How Shocks Travel: The Cross-Border Impact of Natural Disasters in Firm Networks*

Vanessa Alviarez	Brian C. Fujiy	Angel Espinoza E.	Tomasz Swiecki
IDB	U.S. Census	UBC	Opendoor Labs

April, 2026

Abstract

Do the boundaries of the multinational firm shape how shocks propagate through global production networks? We link U.S. Bill of Lading microdata, geocoded natural disaster records, and cross-border ownership data to trace the transmission of exogenous supply disruptions from foreign suppliers to U.S. importers. Exploiting the quasi-random timing and location of natural disasters in a staggered event-study design, we document that disasters abroad generate sharp, persistent export declines at affected suppliers—shocks that propagate downstream, reducing U.S. importers' total purchases by up to 25 percentage points. Transmission, however, is far from uniform: non-MNC buyers contract roughly twice as much as importers belonging to global corporations, and a comparable gap separates arm's-length from intra-firm trade. These patterns point to the internal networks of multinational firms as shock absorbers that attenuate and redirect supply-side disruptions—suggesting that firm boundaries are not merely organizational choices, but also determinants of macroeconomic resilience.

*We give special thanks to Javier Cravino and Verónica Rappoport, editors of the *IMF Economic Review*, for their guidance and encouragement. We are especially grateful to Álvaro García Marín for a thoughtful discussion, and to participants at the Conference on *From Micro to Macro: Leveraging Microdata to Address Macroeconomic Issues* for their valuable feedback. We also thank Keith Head, Matilde Bombardini, Natalia Ramondo, and Andrei Levchenko, as well as the participants of the SEA Annual Conference, UBC VSE-Sauder seminar, and the CEA annual meetings, for very useful comments and suggestions to early versions of this work. Any opinions and conclusions expressed herein are those of the authors and do not represent the views of the U.S. Census Bureau, the Inter-American Development Bank, or its member countries.

Email: valviarezr@iadb.org , brian.c.fujiy@census.gov, aesp@student.ubc.ca, tomasz.swiecki.work@gmail.com.

1 Introduction

Global production networks map the flow of goods—and the pathways through which shocks spread. Firms operate through webs of suppliers, buyers, and affiliates that cross national and corporate boundaries. These linkages bring efficiency but also risk: when disruption occurs—whether due to natural disasters, geopolitical conflict, or logistical breakdowns—the network structure determines how far the shock travels and who bears the burden. This paper studies how large, sudden disruptions to foreign suppliers propagate through the international production network, and how this propagation depends on two key features of global firms: whether trade occurs across firm boundaries or within them, and whether the traded input is easy or difficult to substitute. Using high-resolution data on natural disasters, firm-to-firm U.S. import transactions, and multinational ownership linkages, we show that both the structure of trade relationships and the nature of the traded good jointly shape the exposure, adjustment, and recovery paths following supply shocks.

A striking feature of international trade is the persistent asymmetry in how different trade relationships respond to large-scale disruptions. During the Asian financial crisis, U.S. exports to arm's-length partners in Asia dropped by 26%, while shipments to foreign affiliates fell just 4% (Bernard et al., 2009). A decade later, the Great Recession revealed a similar asymmetry on a larger scale. Between the peak and trough of the crisis, real imports fell by more than 5.6 times the decline in U.S. real GDP (Levchenko et al., 2010). Yet even this steep contraction was uneven: imports from unaffiliated suppliers plummeted by nearly 31%, compared to a 16% drop in imports from foreign affiliates, as shown in Figure 1. When the world stopped, internal trade links bent—but did not break. These patterns suggest that intra-firm trade is more resilient to external disruptions, possibly because multinationals can coordinate production internally, reallocate demand, or bridge financing constraints across affiliates. Yet despite repeated episodes, this asymmetry remains poorly documented—and its underlying mechanisms even less understood. This paper brings both to the micro level.

How far a shock travels depends critically on the structure of the trade relationship. A natural disaster in a supplier's location may affect its direct buyers, but the extent and persistence of the impact depends critically on two features of the relationship: whether the traded input is relationship-specific (Rauch, 1999), and whether the firms involved are bound by ownership ties. The role of input specificity is not straightforward. Highly tailored or certified inputs are harder to replace, which amplifies the immediate costs of disruption. But specificity also creates mutual dependence: a buyer that relies on a given supplier has strong incentives to support, wait for, or reintegrate that supplier after the disruption—

FIGURE 1: Growth in Real U.S. Imports: Total, Intra-Firm, and Arm's-Length

Notes: The figure plots year-on-year growth rates in U.S. goods imports. All series are expressed in real terms using the U.S. GDP deflator. The solid blue line reflects total U.S. imports of goods, the dashed orange line captures intra-firm imports—defined as goods imported by U.S. multinational parents from their foreign affiliates, as well as by foreign affiliates operating in the U.S. from their parent companies abroad—and the dotted black line captures arm's-length imports, defined as trade between unrelated parties. *Sources*: Total imports from the U.S. Census Bureau; intra-firm imports and GDP deflator from the U.S. Bureau of Economic Analysis (BEA).

potentially accelerating recovery. Firm boundaries shape how this dependence plays out. Intra-firm trade may allow for faster internal substitution or reallocation across affiliates, while arm's-length relationships rely more heavily on external re-matching—and may never be reestablished.

We bring these ideas to the data by constructing a panel of U.S. importers linked to foreign suppliers using Panjiva customs shipment records, matched with firm ownership data from Orbis and disaster records from EM-DAT. Our empirical strategy exploits the quasi-random timing and geography of extreme natural disasters that affect the exporter's subnational location as exogenous shocks to the supply capacity of affected foreign exporters. We focus on sourcing relationships that are both recent and economically significant, and compare the dynamic outcomes of exposed and unexposed firms across multiple margins: total imports, imports excluding affected suppliers, the number of active supplier relationships, and– for firms that also sell abroad– exports.

A central challenge in this setting is that we do not observe whether a firm intends to import in a given quarter. Firm-to-firm trade relationships are inherently sparse and intermittent: a firm that does not appear in the data may have stopped importing altogether, or may simply not have intended to buy in that particular quarter. Without knowing the counterfactual, it is difficult to interpret a zero as a disaster-induced disruption rather than a reflection of the inherently intermittent nature of international trade. We address this carefully by restricting the sample to firms with established trading patterns prior to the shock—increasing the likelihood that any subsequent interruption reflects a genuine disruption rather than the natural lumpiness of international sourcing.

Central to our analysis is the ability to observe the organizational structure behind each trade relationship. We exploit two dimensions of this structure. The first distinguishes importers that belong to a multinational group from those that do not, capturing whether the buyer has access to an internal corporate network. The second distinguishes intra-firm trade—transactions between affiliates of the same corporate group, which can only occur within multinational networks—from arm's-length trade, which encompasses transactions between unrelated parties and is conducted by both MNC and non-MNC importers alike. We further stratify trade flows by input specificity, using the Rauch (1999) classification to distinguish differentiated products—typically harder to substitute—from those traded on organized exchanges or with reference prices. These dimensions are central to understanding resilience: organizational form shapes the ability to substitute and reallocate when disruptions occur, while input specificity determines how costly substitution is in the first place. By examining each in turn, we can assess not only who absorbs shocks better, but through which mechanisms.

Our findings paint a clear picture of how supply chain disruptions travel across borders. Natural disasters generate large and persistent shocks to directly affected foreign suppliers, with export growth declining by up to 36 percentage points and remaining depressed for several years. These disruptions do not stop at the supplier's door: U.S. importers connected to affected suppliers see their total imports fall by up to 25 percentage points at the trough, and the contraction extends to purchases from unaffected suppliers, the number of active relationships, and, among firms that also sell abroad, their own export activity. A shock that begins upstream propagates throughout the buyer's entire organizational footprint.

The severity of this propagation, however, depends critically on who is hit and how their relationships are structured. Multinational importers weather supply disruptions more effectively than their non-multinational counterparts—and the advantage is sharpest when the disrupted supplier belongs to the same corporate group. Arm's-length buyers experience a peak decline roughly twice as large as that of multinational importers, and intra-firm trade relationships exhibit smaller initial declines and faster recoveries, consistent with the coordination and reallocation mechanisms that multinational networks make available. Arm's-length trade, by contrast, is more exposed to persistent disruption. Product differentiation plays a more nuanced role: differentiated inputs show a somewhat larger initial response, but the gap relative to homogeneous goods narrows over time as the mutual dependence that makes these inputs hard to replace also creates incentives to restore the relationship.

Our paper speaks to three strands of literature but departs from each in critical ways. A first body of work studies how shocks propagate through production networks, often emphasizing the amplifying role of input linkages in domestic settings (Barrot and Sauvagnat, 2016; Carvalho et al., 2021; Khanna et al., 2022; Freund et al., 2022; Lee and Han, 2022). These papers trace shock transmission across firm-to-firm connections, but typically cannot observe whether the affected links operate within or outside the boundaries of a corporate group. As a result, they remain silent on whether propagation depends on organizational structure. We extend this logic across borders—where ownership can be separately measured from trade—allowing us to disentangle the nature of the link from the structure behind it.

A second strand of research studies how multinational firms organize production across borders—deciding not only where to operate, but also which activities to internalize. This literature has examined patterns of vertical integration across affiliates (Atalay et al., 2014; Li, 2021) and modeled the trade-offs involved in arm's-length versus intra-firm sourcing (Antràs, 2022; Conconi et al., 2022). Cravino and Levchenko (2017) provide empirical evidence on the coordination within multinationals, showing that sales at headquarters and foreign affiliates comove systematically in response to shocks—consistent with internal propagation mechanisms. Irarrazabal et al. (2013) go further by embedding intra-firm trade as a central margin in a structural model and estimating its role in shaping production and trade patterns. Yet across this literature, data rarely allow researchers to observe whether specific trade relationships occur within or across firm boundaries.[1] As a result, we know little about how organizational form shapes exposure and recovery in the face of disruption. We address this gap by combining customs data with firm-level ownership linkages, identifying the structure behind buyer-supplier pairs, and investigating how their exposures and adjustment dynamics differ under supply shocks.

Our paper also contributes to the growing literature on supply chain resilience and input specificity. A foundational insight from Barrot and Sauvagnat (2016) is that shocks propagate more forcefully when upstream inputs are harder to replace—due to differentiation, customization, or contractual frictions. Subsequent work has reinforced this view: specific inputs tend to amplify vulnerability, either because substitution is technologically costly

[1] An exception is Ramondo et al. (2016), who show that a large share of affiliate sales occur without direct trade links to headquarters—highlighting the uneven use of internal trade margins within MNC networks.

(Boehm et al., 2019; Castro-Vincenzi, 2022; Balboni et al., 2023), or because they are embedded in long-term buyer-supplier relationships (Fort et al., 2023; Blaum et al., 2023). Recent studies, such as Blaum et al. (2025) and Heise et al. (2025), also show that firms exposed to recurrent delays or sector-specific risks diversify sourcing strategies but often reduce overall imports. Yet across these papers, the structure of the trade relationship itself—whether the link operates within or across firm boundaries—remains unobserved. We show that input specificity does magnify the initial drop in sourcing after a disruption, but also opens a path to recovery when the buyer and supplier are part of the same multinational group. Intra-firm relationships allow for coordination, support, and reintegration that arm's-length links struggle to replicate.

Finally, our approach connects to recent work that uses firm-to-firm data to study how supply shocks transmit across networks. Papers such as Heise (2023), Huneeus (2023), Kikkawa et al. (2023), and Méjean et al. (2023) document rich heterogeneity in how firms respond to disruptions, depending on network position, linkage intensity, and pre-shock characteristics. We complement and extend this agenda by showing that organizational form—intra-firm vs. arm's-length—is a critical determinant of resilience. This dimension is especially salient when input specificity constrains substitution: under those conditions, ownership boundaries become central to understanding who adjusts, how quickly, and through which mechanisms.

The remainder of the paper is organized as follows. Section 2 describes the construction of our firm-to-firm trade panel, the matching with ownership linkages, and the identification of disaster exposures. Section 3 presents the empirical strategy, including our event-study design, sample definitions, and exposure criteria. Section 4 displays key descriptive patterns and reports the main empirical findings, breaking down responses by product type and ownership structure. Section 5 presents robustness checks and alternative specifications. Section 6 concludes.

2 Data

Our empirical analysis draws on three main data sources: U.S. Bill of Lading records that track firm-to-firm import transactions, geocoded records of global natural disasters, and cross-border ownership linkages. Together, these allow us to trace how supply shocks originating abroad propagate through the network of U.S. importers and their foreign suppliers

2.1 Firm-to-firm trade transactions from Bill of Lading records

We use U.S. shipment-level Bill of Lading (BoL) records compiled by S&P Panjiva as our primary source for firm-to-firm international trade transactions.[2] The dataset contains over 155 million records spanning 2007 to 2022, covering the universe of U.S. maritime imports with the names and addresses of both foreign exporters (shippers) and U.S. importers (consignees), as well as product descriptions and quantities.

To construct firm-level measures of trade activity, we first harmonize the BoL data with official U.S. Census trade statistics. We map BoL product descriptions to six-digit Harmonized System (HS6) codes and merge them with HS10-country-year data from the U.S. Census to recover shipment values. Weighted average unit values at the HS6-country-year level are combined with reported shipment quantities to approximate the value of each transaction. Aggregating across transactions provides a firm-level export measure over time. Transactions that list multiple HS6 product codes but report a single quantity or weight figure are excluded, as the reported quantity cannot be reliably allocated across products.

A key feature of the BoL data is the ability to track firms across time and space. We geocode the physical addresses of all U.S. importers and their foreign partners, enabling us to link each firm to subnational locations and to measure their exposure to natural disasters. Our procedure yields precise coordinates for more than 470,000 U.S. importers and nearly one million foreign exporters. Only a small fraction of records (around 2–3%) lack reliable location information and are excluded. Further details on the harmonization of firm identifiers, geocoding methods, and data cleaning procedures are provided in Appendix A.2.

2.2 Natural Disasters

To quantify the severity of natural disasters, we use data from EM-DAT and measure disaster intensity as deaths per capita at the second administrative level (ADM2, e.g., provinces, counties, or districts depending on the country). To avoid over-weighting small regions—where moderate events can appear disproportionately large in per-capita terms—we winsorize the upper tail of the fatality distribution. Disasters above the 90th percentile of this distribution are classified as extreme and constitute the treated group in our event-study framework; those below the 50th percentile serve as the control group. Events in the intermediate range are excluded from the main analysis to sharpen identification.

[2] A Bill of Lading (BoL) is a legal document issued by a carrier to a shipper that confirms receipt of goods for transport. It includes shipper and consignee names and addresses, product descriptions, vessel and carrier information, ports of loading and unloading, and shipment weights and quantities. Panjiva acquires these records from U.S. Customs and Border Protection (CBP).

FIGURE 2: Distribution of natural disaster events by number of deaths per capita

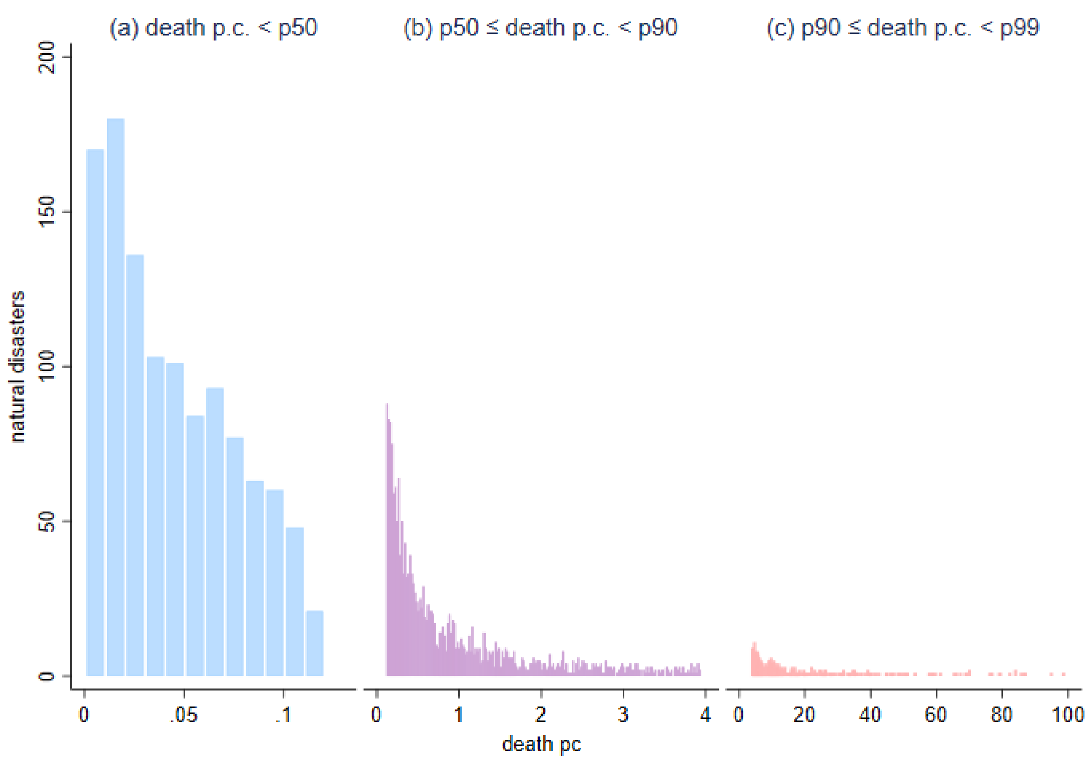

Notes: The histogram groups disaster events into three categories based on their position in the fatality distribution: below the median (p50), between the 50th and 90th percentiles, and between the 90th and 99th percentiles. The x-axis is scaled to accommodate the highly skewed distribution of fatalities across events. For scale purposes, the figure is restricted to natural disasters with at least 3 deaths. The sample covers disaster events worldwide from 2007 to 2022. *Source*: EM-DAT International Disaster Database.

Figure 2 illustrates the highly skewed distribution of natural disaster events by fatality counts. Panel (a) shows that low-fatality events (deaths $<$ p50) are by far the most frequent, with most causing fewer than 5 deaths. Panel (b) covers intermediate-severity disasters (p50$\leq$ deaths $<$ p90), ranging roughly from 13 to 103 deaths. Panel (c) depicts high-fatality events (p90 $\leq$ deaths $<$ p99), which are rare but often catastrophic, starting around 100 deaths and occasionally exceeding 1,000. Across panels, the vast majority of disaster events result in relatively few deaths, whereas severe events are infrequent but disproportionately impactful.

Figure 3 documents the evolution of natural disasters worldwide from 2007 to 2022. Most events fall below the median (p50), while only a small fraction reach the top decile (>p90) or top 5% (>p95) of the fatality distribution. Severe disasters are rare and relatively stable over time. This pattern confirms that our baseline treatment captures genuinely exceptional shocks without being driven by short-term spikes in disaster activity.

Figure 3: Number of Natural Disasters Over Time by Death and Damage Percentiles

(a) Number of natural disasters by death per capita percentile

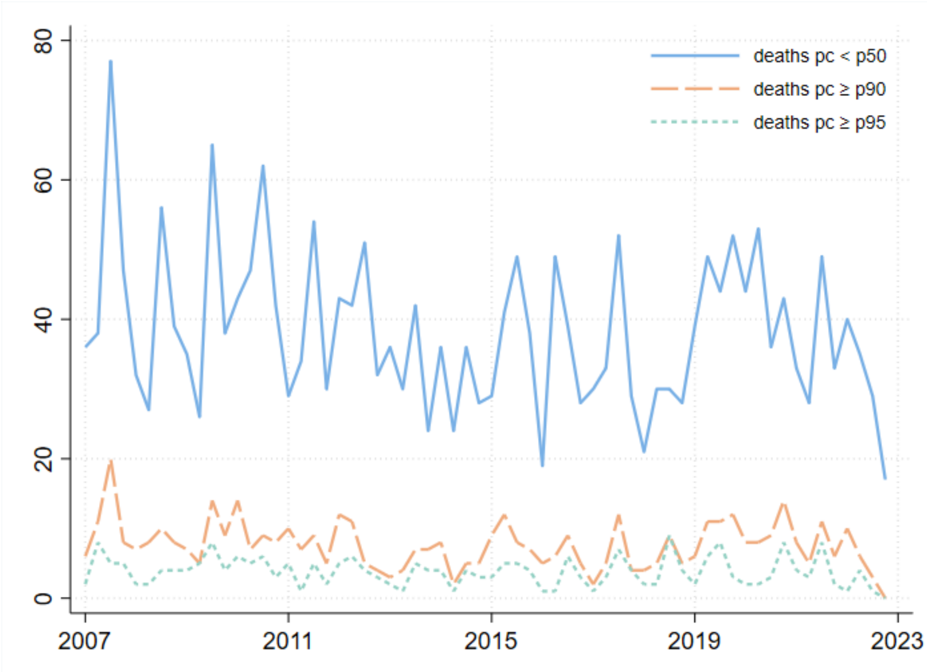

(b) Number of natural disasters by damage percentiles

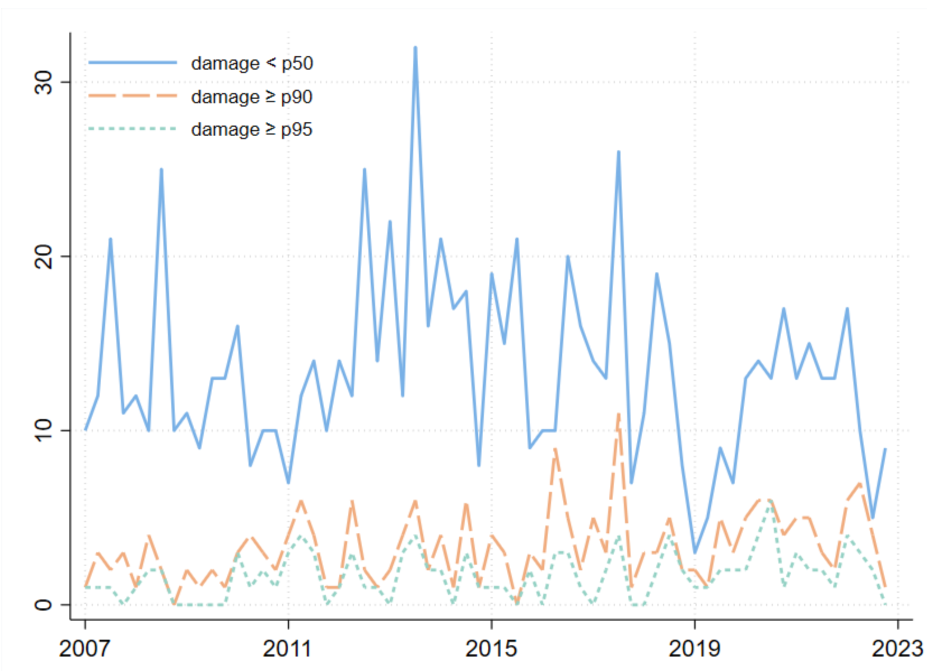

Notes: Panel (a) shows the quarterly number of natural disasters worldwide from 2007 to 2022, classified by the severity of fatalities. The blue line represents disasters with deaths below the median (p50), the orange line corresponds to disasters with deaths in the top 10% (≥p90), and the green line represents the most extreme events in the top 5% (≥p95). Panel (b) shows the corresponding quarterly counts of natural disasters classified by the distribution of reported economic damages (in millions of USD), with the same percentile thresholds applied. *Source*: EM-DAT International Disaster Database.

FIGURE 4: Frequency of Extreme Disaster Events by Mortality Thresholds

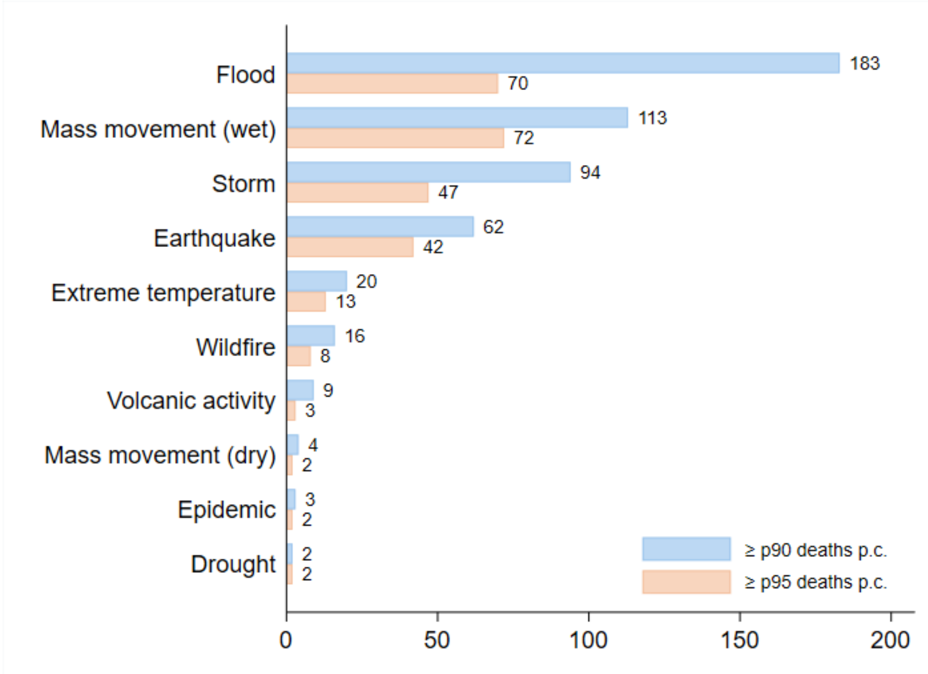

Notes: The figure shows the frequency of extreme disaster events by type, classified by the number of deaths per capita. Blue bars mark events above the 90th percentile (p90) and orange bars mark events above the 95th percentile (p95) of the deaths per capita distribution. Frequencies refer to disasters occurring between 2007 and 2022. *Source*: EM-DAT International Disaster Database.

Figure 4 shows the distribution of natural disasters in our sample by type and severity, measured by deaths per capita. Floods and epidemics dominate the most severe events, while droughts and wildfires rarely reach these thresholds. Our baseline classifies disasters above the 90th percentile of the deaths-per-capita distribution as extreme. As a robustness check, we sequentially exclude each disaster type to verify that no single category drives our results.

To ensure that our estimates reflect supply-chain disruptions rather than domestic shocks, we control for whether the U.S. importer itself was directly affected by a major natural disaster. In our baseline specification, we include an indicator for whether the importer's ADM was affected in the relevant period.

2.3 Global Ownership Linkages

We construct firm-to-firm ownership linkages using Orbis, a global dataset maintained by Bureau van Dijk that provides comprehensive information on corporate structures, including both listed and unlisted firms. Orbis aggregates data from national registries, annual reports, and other country-specific sources, covering firm characteristics such as revenues, assets, and detailed ownership hierarchies

A key advantage of Orbis is the breadth and precision of its ownership records. It identifies each firm's direct and indirect shareholders, subsidiaries, and its global ultimate owner (GUO), allowing us to map the full network of affiliates within a multinational corporation (MNC), even when parent and subsidiary operate in different countries. We define an ownership link when a parent holds at least 50% of an affiliate's equity. This enables us to classify U.S. firms as (i) majority-owned affiliates of foreign MNCs or (ii) U.S. parents with majority-owned foreign operations, providing a complete view of the multinational structures embedded in U.S. trade.

Orbis also provides firm names and physical addresses, which we use to match foreign exporters in the U.S. Bill of Lading data. This match reveals whether a given transaction links affiliates within the same multinational group—flagged as intra-firm trade—or instead occurs at arm's length between unrelated parties. Importantly, Orbis allows us to characterize each party individually, independently of whether they belong to the same group. A transaction between two multinational firms that do not share a corporate parent may still reflect different sourcing strategies, financing constraints, or network positions than one between truly independent firms. Together, these distinctions allow us to study how organizational form—at both the firm and relationship level—shapes the transmission of shocks through global production networks.

3 Empirical Strategy

Understanding how supply disruptions propagate through global production networks requires a framework that can trace shocks from their origin—a supplier hit by a natural disaster—through the web of trade relationships that connect it to downstream buyers. Our empirical strategy is built around this logic. We exploit the quasi-random timing and subnational geography of extreme natural disasters as exogenous shocks to foreign suppliers and trace their effects both directly—on the suppliers themselves—and indirectly, through the network of U.S. importers that depend on them. Our primary interest lies not only in measuring how far shocks travel, but in understanding why some firms absorb them better than others. The answer, we argue, lies in two features of the trading relationship: the organizational form—whether the importer is a multinational and whether trade occurs within or across firm boundaries—and the nature of the traded input. We exploit rich variation in both to characterize the mechanisms behind resilience and recovery. To this end, we integrate event-study specifications within a rich panel of firm-to-firm trade linkages, complemented by a detailed characterization of ownership networks and product specificity.

3.1 Identification and Assumptions

Our identification strategy exploits the quasi-random timing and geographic incidence of extreme natural disasters, which we treat as exogenous shocks to the supply capacity of affected foreign exporters.[3] Conditional on rich fixed effects, the timing of a disaster is assumed to be unrelated to the short-run dynamics of international trade links. We use an event-time specification centered on the first qualifying exposure to make the identifying restriction transparent: absent the disaster, outcomes would have followed comparable paths once the controls are in place. We assess this restriction empirically by inspecting pre-event dynamics in the results. Our estimation sample spans 2012 to 2022, excluding the Great Recession period (2008–2009) to avoid large aggregate demand shocks that could contaminate the identification of supply-side disruptions.[4]

Our identification strategy for importer-level regressions rests on a clear exclusion restriction: conditional on the fixed effects and controls described below, a disaster in a foreign supplier's district affects the U.S. importer only through the supply link with that supplier—not via contemporaneous demand shocks from the affected country, nor via direct physical damage to the importer's own location. This assumption motivates the two controls we include in all specifications. The specification absorbs broader macroeconomic shocks with high-dimensional fixed effects—importer and product-by-calendar-quarter—so the comparison is within product-time cells across importers that are and are not exposed to disaster-affected suppliers. To further sharpen identification, we include two additional controls. First, we account for whether the importer's own location was affected by a domestic disaster, using ADM2-level geographic matching within the United States. Second, we flag whether the affected foreign country is also an important export destination for the importer, helping to separate upstream supply-chain propagation from potential demand-side confounders. Both controls are included in all specifications and are described in more detail in Section 3.3.

A potential concern is that multinational suppliers may systematically locate in districts less exposed to natural disasters. Our exporter fixed effects absorb any time-invariant tendency of MNCs to locate in safer areas, and we find no evidence of differential disaster exposure rates between MNC and non-MNC suppliers in our sample.[5]

[3] Supply disruptions may operate through multiple channels, including direct destruction of productive capacity, damage to local infrastructure, or disruption to logistics and port operations that prevent goods from reaching export markets.

[4] We also exclude 2011 due to the Tōhoku earthquake and tsunami, which generated economy-wide disruptions in Japan—including the Fukushima nuclear disaster and widespread infrastructure damage—that affected suppliers well beyond the directly hit districts, contaminating the within-country control group. Data from 2011 are used solely to construct year-over-year growth rates for observations beginning in 2012.

[5] Disaster exposure rates are comparable across MNC and non-MNC exporters, with 8% of firm-quarters

A second concern relates to spillovers among foreign suppliers within the same country. If a disaster generates local general equilibrium effects—such as input reallocation or infrastructure disruptions—that also affect nearby unaffected suppliers, the control group may be partially contaminated. The direction of such spillovers is theoretically ambiguous, so their net effect on our estimates is unclear. Our country-by-quarter fixed effects absorb aggregate country-level shocks, so identification comes from within-country variation across affected and unaffected districts. Moreover, since EM-DAT records the full geographic extent of each disaster, the treated group already captures a broad affected area, limiting the scope for nearby control units to have been indirectly exposed.

Identifying Firms Exposed to Natural Disasters

The central challenge in our setting is tracing disaster events to the firms that source from affected suppliers. We do this by matching the geocoded subnational location of each foreign exporter to the ADM2-level geographic footprint of each disaster recorded in EM-DAT. This allows us to identify, for any given quarter, which U.S. importers are connected to a foreign supplier located in a disaster-hit district—and therefore potentially exposed to a supply disruption.

Linking disasters to importers is only the first step. Cross-border trade relationships are inherently sparse and irregular on both sides of the market: a foreign exporter may ship intermittently to U.S. buyers, just as a U.S. importer may not place orders in every quarter. A firm that does not appear in the data in a given quarter may have stopped trading altogether or simply not have intended to buy or sell in that period. Without knowing the counterfactual, it is difficult to interpret a zero as a disaster-induced disruption rather than a reflection of the naturally lumpy nature of international sourcing.[6] We address this by working at the quarterly frequency—granular enough to capture near-term disruptions while allowing year-over-year comparisons that smooth out seasonal noise—and by restricting the sample to firms with established trading patterns on both sides.

Specifically, we restrict the sample to *recurrent* firms—those that report at least one transaction per calendar year since their first appearance in the data. This *intensive margin* restriction ensures that both treated and control units face identical sample requirements by construction, and applies equally to foreign exporters in the first stage and to U.S. importers in the second stage. By focusing on firms with consistent trading histories, we increase the

affected among MNCs and 5% among non-MNCs.

[6]This restriction necessarily tilts the sample toward larger and more stable trading relationships. To the extent that smaller or less established firms are more vulnerable to supply disruptions, our estimates likely represent a lower bound on the effects in the broader population of importers.

likelihood that a post-disaster decline in trade reflects genuine disruption rather than the natural lumpiness of international sourcing. In Section 5, we additionally examine the sensitivity of our results to requiring that the affected supplier represent a non-trivial share of the importer's sourcing.

An importer is considered treated if at least one supplier in its import basket in the year prior to the shock is affected by an extreme natural disaster, defined as an event above the 90th percentile of the deaths-per-capita (DPC) distribution. Treatment is absorbing in the baseline specification: once a buyer is first exposed, it remains treated in all subsequent quarters, capturing the persistence of supply chain disruptions. Formally, we define the treatment indicator as:

$$\text{Treated}_{bt_0} = \begin{cases} 1 & \text{if } \exists\, s \in \mathcal{S}_b(t_0) \mid \text{DPC}_{st_0} > p_{90} \\ 0 & \text{otherwise} \end{cases} \tag{3.1}$$

where $\mathcal{S}_b(t_0)$ is the set of suppliers in the import basket of buyer b in the quarter of the disaster t_0, and DPC_{st_0} measures the intensity of disaster exposure of supplier s in that quarter. For treated buyers ($\text{Treated}_{bt_0} = 1$), we define event time as $\tau_{bt} = t - t_0(b)$, where $t_0(b)$ denotes the first quarter in which buyer b is exposed to a disaster-affected supplier; for never-treated buyers, all event-time indicators equal zero by convention.

3.2 Direct Impact on Foreign Suppliers

To estimate the immediate and dynamic effects of extreme natural disasters on foreign suppliers, we use an event-study design at the supplier–quarter level. Our baseline specification is:

$$X_{s,t} = \alpha_s + \delta_{p(s)\,t} + \gamma_{c(s)\,t} + \sum_{k=-K}^{L} \beta_k \cdot \mathbf{1}\{\tau_{st} = k\} + \varepsilon_{st}, \tag{3.2}$$

where $X_{s,t}$ represents two distinct measures of supplier performance: total exports to the U.S. and the total number of transactions each seller engages in across all U.S. buyers and products. Supplier fixed effects α_s control for time-invariant heterogeneity such as average productivity or geographic characteristics. The industry–time fixed effects $\delta_{p(s)\,t}$ capture seasonality and global demand shocks common to suppliers that share the same primary product, where $p(s)$ denotes supplier s's primary export industry at the HS2 level—i.e., the product category that accounts for the largest share of s's exports to the United States, and is therefore a characteristic of the supplier rather than of the individual transaction. Country–time fixed effects $\gamma_{c(s)t}$ absorb macroeconomic conditions, policy shocks, and other country-wide events. Finally, τ_{st} denotes event time relative to the first extreme natural disaster that struck the ADM2 district where supplier s is located, with $\tau_{st} < 0$ for pre-event quarters and $\tau_{st} > 0$ for post-event quarters. The indicator $\mathbf{1}\{\tau_{st} = k\}$ picks out observations k quarters from the event; for never-treated suppliers, all event-time indicators equal zero by convention. The reference period is $\tau_{st} = -1$.

To identify the effects, we restrict the baseline to *recurrent* exporters—foreign firms that report at least one shipment to the United States per calendar year since their first appearance in the data. This intensive margin restriction, consistent with the approach applied to U.S. importers in the second stage, limits the noise from sporadic sellers and isolates the impact of severe natural disasters on firm performance. To ensure clean treatment assignment, we exclude suppliers located in districts where disaster severity falls in an intermediate range ($p_{50} \leq \text{DPC} \leq p_{90}$), retaining only observations at the tails of the severity distribution.

The main variables of interest are measured as quarterly year-on-year midpoint growth rates, smoothed using a two-quarter moving average to reduce high-frequency volatility. Specifically, for each calendar quarter q in year y:

$$\tilde{g}_{q,y} = 2 \cdot \frac{\tilde{X}_{q,y} - \tilde{X}_{q,y-1}}{\tilde{X}_{q,y} + \tilde{X}_{q,y-1}},$$

where $\tilde{X}_{q,y}$ is a two-quarter moving average of trade flows, defined as $\tilde{X}_{q,y} = 0.5 \cdot (X_{q,y} + X_{q-1,y})$. When a firm does not appear in the data in a given quarter, we set $X_{q,y} = 0$, ensuring that firm exit or temporary inactivity is captured as a large negative value rather than a missing observation.[7] Panjiva provides consistent shipment quantities but does not report reliable trade values, which are needed to construct the growth rates described above. To address this, we construct real flows by applying 2010 unit values from the U.S. Census—defined at the HS6-level—to all reported quantities. This approach ensures that we capture changes in physical trade volumes priced at a fixed baseline, isolating the supply-side adjustments that matter for production-based transmission.

The coefficients of interest, β_k, trace the evolution of trade flows before and after a disaster. Leads ($k < 0$) test for pre-trends and potential anticipatory behavior, while lags ($k > 0$) capture the trajectory of post-disaster outcomes. A sharp decline in β_0 or β_1 would indicate an immediate contraction in shipments, while the path of β_k over subsequent quarters reveals the extent and persistence of the disruption—how long it takes for the supplier to recover

[7] Because each observation incorporates information from the adjacent quarter, consecutive observations for the same firm are mechanically correlated. We account for this by clustering standard errors at the ADM2-quarter level for supplier regressions and at the importer level for importer regressions, which allows for arbitrary correlation in the error terms within each firm over time.

its capacity to export to the U.S. Standard errors are clustered at the ADM2-quarter level to account for the fact that multiple suppliers in the same district may be simultaneously affected by the same disaster event.

3.3 Transmission to U.S. Importers

Having established the direct effects of disasters on foreign suppliers, we now turn to the downstream consequences for U.S. importers. The empirical design parallels the supplier-level regressions but shifts the unit of observation to the importer, allowing us to trace how upstream shocks propagate through established trade relationships. We estimate:

$$M_{bt} = \alpha_b + \delta_{p(b)t} + \theta\,\mathbf{1}\{\text{directly hit}_{bt}\} + \sum_{k=-K}^{L} \phi_k \cdot \mathbf{1}\{\tau_{bt} = k\} + X_{bt} + \eta_{bt}, \tag{3.3}$$

where M_{bt} denotes the outcome of importer b in quarter t, and τ_{bt} is the event-time indicator defined in Equation (3.1). The fixed effects α_b and $\delta_{p(b)t}$ control for time-invariant buyer characteristics and product-by-calendar-quarter shocks, respectively, where $p(b)$ is defined as the buyer's primary importing industry based on cumulative import value. The first coefficient, θ, captures the direct effect of a disaster occurring in the buyer's location. The coefficients of interest, ϕ_k, measure how shocks to qualifying foreign suppliers propagate through established trade links, tracing the full dynamic response of U.S. importers before and after the shock. The vector X_{bt} includes a control for whether the affected foreign country is also an important export destination for the importer, helping to separate upstream supply-chain propagation from demand-side confounders.

We estimate this equation separately for several outcomes, grouping them into two groups. The first captures adjustments along the import margin: the growth rate of total imports, the growth rate of imports sourced exclusively from suppliers never affected by a disaster, the number of active supplier relationships, and the number of transactions. The second captures broader firm performance: the growth rate of exports for buyers that also sell abroad and total sales for the subsample of firms matched to Compustat. Together, these outcomes allow us to assess whether supply disruptions are confined to the directly affected trade links or propagate more broadly through the buyer's organizational footprint. As in the supplier-side regressions, all outcomes are expressed as year-on-year midpoint growth rates smoothed using a two-quarter moving average. The reference period is $\tau_{bt} = -1$.[8] Standard errors are clustered at the importer level, allowing for arbitrary serial correlation

[8] In a robustness exercise, we additionally require that the affected supplier account for at least 5% of the importer's total sourcing value. See Section 5.

within each firm over time.

The specification in Equation (3.3) captures the average dynamic adjustment of U.S. importers to shocks in their sourcing network. In the next section, we explore whether these responses vary by firm and product characteristics—for instance, whether the importer is part of a multinational group, whether the affected supplier is an affiliate, or whether the good is differentiated.

3.4 Multinational Supply Networks, Input Specificity, and Recovery Dynamics

We now turn to the central question of the paper: What mechanisms govern the propagation of shocks through global value chains? Multinational corporations play a pivotal role in both absorbing and transmitting these shocks. Several questions guide this analysis: How do shocks propagate through cross-border trade and MNC networks? Does it matter that firms operate globally? Does it matter that trade occurs within firm boundaries? The answers to these questions are critical for understanding the consequences of trade disruptions—and, by extension, trade policy—in a world shaped by complex production networks.

Our analysis centers on the role of multinational ownership. We examine how the effects of foreign shocks vary with the buyer's multinational status and the presence of cross-border shared ownership links. These organizational structures may offer internal margins of adjustment—through coordination, reallocation, or financing—that are unavailable in more fragmented networks. Our estimating equations are:

$$M_{bt} = \alpha_b + \delta_{b(p)t} + \theta\,\mathbf{1}\{\text{directly hit}_{bt}\} + \sum_{k=-K}^{L} \Big(\beta_k + \phi_k \cdot \text{MNC}_s\Big) \cdot \mathbf{1}\{\tau_{bt} = k\} + \eta_{bt}, \tag{3.4}$$

$$M_{bt} = \alpha_b + \delta_{b(p)t} + \theta\,\mathbf{1}\{\text{directly hit}_{bt}\} + \sum_{k=-K}^{L} \Big(\nu_k + \delta_k \cdot \text{Intra-firm}_s\Big) \cdot \mathbf{1}\{\tau_{bt} = k\} + \epsilon_{bt}. \tag{3.5}$$

Here, MNC_s is an indicator for whether the US importer belongs to a U.S. multinational network, and Intra-firm_s denotes whether the trade relationship occurs between a U.S. buyer and a foreign exporter belonging to the same multinational corporate group. These specifications allow us to test whether multinational ownership and intra-firm linkages mediate the severity and persistence of disaster impacts. If MNCs provide an internal insurance channel, suppliers linked by equity ties may benefit from financial, operational, or relational support that accelerates post-disaster recovery. By estimating event-time dynamics separately for MNC ($\beta_k + \phi_k$) and non-MNC (β_k) buyers, and for intra-firm ($\nu_k + \delta_k$) versus arm's-length (ν_k) trade, we assess how cross-border ownership affects the transmission and recovery patterns following extreme shocks.

We also analyze a second source of heterogeneity: product specificity. Suppliers that export highly differentiated or hard-to-substitute goods may be less likely to lose their buyer base after a disruption, enabling faster recovery. We proxy for specificity using measures such as the classification in Rauch (1999), and estimate separate event-time dynamics for high- and low-specificity goods. These parallel specifications help us assess how the nature of the traded input shapes the speed and extent of recovery.

Importantly, the same specificity that amplifies initial disruptions may also accelerate recovery. Buyers—particularly affiliates within MNCs—have strong incentives to preserve relationships with specialized suppliers. During temporary disruptions, they may absorb short-term costs or rely on inventories, rather than switch to alternative suppliers. As a result, once the original supplier resumes operations, the trade link can be quickly restored. Recovery is faster because establishing new supplier relationships involves costly adjustments—especially for highly specific inputs.

4 Results

This section presents the main empirical results on how natural disasters affecting foreign suppliers propagate through international production networks. We first describe the structure of the estimation sample and then estimate the direct effects on foreign exporters—firms located in disaster-hit regions—focusing on their exports and transaction volumes to the United States. We then turn to the U.S. importers connected to those suppliers to examine the extent and timing of downstream disruptions. Finally, we explore heterogeneity in exposure and response, distinguishing firms by ownership structure and the nature of their input relationships.

4.1 Estimation Sample

Before turning to the main results, we briefly describe the structure of the estimation sample. All statistics below reflect the final panel of firm-to-firm relationships used in the main specifications—after applying the filters for importer recurrence and disaster classification described above.

Table 1 presents summary statistics for the estimation sample, which tracks the evolution of international trade relationships between U.S. importers and their foreign suppliers at quarterly frequency. Panel A summarizes the sample of foreign exporters, who serve as potential sources of disruption. These are recurrent firms—those that report at least one shipment

Table 1: Descriptive Statistics

Panel A. Exporters	Obs.	Mean	Std. Dev.	p25	p50	p75
Exports (M USD, 2010)	3,284,657	2.59	88.40	0.03	0.14	0.60
Transactions (count)	3,284,657	16.81	151.35	2.00	5.00	13.00
Total buyers (count)	3,284,657	32.05	119.34	6.00	16.00	35.00
Export growth (y/y)	3,160,324	0.06	0.92	-0.55	0.06	0.67
Transactions growth (y/y)	3,209,334	0.04	0.69	-0.40	0.00	0.47
Active quarters per year	157,168	2.87	0.94	2.00	3.00	3.71
Dummy variables						
Multinational firm	221,825	0.05	0.23			
Natural disaster event	221,825	0.07	0.26			

Panel B. Importers	Obs.	Mean	Std. Dev.	p25	p50	p75
Imports (M USD, 2010)	2,641,032	3.43	81.68	0.04	0.16	0.67
Imports excl. aff. suppliers (M USD, 2010)	2,628,779	3.22	78.53	0.03	0.14	0.62
Transactions (count)	2,641,032	21.00	162.33	2.00	5.00	13.00
Suppliers per quarter (count)	2,641,032	4.15	10.16	1.00	2.00	4.00
Imports growth (y/y)	2,386,869	0.05	0.90	-0.52	0.05	0.64
Transactions growth (y/y)	2,428,214	0.04	0.69	-0.40	0.00	0.44
Suppliers growth (y/y)	2,428,214	0.02	0.54	-0.29	0.00	0.33
Active quarters per year	272,285	1.88	1.10	1.00	1.00	3.00
Dummy variables						
Multinational firm	581,559	0.04	0.19			
Intrafirm trade	581,559	0.01	0.10			
Diff. good (liberal def.)	581,559	0.60	0.49			

Notes: This table reports summary statistics for the estimation sample (2012–2022). Panel A presents foreign exporters—firms located outside the United States that ship to U.S. importers, as recorded in U.S. Bill of Lading data. The unit of observation is the exporter-quarter. Export values are in millions of constant 2010 U.S. dollars, constructed by applying 2010 HS6-by-country unit values from the U.S. Census to quantity data from Panjiva. Growth rates are midpoint year-over-year changes computed over all post-entry quarters and trimmed at $|g| < 2$. *Multinational firm* is an indicator for exporters belonging to a multinational group based on Orbis ownership linkages. *Natural disaster event* equals one if the exporter is located in an ADM2 region affected by a qualifying disaster. Panel B presents statistics for U.S. importers. The unit of observation is the importer-quarter. *Imports excl. aff. suppliers* excludes flows from ever disaster-affected suppliers. *Diff. good (liberal def.)* classifies traded products based on the Rauch (1999) classification.

to the United States per calendar year since their first appearance in the data—matched to natural disaster exposure via geocoded location data. The median foreign exporter ships $14,000 per quarter (2010 USD) across approximately 5 transactions to a total of 16 U.S. buyers. However, the distribution is highly skewed, with exporters at the 99th percentile shipping over $24 million in quarterly flows.[9] Export growth over a four-quarter horizon is 6 percent on average but highly volatile, with a standard deviation of 92. Roughly 5% of exporter-quarters correspond to multinational firms, and 7% coincide with an extreme natural disaster affecting the exporter's subnational location (ADM2), as defined in our identification strategy.

Panel B presents summary statistics for U.S. importers—the central units of analysis—over the sample period 2012–2022. The sample comprises over 2.6 million firm-quarter observations, restricted to *recurrent importers*: U.S. firms with at least one recorded import transaction in each of the four quarters preceding their first appearance in the data. The average importer maintains around 4 active suppliers per quarter. Average quarterly import growth over four quarters is 5%, with considerable dispersion.

Only 4% of importers are affiliated with multinational firms, and just 1% of importer–supplier links are intra-firm, reflecting common ownership across borders. About 60% of transactions involve differentiated goods. The average importer is active in fewer than 2 quarters per year, underscoring the sparse and lumpy nature of international sourcing—and highlighting the need to focus on economically meaningful buyer–supplier relationships when estimating shock transmission.

4.2 Direct Effects on Foreign Exporters

Before tracing how shocks propagate downstream, we first establish that extreme natural disasters generate meaningful disruptions at the source. Table 2 reports event-study coefficients for two key outcomes: total export value to the United States and the number of transactions. Coefficients are expressed as quarterly year-on-year midpoint growth rates, centered on the quarter in which the first qualifying disaster affecting the supplier's ADM2 location occurred. Estimates are based on a panel of recurrent exporters and include exporter, product-by-calendar-quarter, and country-by-quarter fixed effects. Standard errors are clustered at the ADM2-quarter level.

The results reveal a sharp and persistent contraction in trade activity following the disaster. Export value displays flat pre-trends, consistent with the parallel trends assumption, and

[9] For context, this corresponds to roughly 0.005% of total U.S. quarterly goods imports of approximately USD 478 billion, a modest share that reflects the vast number of suppliers serving the U.S. market.

Table 2: Effects of Natural Disasters on Foreign Exporters

	Export Value Growth	Transactions Growth
Disaster hits supplier $(t-4)$	0.006	$0.094^{(c)}$
	(0.055)	(0.051)
Disaster hits supplier $(t-3)$	-0.004	0.047
	(0.054)	(0.047)
Disaster hits supplier $(t-2)$	0.028	0.036
	(0.039)	(0.038)
Reference period: $t-1$		
Disaster hits supplier (t)	-0.051	-0.066
	(0.054)	(0.049)
Disaster hits supplier $(t+1)$	$-0.132^{(c)}$	$-0.129^{(b)}$
	(0.078)	(0.063)
Disaster hits supplier $(t+2)$	$-0.304^{(a)}$	$-0.261^{(a)}$
	(0.090)	(0.067)
Disaster hits supplier $(t+3)$	$-0.359^{(a)}$	$-0.345^{(a)}$
	(0.110)	(0.088)
Disaster hits supplier $(t+4)$	$-0.299^{(a)}$	$-0.324^{(a)}$
	(0.108)	(0.090)
Disaster hits supplier $(t+5)$	-0.096	$-0.172^{(a)}$
	(0.073)	(0.059)
Disaster hits supplier $(t+6)$	0.107	0.015
	(0.089)	(0.086)
Disaster hits supplier $(t+7)$	0.019	-0.123
	(0.092)	(0.089)
Disaster hits supplier $(t+8)$	0.000	-0.088
	(0.128)	(0.134)
Exporter FE	Yes	Yes
Industry × Time FE	Yes	Yes
Country × Time FE	Yes	Yes
Observations	1,580,661	1,699,443
R^2	0.173	0.200

Notes: This table reports event-study estimates of the impact of extreme natural disasters on directly affected foreign exporters. The dependent variables are the midpoint year-on-year growth rates for export value and transaction count, both smoothed with a two-quarter moving average. Estimates are based on the following specification:

$$X_{st} = \alpha_s + \delta_{p(s)t} + \gamma_{c(s)t} + \sum_{k=-K}^{L} \beta_k \cdot \mathbf{1}\Big\{\tau_{st} = k\Big\} + \varepsilon_{st}$$

where X_{st} denotes the outcome of interest for exporter s in calendar quarter t, and τ_{st} denotes the number of quarters elapsed since the first extreme natural disaster affecting supplier s's ADM2 district. The reference period is $t-1$. The sample includes recurrent exporters. All regressions include exporter fixed effects (α_s), industry-calendar-quarter fixed effects ($\delta_{p(s)t}$), and country-calendar-quarter fixed effects ($\gamma_{c(s)t}$). Standard errors are clustered at the ADM2-time level. Significance: $^{(a)}$ 1%; $^{(b)}$ 5%; $^{(c)}$ 10%.

falls by 5.1 percentage points on impact, deepening to 30 percentage points by the second quarter and 36 percentage points by the third. To put this in context, given that average annual export growth in the sample is 6%, the contraction represents roughly a 30% decline in exports. In levels, with average exports in an active quarter of USD 2.6million, this implies an average loss of more than USD 0.8 million per exporter per active quarter.

A similar pattern emerges for the number of transactions: the initial drop is slightly steeper, and recovery occurs around 1.5 years after the event. With the typical exporter conducting nearly 17 transactions per active quarter, the peak decline of 35 percentage points at $t = 3$ implies a loss of more than five shipments per quarter at the height of the disruption. The observed persistence is not driven by outliers or timing heterogeneity—the pre-trends are flat, and the coefficients decline monotonically in the post-event window before partially recovering.

These results establish that extreme natural disasters generate large and lasting disruptions to firms' export capacity. The observed persistence reflects sustained declines in average export volumes among surviving firms—driven by a contraction in both the number of U.S. buyers they serve and their sales per buyer—suggesting that the damage to trade relationships extends well beyond the immediate aftermath of the shock. We next examine how these disruptions propagate downstream to U.S. importers.

Heterogeneity by Ownership at the Supplier Level. To assess whether multinational ownership buffers the effects of disasters on directly hit suppliers, Figure B.1 in Appendix B disaggregates the event-study by exporter status. We find no significant difference in the magnitude or persistence of the impact: both MNC and non-MNC suppliers experience a sharp and enduring fall in trade activity, with no evidence of faster recovery among multinationals. Physical destruction appears to affect all suppliers equally—while multinational status may facilitate coordination and financing, these channels do not mitigate direct export losses once a facility is hit.

Overall, these first-stage results establish that extreme natural disasters generate sizeable and persistent supply shocks at the exporter level, independent of ownership structure. We next examine how these shocks propagate downstream to U.S. importers.

4.3 Propagation to U.S. Importers

Having established that extreme natural disasters generate large and persistent supply shocks at the source, we now examine how these disruptions propagate downstream to U.S. importers. The empirical design parallels the supplier-level regressions but shifts the unit of

Figure 5: Supplier Export Disruptions and Downstream Propagation to U.S. Importers

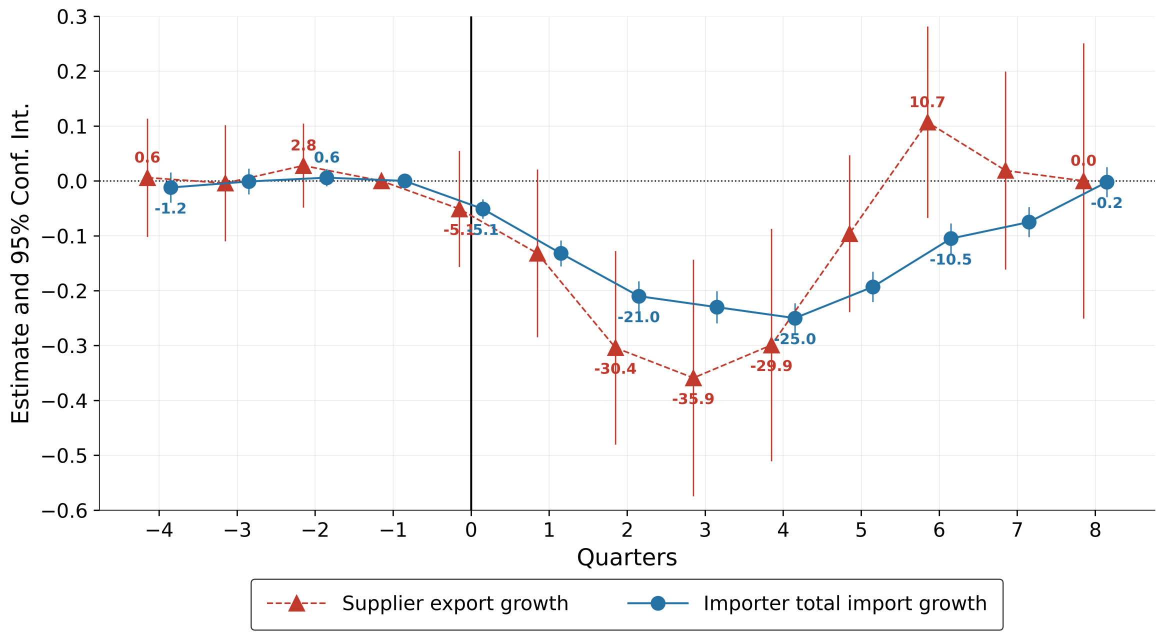

Notes: This figure overlays event-study coefficients from two specifications estimated on directly exposed trade links. The red dashed line plots $\hat{\beta}_k$ from the first-stage supplier regression:

$$X_{st} = \alpha_s + \delta_{p(s)t} + \gamma_{c(s)t} + \sum_{k=-K}^{L} \beta_k \cdot \mathbf{1}\Big\{\tau_{st} = k\Big\} + \varepsilon_{st}$$

where X_{st} is the midpoint year-on-year growth rate of export value for supplier s at calendar quarter t, and τ_{st} denotes quarters elapsed since the first extreme natural disaster affecting supplier s's ADM2 district. The blue solid line plots $\hat{\phi}_k$ from the second-stage importer regression:

$$M_{bt} = \alpha_b + \delta_{p(b)t} + \theta\,\mathbf{1}\{\text{directly hit}_{bt}\} + \sum_{k=-K}^{L} \phi_k \cdot \mathbf{1}\{\tau_{bt} = k\} + X_{bt} + \eta_{bt},$$

where M_{bt} is the midpoint year-on-year growth rate of total import value for importer b at calendar quarter t, τ_{bt} denotes quarters elapsed since the first quarter in which any supplier in buyer b's import basket is hit by an extreme natural disaster, as defined in Section 3, and X_{bt} controls for whether the affected country is also an important export destination for the importer. The reference period is $t-1$. The first-stage sample comprises recurrent exporters with exporter, product-by-calendar-quarter, and country-by-calendar-quarter fixed effects; standard errors are clustered at the ADM2-quarter level. The second-stage sample comprises recurrent importers with importer and product-by-calendar-quarter fixed effects; standard errors are clustered at the importer level. Vertical bars denote 95% confidence intervals.

observation to the importer, allowing us to trace how upstream shocks travel through established trade relationships. Treatment is defined at the importer level and is absorbing: once a firm becomes exposed—through any qualifying foreign supplier hit by an extreme disaster—it remains treated in all subsequent quarters. All specifications include importer fixed effects and product-by-calendar-quarter fixed effects, and standard errors are clustered at the importer level.

Figure 5 juxtaposes the importer response (blue solid line) with the underlying supplier export contraction (red dashed line). Total imports of exposed buyers decline by up to 25 percentage points over the four quarters following the shock. Relative to the 5% baseline growth these firms would otherwise enjoy, this corresponds to a net contraction of 20%—enough to turn a period of steady expansion into one of marked contraction. The importer decline is smaller than the supplier-side contraction, which reaches 35 percentage points at its deepest, as expected, since any given supplier typically accounts for only a fraction of the importer's total sourcing.

The two series also diverge sharply in their dynamics. Supplier exports recover substantially within six quarters as productive capacity is rebuilt, while importer outcomes remain depressed for nearly two years. Were the disruption purely a matter of input availability, buyer imports would track the supplier's recovery. The absence of such co-movement points to a deeper friction: rebuilding a cross-border sourcing relationship takes considerably longer than restoring a supplier's capacity to ship. Reassuringly, neither specification exhibits differential pre-trends, ruling out pre-existing differences between treated and control firms as an alternative explanation.[10]

Table 3 examines how the disruption affects different margins of importer activity. A natural question is whether the decline in total imports simply reflects the mechanical loss of shipments from the affected supplier, or whether it extends to the rest of the importer's sourcing network. The second column addresses this by isolating imports from suppliers that have never been hit by a disaster. This measure effectively captures the importer's ability to sustain trade with its unaffected partners—a closer proxy for the firm's underlying operational capacity than total imports, which are partly driven by the direct supplier link. The results show that even this margin contracts significantly on impact, consistent with production complementarities: when a key input is missing, the firm scales back output and reduces demand for all intermediates, not just the disrupted one. The gradual recovery of this margin

[10]Figure B.2 in Appendix B.2 extends the analysis to the number of transactions—a measure that captures the extensive margin of trade activity. The pattern reinforces the value-based evidence: supplier transactions decline sharply and bottom out near 40 percentage points by $t+3$, while importer transactions follow a more gradual but strikingly persistent path.

Table 3: Effects of Natural Disasters on U.S.-Exposed Importers

	Imports Total	Imports Never-Aff.	Exports	Sales
Disaster hits one supplier $(t-4)$	-0.012	0.043$^{(a)}$	0.030	-0.067
	(0.014)	(0.016)	(0.076)	(0.055)
Disaster hits one supplier $(t-3)$	-0.001	0.030$^{(b)}$	0.027	-0.002
	(0.012)	(0.014)	(0.065)	(0.046)
Disaster hits one supplier $(t-2)$	0.006	0.012	0.033	0.010
	(0.008)	(0.010)	(0.047)	(0.035)
Reference period: $t-1$				
Disaster hits one supplier (t)	-0.051$^{(a)}$	-0.040$^{(a)}$	-0.025	-0.033
	(0.009)	(0.010)	(0.047)	(0.037)
Disaster hits one supplier $(t+1)$	-0.132$^{(a)}$	-0.070$^{(a)}$	-0.195$^{(a)}$	-0.068
	(0.012)	(0.014)	(0.067)	(0.052)
Disaster hits one supplier $(t+2)$	-0.210$^{(a)}$	-0.106$^{(a)}$	-0.246$^{(a)}$	-0.064
	(0.014)	(0.016)	(0.076)	(0.062)
Disaster hits one supplier $(t+3)$	-0.230$^{(a)}$	-0.115$^{(a)}$	-0.269$^{(a)}$	-0.042
	(0.015)	(0.018)	(0.082)	(0.071)
Disaster hits one supplier $(t+4)$	-0.250$^{(a)}$	-0.132$^{(a)}$	-0.405$^{(a)}$	-0.005
	(0.014)	(0.017)	(0.077)	(0.057)
Disaster hits one supplier $(t+5)$	-0.193$^{(a)}$	-0.115$^{(a)}$	-0.406$^{(a)}$	-0.012
	(0.014)	(0.016)	(0.076)	(0.054)
Disaster hits one supplier $(t+6)$	-0.105$^{(a)}$	-0.057$^{(a)}$	-0.439$^{(a)}$	0.030
	(0.014)	(0.016)	(0.078)	(0.052)
Disaster hits one supplier $(t+7)$	-0.075$^{(a)}$	-0.021	-0.396$^{(a)}$	0.022
	(0.014)	(0.016)	(0.080)	(0.051)
Disaster hits one supplier $(t+8)$	-0.002	0.041$^{(b)}$	-0.371$^{(a)}$	0.083
	(0.014)	(0.016)	(0.081)	(0.055)
Importer FE	Yes	Yes	Yes	Yes
Product × Time FE	Yes	Yes	Yes	Yes
Observations	3,092,779	3,016,669	76,849	18,699
R^2	0.454	0.440	0.179	0.717

Notes: This table reports event-study estimates of the impact of extreme natural disasters on U.S. importers. Sales correspond to the subsample of firms matched to Compustat. Estimates are based on the following specification:

$$M_{bt} = \alpha_b + \delta_{p(b)t} + \theta\,\mathbf{1}\{\text{directly hit}_{bt}\} + \sum_{k=-K}^{L} \phi_k \cdot \mathbf{1}\{\tau_{bt} = k\} + X_{bt} + \eta_{bt}$$

where M_{bt} denotes the midpoint year-on-year growth rate of the outcome of interest for importer b in quarter t, smoothed using a two-quarter moving average, τ_{bt} is defined as in Equation (3.1), and X_{bt} controls for whether the affected country is also an important export destination for the importer. The reference period is $t-1$. All regressions include importer fixed effects (α_b) and product-by-quarter fixed effects ($\delta_{p(b)t}$). Standard errors are clustered at the importer level. Significance: $^{(a)}$ 1%; $^{(b)}$ 5%; $^{(c)}$ 10%.

over roughly two years suggests that importers do eventually substitute toward alternative sources, but the speed of this adjustment—measured in quarters, not weeks—reflects the well-documented costs of forming new supplier relationships in international trade.

The consequences extend beyond the import margin. For the subset of buyers that also sell abroad, the third column of Table 3 shows that export growth declines by more than 40 percentage points within four quarters and remains depressed for over two years. This pattern is consistent with the vertical structure of production: when a critical foreign intermediate becomes unavailable, the buyer's own output and export capacity are directly impaired—providing micro-level evidence that localized supply shocks can cascade through multiple layers of a production network. The fourth column reports total sales for firms matched to Compustat. Point estimates are negative and economically meaningful—on the order of 5 to 7 percentage points—but imprecisely estimated, reflecting the limited size of this subsample rather than the absence of a real effect. Table B.1 in the Appendix confirms this interpretation: treated importers experience a significant and sustained erosion of their active supplier base, with no short-run offset through new relationships. The directly affected margin bears the brunt of the adjustment, with declines in import values from disaster-hit suppliers exceeding one full log point over the four quarters.

4.4 Ownership and Input-Specificity as Moderators of Propagation

The baseline results establish that U.S. importers contract sharply when a foreign supplier is hit by a natural disaster, and that part of this contraction spills over to trade with suppliers that were never directly affected. Whether these disruptions propagate uniformly across importers or are shaped by the organizational and contractual features of the trading relationship is a first-order question for understanding—and potentially mitigating—supply-chain fragility. We explore two dimensions: whether the importing firm belongs to a multinational network and whether the disrupted trade occurs within or outside the firm's boundaries. A third dimension, the degree of input specificity, tests whether relationship-specific inputs amplify or dampen the transmission of shocks.

Table 4 interacts the baseline event-study with an indicator for the multinational status of the importer. Both MNC and non-MNC importers experience a significant decline in total imports after the shock, but the magnitudes diverge quickly. Non-MNC importers reach a peak contraction of roughly 30 percentage points within four quarters and remain depressed throughout the two-year horizon, still 12 percentage points below baseline after seven quarters. MNC importers, by contrast, contract by about half as much—14 percentage points at the trough—and return to pre-shock levels within six quarters. The pattern in

Table 4: Effects of Natural Disasters on U.S. Importers: *By Multinational Status*

	Imports Total		Imports Never-Aff.		Exports	
	MNC	Non-MNC	MNC	Non-MNC	MNC	Non-MNC
Disaster hits one supplier ($t-4$)	−0.006	−0.013	0.014	0.059[a]	0.008	0.067
	(0.023)	(0.016)	(0.026)	(0.019)	(0.094)	(0.112)
Disaster hits one supplier ($t-3$)	0.005	−0.003	0.010	0.041[b]	0.015	0.047
	(0.020)	(0.014)	(0.022)	(0.017)	(0.080)	(0.096)
Disaster hits one supplier ($t-2$)	0.018	0.002	0.013	0.012	0.023	0.048
	(0.014)	(0.010)	(0.016)	(0.012)	(0.059)	(0.070)
Reference period: $t-1$						
Disaster hits one supplier (t)	−0.021	−0.064[a]	−0.015	−0.052[a]	0.025	−0.091
	(0.014)	(0.010)	(0.016)	(0.013)	(0.057)	(0.072)
Disaster hits one supplier ($t+1$)	−0.075[a]	−0.157[a]	−0.023	−0.092[a]	−0.142[c]	−0.267[a]
	(0.020)	(0.014)	(0.022)	(0.017)	(0.083)	(0.096)
Disaster hits one supplier ($t+2$)	−0.120[a]	−0.248[a]	−0.054[b]	−0.131[a]	−0.191[b]	−0.321[a]
	(0.023)	(0.016)	(0.026)	(0.020)	(0.095)	(0.109)
Disaster hits one supplier ($t+3$)	−0.127[a]	−0.274[a]	−0.065[b]	−0.139[a]	−0.205[b]	−0.357[a]
	(0.026)	(0.018)	(0.029)	(0.022)	(0.102)	(0.118)
Disaster hits one supplier ($t+4$)	−0.143[a]	−0.296[a]	−0.082[a]	−0.157[a]	−0.328[a]	−0.510[a]
	(0.025)	(0.017)	(0.028)	(0.021)	(0.097)	(0.113)
Disaster hits one supplier ($t+5$)	−0.094[a]	−0.235[a]	−0.069[a]	−0.138[a]	−0.237[b]	−0.639[a]
	(0.024)	(0.017)	(0.027)	(0.020)	(0.094)	(0.112)
Disaster hits one supplier ($t+6$)	−0.018	−0.143[a]	0.004	−0.088[a]	−0.273[a]	−0.671[a]
	(0.024)	(0.017)	(0.027)	(0.020)	(0.098)	(0.113)
Disaster hits one supplier ($t+7$)	0.026	−0.119[a]	0.052[c]	−0.056[a]	−0.281[a]	−0.559[a]
	(0.024)	(0.017)	(0.027)	(0.020)	(0.100)	(0.116)
Disaster hits one supplier ($t+8$)	0.088[a]	−0.041[b]	0.110[a]	0.007	−0.308[a]	−0.464[a]
	(0.024)	(0.017)	(0.027)	(0.020)	(0.103)	(0.117)
Buyer FE	Yes	Yes	Yes	Yes	Yes	Yes
Product × Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3,092,779		3,016,669		76,849	
R^2	0.454		0.440		0.179	

Notes: This table reports event-study estimates of the impact of extreme natural disasters on U.S. importers, disaggregated by the multinational status of the importing firm. Columns (1)–(2) report total import growth, columns (3)–(4) report import growth excluding flows from directly affected suppliers, and columns (5)–(6) report export growth. "MNC" denotes importers that belong to a multinational group; "Non-MNC" denotes all other importers. All dependent variables are midpoint year-on-year growth rates, smoothed using a two-quarter moving average. Estimates are based on:

$$M_{bt} = \alpha_b + \delta_{b(p)t} + \theta\,\mathbf{1}\{\text{directly hit}_{bt}\} + \sum_{k=-K}^{L} \Big(\beta_k + \phi_k \cdot \text{MNC}_s\Big) \cdot \mathbf{1}\{\tau_{bt} = k\} + \eta_{bt},$$

where MNC_s is an indicator for whether the importer belongs to a U.S. multinational network. Reference period: $t-1$. All regressions include buyer-by-exporter, quarter, and ADM1-by-time fixed effects. Standard errors clustered at the buyer-by-exporter level. [a] 1%; [b] 5%; [c] 10%.

imports from never-affected suppliers is qualitatively similar: non-MNC importers cut these flows by 16 percentage points at the peak, whereas MNC importers reduce them by only 8 percentage points and recover fully by the end of the window. Panel A of Figure 6 confirms visually that the gap between the two groups opens immediately after impact and widens over the first year. Appendix Figure B.3 shows that the same ranking holds—and is even more pronounced—for export growth, where non-MNC importers suffer declines exceeding 70 percentage points at their trough.

The MNC comparison, however, conflates organizational structure with the nature of the disrupted sourcing relationship, because multinationals also engage extensively in arm's-length trade. Columns (1)–(2) of Table 5 therefore interact the baseline specification with an indicator for whether the affected supplier-importer pair involves related-party trade. Arm's-length importers reach a peak decline of roughly 26 percentage points within four quarters and remain significantly depressed through two years. Intra-firm importers contract by roughly half that magnitude, and the point estimates lose significance after five quarters, consistent with faster recovery. Panel B of Figure 6 illustrates the divergence. This result runs counter to the intuition that intra-firm trade, which typically involves more specialized inputs and tighter production integration, should be harder to substitute when disrupted. Instead, the evidence points to the advantages of organizational coordination: multinational firms can reallocate orders across affiliates, reroute shipments through internal logistics, or accelerate the restoration of the disrupted link—adjustment margins that are largely unavailable in arm's-length settings.

Columns (3)–(4) of Table 5 interact the baseline specification with an indicator for product differentiation, based on the Rauch (1999) conservative classification. Differentiated products display a somewhat larger initial contraction—about 6 percentage points on impact versus 4 for non-differentiated goods—but the gap narrows quickly and the two categories follow broadly similar trajectories over the remainder of the horizon. By the peak, both groups contract by 23–27 percentage points, and recovery proceeds at comparable rates. This pattern differs from Barrot and Sauvagnat (2016), who find a more persistent role for input specificity in domestic propagation. The difference may reflect the nature of cross-border adjustment: international sourcing relationships are inherently harder to replace regardless of product type, which could compress the differential between specific and non-specific inputs. Appendix Table B.2 confirms that imports from never-affected suppliers behave similarly across the two categories.

The heterogeneity results assign a clear role to organizational structure as a moderator of cross-border shock propagation. The finding that MNC and intra-firm importers recover

Figure 6: Heterogeneous Effects of Natural Disasters on U.S. Importers' Total Imports

Panel A: By Multinational Status

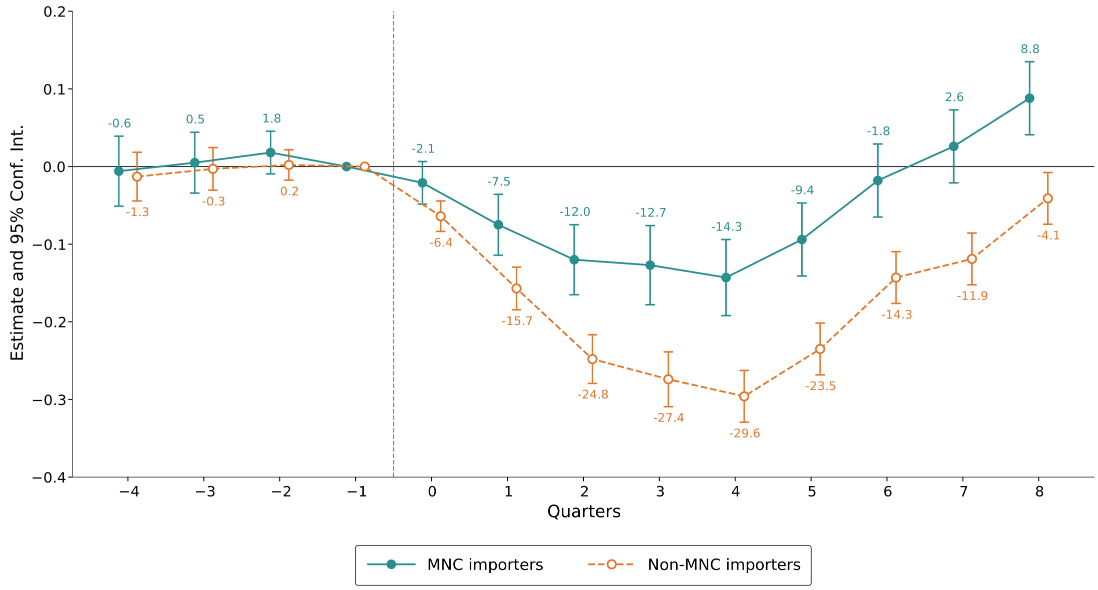

Panel B: By Intra-Firm Trade Status

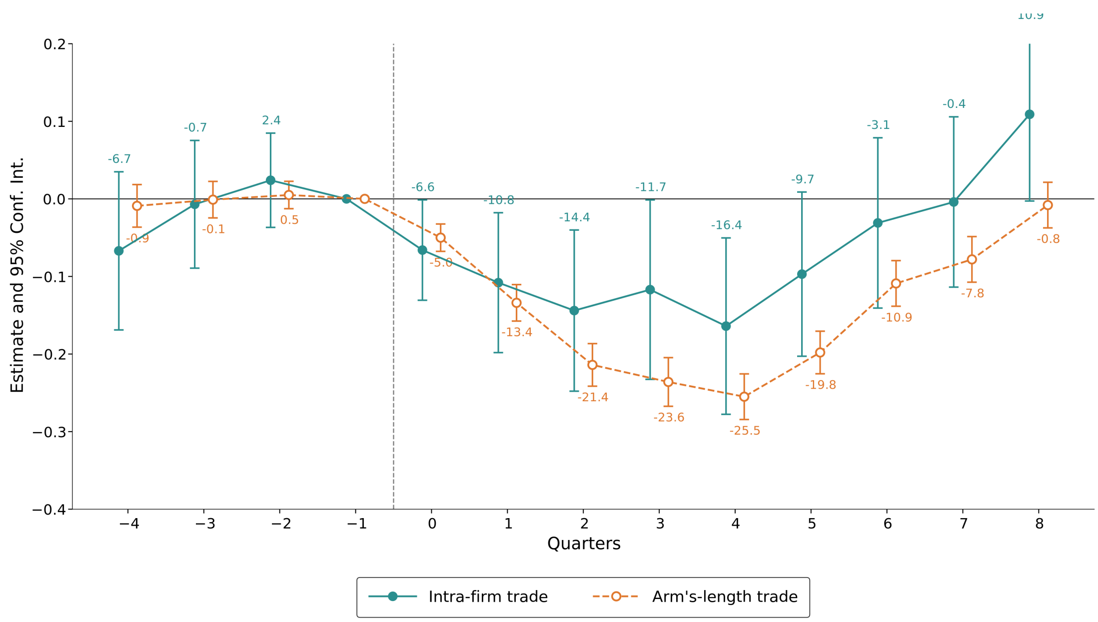

Notes: This figure overlays event-study coefficients from two interaction specifications estimated on directly exposed trade links. The dependent variable in all panels is the midpoint year-on-year growth rate of total import value, smoothed using a two-quarter moving average. Panel A plots $\hat{\beta}_k$ and $\hat{\beta}_k + \hat{\phi}_k$ from:

$$M_{bt} = \alpha_b + \delta_{p(b)t} + \gamma_{ct} + \sum_{k=-K}^{L} (\beta_k + \phi_k \cdot \text{MNC}_b) \, \mathbf{1}\{\tau_{bt} = k\} + \varepsilon_{bt},$$

where MNC_b indicates whether the importer belongs to a U.S. multinational network. The teal solid line plots $\hat{\beta}_k + \hat{\phi}_k$ (MNC importers); the orange dashed line plots $\hat{\beta}_k$ (non-MNC importers). Panel B replaces the interaction with IntraFirm_{bs}, an indicator for whether the importer-supplier pair is classified as related-party trade in U.S. Customs data. All regressions include buyer and product-quarter fixed effects. Standard errors are clustered at the buyer-by-exporter level. Vertical bars denote 95% confidence intervals.

Table 5: Heterogeneous Effects of Natural Disasters on U.S. Importers' Total Imports

	Intra-Firm Trade		Product Differentiation	
	Intra	Arm's-L.	Diff.	Non-Diff.
Disaster hits one supplier $(t-4)$	-0.067	-0.009	-0.006	-0.020
	(0.052)	(0.014)	(0.018)	(0.020)
Disaster hits one supplier $(t-3)$	-0.007	-0.001	0.003	-0.007
	(0.042)	(0.012)	(0.016)	(0.017)
Disaster hits one supplier $(t-2)$	0.024	0.005	0.016	-0.007
	(0.031)	(0.009)	(0.011)	(0.012)
Reference period: $t-1$				
Disaster hits one supplier (t)	$-0.066^{(b)}$	$-0.050^{(a)}$	$-0.059^{(a)}$	$-0.040^{(a)}$
	(0.033)	(0.009)	(0.012)	(0.012)
Disaster hits one supplier $(t+1)$	$-0.108^{(b)}$	$-0.134^{(a)}$	$-0.140^{(a)}$	$-0.123^{(a)}$
	(0.046)	(0.012)	(0.016)	(0.017)
Disaster hits one supplier $(t+2)$	$-0.144^{(a)}$	$-0.214^{(a)}$	$-0.215^{(a)}$	$-0.204^{(a)}$
	(0.053)	(0.014)	(0.018)	(0.020)
Disaster hits one supplier $(t+3)$	$-0.117^{(b)}$	$-0.236^{(a)}$	$-0.224^{(a)}$	$-0.238^{(a)}$
	(0.059)	(0.016)	(0.020)	(0.022)
Disaster hits one supplier $(t+4)$	$-0.164^{(a)}$	$-0.255^{(a)}$	$-0.234^{(a)}$	$-0.269^{(a)}$
	(0.058)	(0.015)	(0.019)	(0.021)
Disaster hits one supplier $(t+5)$	$-0.097^{(c)}$	$-0.198^{(a)}$	$-0.179^{(a)}$	$-0.209^{(a)}$
	(0.054)	(0.014)	(0.019)	(0.020)
Disaster hits one supplier $(t+6)$	-0.031	$-0.109^{(a)}$	$-0.103^{(a)}$	$-0.108^{(a)}$
	(0.056)	(0.015)	(0.019)	(0.021)
Disaster hits one supplier $(t+7)$	-0.004	$-0.078^{(a)}$	$-0.074^{(a)}$	$-0.076^{(a)}$
	(0.056)	(0.015)	(0.019)	(0.021)
Disaster hits one supplier $(t+8)$	$0.109^{(c)}$	-0.008	-0.002	-0.002
	(0.057)	(0.015)	(0.019)	(0.021)
Buyer FE	Yes	Yes	Yes	Yes
Product × Time FE	Yes	Yes	Yes	Yes
Observations	3,092,779		3,092,779	
R^2	0.454		0.454	

Notes: This table reports event-study estimates of the impact of extreme natural disasters on U.S. importers' total import growth, disaggregated by relationship type and product characteristics. Columns (1)–(2) split the sample by whether the importer-supplier pair involves intra-firm (related-party) trade; columns (3)–(4) split by whether the HS-6 product is classified as differentiated under the Rauch (1999) conservative classification. All dependent variables are midpoint year-on-year growth rates, smoothed using a two-quarter moving average. Standard errors clustered at the buyer-by-exporter level. $^{(a)}$ 1%; $^{(b)}$ 5%; $^{(c)}$ 10%.

faster is consistent with a growing literature documenting the resilience advantages of vertical integration in the face of supply-chain disruptions, and suggests that multinational networks facilitate adjustment not only through geographic diversification but also through internal reallocation across affiliates. Pre-trends in all interaction specifications are close to zero and statistically insignificant, indicating that the differential responses are not driven by pre-existing divergences across importer types.

5 Robustness

In this section, we assess the robustness of our findings along three dimensions. First, we address the well-known limitations of two-way fixed effects estimators under staggered adoption by replicating our baseline using alternative estimators designed to accommodate heterogeneous treatment effects. Second, we implement a placebo test on never-treated units to verify that the baseline estimates reflect the causal effect of the disaster rather than a spurious correlation with its timing. Third, we examine whether any single disaster category drives our results by re-estimating the baseline specification with one disaster type excluded at a time. Our main conclusions remain stable throughout.

5.1 Robustness to Alternative Estimators under Staggered Exposure

To assess the robustness of our results to the known limitations of two-way fixed effects (TWFE) estimators in staggered adoption settings, we replicate our baseline specifications using recently proposed methods that accommodate heterogeneous treatment effects. As highlighted by Goodman-Bacon (2021), the standard TWFE approach aggregates effects across multiple treatment timings by constructing weighted averages of pairwise comparisons between earlier- and later-treated units, which can yield biased estimates when treatment effects vary across cohorts or over time. This concern is particularly relevant in our context, where both the timing of exposure and the severity of disruption vary across disaster events.

We therefore re-estimate our event-study specifications using the approaches developed by Borusyak et al. (2021), and Callaway and Sant'Anna (2021). These estimators restrict comparisons to not-yet-treated or never-treated units, and are designed to recover consistent average treatment effects in the presence of heterogeneous dynamics.

Figure 7 plots the estimated event-time coefficients across methods. As shown, the patterns are broadly consistent across specifications. While the alternative estimators tend to produce wider confidence intervals, they closely track the TWFE estimates both before and after the event, supporting the robustness of our findings.

Figure 7: Robustness to Staggered Treatment Timing: Disruption to U.S. Importers

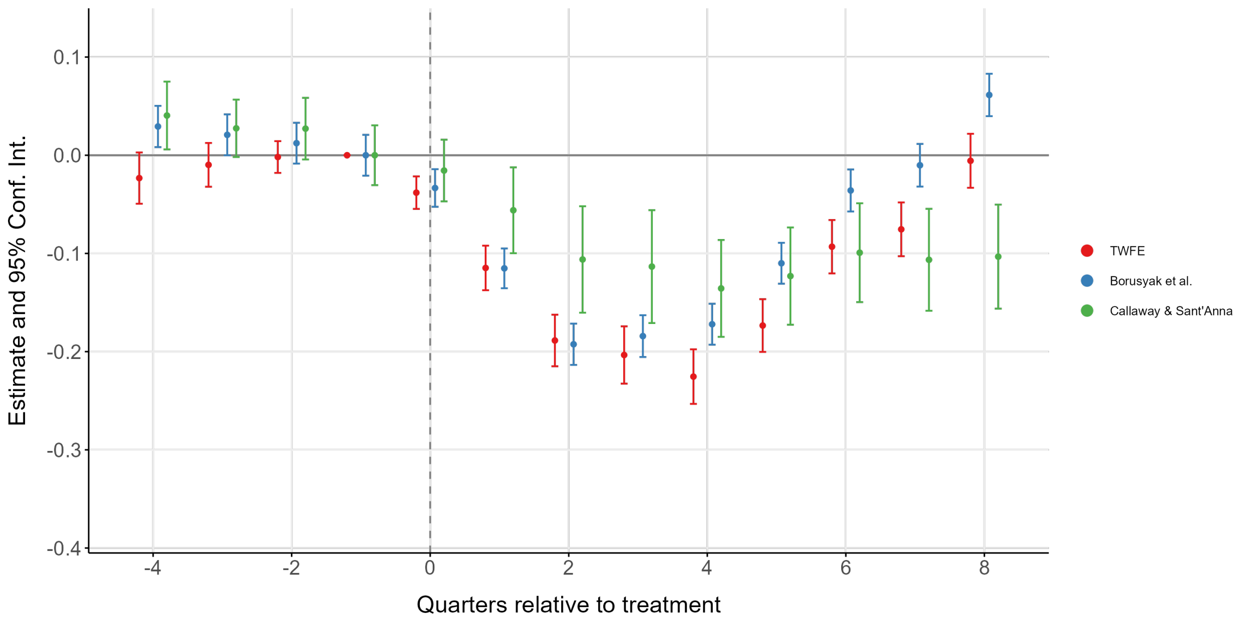

Notes: This figure compares event-study estimates of the effect of extreme natural disasters on U.S. importers exposed to affected foreign suppliers using three estimators that account for heterogeneous treatment effects under staggered adoption: the imputation estimator of Borusyak et al. (2021), and the doubly-robust estimator of Callaway and Sant'Anna (2021). The dependent variable is the midpoint year-on-year growth rate of total imports, smoothed with a two-quarter moving average. Treatment is defined as the first quarter in which any qualifying supplier is hit by a disaster at the ADM1 level. Dots report point estimates; vertical bars denote 95% confidence intervals. The dashed vertical line marks the onset of the disaster. All regressions include buyer × exporter, quarter, and ADM1 × time fixed effects. Standard errors are clustered at the buyer-by-exporter level.

5.2 Placebo Test on Never-Treated Units

To assess whether our baseline estimates reflect the causal effect of the disaster rather than a concurrent shock correlated with its timing and geography, we implement a placebo test on never-treated units. We restrict the sample to units with a treatment indicator equal to zero throughout the panel. For each remaining unit, we draw a pseudo-treatment quarter from the empirical distribution of actual disaster onsets, truncated to the quarters in which the unit is observed with a valid pre- and post-treatment window. We then randomly label a subset of these units as placebo-treated and use the remaining units as controls. On this constructed sample, we estimate the same event-study specification as in the baseline, with identical fixed effects, controls, clustering, and event window.

Under the null that the baseline estimates capture a genuine causal effect rather than a confounder, the placebo coefficients should be statistically indistinguishable from zero across all leads and lags; any systematic pattern would indicate a mechanical component in the baseline results. Figure 8 confirms this prediction: the placebo estimates are tightly centered around zero with no discernible dynamics, ruling out that the baseline patterns are an artifact of the research design. Random draws are seeded for reproducibility.

Figure 8: Placebo Test on Never-Treated Units: Total Imports

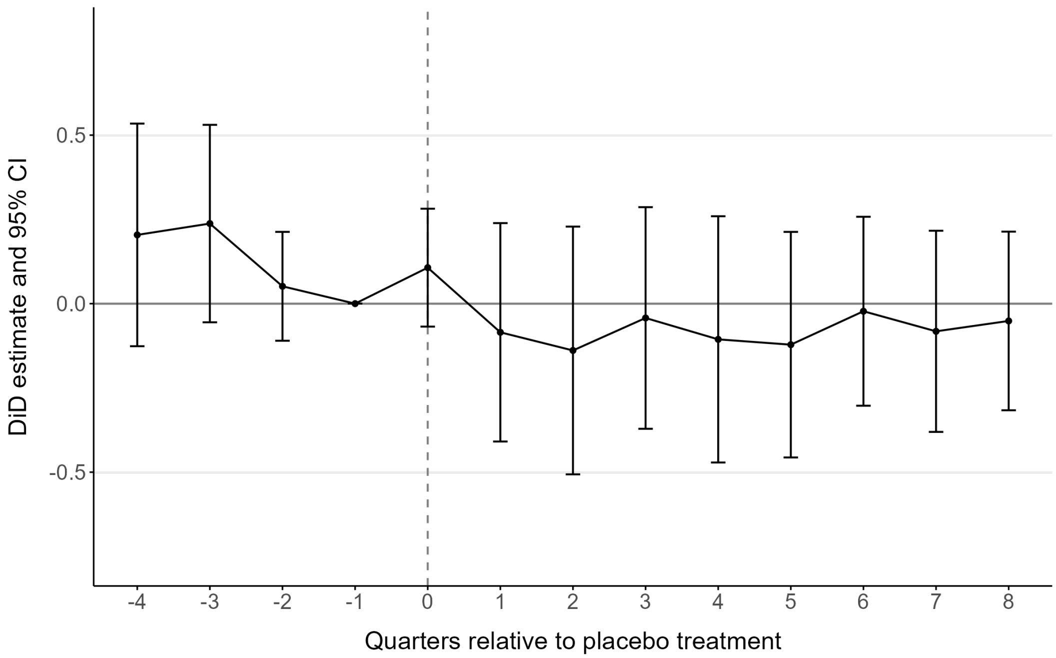

Notes: This figure reports event-study estimates from a placebo exercise in which pseudo-treatment is randomly assigned to units that are never treated in the actual sample. For each unit, a placebo disaster onset is drawn from the empirical distribution of actual disaster timings, conditional on the unit being observed with a valid pre- and post-treatment window. The dependent variable is the midpoint year-on-year growth rate of total imports, smoothed with a two-quarter moving average. The specification mirrors the baseline, with identical fixed effects, controls, clustering, and event window. Dots report point estimates; vertical bars denote 95% confidence intervals. The dashed vertical line marks the onset of the placebo. Random draws are seeded for reproducibility.

5.3 Sensitivity to Individual Disaster Events

To further probe the credibility of our results, we implement a series of additional robustness checks. First, we assess whether specific disaster types drive our findings by sequentially excluding one disaster category at a time—floods, earthquakes, epidemics, extreme temperatures, and storms—and re-estimating the main event-study specification. This leave-one-out approach verifies that no single type of shock disproportionately influences our results. As shown in Figure 9, the estimates are remarkably stable across all five permutations, both in magnitude and dynamics, lending confidence that our findings reflect a general response to supply disruptions rather than the particular features of any one disaster type.

Figure 9: Leave-One-Out by Disaster Type—Total Imports

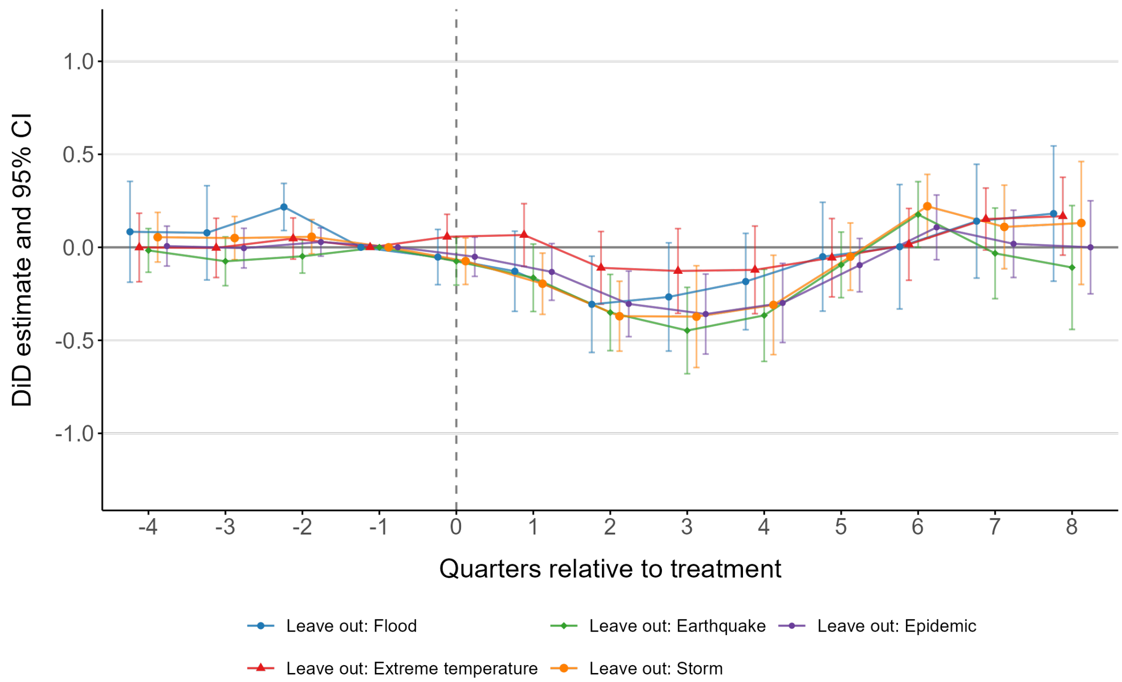

Notes: This figure reports event-study estimates from a leave-one-out exercise in which one disaster category is excluded at a time. Each line corresponds to a separate re-estimation of the baseline importer specification, omitting in turn: floods, earthquakes, epidemics, extreme temperatures, and storms. The dependent variable is the midpoint year-on-year growth rate of total import value. The specification follows the second stage described above, with importer and product-by-calendar-quarter fixed effects and standard errors clustered at the importer level. Vertical bars denote 95% confidence intervals.

6 Conclusions

This paper studies how natural disasters affecting foreign suppliers propagate through global production networks. Integrating U.S. Bill of Lading microdata, geocoded disaster records, and cross-border ownership linkages, we trace these shocks from their origin through firm-to-firm trade relationships and into the operations of downstream buyers.

We find that extreme natural disasters generate large and persistent disruptions. Exporters in affected regions experience a sharp decline in shipments to the United States, accompanied by a sustained drop in transaction volume over several years. On the buyer side, U.S. importers exposed to affected suppliers experience substantial declines in total import value — reaching up to 25 percentage points at the trough — with limited recovery over two years. The contraction extends beyond directly hit partners: imports from suppliers never exposed to a disaster also fall, and importers reduce the number of active suppliers. Among firms that also export, we observe a decline in their own outward shipments, indicating that the shock propagates not just along the supply chain but across the buyer's broader operations.

The magnitude and persistence of these effects vary with organizational structure. Arm's-length buyers experience a decline roughly twice as large as that of importers linked to their suppliers through common ownership, and a similar gap separates arm's-length from intra-firm trade. Multinational importers exhibit smaller declines and faster recoveries, particularly when the disrupted relationship is within the same corporate group, pointing to the role of internal coordination and redeployment across affiliates. These patterns suggest that the organizational boundaries of firms—not just the geographic spread of their operations—play a central role in shaping supply chain resilience.

Our findings highlight that global production networks simultaneously enable access to diverse foreign suppliers and serve as conduits for the transmission of shocks. The heterogeneity we document implies that firms' ability to withstand disruptions depends critically on how their supply relationships are structured and embedded within corporate networks—a dimension that deserves further attention in both firm strategy and trade policy.

References

Antràs, P. (2022): "Global Sourcing and Firm Organization," *Annual Review of Economics*, 14, 137–168.

Atalay, E., A. Hortacsu, and C. Syverson (2014): "Vertical integration and input flows," *American Economic Review*, 104, 1120–1148.

Balboni, C. et al. (2023): "Extreme Weather and Firm Performance in Pakistan," *Working Paper.*

Barrot, J.-N. and J. Sauvagnat (2016): "Input specificity and the propagation of idiosyncratic shocks in production networks," *The Quarterly Journal of Economics*, 131, 1543–1592.

Bernard, A. B., J. B. Jensen, S. J. Redding, and P. K. Schott (2009): "The Margins of U.S. Trade," *American Economic Review*, 99, 487–493.

Blaum, J., F. Esposito, and S. Heise (2023): "Input Sourcing under Supply Chain Risk," *Working Paper.*

——— (2025): "Input Sourcing Under Supply Chain Risk: Evidence from U.S. Manufacturing Firms," Tech. Rep. 1141, Federal Reserve Bank of New York Staff Reports.

Boehm, C. E., A. Flaaen, and N. Pandalai-Nayar (2019): "Input linkages and the transmission of shocks: Firm-level evidence from the 2011 Tōhoku earthquake," *Review of Economics and Statistics*, 101, 60–75.

Borusyak, K., X. Jaravel, and J. Spiess (2021): "Revisiting Event Study Designs," *arXiv preprint arXiv:2108.12419.*

Callaway, B. and P. H. C. Sant'Anna (2021): "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics.*

Carvalho, V. M., M. Nirei, Y. U. Saito, and A. Tahbaz-Salehi (2021): "Supply Chain Disruptions: Evidence from the Great East Japan Earthquake," *The Quarterly Journal of Economics*, 136, 1255–1321.

Castro-Vincenzi, A. (2022): "Climate shocks and supply chains," *Working Paper.*

Conconi, P. et al. (2022): "Multinational firms and the international transmission of shocks: Micro evidence from the Covid-19 crisis," *Working Paper.*

Cravino, J. and A. A. Levchenko (2017): "Multinational firms and international business cycle transmission," *The Quarterly Journal of Economics*, 132, 921–962.

Fort, T. C. et al. (2023): "Changing supply chains: Firm responses to the COVID-19 pandemic," *Working Paper.*

Freund, C. et al. (2022): "Natural Disasters and the Global Supply Chain," *Working Paper.*

Goodman-Bacon, A. (2021): "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*.

Heise, S. (2023): "Firm-to-Firm Relationships and the Pass-Through of Shocks: Theory and Evidence," *Review of Economics and Statistics*.

Heise, S., J. R. Pierce, G. Schaur, and P. K. Schott (2025): "How Do Firms in Different Sectors Organize Their Supply Chains? Evidence from Transaction-Level Import Data," *AEA Papers and Proceedings*, 115, 177–181.

Huneeus, F. (2023): "Production Network Dynamics and the Propagation of Shocks," *Working Paper*.

Irarrazabal, A., A. Moxnes, and L. D. Opromolla (2013): "The margins of multinational production and the role of intrafirm trade," *Journal of Political Economy*, 121, 74–126.

Khanna, G. et al. (2022): "Supply Chain Resilience: Evidence from Indian Manufacturing," *Working Paper*.

Kikkawa, K. et al. (2023): "Shocks and the Structure of Global Value Chains," *Working Paper*.

Lee, S. and J. Han (2022): "Global Supply Chain Shocks," *Working Paper*.

Levchenko, A. A., L. T. Lewis, and L. L. Tesar (2010): "The Collapse of International Trade during the 2008–09 Crisis: In Search of the Smoking Gun," *IMF Economic Review*, 58, 214–253.

Li, X. (2021): "The effect of input specificity on firm boundaries: Evidence from buyer–supplier relationships," *Journal of International Economics*, 130, 103444.

Méjean, I. et al. (2023): "Networks, Shocks, and Firm Responses: Evidence from International Trade," *Working Paper*.

Ramondo, N., V. Rappoport, and K. J. Ruhl (2016): "Intrafirm trade and vertical fragmentation in US multinational corporations," *Review of Economics and Statistics*, 98, 701–714.

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

Online Appendix

"How Shocks Travel: The Cross-Border Impact of Natural Disasters in Firm Networks"

Vanessa Alviarez, Brian C. Fujiy, Angel Espinoza E., Tomasz Swiecki

This Appendix provides supplementary information on the construction of the firm-to-firm Bill of Lading (BoL) dataset and the empirical procedures used in the main analysis. We begin by documenting the processing of Panjiva's raw shipment-level records, including product code harmonization, firm identification, and the estimation of shipment values. We then describe the geocoding of shipper and consignee addresses, which enables spatial analyses of trade linkages, and we outline the quality checks implemented to ensure data consistency. Additional tables and figures illustrate the robustness of our procedures and support the results reported in the main text.

A Data Construction

A.1 Construction of Firm-to-Firm BoL Data

The raw Panjiva dataset contains over 155 million shipment-level records of U.S. maritime imports since 2007, sourced from U.S. Customs and Border Protection (CBP). Each Bill of Lading (BoL) reports the shipment date, port of arrival, shipper and consignee names and addresses, product descriptions, shipment weight and quantity, and carrier information. We process these transaction records to construct a panel of firm-to-firm trade flows suitable for analysis.

Data harmonization and firm identifiers. Panjiva assigns six-digit Harmonized System (HS6) product codes to BoL product descriptions using a proprietary text-processing algorithm. These codes are harmonized to U.S. Census HS10 classifications to recover shipment values. Panjiva also provides numeric identifiers for foreign shippers and U.S. consignees, but the same legal entity may appear under multiple IDs due to minor name or address differences. To create temporally consistent firm identifiers, we standardize company names and addresses, geocode all addresses to detect duplicates within a 100-meter radius, and reconcile firm IDs over time to ensure longitudinal consistency. Where possible, we link Panjiva firms to Capital IQ using S&P Global's CUSIP/CIK crosswalk, allowing integration with Compustat.

Estimating shipment values. Because Panjiva does not report shipment values, we impute them by combining BoL quantities with official U.S. Census unit values. Weighted average unit values are computed at the HS6–origin-country–year level and multiplied by reported quantities to generate transaction-level value proxies. We drop shipments listing multiple HS6 codes but a single quantity or weight, as they cannot be allocated reliably. Aggregating these proxies yields a panel of firm-to-firm U.S. import flows with consistent value measures.

A.1.1 Geocoding and spatial linking

To measure firms' exposure to local shocks, we geocode all shipper and consignee addresses. After cleaning and standardizing addresses (removing extraneous characters and PO boxes), we obtain latitude and longitude coordinates using Google and OpenStreetMap APIs. Ambiguous results are manually reviewed, and about 2–3% of addresses lacking reliable coordinates are excluded from spatial analyses. This procedure produces precise locations for over 470,000 U.S. importers and nearly one million foreign exporters, enabling subnational mapping of trade networks.

Data quality checks. We validate the resulting panel by verifying that shipment counts by HS6–country–year align with U.S. Census totals, cross-checking geocoded firm locations with port-of-entry information, and ensuring that longitudinal firm IDs do not generate spurious entry or exit due to misspellings or address noise. These steps produce a clean, high-quality dataset suitable for linking with other economic and environmental datasets.

A.2 Geocoding firm addresses

To measure the distance between firms and the locations of natural disasters, we geocode the physical addresses of all U.S. importers and their corresponding foreign partners, converting them into geographic coordinates (latitude and longitude). To improve consistency, we first harmonize the addresses of firms with the same Panjiva ID but different recorded spellings. This process yields geocodes for 470,435 unique U.S. importer addresses and 995,568 unique foreign exporter addresses. [11]

[11] The full dataset includes 2,586,903 unique addresses for U.S. importers and 2,782,417 for foreign exporters. We prioritize geocoding a subset of firms based on three criteria: (1) U.S. importers that are also

We geocode each firm using the most complete address available in the Panjiva Bills of Lading (BoL) data. Our preferred address format is constructed by concatenating structured components of the firm's physical address—including street, city, region/state, postal code, and country. When only region/state and country (or country alone) are available, we rely on the unparsed full address reported in the BoL. We favor the parsed-and-concatenated version whenever possible, as it is typically more structured and thus easier to geocode using API services. In contrast, full addresses as recorded in the BoL often lack punctuation or consistent formatting, which can hinder parsing and lower geocoding precision. [12]

Table A.1 summarizes the address types used in the geocoding process. For 59% of U.S. importers, geocodes are based on concatenated fields that include all address components. An additional 4.6% are based on addresses missing only the street component. For 25.5%, we rely on the full unstructured address as reported in the BoL. In contrast, for foreign exporters, the full BoL address is used in approximately 60% of cases, reflecting more limited standardization in the original data.

Table A.1: Distribution of Geocoded Firms by Address Type (%)

	U.S. Importers	Foreign Partners	Total
Full address	25.5	37.3	33.5
Route/City/Region/Postal Code/Country	59.0	21.1	33.3
City/Region/Postal Code/Country	4.6	7.7	6.7
City/Region/Country	3.5	11.7	9.1
Route/City/Region/Country	2.7	6.6	5.3
Route/Region/Postal Code/Country	1.8	3.2	2.8
Other	2.9	12.4	9.3

Notes: This table shows the information from the firm's physical address used in the geocode process for U.S. importers and foreign exporters. The first row shows the percentage of addresses that were geocoded using the full or complete address directly reported by the Panjiva BoL. The next five rows show the percentage of geocoded addresses that were constructed by concatenating the individual components of the firm's physical address, including route, city, region/state, postal code, and country. The last row shows the percentage of firms that use a different combination of route, city, region/state, postal code, and country, not listed in the previous rows.

We geocode firm addresses using the Geopy Python library, a client for several major geocoding web services. Specifically, we employ the Bing Maps Location API to retrieve latitude

exporters, (2) U.S. importers that appear in both Panjiva and Compustat, and (3) U.S. importers with more than one foreign supplier. After applying these filters, we geocode 470,435 U.S. addresses and 995,568 foreign addresses.

[12] Addresses constructed from individual components tend to be of higher quality and are less prone to input errors.

and longitude coordinates for each address. The API provides several metadata fields that allow us to assess the quality of the geocoding output: (1) confidence, (2) match code, (3) inland, and (4) country. A geocode is classified as having medium or low confidence when only a subset of address components is matched (e.g., when only the postal code matches). The match code distinguishes between good, ambiguous, and up-hierarchy matches, depending on the uniqueness and precision of the returned location.

To further evaluate reliability, we construct two additional indicators. First, we generate an inland dummy equal to one if the coordinates fall on land—based on satellite data from the International Space Station—and zero otherwise. Second, we create a dummy that compares the country reported by the firm to the country returned by Bing's geocoder, allowing us to flag inconsistencies.

After obtaining geocodes, we match firms to administrative units using shapefiles from the Global Administrative Unit Layers (GAUL) maintained by the Food and Agriculture Organization of the United Nations (FAO). These shapefiles cover three administrative levels: national, provincial, and district. [13]

B Additional Results

B.1 Direct Impact on Foreign Suppliers by MNC Status

This appendix reports first-stage estimates disaggregated by the multinational status of the foreign supplier. Specifically, we re-estimate the supplier-level event study separately for exporters that belong to a multinational group and those that do not. This decomposition allows us to assess whether the direct impact of a natural disaster on foreign suppliers' export performance differs across organizational structures, providing context for the heterogeneous propagation patterns documented in the main text.

[13] Shapefiles were obtained by contacting `GeoNetwork@fao.org` . We thank Nelson Rosas Ribeiro Filho for sharing the maps. Minor adjustments were made, including the computation of polygon centroids.

Figure B.1: Effects of Natural Disasters on Foreign Exporters by MNC Status

(a) Export value to the U.S. (YoY midpoint growth)

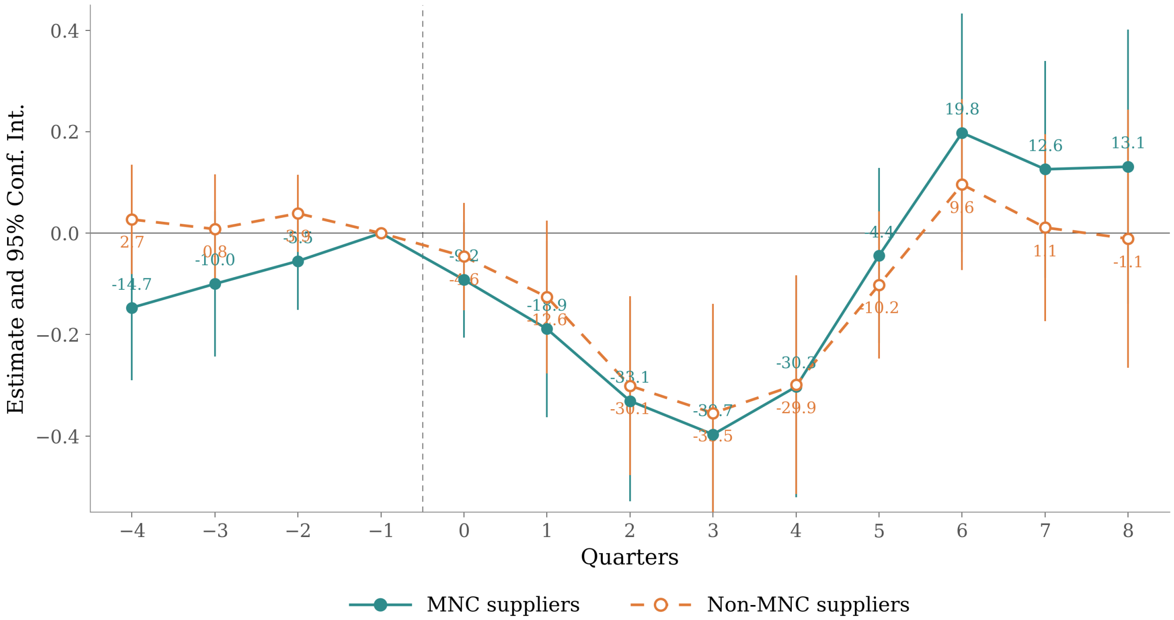

(b) Number of transactions to the U.S. (YoY midpoint growth)

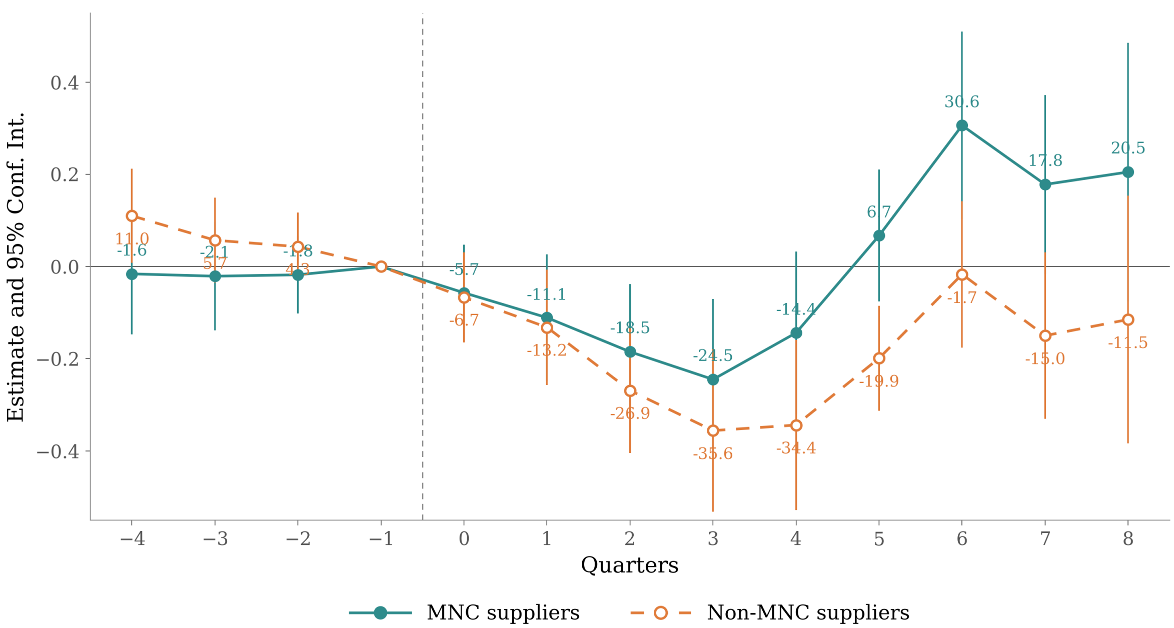

Notes: This figure plots the estimated effects of extreme natural disasters on directly affected foreign exporters, disaggregated by multinational status. Each panel reports coefficients from an event-study regression using quarterly year-on-year midpoint growth as the outcome. Panel A shows total export value to the U.S.; Panel B shows the number of transactions. The sample includes recurrent exporters, and treatment is defined at the ADM2 level based on the first qualifying disaster. *Multinational* is defined based on Orbis ownership links. All regressions include exporter fixed effects, product-by-calendar-quarter fixed effects, and country-by-quarter fixed effects. Standard errors are clustered at the ADM2 level.

B.2 Propagation to U.S. Importers: Effect on Sales Abroad

This appendix presents supplementary results on the downstream propagation of foreign supplier shocks to U.S. importers. We report additional outcomes beyond total import value, including the number of transactions and active suppliers, to characterize the breadth of the disruption across the buyer's sourcing margin. We also examine whether the effects extend to importers' own export activity, testing whether supply-side shocks propagate beyond the import channel and into the buyer's operations abroad.

To complement the heterogeneity analysis in the main text, we replicate the interaction specifications for imports sourced exclusively from suppliers never directly exposed to a disaster. This exercise isolates the indirect component of the propagation—the extent to which buyers adjust their purchases from unaffected partners in response to a shock elsewhere in their supplier network. We further report the corresponding heterogeneity results for exports, disaggregated by multinational status and intra-firm linkages. Tables and figures are referenced in the main text.

Figure B.2: Transactions Growth: Supplier Disruptions and Downstream Propagation

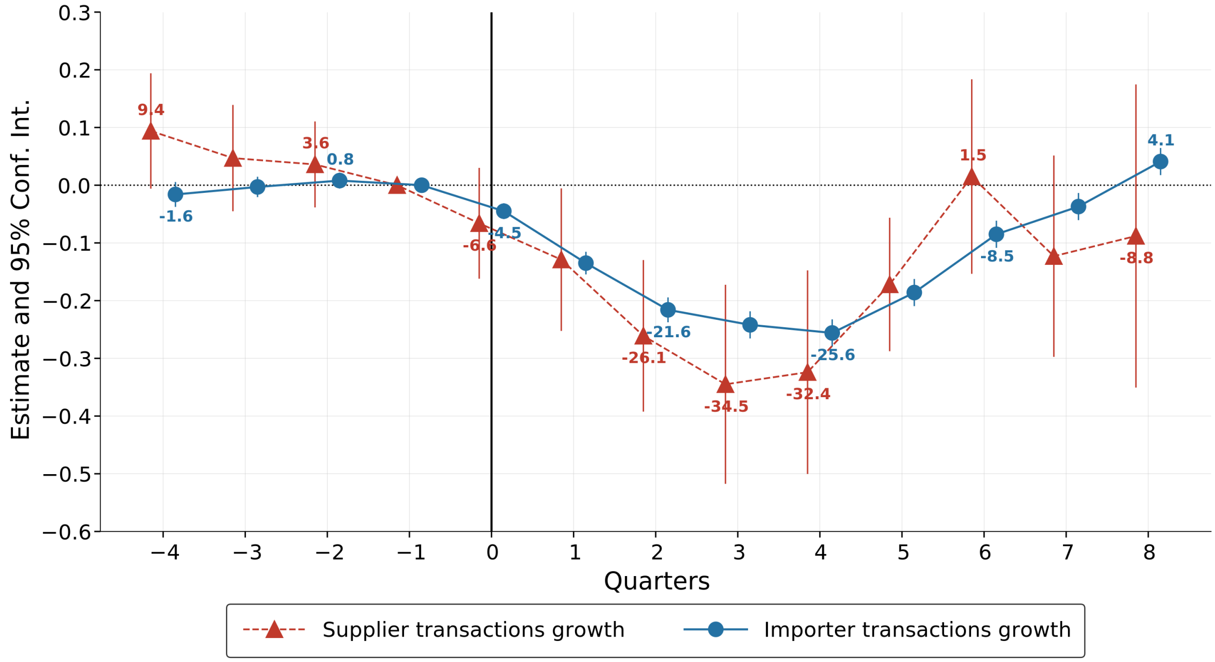

Notes: This figure overlays event-study coefficients from two specifications estimated on directly exposed trade links. The red dashed line plots $\hat{\beta}_k$ from the first-stage supplier regression:

$$X_{st} = \alpha_s + \delta_{p(s)t} + \gamma_{c(s)t} + \sum_{k=-K}^{L} \beta_k \cdot \mathbf{1}\Big\{\tau_{st} = k\Big\} + \varepsilon_{st}$$

where X_{st} is the midpoint year-on-year growth rate of the number of transactions for supplier s at calendar quarter t, and τ_{st} denotes the number of quarters elapsed since the first extreme natural disaster affecting supplier s's ADM2 district. The blue solid line plots $\hat{\phi}_k$ from the second-stage importer regression:

$$M_{bt} = \alpha_b + \delta_{p(b)t} + \theta\, \mathbf{1}\{\text{directly hit}_{bt}\} + \sum_{k=-K}^{L} \phi_k \cdot \mathbf{1}\{\tau_{bt} = k\} + X_{bt} + \eta_{bt},$$

where M_{bt} is the midpoint year-on-year growth rate of the number of transactions for importer b at quarter t, τ_{bt} denotes quarters elapsed since the first quarter in which any supplier in buyer b's import basket is hit by an extreme natural disaster, as defined in Section 3, and X_{bt} controls for whether the affected country is also an important export destination for the importer. The reference period is $t-1$. The first-stage sample comprises recurrent exporters with exporter, product-by-calendar-quarter, and country-by-quarter fixed effects; standard errors are clustered at the ADM2-quarter level. The second-stage sample comprises recurrent importers with importer and product-by-calendar-quarter fixed effects; standard errors are clustered at the importer level. Vertical bars denote 95% confidence intervals.

TABLE B.1: Effects of Natural Disasters on U.S. Importers: Additional Outcomes

	Transactions	Imports Affected	No. of Active Suppliers
Disaster hits one supplier $(t-4)$	-0.016	$-0.206^{(a)}$	$-0.020^{(b)}$
	(0.011)	(0.023)	(0.010)
Disaster hits one supplier $(t-3)$	-0.003	0.002	-0.008
	(0.009)	(0.019)	(0.009)
Disaster hits one supplier $(t-2)$	0.008	$0.130^{(a)}$	0.003
	(0.007)	(0.013)	(0.006)
Reference period: $t-1$			
Disaster hits one supplier (t)	$-0.045^{(a)}$	$-0.168^{(a)}$	$-0.032^{(a)}$
	(0.007)	(0.014)	(0.006)
Disaster hits one supplier $(t+1)$	$-0.135^{(a)}$	$-0.708^{(a)}$	$-0.116^{(a)}$
	(0.010)	(0.019)	(0.009)
Disaster hits one supplier $(t+2)$	$-0.216^{(a)}$	$-1.250^{(a)}$	$-0.205^{(a)}$
	(0.011)	(0.020)	(0.010)
Disaster hits one supplier $(t+3)$	$-0.242^{(a)}$	$-1.420^{(a)}$	$-0.236^{(a)}$
	(0.012)	(0.023)	(0.011)
Disaster hits one supplier $(t+4)$	$-0.256^{(a)}$	$-1.630^{(a)}$	$-0.261^{(a)}$
	(0.012)	(0.022)	(0.010)
Disaster hits one supplier $(t+5)$	$-0.186^{(a)}$	$-1.460^{(a)}$	$-0.196^{(a)}$
	(0.012)	(0.022)	(0.010)
Disaster hits one supplier $(t+6)$	$-0.085^{(a)}$	$-1.150^{(a)}$	$-0.088^{(a)}$
	(0.012)	(0.024)	(0.010)
Disaster hits one supplier $(t+7)$	$-0.037^{(a)}$	$-1.110^{(a)}$	$-0.034^{(a)}$
	(0.012)	(0.024)	(0.011)
Disaster hits one supplier $(t+8)$	$0.041^{(a)}$	$-0.993^{(a)}$	$0.044^{(a)}$
	(0.012)	(0.025)	(0.011)
Importer FE	Yes	Yes	Yes
Product × Time FE	Yes	Yes	Yes
Observations	3,468,296	3,054,393	3,468,296
R^2	0.490	0.454	0.512

Notes: This table reports event-study estimates of the impact of extreme natural disasters on U.S. importers, complementing Table 3. All dependent variables are midpoint year-on-year growth rates, smoothed using a two-quarter moving average. Transactions refers to the total number of shipments received by the importer. Affected Imports refers to imports from suppliers directly hit by a disaster. No. Suppliers is the number of active suppliers. Estimates are based on the following specification:

$$M_{bt} = \alpha_b + \delta_{p(b)t} + \theta\,\mathbf{1}\{\text{directly hit}_{bt}\} + \sum_{k=-K}^{L} \phi_k \cdot \mathbf{1}\{\tau_{bt} = k\} + X_{bt} + \eta_{bt},$$

where M_{bt} denotes the midpoint year-on-year growth rate of the outcome of interest for importer b in calendar-quarter t, smoothed using a two-quarter moving average, and τ_{bt} is defined as in Section 3. The reference period is $t-1$. Standard errors are clustered at the importer level. Significance: $^{(a)}$ 1%; $^{(b)}$ 5%; $^{(c)}$ 10%.

Table B.2: Heterogeneous Effects on U.S. Imports from Never-Affected Suppliers

	Intra-Firm Trade		Product Differentiation	
	Intra	Arm's-L.	Diff.	Non-Diff.
Disaster hits one supplier ($t-4$)	-0.102	$0.051^{(a)}$	$0.064^{(a)}$	0.018
	(0.066)	(0.016)	(0.021)	(0.024)
Disaster hits one supplier ($t-3$)	-0.018	$0.033^{(b)}$	$0.041^{(b)}$	0.017
	(0.055)	(0.014)	(0.018)	(0.020)
Disaster hits one supplier ($t-2$)	-0.017	0.014	$0.024^{(c)}$	-0.003
	(0.042)	(0.010)	(0.013)	(0.015)
Reference period: $t-1$				
Disaster hits one supplier (t)	-0.058	$-0.039^{(a)}$	$-0.040^{(a)}$	$-0.040^{(a)}$
	(0.040)	(0.010)	(0.014)	(0.015)
Disaster hits one supplier ($t+1$)	-0.080	$-0.069^{(a)}$	$-0.058^{(a)}$	$-0.083^{(a)}$
	(0.054)	(0.014)	(0.018)	(0.020)
Disaster hits one supplier ($t+2$)	$-0.139^{(b)}$	$-0.104^{(a)}$	$-0.102^{(a)}$	$-0.110^{(a)}$
	(0.065)	(0.016)	(0.022)	(0.024)
Disaster hits one supplier ($t+3$)	-0.120	$-0.115^{(a)}$	$-0.103^{(a)}$	$-0.129^{(a)}$
	(0.076)	(0.018)	(0.024)	(0.026)
Disaster hits one supplier ($t+4$)	$-0.133^{(c)}$	$-0.132^{(a)}$	$-0.111^{(a)}$	$-0.158^{(a)}$
	(0.072)	(0.017)	(0.022)	(0.025)
Disaster hits one supplier ($t+5$)	-0.093	$-0.116^{(a)}$	$-0.105^{(a)}$	$-0.127^{(a)}$
	(0.069)	(0.016)	(0.021)	(0.024)
Disaster hits one supplier ($t+6$)	-0.009	$-0.060^{(a)}$	$-0.047^{(b)}$	$-0.070^{(a)}$
	(0.069)	(0.017)	(0.022)	(0.024)
Disaster hits one supplier ($t+7$)	0.020	-0.023	0.001	$-0.047^{(c)}$
	(0.071)	(0.017)	(0.022)	(0.024)
Disaster hits one supplier ($t+8$)	0.080	$0.039^{(b)}$	$0.059^{(a)}$	0.019
	(0.070)	(0.017)	(0.022)	(0.024)
Buyer FE	Yes	Yes	Yes	Yes
Product × Time FE	Yes	Yes	Yes	Yes
Observations	3,016,669		3,016,669	
R^2	0.440		0.440	

Notes: This table reports event-study estimates of the impact of extreme natural disasters on U.S. importers' import growth from never-affected suppliers, disaggregated by relationship type and product characteristics. Columns (1)–(2) split the sample by whether the importer-supplier pair involves intra-firm (related-party) trade; columns (3)–(4) split by whether the HS-6 product is classified as differentiated under the Rauch (1999) conservative classification. All dependent variables are midpoint year-on-year growth rates, smoothed using a two-quarter moving average. Standard errors clustered at the buyer-by-exporter level. $^{(a)}$ 1%; $^{(b)}$ 5%; $^{(c)}$ 10%.

FIGURE B.3: Heterogeneous Effects of Natural Disasters on U.S. Importers' Exports

Panel A: By Multinational Status

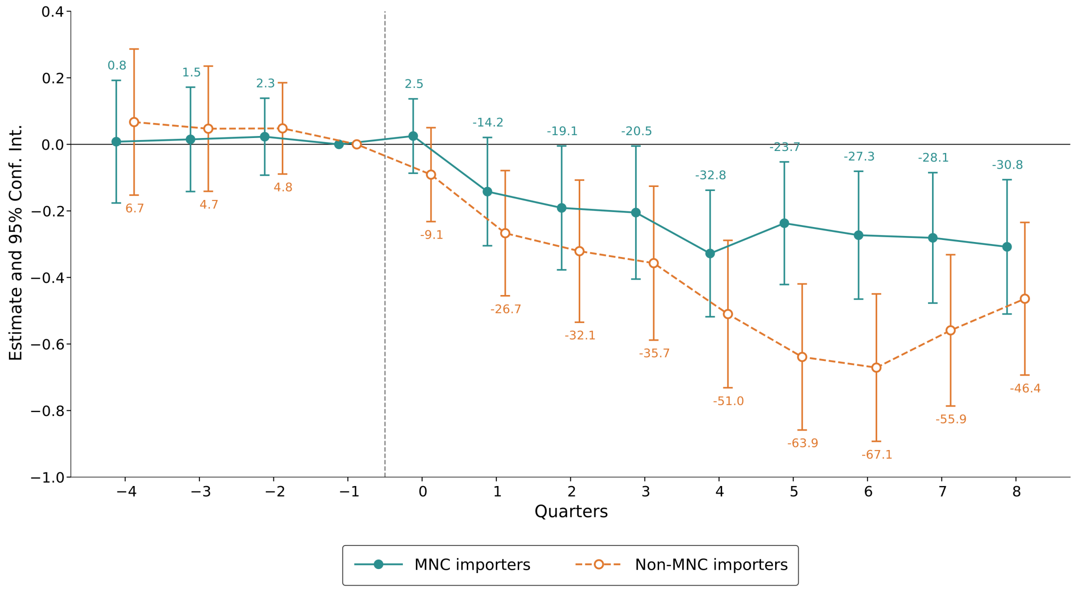

Panel B: By Intra-Firm Trade Status

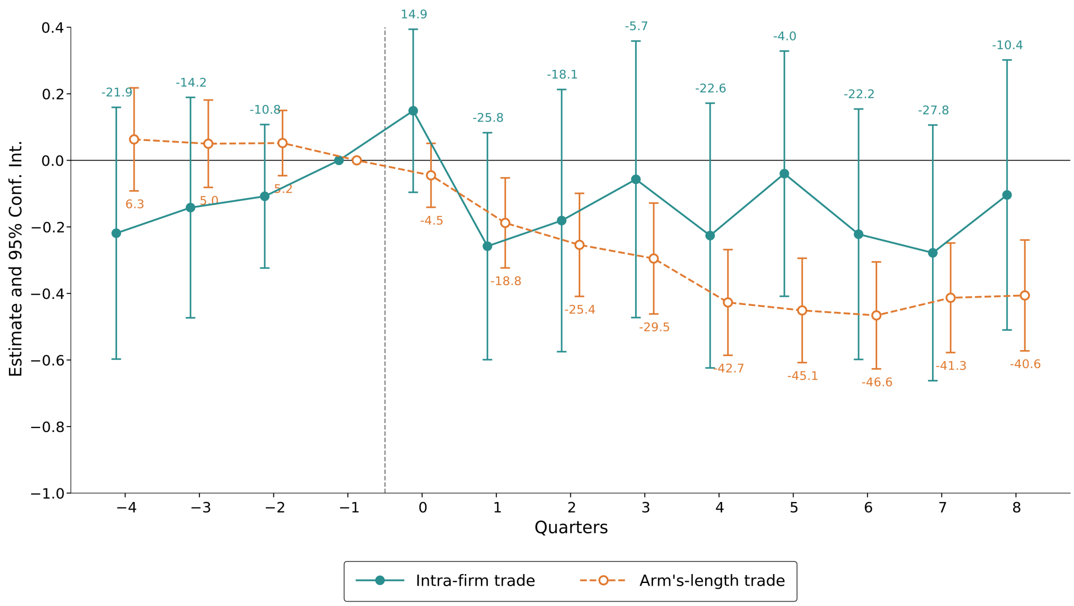

Notes: This figure overlays event-study coefficients from two interaction specifications estimated on directly exposed trade links. The dependent variable in all panels is the midpoint year-on-year growth rate of export value, smoothed using a two-quarter moving average. Panel A plots $\hat{\beta}_k$ and $\hat{\beta}_k + \hat{\phi}_k$ from:

$$M_{bt} = \alpha_b + \delta_{p(b)t} + \gamma_{ct} + \sum_{k=-K}^{L} (\beta_k + \phi_k \cdot \text{MNC}_b)\, \mathbf{1}\{\tau_{bt} = k\} + \varepsilon_{bt},$$

where MNC_b indicates whether the importer belongs to a U.S. multinational network. The teal solid line plots $\hat{\beta}_k + \hat{\phi}_k$ (MNC importers); the orange dashed line plots $\hat{\beta}_k$ (non-MNC importers). Panel B replaces the interaction with IntraFirm_{bs}, an indicator for whether the importer-supplier pair is classified as related-party trade in U.S. Customs data. All regressions include buyer and product-quarter fixed effects. Standard errors are clustered at the buyer-by-exporter level. Vertical bars denote 95% confidence intervals.