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

Vanessa Alviarez
IADB

Brian Fujiy
U.S. Census

Angel Espinoza
UBC

Tomasz Swiecki
Opendoor Labs

October, 2025

Abstract

We study how disruptions in global production networks propagate across borders through firm-to-firm linkages. Using a novel dataset that combines U.S. Bill of Lading microdata, geocoded records of global natural disasters, and cross-border ownership linkages, we trace the effects of exogenous shocks to foreign input suppliers on U.S.-based firms. Our empirical strategy is an event study with staggered exposure, exploiting the exogeneity of natural disasters to estimate causal impacts on trade flows and firm performance. We first document that natural disasters abroad lead to a decline in the exports of affected foreign suppliers from which U.S. firms source intermediate inputs. Building on this first stage, we study how such shocks propagate through global value chains and how the strength of the transmission varies with the nature of the firm-to-firm relationship, particularly the presence of ownership links. This approach allows us to characterize the role of multinational networks and intra-firm linkages in shaping the cross-border transmission of supply chain disruptions.

*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 motivating fact for this paper is the persistent asymmetry in how different trade relationships respond to recent large-scale trade shocks. 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 at a larger scale. Between the peak and trough of the crisis, U.S. real GDP declined by 3.9%, while real imports fell by 18.6% (Levchenko et al., 2010). Yet even this steep contraction was uneven: imports from unaffiliated suppliers plummeted by nearly 26%, 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.

Our analysis begins with the idea that not all shocks propagate equally. A natural disaster in a supplier's location may affect its direct buyers, but the extent and persistence of the impact depends on whether (i) the traded-good is relationship-specific (Rauch, 1999), (ii) the firms involved belong to a multinational group, and (iii) the trading partner is part of the same corporate group. The role of input specificity is not obvious: it could amplify the immediate costs of disruption, as highly tailored or certified inputs are harder to replace—but it could also foster recovery. When a buyer depends on a given supplier, it has stronger incentives to support, wait for, or reintegrate that supplier after the disruption. Firm boundaries shape

Figure 1: Growth in US Imports: Total vs. Intra-Firm

Notes: The figure plots year-on-year growth rates in U.S. goods imports using data from the U.S. Bureau of Economic Analysis (BEA). The solid blue line reflects total U.S. imports of goods, while 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.

how this dependence plays out: intra-firm trade may allow for faster internal substitution or relocation, while arm's-length relationships may rely more heavily on external re-matching and may never be reestablished.

We apply this logic to the data by constructing a panel of U.S. importers linked to foreign suppliers using customs shipment records from Panjiva, matched with firm ownership data from Orbis, and disaster records from EM-DAT. Our event-study design exploits the quasi-random timing and geography of extreme natural disasters that affect the exporter's subnational location. We define importer exposure based on active sourcing ties to affected suppliers, requiring that a supplier represents at least 10% of the importer's sourcing value over the prior six quarters. We then compare the dynamic outcomes of exposed vs. unexposed firms across three margins: total imports, imports excluding affected suppliers, and exports, restricting the last margin to importers that also engage in foreign sales. These variables are measured at quarterly frequency, expressed in year-over-year midpoint growth rates, and smoothed to filter short-term volatility.

Importantly, we distinguish between arm's-length and intra-firm relationships using ownership data. We classify a buyer–supplier pair as intra-firm if both entities belong to the same multinational group, and arm's-length otherwise. This allows us to examine whether supply disruptions are transmitted differently depending on the organizational form of the relationship. We also stratify trade flows by input specificity, classifying goods based on their economic use (intermediate, consumption, or capital) and on the Rauch (1999) classification, which distinguishes differentiated products—typically harder to substitute—from those traded on organized exchanges or with reference prices. This allows us to test whether the resilience of trade relationships depends jointly on product characteristics and ownership structure.

The paper implements an event-study and staggered difference-in-differences design to estimate the dynamic effects of extreme natural disasters on international trade relationships. Treatment is defined at the firm level, based on exposure to geographically localized shocks affecting foreign suppliers. The timing of treatment varies across firms and quarters, and we allow for absorbing exposure as well as repeated events. The design flexibly identifies deviations in trade flows relative to pre-shock trends, while conditioning on high-dimensional fixed effects to absorb seasonality, macroeconomic shocks, and importer-supplier heterogeneity. We estimate responses both for directly affected exporters and for U.S. importers linked to them through active, economically significant trade relationships. By interacting exposure with product specificity and organizational structure, we test whether intra-firm and multinational links modulate the intensity and persistence of disruptions.

Our findings show that both intra-firm and relationship-specific trade relationships are more resilient to supply disruptions. These results suggest that multinational networks and relationship-specific trading contracts provide an internal insurance margin that is activated precisely when switching is difficult—offering coordination, support, and reallocation mechanisms unavailable to independent firms.

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 (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 also 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 position in the network, intensity of linkages, 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

[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.

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 and exposure definitions. Section 4 then displays key descriptive patterns—showing the prevalence of intra-firm links, input specificity distributions, and baseline sourcing behavior. Section 5 reports the main empirical findings, breaking down responses by product types and ownership forms. Section 6 presents robustness checks and alternative specifications. Section 7 concludes with a discussion of implications, caveats, and paths for further work.

2 Data

In this section, we explain the three main sources of data we use: (1) bill of lading records recording international transactions of US importers, (2) Natural disasters, (3) ownership linkages

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 since 2007, 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 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 with ambiguous product-quantity assignments are excluded to preserve measurement accuracy.

A key feature of the BoL data is the ability to track firms across time and space. We

[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).

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 C.

2.2 Natural Disasters

To quantify the severity of natural disasters, we construct an Extreme Disaster Index (EDI) using EM-DAT and SHELDUS, which provide detailed information on events in the U.S. and globally. For each event, we compute deaths per capita, affected population per capita, and economic damage as a share of regional GDP, using population and economic data at the second administrative level (ADM2). These indicators are standardized (z-scores) and combined into a composite index that captures both intensity and human impact. To avoid over-weighting small countries—where moderate events can look disproportionately large in per-capita terms—we winsorize the tails and weight components using both per-capita and absolute magnitudes. Events in the top decile of the EDI are classified as "extreme disasters" and serve as the treatment in our event-study framework.

Our baseline severity proxy uses fatalities, not reported monetary damages. Death counts are more comparable across places because monetary losses reflect local wealth, asset density, and insurance coverage; a physically similar shock can seem larger in high-income regions simply due to higher asset values. Fatalities better track underlying intensity and are less sensitive to differences in reporting. That said, deaths may not fully align with the economic channels we study—for example, a low-casualty event can still disrupt production if it hits industrial or logistics hubs. As a robustness check, we replicate the analysis using reported damages as an alternative severity measure to verify that results do not hinge on the metric and to compare how different dimensions of intensity map into cross-border supply-chain disruptions. By considering both fatalities and monetary damages as alternative severity measures—and both single- and multi-shock exposure definitions—we can assess whether the propagation of disruptions through global value chains is robust to how shock intensity and exposure are defined. This is what we chose in the baseline, but we performed robustness with alternative specifications.

Figure 2 illustrates the distribution of natural disaster events by fatality counts, classified into three percentile-based groups. Panel (a) shows low-fatality events (deaths $<$ p50), which are most frequent and highly concentrated at very low counts, with most causing

Figure 2: Distribution of natural disaster events by number of deaths

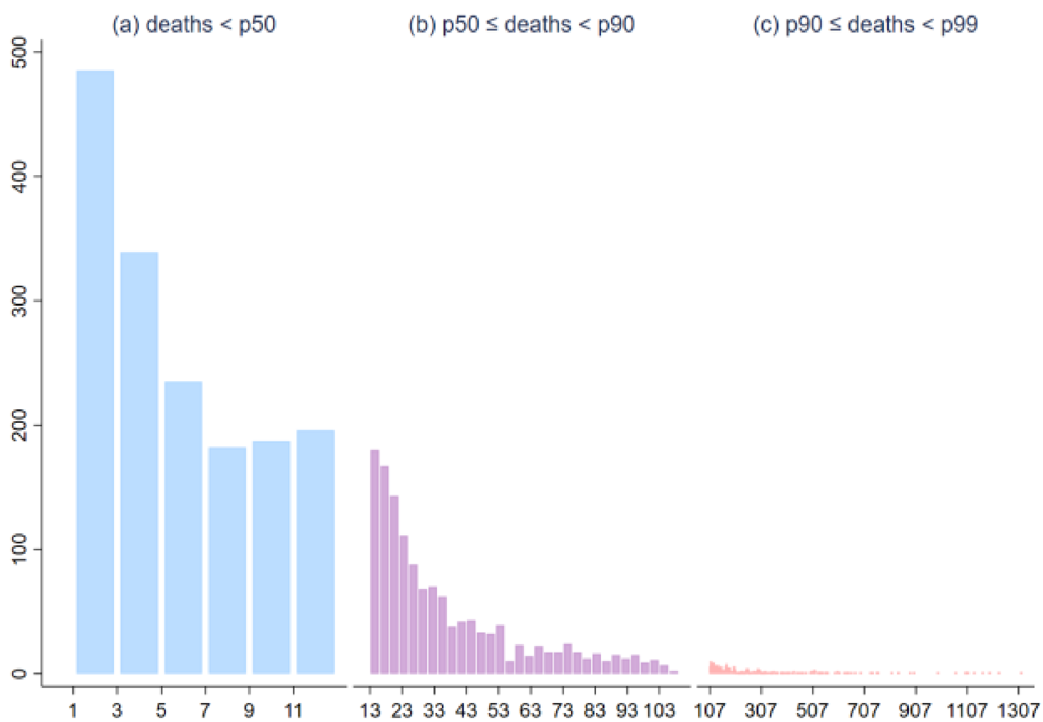

Notes: The histogram groups disaster events into three categories based on death percentiles: below the median (p50), between the 50th and 90th percentiles, and between the 90th and 99th percentiles. The x-axis is scaled to accommodate the highly skewed distribution of fatalities.

fewer than 5 deaths. Panel (b) covers intermediate-severity disasters (p50$\leq$ deaths $<$ p90), which are less frequent and exhibit a right-skewed distribution 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 distributions highlight that the vast majority of disaster events result in relatively few deaths, whereas severe events are infrequent but disproportionately impactful.

Figure 3 illustrates the evolution of natural disasters worldwide from 2007 to 2018, categorized by severity, with data on deaths (Figure 3a) and economic damages (Figure 3b). Most events fall below the median (p50), while only a small fraction reach the top decile (>p90) or top 5% (>p95). Severe disasters—particularly those above p95—are rare and relatively stable over time, whereas moderate events fluctuate more widely. This pattern highlights that our baseline definition of extreme disasters, the top 10% of the mortality distribution, focuses on genuinely exceptional shocks without being driven by short-term spikes in more frequent, lower-severity events.

Figure 4 shows the distribution of natural disasters in our sample by type and severity, measured by deaths in the affected district. Blue bars mark events in the top decile (>p90 deaths), and orange bars the top 10% (>p95 deaths). Floods and epidemics dominate high-fatality events, while droughts and wildfires rarely reach these extremes. Our baseline uses the top 5% threshold to capture truly severe shocks; as visible in the figure, these events are also more evenly spread across disaster types than those above p90, which are heavily skewed toward floods. As a robustness check, we sequentially leave out one disaster type at a time to ensure that no single category—such as floods or epidemics—drives the results.

To ensure that our estimates reflect supply-chain disruptions rather than local disruptions, we control for whether the U.S. importer itself was directly affected by a major natural disaster. We construct this measure using the SHELDUS (Spatial Hazard Events and Losses Database for the United States), which reports the start and end dates and the FIPS codes of all affected counties for each disaster. Through GIS mapping, we assign each FIPS code to our ADM geographic units and aggregate disaster exposure at the ADM–quarter level. In our baseline specification, we include an indicator for whether the importer's ADM was affected in the relevant period. As a robustness check, we re-estimate the event-study restricting the sample to importers that were never directly exposed to a major disaster, confirming that the results are driven by foreign supplier shocks rather than local disruptions.

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

(a) Number of natural disasters by death percentile

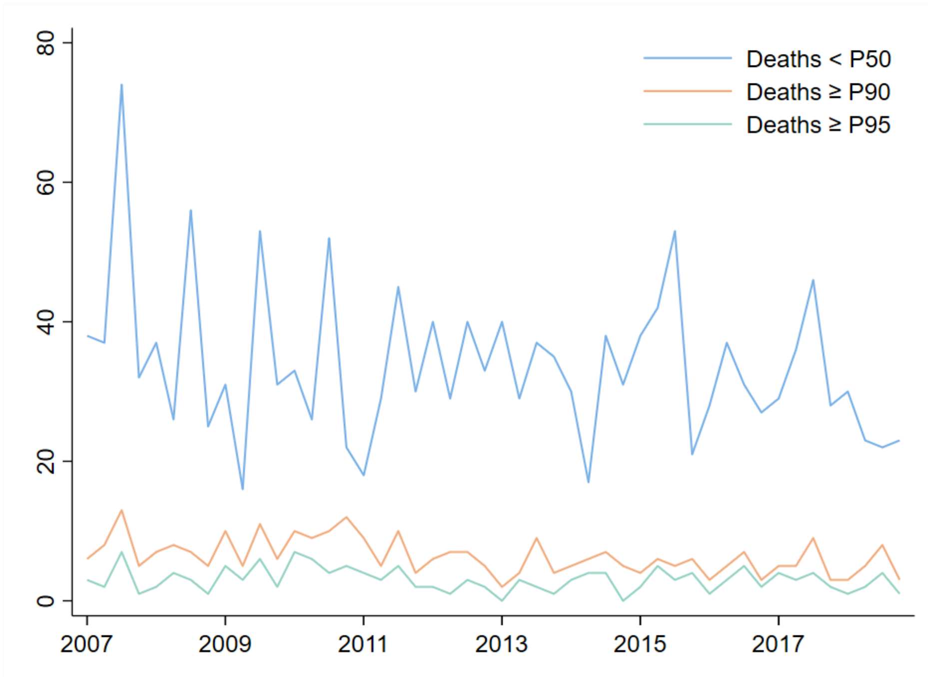

(b) Number of natural disasters by damage percentiles

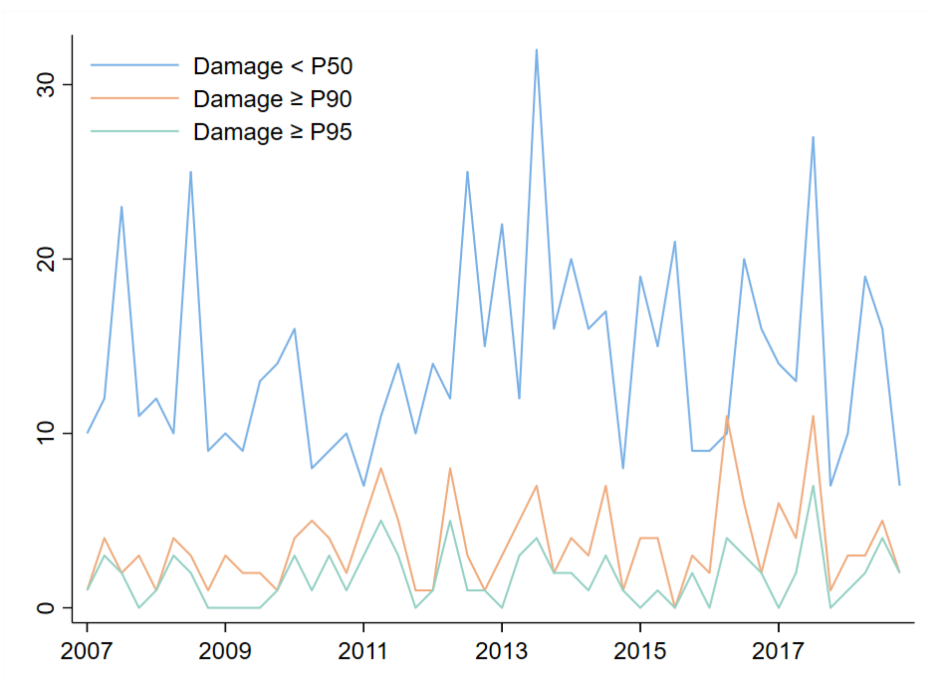

Notes: Panel (a) shows the quarterly number of natural disasters worldwide from 2007 to 2018, classified by the severity of fatalities as measured by death percentiles. The blue line represents disasters with deaths below the median (P50), the orange line corresponds to disasters with deaths in the top 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 to disaster damages.

FIGURE 4: Frequency of Extreme Disaster Events by Mortality Thresholds

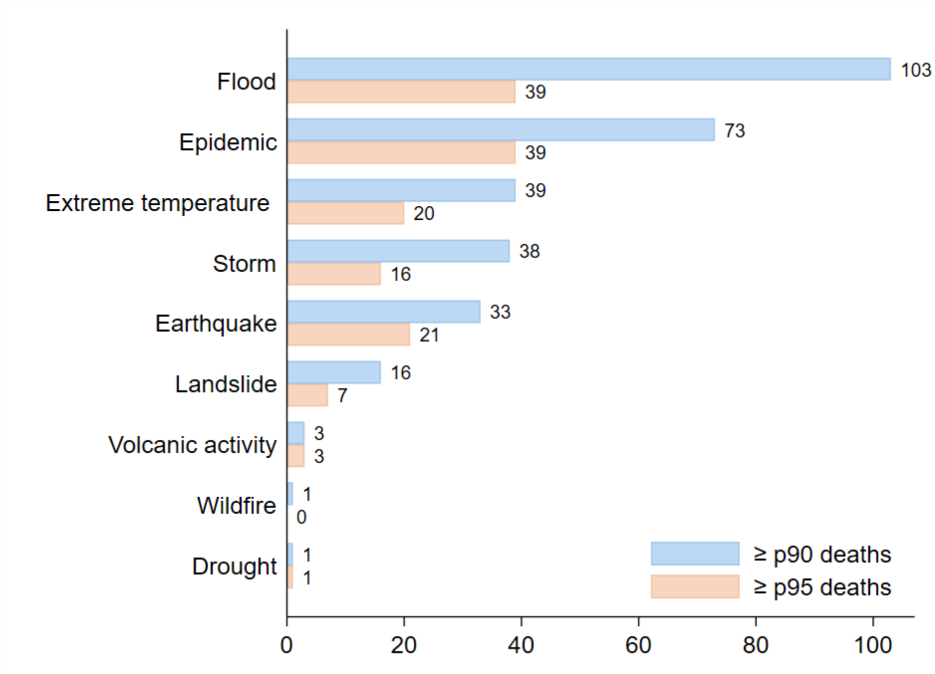

Notes: Severity is classified using historical mortality percentiles from the EM-DAT database. Disasters above the 90th percentile (p90) represent the deadliest 10%, while those above the 95th percentile (p95) mark the top 5% of events by death toll. Frequencies refer exclusively to disasters occurring between 2007 and 2018.

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—crucially for our purposes—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.

Crucially for our analysis, Orbis provides both firm names and physical addresses, which we use to match foreign exporters in the U.S. Bill of Lading (BoL) data. This match reveals whether a given transaction links affiliates within the same multinational group or instead occurs at arm's length between unrelated parties. Identifying these intra-firm and inter-firm links allows us to study how ownership structures shape the transmission of shocks through global production networks.

But beyond matching pairs, Orbis also lets us characterize each party individually. We can test whether the buyer or the seller is part of a multinational, regardless of whether they belong to the same group. This distinction is key. A transaction between two multinational firms that do not belong to the same corporate family may still reflect different sourcing strategies, financing constraints, or network positions than one between truly independent firms. And of course, when both parties do belong to the same corporate group, we can explicitly flag the trade as occurring within the firm.

3 Empirical Strategy

The empirical strategy is designed to analyze how extreme natural disaster shocks propagate through the international production network and affect firms along multiple margins. Our primary interest lies in understanding the mechanisms by which shocks to foreign suppliers can transmit to U.S. importers and eventually influence other suppliers that share these buyers, either amplifying or dampening the overall impact on global trade. This approach requires a framework that not only identifies the direct effects of disasters on exposed suppliers but also uncovers the second-order and network-mediated consequences that follow. To achieve this, we integrate event-study and difference-in-differences methods within a rich panel of firm-to-firm trade linkages, complemented by a detailed characterization of product specificity and ownership networks.

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. 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.

For importer-level regressions, identification relies on exposure operating only through an importer's directly linked foreign suppliers. A disaster in a supplier's district affects the importer through supply-chain disruptions, not through a contemporaneous collapse in final demand from that country. To reinforce this assumption, we indicate whether the affected country is also a destination for the importer's exports. The specification absorbs broader shocks and bilateral trends with high-dimensional fixed effects—importer-firm and product-by-calendar-quarter—so the comparison is within product-time cells across links that do or do not sit in the struck ADM2.

Identifying Firms Exposed to Natural Disasters

How do we identify firms that are exposed to natural disasters through their network of foreign suppliers? We integrate disaster severity data—measured at the ADM2 level and drawn from EM-DAT—into our empirical framework by linking each event to the geocoded location of foreign exporters. This assignment builds on location identifiers already established in our dataset, where each foreign supplier is matched to a U.S. importer using the universe of Bills of Lading. As a result, for any given quarter, we can identify which U.S. importers are connected to a foreign supplier hit by an extreme natural disaster.

Studying the propagation of shocks through international trade networks poses distinct empirical challenges. Cross-border relationships tend to be sparse and irregular, and trade within a given buyer–supplier pair may be intermittent. To address this, we work at the quarterly frequency—a level granular enough to capture near-term disruptions while still allowing year-over-year comparisons that help smooth out seasonal noise.

Moreover, we restrict the sample to *recurrent importers*—firms that appear in the data as importers every year (2007-2022) prior to being indirectly exposed to a severe natural disaster. This definition helps ensure that exposure is not mistakenly assigned to firms that exit or re-enter trade relationships for reasons unrelated to the shock. To determine whether a given importer is exposed, we begin by identifying which of its foreign suppliers are eligible to transmit a shock in a given quarter. We look back over the six quarters preceding quarter t, and calculate each supplier's share of the importer's total import value during that window: $w_{bs}(t) = \sum_{\tau=t-6}^{t-1} v_{bs\tau} / \sum_{\tau=t-6}^{t-1} \sum_{s'} v_{bs'\tau}$, where $v_{is\tau}$ is total import value of importer b of goods from supplier s in year τ.

A supplier s qualifies for importer b in quarter t if $w_{is}(t) \geq 0.10$ and i received at least one shipment from s during the six-quarter window. We denote the set of such qualifying suppliers by $\mathcal{S}_b(t)$. These thresholds serve to isolate economically meaningful and recently active relationships—those most likely to matter for sourcing continuity.

This approach is central to our identification strategy. We are interested in disruptions to active supplier relationships—not sporadic or incidental links. Without these thresholds, we risk mistaking the natural lags and lumpiness in international trade for disaster-induced substitution. By focusing on relationships that were both recent and non-trivial in value, we increase the likelihood that a post-disaster drop in trade reflects genuine disruption rather than a buyer simply discontinuing a marginal tie.

Therefore, an importer is considered treated if at least one of its qualifying foreign suppliers experiences an extreme natural disaster. In our baseline specification, a disaster is classified as extreme if its death toll lies above the 90th percentile of the global distribution. Events below the median are considered non-extreme, and those in the intermediate range are excluded from the main analysis to sharpen identification. In the baseline specification, treatment is absorbing: once a buyer is first exposed, it remains treated in all subsequent quarters. This design captures persistent disruptions. In robustness checks, we allow for repeated exposure: up to three separate qualifying suppliers may be hit over time, and we track each episode to study dynamic responses. We also flag whether the U.S. importer itself is affected by a disaster at home, using ADM2-level matching for U.S. locations.

Formally, we define the treatment indicator as:

$$\text{Treated}_{bt} = \mathbf{1}\left\{ \sum_{s \in \mathcal{S}_b(t)} w_{bs}(t) \cdot D_{st} > 0 \right\} \tag{3.1}$$

where $\mathcal{S}_b(t)$ is the set of qualifying suppliers for importer b in quarter t (defined by the 10% threshold and recent activity), $w_{bs}(t)$ is the share of total import value accounted for by supplier s, and D_{st} is an indicator equal to 1 if supplier s experiences an extreme natural disaster in quarter t. The indicator Treated_{bt} is equal to one if any qualifying supplier of buyer b is hit by a severe disaster in that period.

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–product-time level. Our baseline specification is:

$$X_{spt} = \alpha_s + \delta_{pt} + \gamma_{ct} + \sum_{k=-K}^{L} \beta_k \cdot \mathbf{1}\{\tau_{sct} = k\} + \varepsilon_{spt}, \tag{3.2}$$

where X_{spt} measures of supplier performance: total exports to the U.S., and as a complementary measure of network intensity, we use the total number of transactions each seller engages in across all US buyers and products. Supplier fixed effects, α_s, are controlling for time-invariant heterogeneity such as average productivity or geographic characteristics; δ_{pt} are product-time fixed effects capturing seasonality and global demand shocks at the HS4 level; and γ_{ct} are country-time fixed effects that absorb macroeconomic conditions, policy shocks, and other country-wide events.

To identify the effects, we restrict the baseline to recurrent exporters, defined as foreign firms that report positive exports to the United States in each of the previous six quarters since their first appearance. For treated firms, we relax this requirement after the shock, allowing export spells to become intermittent or to cease altogether. Focusing on recurrent exporters limits the noise from sporadic sellers and isolates the impact of severe natural disasters on firm performance, particularly on exports.

The main variables of interest are measured as quarterly year-on-year midpoint growth rates. For each calendar quarter q in year y, we compute $g_{q,y} = 2 \cdot \frac{X_{q,y} - X_{q,y-1}}{X_{q,y} + X_{q,y-1}}$, and smooth $g_{q,y}$ using a two-quarter moving average to reduce high-frequency volatility. Panjiva provides consistent quantities but does not report reliable trade values. To address this, we construct real flows by applying 2010 unit values from the U.S. Census—defined at the HS6–by–origin-country 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. As a robustness check, we also compute flows using current prices lagged one quarter to mitigate endogeneity from contemporaneous price movements.

The coefficients of interest, β_k, trace the evolution of trade flows before and after a disaster. Leads ($\kappa < 0$) test for pre-trends and potential anticipatory behavior, while lags ($\kappa > 0$) capture the trajectory of post-disaster outcomes. We normalize $\kappa = -1$ as the reference period to avoid collinearity. This event-study design allows us to assess both the timing and persistence of the disruption. 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 to which the disaster impairs the capacity and ability of the foreign supplier to continue exporting to the U.S.

Standard errors are clustered at the ADM2 level for supplier regressions and at the importer level for importer regressions to account for serial and cross-sectional correlation. Robustness checks vary the severity threshold used to define extreme events, drop quarters spanning global crises (e.g., 2008–09), and implement placebos that randomly assign treatment to firms that never experience a qualifying shock. Across these exercises, the estimated responses align with localized supply-side disruptions propagating through global value chains, not with demand conditions or macroeconomic confounders.

3.3 Transmission to U.S. Importers

This section examines the second stage of transmission: the response of U.S. importers to supply disruptions affecting their foreign suppliers. The design parallels the supplier-level regressions but shifts the unit of observation to the importer, allowing us to trace downstream effects once an upstream shock occurs. We estimate the following specification:

$$\begin{aligned} M_{bt} = \alpha_b + \delta_{p(b)t} + \theta^{(1)}\, \mathbf{1}\{\text{ND hits buyer itself}_{bt}\} \\ + \sum_{k=-K}^{L} \theta_k^{(2)}\, \text{Treated}_{bt} + X_{bt} + \eta_{bt}. \end{aligned} \tag{3.3}$$

where M_{bt} are outcome variables for importer b in year t. We estimate this equation separately for three dependent variables: (i) the growth rate of total imports; (ii) the growth rate of imports excluding those from affected suppliers; and (iii) the growth rate of exports, for buyers that also sell abroad. As in the supplier-side regressions, all outcomes are expressed in year-on-year midpoint growth rates and smoothed using a two-quarter moving average.

The fixed effects α_b and $\delta_{p(b)t}$ control for time-invariant buyer characteristics and product-by-calendar-quarter shocks, respectively. The product $p(b)$ is defined as the buyer's primary importing industry, based on cumulative import value. The first coefficient, $\theta^{(1)}$, captures the direct effect of a disaster striking the buyer's own ADM2 location. The second set, $\theta_k^{(2)}$, measures how shocks to qualifying foreign suppliers propagate through established trade links.

The definition of Treated_{bt} follows the definition in Equation (3.1). A buyer is considered treated if, in a given quarter, at least one of its *qualified* foreign suppliers—defined as those accounting for at least 10% of past import value and active within the previous six quarters—is located in a district hit by an extreme natural disaster.[3]

To address potential confounding factors, we include two additional controls. First, we account for whether the importer itself is affected by a domestic disaster, using ADM2-level geographic matching within the United States. Second, for importers that also export, we flag whether the affected foreign country is also a destination for their sales, X_{bt}. These adjustments help separate upstream propagation through supplier links from confounding demand-side shifts.

The specification in 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 multinational status of the buyer and the presence of shared ownership links across borders. These organizational structures may offer internal margins of

[3]The minimum threshold requirement is tightened in robustness exercises; see Section 6.

adjustment—through coordination, reallocation, or financing—that are unavailable in more fragmented networks. Our estimating equations are:

$$M_{bt} = \alpha_s + \delta_{b(p)t} + \gamma_{ct} + \sum_{k=-K}^{L} \Big(\beta_k + \phi_k \cdot \text{MNE}_s\Big) \mathbf{1}\{\tau_{sct} = k\} + \varepsilon_{bt}. \tag{3.4}$$

$$M_{bt} = \alpha_s + \delta_{b(p)t} + \gamma_{ct} + \sum_{k=-K}^{L} \Big(\nu_k + \delta_k \cdot \text{Intra-firm}_s\Big) \mathbf{1}\{\tau_{sct} = k\} + \epsilon_{bt}. \tag{3.5}$$

Here, MNE_s is an indicator for whether the supplier belongs to a U.S. multinational network, and Intra-firm_s denotes whether the trade relationship occurs between a US buyer and a foreign exporter belonging to the same multinational corporate group. These specifications allow us to test whether multinational ownership and intrafirm 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 (β_k) and non-MNC (ϕ_k) suppliers, and for intra- (ν_k) versus arm's-length (δ_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 Rauch (1999) classification, 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 Descriptive Statistics

Before turning to the 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 supplier relevance, importer recurrence, and disaster classification described above.

Table 1 presents summary statistics for our main firm-level panel, 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 firms that report positive shipments to the U.S. in each of the six quarters prior to their first shock, and are matched to natural disaster exposure via geocoded location data. On average, a foreign exporter ships $170,000 per quarter (2010 USD) across approximately 16 transactions and 24 U.S. buyers. However, the distribution is highly skewed, with the top 1% of exporters accounting for over $41 million in quarterly flows. Export growth over a four-quarter horizon is near zero on average but highly volatile, with a standard deviation above 1.

Roughly 18% of exporter-quarters correspond to multinational firms, and 10% coincide with an extreme natural disaster affecting the exporter's subnational location (ADM2), as defined in our identification strategy. These events represent the source of exogenous shocks in the empirical design.

Panel B presents summary statistics for U.S. importers—the central units of analysis in our study. We restrict attention to importer–supplier links that are non-trivial in value, requiring that the supplier accounts for at least 10% of the importer's total import value over the relevant horizon. This sample includes over 8.8 million firm-quarter observations, restricted to *recurrent importers*—U.S. firms that, upon first appearance in the data, have at least one recorded import transaction in each of the four preceding quarters. The average importer reports fewer than one active supplier per quarter, reflecting the sparse and intermittent nature of international sourcing relationships. Average import growth over four quarters is 38%, with considerable dispersion.

Only 7% of importer-quarters belong to multinational firms, and just 2% of importer–supplier links are intra-firm, based on common ownership across borders. About 65% of transactions involve differentiated goods, depending on the definition used. About 3% of importer-quarters also report positive exports, underscoring the presence of dual-role firms. These patterns confirm the sparse, lumpy nature of international sourcing—and highlight the need to focus on economically meaningful buyer–supplier relationships when estimating shock

Table 1: Descriptive Statistics

Panel A. Exporters	Obs.	Mean	Std. Dev.	p25	p50	p75
Transactions (count)	1,079,351	15.96	110.31	1.00	4.00	12.00
Buyers (count) ncon_tinvariant	1,087,548	23.52	58.39	7.00	13.00	25.00
Exports growth (y/y)	1,030,609	0.03	1.03	-0.62	0.03	0.69
Transactions growth (y/y)	1,031,607	0.02	0.83	-0.40	0.00	0.48
Dummy variables						
Multinational firm	1,087,548	0.18	0.39			
Natural disaster event	1,087,548	0.10	0.30			

Panel B. Importers	Obs.	Mean	Std. Dev.	p25	p50	p75
Transactions (count)	4,049,573	9.59	97.28	0.00	1.00	4.00
Suppliers (count) nshp_tvariant	8,800,764	0.95	4.93	0.00	0.00	1.00
Imports growth (y/y)	2,583,151	0.38	1.24	-0.45	0.29	1.69
Imports growth excl. aff. (y/y)	2,549,420	0.39	1.23	-0.45	0.30	1.70
Exports growth (y/y)	92,998	0.16	1.64	-1.94	0.26	2.00
Transactions growth (y/y)	2,912,832	0.34	1.23	-0.40	0.18	1.64
Suppliers growth (y/y)	2,912,832	0.33	1.19	-0.29	0.00	1.33
Dummy variables						
Multinational firm	8,800,764	0.07	0.26			
Intrafirm trade	8,800,764	0.02	0.13			
Capital good	8,800,764	0.01	0.11			
Diff. good (liberal def.)	8,800,764	0.63	0.48			
Exporter	8,800,764	0.03	0.16			

Notes: This table reports summary statistics for the estimation sample used in the analysis. Panel A presents foreign exporters—firms located outside the United States that ship to U.S. importers, as recorded in US Bill of Lading data. Exporters are required to be recurrent: they must report at least one shipment to the U.S. in each of the six quarters preceding their first observation in the panel. The unit of observation is the exporter-quarter. Export values are expressed in millions of constant 2010 U.S. dollars and are constructed by applying 2010 HS6-by-country unit values from the U.S. Census to quantity data from Panjiva. The variables *Export growth* and *Transactions growth* are defined as midpoint year-over-year growth rates, smoothed with a two-quarter moving average. *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 primary units of analysis in the paper. Importers are classified as recurrent if they appear in at least one quarter of every calendar year in the sample. The unit of observation is the importer-quarter. All monetary variables are in millions of 2010 U.S. dollars. *Imports excl. affected* excludes flows from disaster-affected suppliers, based on the importer's qualifying supplier set (defined using a six-quarter lookback and a 5% value share threshold). *Transactions* and *Suppliers* count the number of distinct importer–supplier links per quarter. Growth variables are defined as midpoint year-over-year changes and smoothed over two quarters. *Multinational firm* is an indicator for U.S. firms that belong to a multinational group. *Intrafirm trade* identifies importer–supplier pairs that share common ownership across borders. *Diff. good* and *Diff. good (liberal def.)* classify traded products based on the Rauch (1999) classification. *Exporter* indicates whether the U.S. importer also exports in a given quarter.

transmission.

5 Results

This section presents the main empirical results on how natural disasters affecting foreign suppliers propagate through international production networks. We begin by estimating the direct effects on foreign exporters—firms located in disaster-hit regions—focusing on their sales 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.

5.1 Direct Effects on Foreign Exporters

We begin by estimating the direct effects of extreme natural disasters on foreign exporters. Figure 5 plots event-study coefficients for two key outcomes: total export value to the United States (Figure 5a) and the number of transactions (Figure 5b). Each panel reports quarterly year-on-year midpoint growth, centered on the quarter of the first qualifying disaster affecting the supplier's ADM2 location. Estimates are based on a panel of recurrent exporters and include exporter fixed effects, product-by-calendar-quarter fixed effects, and country-by-quarter fixed effects. Standard errors are clustered at the ADM2 level.

The results indicate a sharp and persistent contraction in trade activity following the disaster. Export value declines by roughly 15 percentage points on impact (quarter $t = 0$) and continues falling to nearly 20% by quarter $t = 2$, before stabilizing. Given that average annual export growth in the sample is just 3%, this represents a fall more than six times the typical growth rate. In levels, with average quarterly exports of USD 5.37 million, the shock corresponds to a loss of more than USD 1 million per exporter on average. The effect remains statistically and economically significant for at least eight quarters after the shock.

A similar pattern emerges for the number of transactions: the initial drop is slightly steeper, around 18 percentage points, and recovery is limited even three years after the event. With the typical exporter conducting nearly 16 transactions per quarter, the reduction implies almost three shipments less per quarter immediately after the disaster. Both outcomes show flat pre-trends, supporting the assumption that parallel dynamics exist in the absence of treatment.[4]

[4] When we decompose total exports to the U.S. into the number of U.S. clients (extensive margin) and average sales per client (intensive margin), we find that both margins are negatively and significantly affected by the shock (TBA).

FIGURE 5: Effects of c Disasters on Foreign Exporters

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

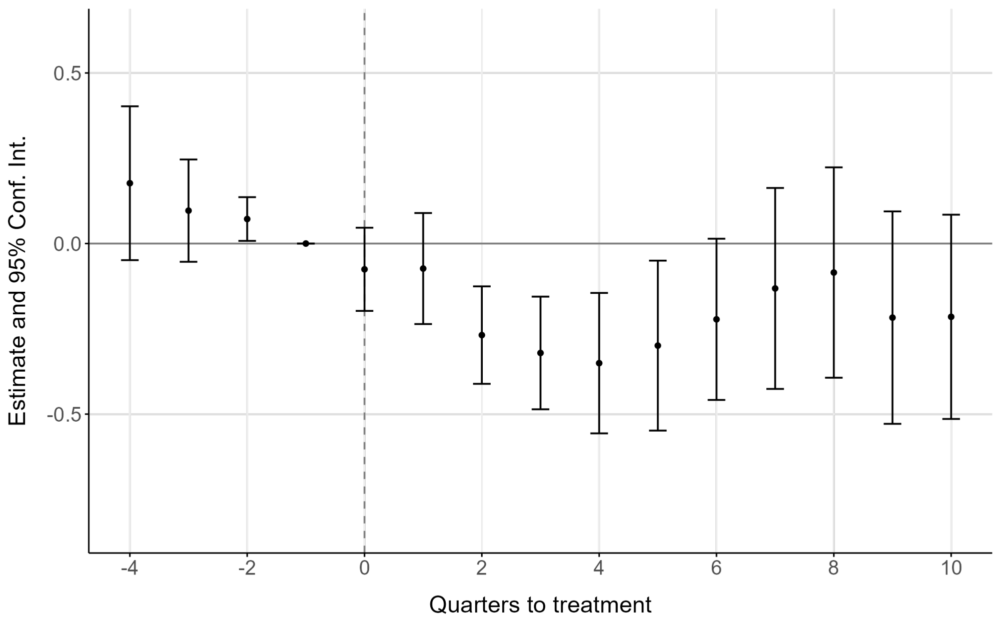

(B) Number of transactions to the U.S. (YoY midpoint growth)

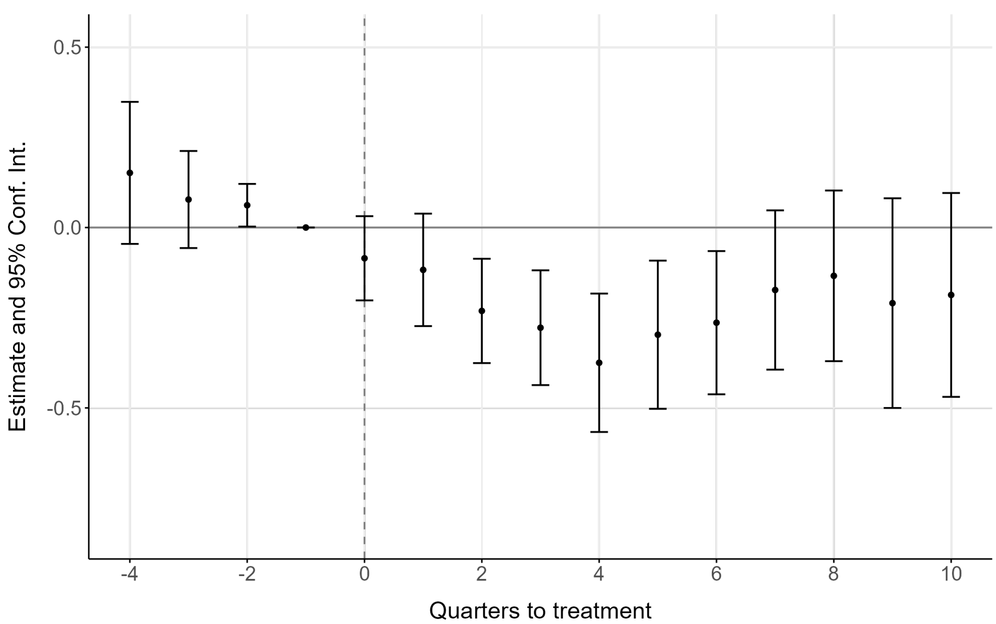

Notes: This figure plots the estimated effects of extreme natural disasters on directly affected foreign exporters. Each panel reports coefficients from an event-study regression using quarterly year-on-year midpoint growth as the outcome. Panel A shows the response of total export value to the U.S.; 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. 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.

Table 2 summarizes these dynamic responses using a difference-in-differences specification that aggregates the post-treatment periods into two windows: quarters 0–5 and quarters 6–10. The estimates confirm a large and persistent decline in both exports and transactions immediately after the shock (–22% and –19%, respectively), followed by only partial recovery in later quarters (–12% and –11%). The DiD results align closely with the event-study trajectories, indicating that the observed persistence is not driven by outliers or timing heterogeneity.

These results show that disasters cause immediate and lasting disruptions to firms' export capacity. The observed persistence is partly driven by supplier exit, as well as by sustained declines in average export volumes among surviving firms.

TABLE 2: Difference-in-Differences Estimates: Exporter Response to Foreign Natural Disasters

	Export growth	Transactions growth
Quarter 0–5 after shock	-0.223$^{(a)}$	-0.192$^{(a)}$
	(0.068)	(0.064)
Quarter 6–10 after shock	-0.123$^{(b)}$	-0.117$^{(b)}$
	(0.058)	(0.053)
Fixed Effects		
Exporter	Yes	Yes
Industry × Quarter	Yes	Yes
Country × Quarter	Yes	Yes
Observations	1,030,609	1,031,607
R^2	0.065	0.077

Notes: This table reports difference-in-differences estimates of the impact of extreme natural disasters on foreign exporters. The dependent variables are the midpoint year-on-year growth in export value and transaction count. Estimates are based on the following specification:

$$y_{st} = \beta_1 \cdot \text{Post}_{[0,5]st} + \beta_2 \cdot \text{Post}_{[6,10]st} + \alpha_s + \delta_{p(s)t} + \theta_{c(s)t} + \varepsilon_{st}$$

where y_{st} denotes the outcome of interest for exporter s from country c, in quarter t. The indicators $\text{Post}_{[0,5]}$ and $\text{Post}_{[6,10]}$ equal one if quarter t falls within the first 6 or subsequent 5 quarters after the first qualifying disaster. The regression includes exporter fixed effects (α_s), industry of the seller-by-quarter fixed effects ($\delta_{p(s)t}$), and exporter country-by-quarter fixed effects ($\theta_{c(s)t}$). Standard errors clustered at the ADM2 level. Significance: $^{(a)}$ 1%; $^{(b)}$ 5%; $^{(c)}$ 10%.

Exporters by Ownership: to assess whether multinational ownership buffers the effects of disasters on directly hit suppliers, Figure A.1 disaggregates the event-study by exporter status. Figure A.1a shows the response of export value to the U.S., and Figure A.1b shows the number of transactions. We find no significant difference in the magnitude or persistence of the impact across groups: both multinational and non-multinational suppliers experience a sharp and enduring fall in trade activity, with no evidence of faster recovery among multinationals. Table A.1 reports the corresponding DiD estimates, which corroborate the graphical evidence—coefficients are virtually identical across ownership types. This suggests that while multinational status may facilitate coordination and financing, these channels do not mitigate the direct export losses once a facility is physically hit.

Overall, the first-stage results establish that extreme natural disasters generate sizeable, persistent supply shocks at the exporter level, independent of multinational affiliation. We next examine how these shocks propagate downstream to U.S. importers. In what follows, we shift focus to the downstream margin—examining whether the importer response varies depending on organizational form and product specificity.

5.2 Propagation to U.S. Importers

We now turn to the downstream effects of foreign shocks on U.S. importers. Figure 6 presents event-study estimates for treated importers, reporting YoY midpoint growth in total import value (Figure 6a) and in import value excluding directly affected suppliers (Figure 6b). Treatment is defined at the importer level and is absorbing: once a firm becomes exposed—through any qualifying foreign supplier hit by a disaster—it remains treated in all subsequent quarters. Specifications include importer fixed effects, product-by-quarter fixed effects, and exporter-country–by–quarter fixed effects. Standard errors are clustered at the ADM2 level of the affected supplier.

The figure reveals two distinct but complementary dynamics. Figure 6a shows a dramatic drop in total import value on impact—exceeding one standard deviation—followed by a slow but steady recovery. This aggregate response masks an important pattern: Figure 6b excludes trade with directly affected suppliers and isolates the behavior of flows from unaffected sources. Although the initial contraction is more modest, imports from these suppliers also decline significantly on impact. Over time, this margin gradually rebounds—beginning around quarter 5 and closing much of the gap by the end of the horizon—indicating some substitution toward unaffected suppliers, albeit with frictions that delay the adjustment.

The fact that both margins decline on impact indicates that importers do not immediately reallocate demand to unaffected suppliers. In many cases, the shock involves the loss of a critical input for which no close substitute exists in the short run. This limits the firm's ability to continue producing at previous levels. Even when substitution occurs—whether toward domestic suppliers or lower-quality foreign matches—it may not fully restore operational capacity. Imported intermediates are often selected precisely for their specificity

FIGURE 6: Disruption to U.S. Importers Following Foreign Supplier Shocks

(A) Total import value (YoY midpoint growth)

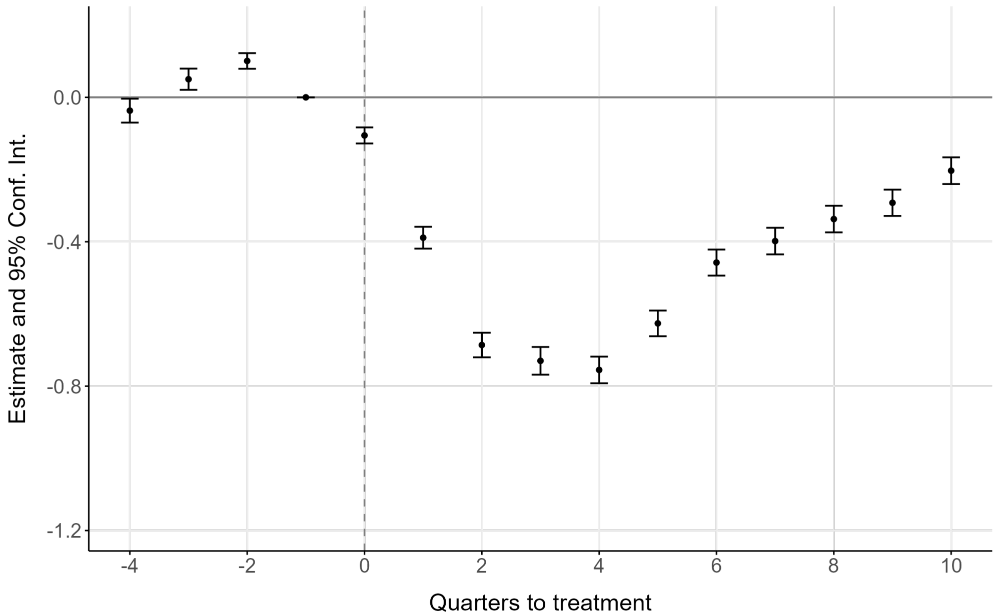

(B) Import value excluding directly affected suppliers (YoY midpoint growth)

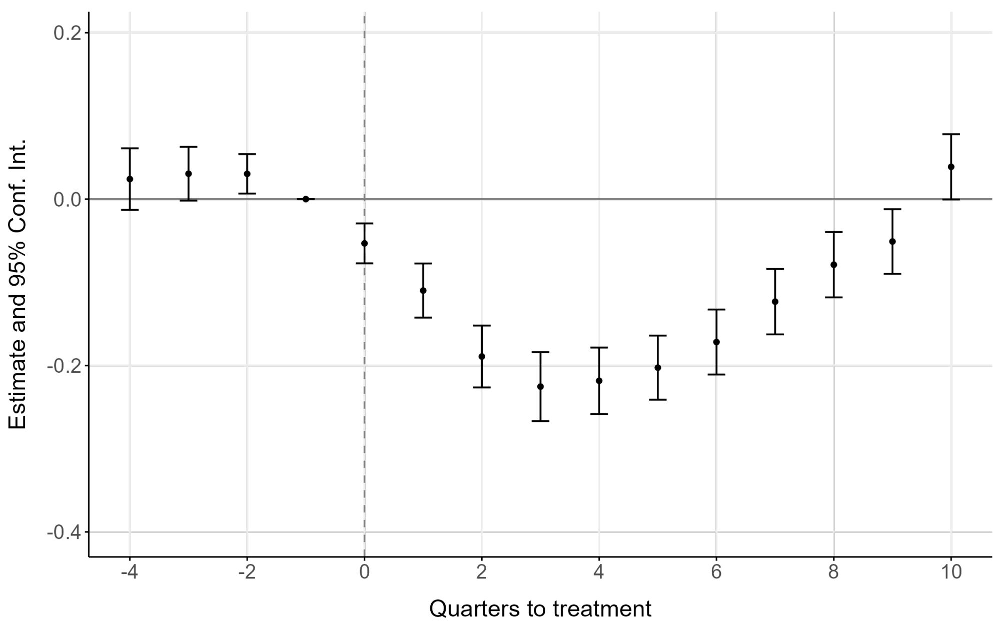

Notes: This figure plots the estimated effects of extreme natural disasters on U.S. importers exposed to affected foreign suppliers. Each panel reports coefficients from an event-study regression using quarterly year-on-year midpoint growth as the outcome. Panel (a) shows total import value; Panel (b) excludes shipments from affected suppliers. The sample includes recurrent importers, and treatment is defined as the first quarter in which any qualifying supplier is hit by a disaster at the ADM2 level. All regressions include importer fixed effects, product-by-calendar-quarter fixed effects, and exporter-country-by-quarter fixed effects. Standard errors are clustered at the ADM2 level of the foreign supplier.

Table 3: Difference-in-Differences Estimates: U.S. Importer Exposure to Foreign Shocks

	Import value	Imports from non-affected suppliers	Transaction count	Export value	Number of suppliers
Quarter 0–5 after shock	-0.531$^{(a)}$	-0.266$^{(a)}$	-0.501$^{(a)}$	-0.073$^{(c)}$	-0.495$^{(a)}$
	(0.0095)	(0.0100)	(0.0087)	(0.0408)	(0.0082)
Quarter 6–10 after shock	-0.324$^{(a)}$	-0.190$^{(a)}$	-0.274$^{(a)}$	0.036	-0.264$^{(a)}$
	(0.0107)	(0.0110)	(0.0098)	(0.0471)	(0.0092)
Fixed Effects					
Importer	Yes	Yes	Yes	Yes	
Industry × Quarter	Yes	Yes	Yes	Yes	
Observations	2,583,151	2,549,420	2,912,832	92,998	2,912,832
R^2	0.303	0.302	0.323	0.219	0.337

Notes: This table reports difference-in-differences estimates of the impact of foreign natural disasters on U.S. importers. The dependent variables are expressed as year-on-year (YoY) midpoint growth rates. Each column corresponds to a different outcome measured at the importer–quarter level. The treatment indicator equals one starting in the first quarter in which any qualified supplier is affected by a natural disaster, identified via an ADM2-level geographic match. The regression specification is as follows:

$$y_{bt} = \beta_1 \cdot \text{Post}_{[0,5]bt} + \beta_2 \cdot \text{Post}_{[6,10]bt} + \alpha_b + \delta_{p(b)t} + \varepsilon_{bt}$$

where y_{bt} denotes the outcome of interest for importer b, and quarter t. Fixed effects include importer, industry-by-quarter. "Non-affected suppliers" exclude trade with directly impacted exporters. Standard errors are clustered at the ADM2 level. Significance: $^{(a)}$ 1%; $^{(b)}$ 5%; $^{(c)}$ 10%.

or reliability, and losing access to them can weaken the performance of downstream activities. For importing firms that also export, these disruptions may translate into lower sales abroad.[5]

Table 3 complements these dynamic patterns with average estimates across multiple margins. In addition to total imports and imports from unaffected suppliers, we observe sizable declines in the number of suppliers and in the importer's own export activity—suggesting both upstream and downstream consequences. The extensive margin is especially affected: treated importers shrink their supplier base substantially in the aftermath of the shock, with no short-run offset via new matches.

Taken together, the evidence reveals both the limits and possibilities of adjustment. The initial shock triggers widespread disruption across the supply base, but some substitution toward unaffected suppliers emerges gradually. The slow pace of this recovery, however,

[5]See Figure A.2 for the effect on foreign sales, disaggregated by MNC status.

suggests that reconfiguring international input chains is costly—particularly when trade relationships are relationship-specific and sticky.

5.3 Ownership and Input-Specificity as Moderators of Propagation

Figures 7 and 8, and Table 4 explore how organizational structure and input-specificity shapes the response of U.S. importers to foreign supply disruptions. We begin by comparing importers that are part of a multinational corporation to those that are not. Figure 7a shows that both MNCs and non-MNCs experience a sharp contraction in total import value following the shock, but the magnitude and persistence differ markedly. Non-MNCs exhibit a peak decline of approximately –0.96 in YoY midpoint growth—around three times the average annual import growth in the sample (0.38)—and remain substantially below baseline throughout the horizon. MNCs, by contrast, experience a smaller initial drop of about –0.3 and begin recovering after four quarters, returning close to pre-shock levels by the end of the sample window.

Figure 7b, which excludes imports from directly affected suppliers, reveals broadly similar initial responses across importer types. Both MNCs and non-MNCs reduce imports from unaffected suppliers following the shock, but these initial declines are not statistically different on impact. However, the trajectories begin to diverge meaningfully after a few quarters. While MNCs exhibit a gradual but sustained rebound—returning to baseline around quarter 6—non-MNCs show limited recovery and remain well below pre-shock levels throughout the horizon.

However, because MNCs also engage in arm's-length trade, the comparison does not isolate the role of intra-firm sourcing. Figure 8 therefore examines whether the nature of the disrupted trade relationship—specifically, whether the affected supplier belongs to the same corporate group as the importer (intra-firm) or not—modulates the severity of disruption. Figure 8a shows that importers exposed to arm's-length shocks experience a peak decline of nearly –0.8 in total import value, while those with intra-firm disruptions contract by about –0.4. The difference is statistically significant over multiple quarters, and recovery is faster and more complete for intra-firm trade.

This finding challenges the notion that intra-firm trade is inherently more exposed to disruption due to input specificity or tight technological links. In practice, the organizational structure of multinational firms appears to offer a buffer that more than compensates for these features when shocks occur. Intra-firm relationships enable faster substitution across affiliates and more coordinated reallocation within corporate networks—adjustment channels rarely accessible in arm's-length trade.

Figure 7: Effects of Foreign Shocks on U.S. Importers, by MNC Status

(a) Total import value (YoY midpoint growth)

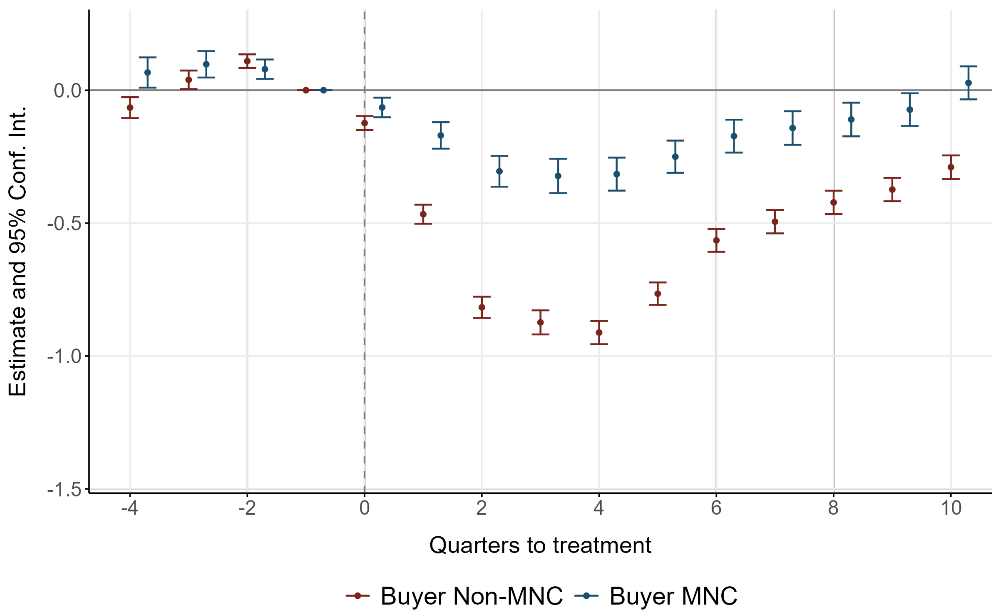

(b) Import value excluding affected suppliers (YoY midpoint growth)

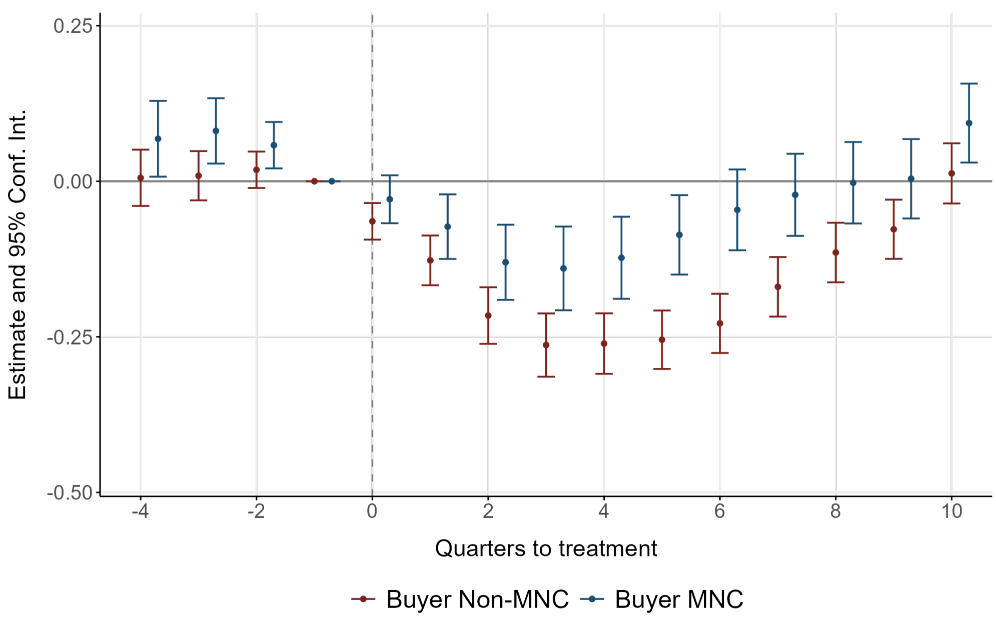

Notes: This figure plots the estimated effects of extreme natural disasters on U.S. importers, disaggregated by multinational status. Each panel shows event-study coefficients using quarterly YoY midpoint growth as the outcome. Panel (a) reports total import value; Panel (b) excludes imports from directly affected suppliers. The sample includes recurrent importers, and treatment is defined based on exposure to a qualifying supplier affected by a disaster (ADM2-level match). Multinational status is defined using Orbis ownership links. All regressions include importer fixed effects, product-by-quarter fixed effects, and exporter-country–by–quarter fixed effects. Standard errors are clustered at the ADM2 level of the supplier.

Figure 8b focuses on imports from unaffected suppliers. Here, the distinction between intra-firm and arm's-length cases is not statistically significant. Point estimates for intra-firm links are slightly more negative, but confidence intervals are wide throughout. This imprecision likely reflects sample selection: when an importer relies exclusively on a disrupted intra-firm supplier, it exits the sample in Panel (b), which excludes affected flows. As a result, we lose precisely those firm–product pairs that may be most exposed, leading to attenuation and noise.

Even so, the absence of a clear recovery in this margin suggests that reallocation toward unaffected suppliers is limited. If intra-firm importers do recover faster in the aggregate—as seen in Panel (a)—the adjustment is unlikely to come from outside the corporate group. Rather, the recovery may reflect the restoration of the disrupted internal link itself, or substitution across affiliates within the same firm network—adjustments that are not visible in Panel (b) by construction.

These dynamics provide suggestive evidence that intra-firm links may be less vulnerable—not because the inputs involved are easier to replace, but because multinational firms can mobilize internal networks more quickly. While intra-firm trade is often assumed to involve more tailored or integrated production processes, the relative resilience of these flows points instead to the advantages of organizational coordination: firms may be able to substitute the disrupted affiliate with another within the same corporate group, or to reconfigure production through internal logistics channels unavailable in arm's-length settings.

Intra-firm trade, by definition, occurs within multinational networks and represents a narrower—but more flexible—sourcing channel. The evidence suggests that this internal margin plays a central role in the aggregate resilience of MNCs. In contrast, arm's-length trade spans both non-MNCs and MNCs sourcing from unaffiliated suppliers and appears more exposed to persistent disruption. This distinction could help explain why MNCs, as a group, perform better in the face of shocks and underscores organizational integration as a key moderator of shock transmission.

Finally, panel C in Table **??**tab:appendix$_{i}nteractions]showstheroleofinput-specificityinthetransmission$
$bordershocks.Asdiscussedintheintroduction, theeffectofimportingrelationship-specificinputsonthetra$
$specificitymightfosterrecoveryduetomutualinterestfromthetradingpartnerstomaintaintheirrelationsh$

This pattern aligns with macro-level observations from the Great Recession, during which intra-firm trade flows tended to contract less than arm's-length transactions. Our micro-level findings are consistent with this broader trend. Multinational networks appear to offer not only geographic diversification but also structural flexibility, giving firms more room to adapt

Figure 8: Effects of Foreign Shocks on U.S. Importers, by Intra-Firm vs. Arm's-Length Trade

(a) Total import value (YoY midpoint growth)

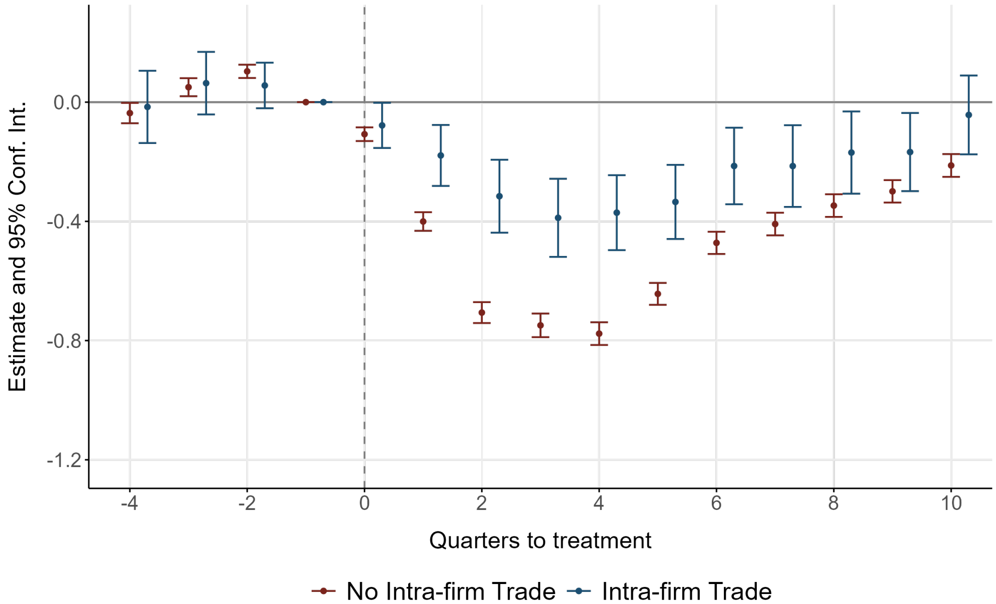

(b) Import value excluding affected suppliers (YoY midpoint growth)

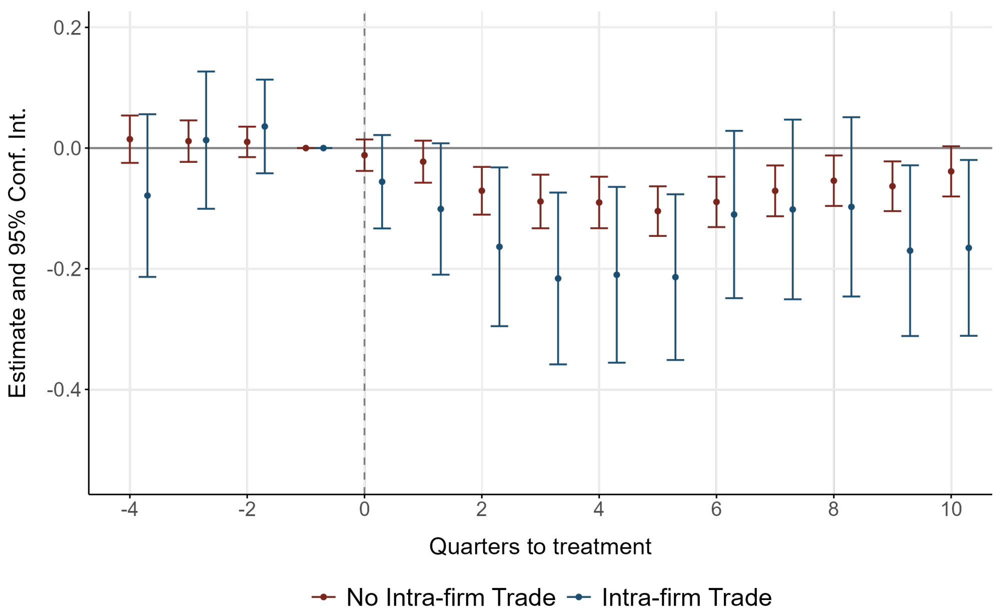

Notes: This figure plots the estimated effects of extreme natural disasters on U.S. importers, disaggregated by whether the affected supplier belongs to the same corporate group (intra-firm) or not (arm's-length). Each panel shows event-study coefficients using quarterly YoY midpoint growth as the outcome. Panel (a) reports total import value; Panel (b) excludes imports from directly affected suppliers. The sample includes recurrent importers, and treatment is defined based on exposure to a qualifying supplier affected by a disaster (ADM2-level match). Corporate group affiliation is identified using Orbis ownership links. All regressions include importer fixed effects, product-by-quarter fixed effects, and exporter-country–by–quarter fixed effects. Standard errors are clustered at the ADM2 level of the supplier.

when disruptions occur. These internal margins of adjustment, such as shifting production across affiliates or reallocating orders internally, are typically less available to firms that rely on a more fragmented network of external suppliers.

6 Robustness

In this section, we assess the robustness of our findings to alternative estimation strategies and sample definitions. We focus on two main concerns. First, when treatment is rolled out at different times—as in our setting, where different suppliers are hit by disasters in different quarters—standard fixed effects models can produce misleading results. These models average across units treated at different times, even when the shocks they experience differ in both timing and severity. By pooling these effects without accounting for such differences, the estimates may obscure the true underlying dynamics. Second, we examine whether our results are sensitive to the inclusion of specific types of disasters, to local disruptions affecting U.S. buyers, or to other sample-related features. Across exercises, our main conclusions remain stable.

6.1 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), Callaway and Sant'Anna (2021), and Sun and Abraham (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 9 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 9: Staggered DiD Robustness: Disruption to U.S. Importers

(A) Total import value (YoY midpoint growth)

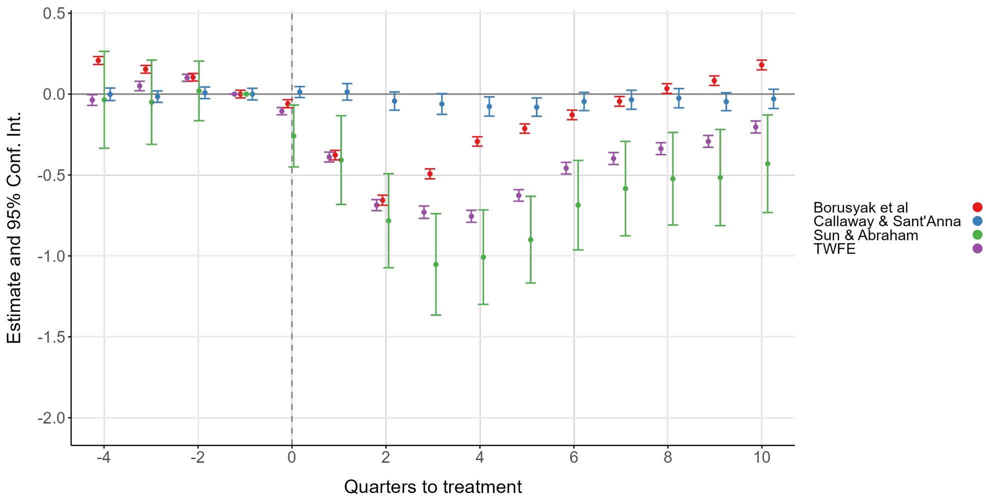

(B) Import value excluding directly affected suppliers (YoY midpoint growth)

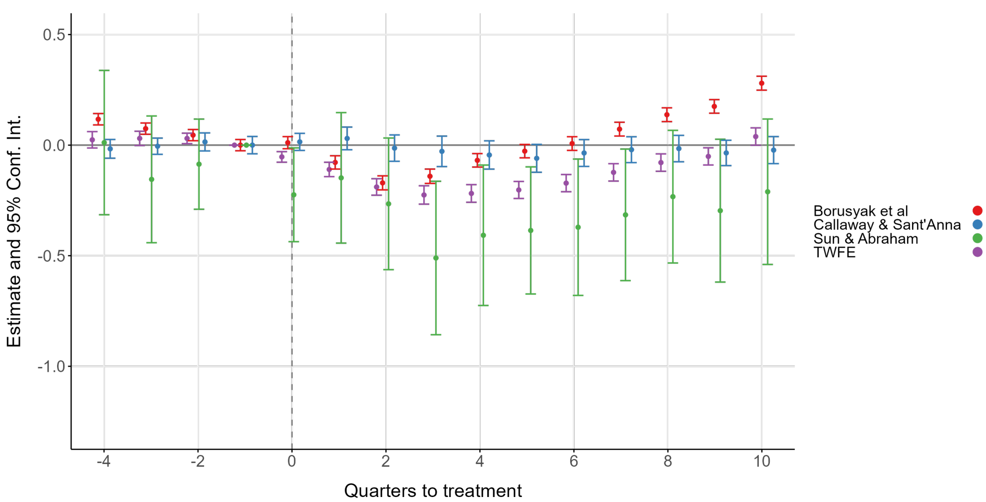

Notes: This figure plots the estimated effects of extreme natural disasters on U.S. importers exposed to affected foreign suppliers, using methods that account for heterogeneous treatment effects. Each panel reports coefficients from an event-study regression using quarterly year-on-year midpoint growth as the outcome. Panel (a) shows total import value; Panel (b) excludes shipments from affected suppliers. The sample includes recurrent importers, and treatment is defined as the first quarter in which any qualifying supplier is hit by a disaster at the ADM2 level. All regressions include importer fixed effects, product-by-calendar-quarter fixed effects, and exporter-country-by-quarter fixed effects. Standard errors are clustered at the ADM2 level of the foreign supplier.

6.2 Placebo and Other Robustness Checks

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. To do so, we sequentially exclude one category of disasters at a time—such as floods, earthquakes, or epidemics—and re-estimate the main event-study specification. This jackknife-style approach helps verify that a particular type of shock does not disproportionately influence our results.

Second, we address the concern that effects might be confounded by local disruptions affecting U.S. importers directly. We restrict the sample to buyers who have never experienced a disaster at their own location and re-estimate the baseline specification. This allows us to isolate the propagation of shocks through supplier linkages, independently of any concurrent domestic events.

We implement additional placebo tests. One strategy involves artificially shifting the timing of disasters—e.g., assigning placebo events several quarters before the actual shock—to ensure that the estimated effects do not reflect pre-existing trends. Another approach uses pseudo-treatment assignments across randomly selected importer–exporter pairs that were never actually affected, to validate that the estimated dynamics are not observed in unaffected matches. These exercises strengthen the case that the observed patterns are not artifacts of specification choices or sample composition.

7 Conclusions

Understanding how supply chain shocks travel across international networks requires careful attention to both the origin of the shock and the pathways of its diffusion. In this paper, we focus on the consequences of natural disasters affecting foreign suppliers and investigate how these events ripple through the web of global production. Specifically, we analyze the extent to which such shocks impair the operations of directly affected suppliers, how they subsequently influence U.S. importers that rely on those suppliers, and whether these disruptions cascade further to other suppliers connected through shared customers. Our approach integrates firm-level trade data, detailed disaster records, and ownership information to capture the multifaceted nature of supply chain propagation.

We find that extreme natural disasters generate large and persistent disruptions in trade relationships. Exporters located in affected regions experience a sharp decline in shipments to the United States—on the order of 15–20% in year-over-year growth—accompanied by a drop in the number of transactions that lasts several years. These effects reflect not only immediate physical damage, but also the difficulty of re-establishing disrupted connections in international production networks.

On the buyer side, U.S. importers exposed to affected suppliers suffer substantial declines in import value, both from directly hit partners and from unaffected suppliers. The contraction in imports reaches over 50% in the short run, with limited recovery over two and a half years. This contraction is not confined to trade volumes: importers reduce their number of suppliers and, among firms that also export, we observe a decline in their own export activity. The shock thus propagates not just along the supply chain, but also across the buyer's organizational footprint.

Importantly, we show that the magnitude and persistence of these effects vary sharply by organizational structure. Multinational importers are better able to absorb and recover from disruptions—particularly when the affected link is within the same corporate group. Intra-firm trade relationships exhibit smaller declines and faster rebounds, pointing to the importance of internal coordination and redeployment across affiliates. By contrast, arm's-length trade—whether conducted by MNCs or non-MNCs—is more exposed to persistent disruption. These results suggest that the organizational boundaries of firms—not just the geographic spread of their operations—play a central role in shaping supply chain resilience.

Viewed together, our findings highlight the dual nature of global production networks: while they allow firms to access diverse suppliers across borders, they also transmit shocks across organizational and transactional links. The heterogeneity in responses we document implies that firms' ability to navigate disruptions depends not only on the origin of the shock, but on how supply relationships are structured and embedded within broader corporate networks. These insights are especially relevant as firms and policymakers re-evaluate supply chain risks in a post-pandemic world.

Table 4: Difference-in-Differences Estimates by Importer Characteristics and Trade Relationship

	Import value	Export value	Import value excl. affected
Panel A: Importer is MNC			
Quarter 0–5 after shock — Non-MNC	-0.602$^{(a)}$	-0.207$^{(a)}$	-0.277$^{(a)}$
	(0.011)	(0.062)	(0.012)
Quarter 6–10 after shock — Non-MNC	-0.373$^{(a)}$	-0.089	-0.209$^{(a)}$
	(0.013)	(0.073)	(0.013)
Quarter 0–5 after shock — MNC	-0.350$^{(a)}$	0.020	-0.243$^{(a)}$
	(0.016)	(0.050)	(0.017)
Quarter 6–10 after shock — MNC	-0.210$^{(a)}$	0.117$^{(b)}$	-0.149$^{(a)}$
	(0.019)	(0.057)	(0.019)
Panel B: Disrupted Supplier is Intra-Firm			
Quarter 0–5 — Arm's-length	-0.540$^{(a)}$	-0.089$^{(b)}$	-0.265$^{(a)}$
	(0.010)	(0.043)	(0.010)
Quarter 6–10 — Arm's-length	-0.328$^{(a)}$	0.021	-0.189$^{(a)}$
	(0.011)	(0.050)	(0.011)
Quarter 0–5 — Intra-firm	-0.375$^{(a)}$	0.034	-0.282$^{(a)}$
	(0.034)	(0.100)	(0.036)
Quarter 6–10 — Intra-firm	-0.264$^{(a)}$	0.128	-0.206$^{(a)}$
	(0.041)	(0.110)	(0.042)
Panel C: Input Type (Differentiated vs. Homogeneous)			
Quarter 0–5 — Homogeneous input	-0.631$^{(a)}$	-0.103$^{(b)}$	-0.312$^{(a)}$
	(0.011)	(0.046)	(0.012)
Quarter 6–10 — Homogeneous input	-0.326$^{(a)}$	0.020	-0.189$^{(a)}$
	(0.013)	(0.054)	(0.013)
Quarter 0–5 — Differentiated input	-0.249$^{(a)}$	0.004	-0.142$^{(a)}$
	(0.016)	(0.076)	(0.018)
Quarter 6–10 — Differentiated input	-0.311$^{(a)}$	0.077	-0.192$^{(a)}$
	(0.019)	(0.085)	(0.020)
Observations	2,583,151	92,998	2,549,420
R^2	0.303	0.219	0.302
Fixed effects	Importer, industry × quarter		

Notes: This table reports difference-in-differences estimates of importer exposure to foreign supplier shocks, disaggregated by three dimensions: (A) whether the importer is part of a multinational group; (B) whether the disrupted supplier belongs to the same corporate group (intra-firm) or not (arm's-length); and (C) whether the traded input is classified as differentiated or homogeneous (Rauch (1999) liberal classification). Outcomes are measured as YoY midpoint growth. The sample includes recurrent importers. Standard errors clustered at the ADM2 level. Significance: $^{(a)}$ 1%; $^{(b)}$ 5%; $^{(c)}$ 10%.

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.

Sun, L. and S. Abraham (2021): "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics.*

Online Appendix

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

Vanessa Alviarez, Brian Fujiy, Anegel Espinoza, 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 Additional Results

A.1 Direct Impact on Foreign Suppliers by MNC Status

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

B 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

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

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

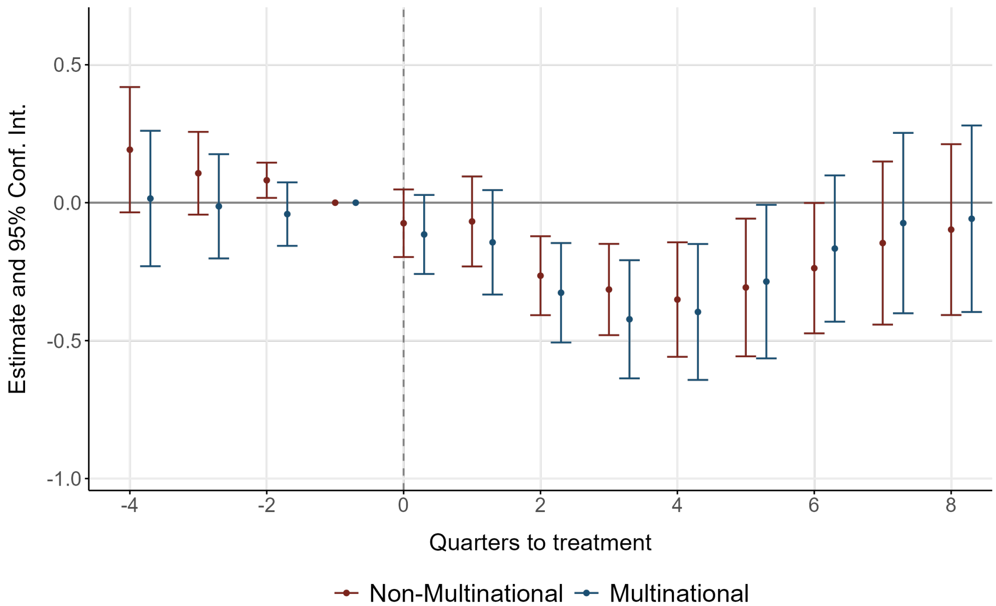

(b) Number of transactions to the U.S. (YoY midpoint growth)

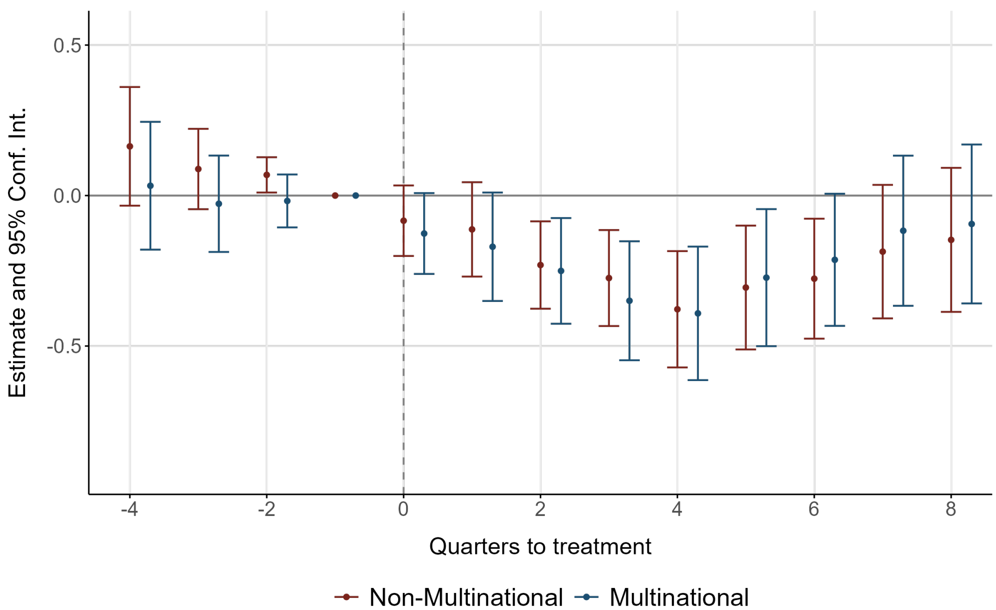

Notes: This figure plots the estimated effects of extreme natural disasters on directly affected foreign exporters, disaggregated by multinational status. Each panel reports coefficients from an event-study regression using quarterly year-on-year midpoint growth as the outcome. Panel A shows total export 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.

TABLE A.1: Difference-in-Differences Estimates by Exporter Ownership (MNC vs. Non-MNC)

	Export growth	Transactions growth
Quarter 0–5 × Non-MNC exporter	-0.227$^{(a)}$	-0.196$^{(a)}$
	(0.068)	(0.065)
Quarter 6–10 × Non-MNC exporter	-0.136$^{(b)}$	-0.127$^{(b)}$
	(0.058)	(0.053)
Quarter 0–5 × MNC exporter	-0.192$^{(a)}$	-0.170$^{(b)}$
	(0.071)	(0.066)
Quarter 6–10 × MNC exporter	-0.029	-0.041
	(0.067)	(0.061)
Fixed Effects		
Exporter	Yes	Yes
Industry × Quarter	Yes	Yes
Country × Quarter	Yes	Yes
Observations	1,030,609	1,031,607
R^2	0.065	0.077

Notes: This table shows heterogeneous difference-in-differences estimates by exporter ownership. Outcomes are defined as midpoint year-on-year growth in export value and transaction count. Multinational status is based on Orbis ownership links. The regression specification is:

$$y_{iqt} = \sum_{m \in \{\text{MNC, Non-MNC}\}} \left[\beta_{1m} \cdot \text{Post}^m_{[0,5]iqt} + \beta_{2m} \cdot \text{Post}^m_{[6,10]iqt}\right] + \alpha_i + \delta_{jt} + \theta_{ct} + \varepsilon_{iqt}$$

where $\text{Post}^m_{[0,5]}$ and $\text{Post}^m_{[6,10]}$ are post-shock indicators interacted with firm type $m \in \{\text{MNC}, \text{Non-MNC}\}$. The sample includes exporters with and without multinational ownership. All regressions include exporter fixed effects, industry-by-quarter fixed effects (2-digit HS), and exporter country-by-quarter fixed effects. Standard errors are clustered at the ADM2 level. Significance: $^{(a)}$ 1%; $^{(b)}$ 5%; $^{(c)}$ 10%.

Figure A.2: Effects of Foreign Shocks on Exports of U.S. Importers, by MNCs status

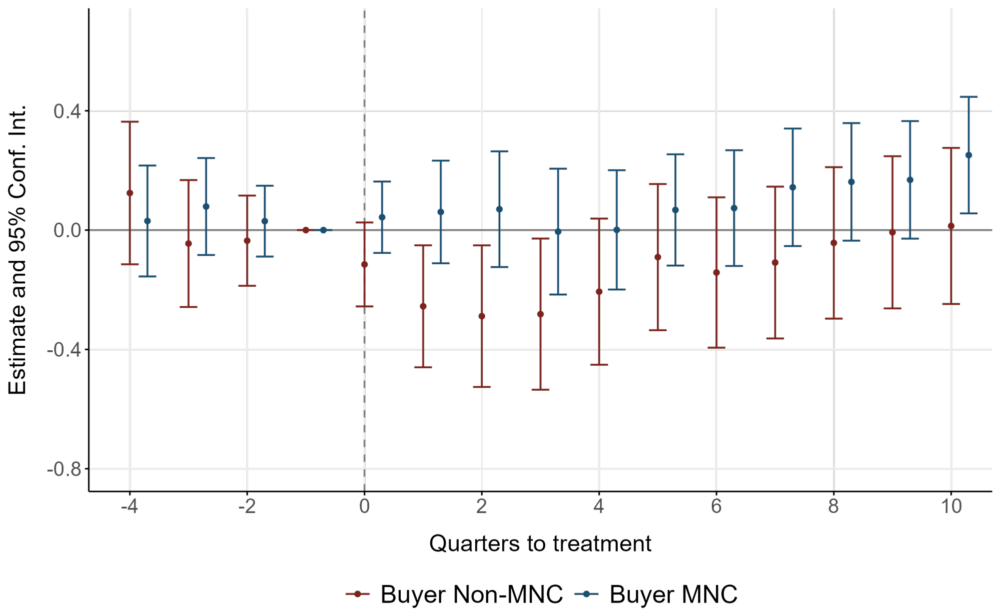

Notes: This figure plots the estimated effects of extreme natural disasters on export sales of U.S. importers, disaggregated by whether the affected supplier belongs to the same corporate group (intra-firm) or not (arm's-length). The sample includes recurrent importers, and treatment is defined based on exposure to a qualifying supplier affected by a disaster (ADM2-level match). Corporate group affiliation is identified using Orbis ownership links. All regressions include importer fixed effects, product-by-quarter fixed effects, and exporter-country–by–quarter fixed effects. Standard errors are clustered at the ADM2 level of the supplier.

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.

B.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.

C 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. [6]

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. [7]

Table C.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 C.1: Distribution of geocoded firms by address type (%)

	US Importers	Foreign Partners	Total
Full address	25.5	37.3	33.5
Route/City/Region/Postal Code/Country	59	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 contained in the firm's physical address used in the geocode process, for the US importers and the 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.

[6] 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 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.

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

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 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. [8]

[8] 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.