construct	Affected mean (SD)	Comparison mean (SD)	25th percentile	75th percentile	Analysis N
Table 1 (Continued)					
Portfolio breadth, same common-outcome sample	3.91 (2.87)	3.46 (2.61)	2	5	259,372

Note. Patent counts are divided fractionally across legal assignees, inventor-country pairs, and normalized IPC subclasses. Affected and comparison counts are retained model-specific estimation observations after eligibility restrictions and, for PPML models, separation. Content-model Ns are two-category outcome-expanded rows; share Ns are underlying base rows. Slash-separated values distinguish paired source and destination models or citation-rate and count models. The 13,402 value is the number of eligible comparison pairs appearing across the external stacks, not the number of distinct observed country pairs. The exact-IPC row combines the retained support of the endpoint-retaining and endpoint-free regressions in Table 4, Panel C; 290,376 is combined support across two regressions rather than a single-regression N. Year −1 is omitted. Country-pair models report endpoint countries because dyadic-robust inference does not use conventional clusters. Online Appendix Tables A1–A3 and A7–A9 report the patent reconciliation, model-specific sample flows, citation, prominence, capacity, weighting, and diagnostic support.

Table 2. Relationship Continuity, Total Co-Invention, and Technological Content

Panel A. Relationship continuity and total co-invention

	(1) Durable dissolution	(2) Total, within firm	(3) Total, country pair
	0.007	−0.018	−0.026
Post-shock term, years 0–4	(0.003)	(0.021)	(0.025)
	[.020]	[.391]	[.298]
Transformed magnitude	0.7 pp	−1.8%	−2.6%
		−0.020	−0.029
Year 4	n.a.	(0.036)	(0.042)
		[.579]	[.490]
Transformed year-4 magnitude	n.a.	−2.0%	−2.9%
Prediction without / with exposure	11.6% / 12.3%	n.a.	n.a.
Pooled conventional 95% confidence interval	n.a.	[−5.8%, 2.3%]	[−7.3%, 2.3%]
Pooled TOST p, 5% / 7.5% / 10%	n.a.	[.056] / [.002] / [<.001]	[.156] / [.020] / [.0008]
Year-4 90% equivalence interval	n.a.	[−7.6%, 4.0%]	[−9.3%, 4.1%]
Year-4 TOST p, 5% / 7.5% / 10%	n.a.	[.192] / [.054] / [.009]	[.298] / [.122] / [.035]
Observations, affected / comparison / total	16,502 / 62,916 / 79,418	190,256 / 1,256,664 / 1,446,920	4,128 / 84,636 / 88,764
Firms, pairs, or inference units	5,984 firms / 541 pairs	6,501 firms / 590 pairs	529 affected / 13,402 eligible comparison pairs; 166 countries
PPML separation losses	n.a.	65,884	3,422

Panel B. Technological content within firms

	(1) Sensitive	(2) Other classified	(3) Sensitive × exposed	(4) Sensitive share
	−0.130	0.058	−0.188	
Post, years 0–4	(0.028)	(0.022)	(0.032)	
	[<.001]	[.008]	[<.001]	
Transformed magnitude	−12.2%	6.0%	−17.2%	
	−0.105			
Sensitive co-invention, year 0	(0.032)			
	[.001]			
Transformed year-0 magnitude	−10.0%			
	−0.139			
Sensitive co-invention, year 4	(0.036)			
	[<.001]			
Transformed year-4 magnitude	−13.0%			
				−4.1 pp
Change in sensitive share				(0.7 pp)
				[<.001]
Sample support	Content-expanded N A/C/T=253,376/1,293,352/1,546,728; firms/clusters=6,408/588; events=15,836; separated=74,708. Share/base N A/C/T=126,688/646,676/773,364; separated=37,354.			

Panel C. Technological content across country pairs

	(1) Sensitive	(2) Other classified	(3) Sensitive × exposed	(4) Sensitive share
	−0.103	0.012	−0.115	
Post, years 0–4	(0.0415)	(0.0208)	(0.0438)	
	[.013]	[.566]	[.009]	
Transformed magnitude	−9.8%	1.2%	−10.9%	
				−2.6 pp
Change in sensitive share				(0.9 pp)
				[.004]
Sample support	Content-expanded N A/C/T=8,256/169,272/177,528; 529 affected pairs; countries=162; separated=9,296. Share/base N A/C/T=4,128/84,636/88,764; separated=4,648.			

Table 2 (Continued)

Note. Panel A uses one complete-follow-up observation per relationship for durable dissolution. The total and content models use distinct samples. Panels B and C estimate sensitive and other-classified co-invention on identical observations within each joint model. Content Ns are two-category outcome-expanded rows; share Ns are underlying base rows. Coefficients are followed by standard errors in parentheses and exact p-values in brackets. Percentage changes equal $100[\exp(\beta)-1]$. The 10% equivalence bounds are $\ln(0.90)$ and $\ln(1.10)$, and each TOST entry reports the larger one-sided p-value. The dissolution model is an unweighted LPM with firm-event fixed effects and pretreatment activity-history controls. All other coefficient models are unweighted PPML. Within-firm total models include firm-relationship-event and firm-event-year fixed effects; within-firm content models allow both sets to vary by content category. Country-pair models include pair-by-stack and endpoint-country-by-year-by-stack fixed effects; country-pair content models allow both sets to vary by content category. Share changes are model-based predictions with delta-method inference. Firm models use two-way clustering by firm and affected country pair; country-pair models use dyadic-robust inference. The content estimates do not decompose the total estimate.

Table 3. Technological Prominence on Affected Relationships

Panel A. Continuous model and adjusted predictions

Result	Continuous model	Prediction at mean prominence	At 0.03	At 0.27	0.27 − 0.03
Exposure at prominence zero	−0.075 (0.029) [.009]				
Mean eligible prominence (unique firm-event-subclass units)		0.160			
Exposure × prominence	−0.281 (0.096) [.003]				
Predicted post-shock change		−0.120 (0.029) [<.001]	−0.083 (0.028) [.003]	−0.151 (0.034) [<.001]	−0.067 (0.023) [.003]
Transformed magnitude	−7.2% at zero	−11.3%	−8.0%	−14.0%	−0.067 log points
Sample support	N=412,568; firms=2,020; firm-subclass pairs=7,829; affected-pair clusters=331; retained zeros=284,116; separated=18,932				

Panel B. Common-support reconciliation with Table 2

Result	Full Table 2 sample	Prominence-eligible support	Reconciliation	N
Sensitive effect	−0.130 (0.028) [<.001]	−0.154 (0.034) [<.001]	−14.3%	262,368
Other-classified effect	0.058 (0.022) [.008]	−0.087 (0.029) [.003]	−8.3%	262,368
Sensitive − other classified	−0.188 (0.032) [<.001]	−0.067 (0.031) [.031]	−6.5%	524,736 expanded
Subclass-share-weighted reconstruction	n.a.	−0.121 (0.030) [<.001]	Category reconstruction	524,736 expanded
All-classified prediction at mean prominence	n.a.	−0.120 (0.029) [<.001]	Prominence reconstruction	412,568
Difference between reconstructions	n.a.	0.001 (0.007) [.886]	Statistically indistinguishable	Common support
Sensitive-only interaction / prediction at sample mean (0.143)	n.a.	−0.245 (0.121) [.043] −0.154 (0.037) [<.001]	Per unit / −14.3%	440,280
Model support	Common-support model: A/C/T=117,600/407,136/524,736 expanded; firms=2,020; clusters=382; separated=18,592. Sensitive-only model: A/C/T=103,288/336,992/440,280; firms=1,784; clusters=356; separated=22,184.			

Note. Prominence is the focal normalized IPC subclass's share of the firm's classified patenting during the twenty complete quarters before treatment. Values 0.03 and 0.27 are presentation points, not groups. The exposure coefficient at prominence zero is the interacted model's main effect; the population-average entry evaluates that model at mean eligible prominence. Predictions and contrasts use identical observations and the joint covariance matrix. All models use unweighted PPML with firm-relationship-event-subclass and firm-event-subclass-year fixed effects and two-way clustering by firm and affected country pair. In Panel B, N=262,368 is the per-category base-row count and N=524,736 is the two-category outcome-expanded count; affected and comparison counts and separation losses in the support statement are outcome-expanded. The 412,568 entry for the all-classified prediction reproduces the full prominence-model support and is distinct from the two-category common-support reconstruction. The broader Table 2 sample remains the estimand for the paper's sensitive-versus-other comparison. n.a. indicates that the statistic does not apply to the specification.

Table 4. Expansion of Geopolitically Sensitive Co-Invention on Other Country Relationships

Panel A. Average expansion

	(1) Geopolitically sensitive	(2) Other classified
Pooled years 0–4	0.103 (0.031) [.001]	0.075 (0.026) [.004]
Transformed magnitude	10.8%	7.8%
Sensitive co-invention, year 4	0.130 (0.046) [.005]	
Transformed year-4 magnitude	13.9%	
Sample support	N=653,612; treated/comparison observations=16,008/637,604; treated/total firms=825/8,022; affected pairs=142; events=1,006; treated/comparison cells=3,085/116,835; separated=5,350	

Panel B. Geographic comparisons and split support

	(1) Retains one country	(2) Retains neither	(3) Direct endpoint	(4) Sensitive − other
Sensitive co-invention	0.140 (0.040) [<.001]	0.077 (0.035) [.028]	0.061 (0.024) [.011]	
Transformed magnitude	15.0%	8.0%	6.3%	
Other-classified endpoint contrast			0.015 (0.019) [.430]	
Difference between content contrasts				0.046 (0.021) [.028]
Affected firms / firm-events / firm-event-alternative-relationship cells	1,632 / 1,798 / 4,256	1,856 / 2,214 / 6,057	1,287 / 1,382 / 17,355	1,287 / 1,382 / 34,710 expanded
Total firms / affected-pair clusters	5,604 / 241	6,421 / 249	1,287 / 241	1,287 / 241
Retained affected / comparison observations	19,793 / 70,291	28,143 / 87,983	16,970 / n.a.	33,940 / n.a.
Kish ESS / separation losses	54,201 / 2,164	68,997 / 2,698	n.a. / 385	n.a. / 770

Panel C. Same-IPC geographic comparisons

	(1) Retains one country	(2) Retains neither	(3) Direct endpoint
Sensitive co-invention	0.174 (0.052) [.001]	0.086 (0.043) [.045]	0.084 (0.034) [.013]
Transformed magnitude	19.0%	9.0%	8.8%
Affected firms / firm-events / firm-event-alternative-relationship-subclass cells	1,012 / 1,104 / 5,282	1,281 / 1,465 / 10,562	780 / 782 / 27,582
Total firms / affected-pair clusters	4,820 / 193	7,560 / 201	780 / 193
Retained affected / comparison observations	23,920 / 73,267	47,762 / 145,427	27,000 / n.a.
Kish ESS / separation losses	56,552 / 2,442	112,393 / 4,856	n.a. / 582

Table 4 (Continued)

Note. Across-firm models compare affected firms with unexposed firms previously active on the same alternative country relationship. Panels B and C divide alternatives by whether they retain one endpoint of the shocked pair. Counts of 4,256 versus 6,057 in Panel B and 5,282 versus 10,562 in Panel C are distinct firm-event-alternative-relationship cells, with Panel C additionally indexed by subclass; 241 versus 249 and 193 versus 201 are the corresponding affected shock-pair clusters. Direct endpoint contrasts compare the two relationship types within affected firms on identical observations. Across-firm models use ATT inverse-odds weights normalized within relationship-by-stack strata and include firm-relationship-stack and relationship-year-stack fixed effects. Direct endpoint models are unweighted and include firm-relationship-event and firm-event-year fixed effects. All models use PPML, retain later zeros, right-censor subsequent shocks, and cluster by firm and affected country pair.

Table 5. Technology-Matched Relational Capacity and Portfolio Reconfiguration

Panel A. Conditional post-shock predictions

	(1) Source C=0, B=2	(2) Source C=0, B=5	(3) Source C=.60, B=5	(4) Other rels. C=0, B=2	(5) Other rels. C=0, B=5	(6) Other rels. C=.60, B=5
	−0.021	−0.134	−0.163	−0.007	0.028	0.157
Predicted post-shock change	(0.027)	(0.042)	(0.043)	(0.033)	(0.034)	(0.041)
	[.432]	[.002]	[<.001]	[.831]	[.410]	[<.001]
Transformed magnitude	−2.1%	−12.5%	−15.0%	−0.7%	2.8%	17.0%
Sample support	Source: N=286,744; firms=3,918; affected-pair clusters=512; separated=12,604			Destination: N=402,516; total firms=6,812; affected-pair clusters=134; separated=19,340		

Panel B. Direct contrasts in common units

Contrast	Source log points	Source proportional change	Destination log points	Destination proportional change
	−0.112		0.035	
Breadth difference at C=0	(0.039)	−10.6%	(0.032)	3.6%
	[.004]		[.279]	
	−0.029		0.129	
Capacity difference at B=5	(0.014)	−2.9%	(0.039)	13.8%
	[.038]		[.001]	
			0.095	
Capacity contrast − breadth contrast			(0.043)	10.0%
			[.028]	

Panel C. Destination falsification

Result	Standard affected bundle	Leave-one-destination-out	Outside affected bundle	Matched placebo bundle	Leave-one-out affected − placebo
	0.129	0.118	0.026	0.014	0.104
Capacity contrast at B=5	(0.039)	(0.041)	(0.034)	(0.036)	(0.045)
	[.001]	[.004]	[.444]	[.697]	[.021]
Proportional change	13.8%	12.5%	2.6%	1.4%	11.0%
	0.035	0.032	0.030	0.021	0.011
Breadth contrast at C=0	(0.032)	(0.033)	(0.030)	(0.031)	(0.038)
	[.279]	[.332]	[.317]	[.498]	[.772]
	0.095	0.086	−0.004	−0.007	0.093
Capacity contrast − breadth contrast	(0.043)	(0.044)	(0.041)	(0.043)	(0.052)
	[.028]	[.051]	[.922]	[.871]	[.074]
Leave-one-out affected bundle − outside-bundle capacity contrast	n.a.	0.092 (0.042) [.029] 9.6%	Reference outcome	n.a.	n.a.

Table 5 (Continued)

Panel D. Falsification sample support

Sample	Candidate N	Separated	Retained A / C / total	Firms / pair clusters
Standard affected bundle	276,344	16,972	46,127 / 213,245 / 259,372	9,730 / 249
Leave-one-destination-out	276,344	16,972	46,127 / 213,245 / 259,372	9,730 / 249
Outside affected bundle	276,344	21,888	44,792 / 209,664 / 254,456	9,684 / 249
Matched placebo bundle	270,116	22,614	43,628 / 203,874 / 247,502	9,512 / 247
Common true-placebo support	265,804	21,898	43,182 / 200,724 / 243,906	9,406 / 244

Panel E. Capacity support

Sample	Zero capacity, retained / separated	Positive capacity, retained / separated	Total retained
Affected relationship	26,029 / 1,021	39,899 / 466	65,928
Destination relationships	100,883 / 10,946	158,489 / 6,026	259,372
Leave-one-destination-out	116,180 / 12,226	143,192 / 4,746	259,372

Note. C is pre-shock technology-matched relational capacity and B is alternative-relationship breadth; both enter continuously and jointly. C=.60 is the 75th percentile; B=2 and B=5 are the 25th and 75th percentiles. Predictions and contrasts use the model's joint covariance matrix. All models use PPML and two-way firm/pair clustering. Source models are unweighted with firm-relationship-event and firm-event-year fixed effects. Destination models use ATT inverse-odds weights with firm-relationship-stack and relationship-year-stack fixed effects. Proportional changes equal 100[exp(β)−1]. Leave-one-out capacity removes the focal destination. Panel C's last column contrasts that estimate with the matched placebo; outside and placebo bundles test generic growth. Panels C–E use design-specific falsification and common-outcome samples, which are distinct from the source and destination prediction samples in Panel A. n.a. indicates that the statistic does not apply to the specification.

Table 6. Citations and Continuing Technological Relevance

	(1) Rate Sensitive	(2) Rate Other	(3) Rate Direct diff.	(4) Count Sensitive	(5) Count Other	(6) Count Direct diff.
Post, years 0–4	0.077 (0.020) [<.001]	0.030 (0.016) [.065]	0.048 (0.020) [.019]	0.088 (0.026) [.001]	0.037 (0.021) [.078]	0.051 (0.025) [.041]
Transformed magnitude	8.0%	3.0%	4.9%	9.2%	3.8%	5.2%
Eligible citing-patent exposure	Yes	Yes	Yes	No	No	No
Observations	594,116	594,116	594,116	578,902	578,902	578,902
Firms / affected-pair clusters	5,218 / 487	5,218 / 487	5,218 / 487	5,218 / 487	5,218 / 487	5,218 / 487
Directed-link events	26,114	26,114	26,114	26,114	26,114	26,114
Retained zero rows	386,904	386,904	386,904	379,812	379,812	379,812
PPML separation losses	27,332	27,332	27,332	28,382	28,382	28,382
Pseudo-R^2	.907	.907	.907	.899	.899	.899

Note. Columns (1)–(3) use the citation-rate sample and the logarithm of eligible citing patents as the PPML exposure. The retained rate-model sample contains 2,806,417 fractional eligible citing-patent equivalents across stacked observations. Columns (4)–(6) use the companion count model. Models are unweighted. The numerator retains applicant-disclosed citations and excludes examiner-added and same-family citations. Sensitive and other outcomes use identical observations, and direct differences are estimated in joint models. Models include firm-directed-pair-event, firm-origin-country-event-year, and cited-country-year-stack fixed effects and cluster by firm and affected country pair.

Table 7. Selected Robustness, Comparison-Set, and Inference Checks

Panel A. Design and construction

Specification	(1) Affected rel. sensitive vs. other	(2) External pair sensitive vs. other	(3) Destination expansion	(4) Endpoint contrast	(5) Prominence contrast	(6) Destination capacity / breadth
Independent conflict-composition measure	−0.173 (0.063) [.006]	−0.102 (0.055) [.064]	0.090 (0.042) [.031]	0.053 (0.030) [.082]	n.a.	n.a.
First complete post-shock year	−0.195 (0.034) [<.001]	−0.121 (0.048) [.012]	0.110 (0.042) [.009]	0.065 (0.031) [.038]	−0.070 (0.025) [.005]	0.089 (0.046) [.053]
Common-PPML-support earlier qualification window	−0.171 (0.059) [.004]	−0.107 (0.051) [.036]	0.088 (0.044) [.044]	0.050 (0.030) [.096]	−0.061 (0.025) [.015]	0.082 (0.047) [.081]
Inverse control-reuse weights	−0.185 (0.044) [<.001]	−0.111 (0.047) [.018]	0.101 (0.040) [.012]	0.060 (0.030) [.047]	−0.065 (0.024) [.007]	0.090 (0.045) [.046]
Majority-sensitive patent classification	−0.176 (0.049) [<.001]	−0.096 (0.054) [.075]	0.094 (0.043) [.029]	0.055 (0.030) [.067]	n.a.	0.091 (0.049) [.063]
Whole counting	−0.181 (0.049) [<.001]	−0.108 (0.043) [.012]	0.101 (0.039) [.010]	0.059 (0.028) [.035]	−0.063 (0.022) [.004]	0.092 (0.046) [.046]
Exclude patents with more than two inventor countries	−0.169 (0.037) [<.001]	−0.104 (0.047) [.027]	0.093 (0.035) [.008]	0.053 (0.027) [.050]	−0.061 (0.024) [.011]	0.084 (0.046) [.068]
Balanced follow-up through year 4	−0.180 (0.040) [<.001]	−0.112 (0.052) [.031]	0.098 (0.039) [.012]	0.058 (0.030) [.053]	−0.065 (0.027) [.016]	0.091 (0.052) [.080]
Mature grant cohort	−0.181 (0.061) [.003]	−0.109 (0.053) [.040]	0.096 (0.049) [.048]	0.049 (0.033) [.134]	−0.061 (0.026) [.019]	0.092 (0.045) [.041]
Prominence denominator, five / twenty patent equivalents	n.a.	n.a.	n.a.	n.a.	−0.061 (0.021) [.004] / −0.071 (0.028) [.011]	n.a.
Lagged prominence window ending before year −3	n.a.	n.a.	n.a.	n.a.	−0.064 (0.024) [.008]	n.a.
Normalized IPC-subclass-by-year fixed effects	−0.163 (0.052) [.002]	−0.108 (0.049) [.028]	0.094 (0.045) [.038]	0.054 (0.032) [.091]	−0.062 (0.025) [.013]	0.087 (0.047) [.064]
Leave out advanced electronics and semiconductors	−0.201 (0.074) [.007]	−0.105 (0.047) [.025]	0.112 (0.054) [.038]	0.065 (0.036) [.071]	n.a.	0.101 (0.060) [.092]
Leave out telecommunications and information security	−0.157 (0.068) [.021]	−0.121 (0.050) [.016]	0.086 (0.046) [.061]	0.049 (0.031) [.114]	n.a.	0.078 (0.055) [.156]
Leave out computing	−0.186 (0.055) [.001]	−0.110 (0.046) [.017]	0.105 (0.047) [.025]	0.060 (0.032) [.061]	n.a.	0.097 (0.052) [.062]
Leave out navigation and sensing	−0.173 (0.050) [.001]	−0.126 (0.054) [.020]	0.097 (0.044) [.028]	0.056 (0.030) [.062]	n.a.	0.089 (0.049) [.069]
Leave out aerospace	−0.190 (0.047) [<.001]	−0.113 (0.047) [.016]	0.108 (0.042) [.010]	0.063 (0.029) [.030]	n.a.	0.099 (0.047) [.035]
Leave out nuclear technologies	−0.184 (0.048) [<.001]	−0.114 (0.044) [.010]	0.102 (0.041) [.013]	0.059 (0.029) [.042]	n.a.	0.094 (0.046) [.041]

Table 7 (Continued)

Panel B. Comparison-set and inference checks

Specification	(1) Affected rel. sensitive vs. other	(2) External pair sensitive vs. other	(3) Destination expansion	(4) Endpoint contrast	(5) Prominence contrast	(6) Destination capacity / breadth
Exclude shocked-pair firms' contributions; N=132,444	n.a.	−0.103 (0.047) [.028]	n.a.	n.a.	n.a.	n.a.
Endpoint-disjoint external benchmark; N=68,422	n.a.	−0.095 (0.053) [.073]	n.a.	n.a.	n.a.	n.a.
Endpoint-disjoint, no endpoint-country-year effects; N=68,422	n.a.	−0.089 (0.056) [.112]	n.a.	n.a.	n.a.	n.a.
Within-firm controls share neither endpoint nor affected subclasses; N=26,126	−0.142 (0.061) [.020]	n.a.	n.a.	n.a.	n.a.	n.a.
Constrained pseudo-events, p [exceedances/1,999]	.003 [5/1,999]	.022 [43/1,999]	.006 [11/1,999]	.061 [121/1,999]	.018 [35/1,999]	.071 [141/1,999]
Alternative dependence-preserving inference, two-sided p	.003	.032	.006	.061	.018	.071
Bootstrap successful / requested resamples	9,999 / 9,999	1,966 / 1,999	1,987 / 1,999	1,987 / 1,999	9,999 / 9,999	1,987 / 1,999
Baseline observations	1,546,728	177,528	653,612	16,970	412,568	402,516
Firms / clusters or dyadic units	6,408 / 588	529 affected pairs; 162 endpoint countries; dyadic	8,022 / 142	1,287 / 241	2,020 / 331	6,812 / 134
PPML separation losses	74,708	9,296	5,350	385	18,932	19,340
Estimator / baseline fixed effects	PPML; firm-rel.-event×content; firm-event-year×content	Joint PPML; pair-stack×content; endpoint-country-year-stack×content	PPML; firm-rel.-stack; rel.-year-stack	PPML; firm-rel.-event; firm-event-year	PPML; firm-rel.-event-subclass; firm-event-subclass-year	PPML; firm-rel.-stack; rel.-year-stack
Weighting / inference	Unweighted; two-way firm/pair	Unweighted; dyadic robust	ATT inverse-odds; two-way firm/pair	Unweighted; two-way firm/pair	Unweighted; two-way firm/pair	ATT inverse-odds; two-way firm/pair

Note. Cells report coefficients, standard errors in parentheses, and exact p-values in brackets. Column (1) is the within-firm sensitive-versus-other comparison; column (2) is the external country-pair comparison; column (3) is the ATT-weighted destination estimate; column (4) is the within-firm sensitive endpoint contrast; column (5) contrasts prominence at 0.27 and 0.03; column (6) is the destination capacity contrast minus the breadth contrast. Baseline Ns in columns (1) and (2) are two-category outcome-expanded rows; column (4) reports base-row support for the sensitive endpoint contrast. Unless stated otherwise, each check retains the corresponding baseline comparison set, weighting, fixed effects, and inference. Source treatment-definition sensitivities labeled as earlier-window diagnostics use that pre-event frame; comparison-set repairs and opportunity specifications use the samples and diagnostics reported in the corresponding appendix panels. The common-PPML-support earlier-qualification row restricts each estimand to regression rows that survive the final PPML fixed point under both its primary and earlier-qualification rules. This is stricter than the pre-estimation common-eligible firm-event panels in Online Appendix Tables A5 and A6, which intersect eligibility before each model applies its own separation and fixed-effect support. The normalized IPC-subclass-by-year row adds subclass-year fixed effects. Pseudo-event cells report empirical p-values and exceedances among 1,999 successful assignments. Bootstrap rows identify the dependence-preserving procedure reported in Online Appendix Table A18. Endpoint-disjoint models exclude comparison pairs sharing either affected endpoint. Repair-specification A/C/T support is 7,846/124,598/132,444 for shocked-firm exclusion, 7,690/60,732/68,422 for endpoint-disjoint external models, and 10,944/15,182/26,126 for within-firm disjoint controls; corresponding separation losses are 15,976, 5,816, and 2,198. Online Appendix Tables A10–A19 report the treatment, classification, weighting, restriction, opportunity, news-volume, ownership, inference, and pseudo-event checks. Each column retains its design-specific baseline sample; observation totals are not intended to be common across outcomes. n.a. indicates that the statistic does not apply to the specification.

Figure 1. Event Paths for Total Co-Invention and Technological Content

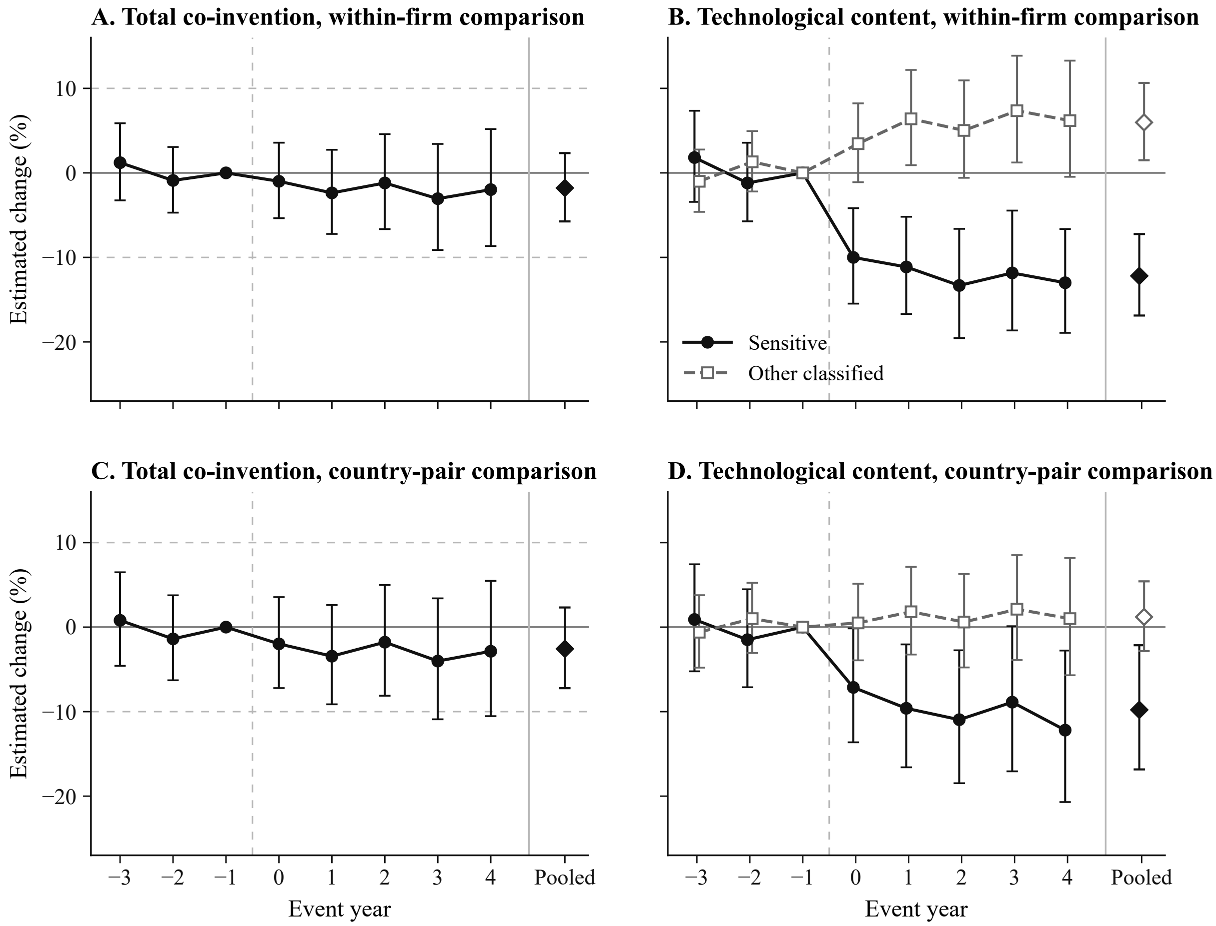

Notes. Points report proportional changes, calculated as 100 × [exp(coefficient) − 1]; whiskers are 95% confidence intervals. Year −1 is normalized to zero and has no confidence interval. Diamonds report the pooled estimate for years 0–4 and are separated from the annual path because they are summary estimates. The dotted horizontal lines in Panels A and C mark the ±10% equivalence boundaries. Panels A and B use the recent-activity sample and compare the affected relationship with the same firm's fixed set of eligible relationships, with firm-relationship-event and firm-event-year fixed effects. The displayed leads in these panels are conditional on recent-window sample construction. The Online Appendix reports the earlier-window diagnostic, for which eligibility is fixed before the displayed leads. Panels C and D compare affected country pairs with eligible never-treated and not-yet-treated pairs, with pair-by-stack and endpoint-country-year-by-stack fixed effects. Firm-level models cluster by firm and affected country pair; country-pair models use dyadic-robust inference. Sensitive and other classified co-invention use identical observations within each content model. Joint tests of the displayed leads (years −3 and −2) are: Panel A, $\chi^2(2) = 0.73$, p = .695; Panel B, sensitive relative to other classified, $\chi^2(2) = 2.68$, p = .262; Panel C, $\chi^2(2) = 0.55$, p = .761; and Panel D, sensitive relative to other classified, $\chi^2(2) = 1.05$, p = .593. The broader firm-level construction contains 6,645 firms and 25,841 firm-by-affected-pair shock events; the external benchmark contains 593 affected pairs. Regression observations, clusters, separation losses, and the earlier-window diagnostic are reported in the Online Appendix.

Figure 2. Technological Prominence and Co-Invention on Affected Relationships

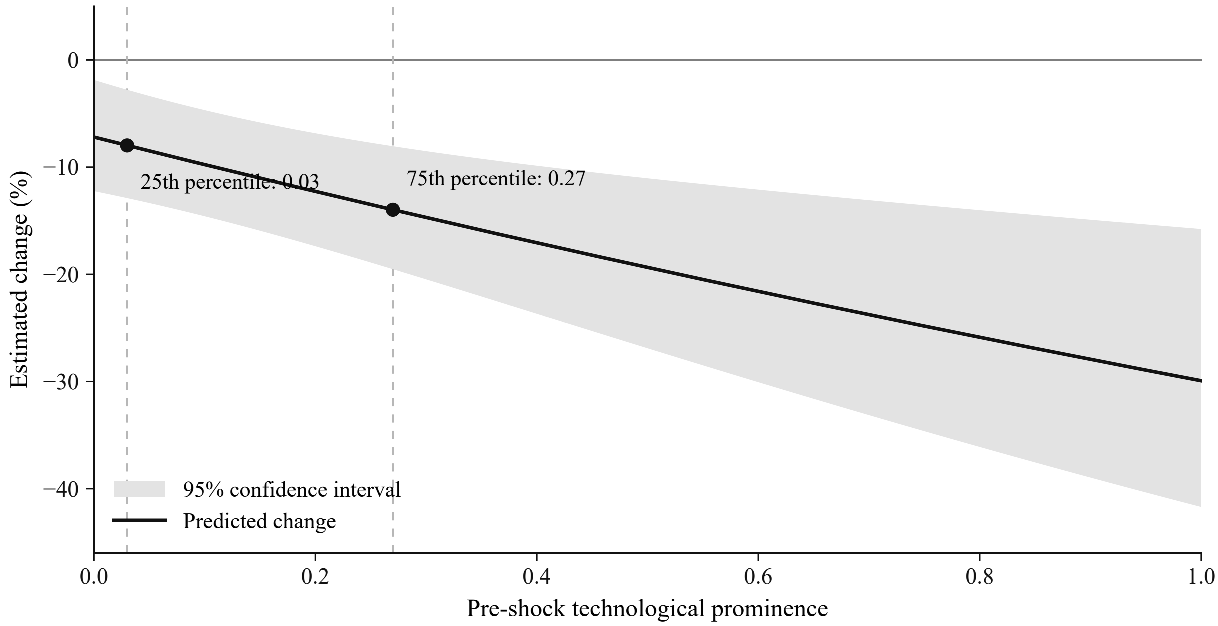

Notes. The line reports the predicted pooled change in co-invention on the affected relationship during years 0–4, calculated as 100 × [exp(prediction) − 1]; the shaded region is the 95% confidence interval derived from the joint covariance matrix. The unit of observation is a firm-relationship-event-year-IPC-subclass cell in the eligible all-classified sample. Technological prominence is continuous and equals the IPC subclass's share of the firm's classified patenting during the twenty complete quarters before the shock. Values 0.03 and 0.27 are the 25th and 75th percentiles and are presentation points, not groups. The PPML model includes firm-relationship-event-subclass and firm-event-subclass-year fixed effects and clusters by firm and affected country pair. The prediction is −8.0% at 0.03 and −14.0% at 0.27; their direct difference is −0.067 log points ($p = .003$). Coverage includes 2,020 firms and 7,829 legal-assignee-by-IPC-subclass pairs. Annual regression observations, cluster counts, separation losses, and fit statistics are reported in the Online Appendix.

Figure 3. Geopolitically Sensitive Co-Invention on Other Country Relationships

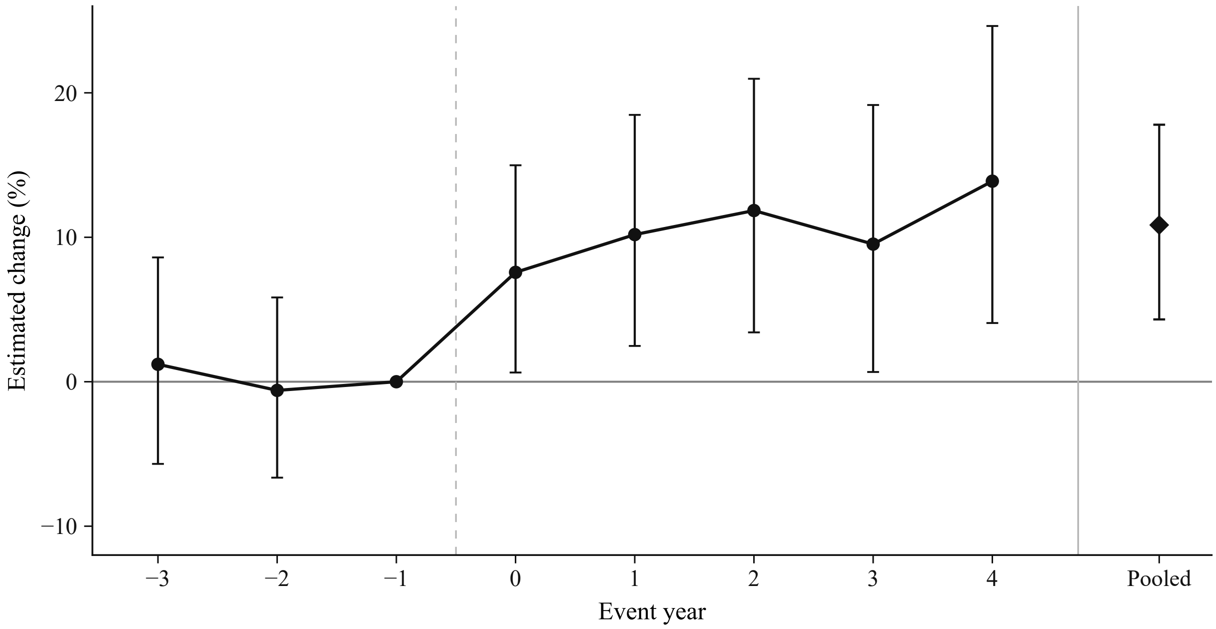

Notes. The outcome is geopolitically sensitive co-invention in a firm-event-other-relationship-year cell. Points report proportional changes, calculated as 100 × [exp(coefficient) − 1]; whiskers are 95% confidence intervals. Year −1 is normalized to zero and has no confidence interval. The diamond reports the pooled estimate for years 0–4 and is separated from the annual path because it is a summary estimate. The earlier-window sample fixes activity-based eligibility in years −6 through −4 and uses an across-firm comparison between affected firms and unexposed firms previously active on the same country relationship. ATT inverse-odds weights impose common support and are normalized within relationship-by-stack strata; cells are right-censored at later shocks. The PPML model includes firm-relationship-stack and relationship-year-stack fixed effects and clusters by firm and affected country pair. The joint test of the displayed leads is $\chi^2(2) = 0.22$, $p = .897$. The estimation sample covers 825 affected firms, 8,022 firms in total, 142 affected country pairs, 1,006 firm-events, 3,085 treated cells, 116,835 comparison cells, and 653,612 annual observations. Comparison observations, effective sample size, balance, cluster counts, and separation losses are reported in the Online Appendix.

Figure 4. Predicted Source and Destination Changes Across Relational Capacity

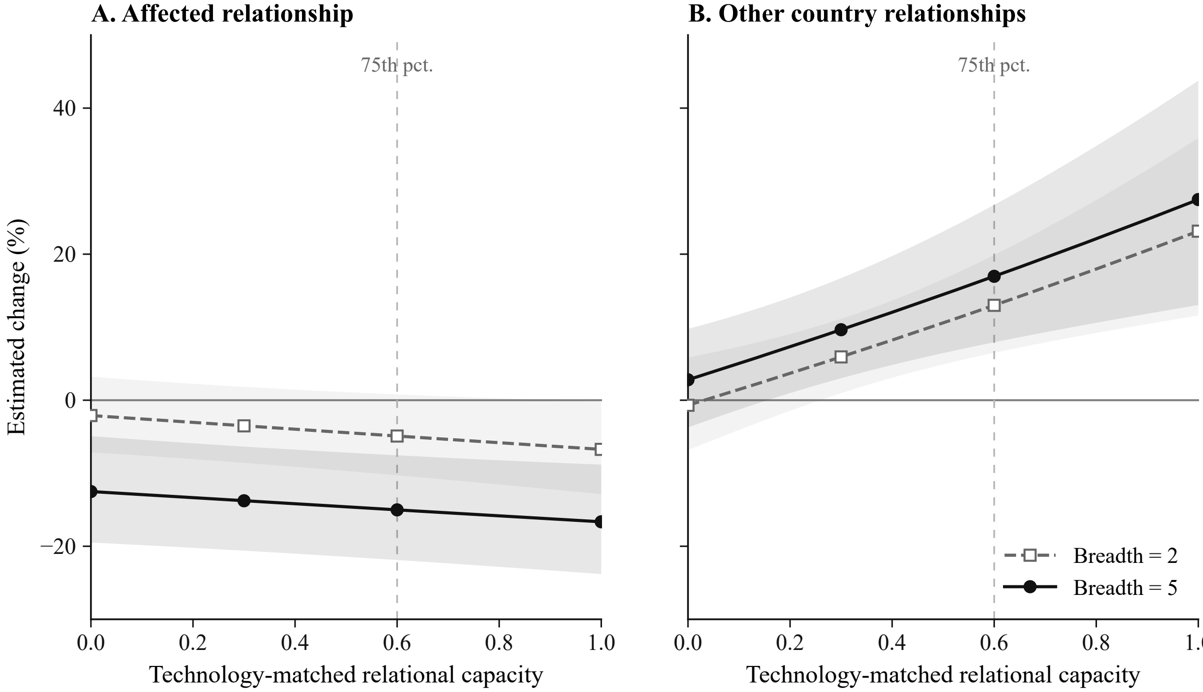

Notes. Lines report predicted changes in geopolitically sensitive co-invention during years 0–4, calculated as 100 × [exp(prediction) − 1]; shaded regions are 95% confidence intervals derived from each model's full joint covariance matrix. Technology-matched relational capacity is continuous and measures the share of the affected sensitive technology bundle already present on another eligible relationship. Breadth enters continuously; values two and five are the 25th and 75th percentiles. The vertical line marks capacity 0.60, the 75th percentile. Panel A uses firm-relationship-event-year observations and a within-firm comparison between the affected relationship and the same firm's eligible relationships, with firm-relationship-event and firm-event-year fixed effects. Panel B uses firm-relationship-stack-year observations and an across-firm comparison between affected firms and unexposed firms previously active on the same country relationship, with firm-relationship-stack and relationship-year-stack fixed effects and ATT inverse-odds weights. All models cluster by firm and affected country pair. Within each panel, predictions use identical observations. The source breadth and capacity contrasts are −0.112 ($p = .004$) and −0.029 ($p = .038$); the destination contrasts are 0.035 ($p = .279$) and 0.129 ($p = .001$). Regression observations, cluster counts, separation losses, and fit statistics are reported in the Online Appendix.