Where Imported Physical Climate Risk Stops: Supply Chains, Overseas Plants, and the Manufacturing Balance Sheet

Hyun Hak Kim*

September 1, 2026

Abstract

Physical climate risk realised abroad reaches a manufacturing economy through two channels: the inputs its firms buy and the plants they own. We measure both for Korean manufacturing and ask how far a shock travels. For the ownership channel we geocode 1,399 overseas plants and 1,844 sales offices of 474 listed manufacturers and evaluate daily ERA5 temperature, precipitation and wind at each site; for the trade channel we weight supplier-country hazards by embodied value added from international input-output tables. National-aggregate exposure measures overstate the heat that plants actually face by a factor of three to four in the hottest supplier economies and understate their flood exposure, because they average over territory with no production. Tracing the shock step by step, we find that severe floods and droughts at a firm's principal plant leave a trace on consolidated asset growth of about two percent, consistent with write-downs; they do not move profitability, capital expenditure, or plant survival, and each null is bounded. Through the trade channel a severe drought in the largest supplier reaches the input bill as 0.2 percent, trade with the stressed partner falls symmetrically in both directions, and the dispersion of marginal revenue products of capital across domestic establishments is unmoved, with effects above two to five percent ruled out. An earthquake placebo, the contrast between consolidated and separate statements, and the contrast between plants and sales offices separate a climate effect from disaster exposure in general, an overseas effect from parent-side shocks, and a production effect from mere presence abroad. Extreme heat is the one hazard on which the evidence divides, and plant-level tests do not resolve it.

Keywords: physical climate risk, global value chains, multinational firms, event study, capital misallocation, exposure measurement

JEL: F23, F64, Q54, D24, G32

*Professor, Department of Economics, Kookmin University. E-mail: hyunhak.kim@kookmin.ac.kr. The research was conducted as a component of the Ministry of Climate, Energy and Environment (MCEE)-funded Graduate School Specialized in Climate Change.

1 Introduction

A manufacturing economy is exposed to physical climate risk realised abroad in two ways. Its firms buy intermediate inputs from countries that are getting hotter, drier and more frequently flooded, and they own production plants in those same countries. Supervisors treat both as channels through which foreign disasters could impair domestic balance sheets, and both are now written into stress-testing frameworks. Whether either channel operates at a magnitude that matters is an empirical question, and answering it requires measuring exposure at the level where the shock actually lands and then following it, step by step, to see how far it travels.

This paper does that for Korean manufacturing over 2015–2024, and the answer is that the shock stops early. Physical damage at a firm's principal overseas plant leaves a trace on the consolidated balance sheet: after a severe flood or drought, consolidated asset growth is about two percentage points lower than separate-statement growth, a difference that appears in the consolidated accounts and not in the parent's own. It goes no further. Profitability, capital expenditure and plant survival do not respond, and each null carries a stated bound. Through the trade channel the same hazards reach the input bill diluted to a fifth of a percent, trade with the stressed partner contracts symmetrically in both directions, and the allocation of capital across domestic establishments is unmoved to within two to five percent of its dispersion.

Three things make these bounds informative rather than merely negative. The first is measurement. We geocode 1,399 overseas manufacturing plants and 1,844 sales offices from firms' subsidiary disclosures and evaluate daily reanalysis weather at each site. Comparing what plants experience with what national averages report, the averages overstate extreme heat by a factor of 3.6 in Thailand and 4.2 in Brazil and understate five-day rainfall extremes in China by more than half, because they weight territory with no production equally with the coast and the river basins where plants sit. Every study that measures foreign exposure at the country-year level, including the trade channel of this paper, inherits that error; our plant-level nulls hold at the resolution where it is absent.

The second is a placebo architecture that separates three confounds which usually travel together. Earthquakes destroy productive capacity exactly as storms do but have nothing to do with climate; consolidated statements include overseas subsidiaries while separate statements exclude them; sales offices sit in the same countries as plants but produce nothing. A climate effect must survive the earthquake comparison, an overseas effect must appear in consolidated and not separate accounts, and a production effect must appear at plants and not at sales offices. The asset-growth trace passes the second test; the null on profits and investment passes all three.

The third is arithmetic. Combining the observed foreign value-added share of Korean manufacturing (32.8 percent), the observed concentration of sourcing (the largest supplier holds 23.5 percent of the foreign share) and the observed trade response to a hazard (2.8 percent per standard deviation of drought), a severe drought in an industry's largest supplier moves its input bill by 0.22 percent. The null on capital allocation is not a failure to detect; it is what that number implies, and the production-network literature's amplification mechanisms would have had to be large to overturn it.

Extreme heat is the exception, and we report it as such. Its coefficient in the trade channel is positive and confined to the capital margin, but also confined to Chinese suppliers and sensitive to how the hazard is standardised; at plant level, where no standardisation is needed, it is absent, but Korean plants sit where days above 40°C are rare, so the plant-level test is weak by construction. We do not adjudicate.

The paper proceeds as follows. Section 2 places the contribution. Section 3 builds the two exposure measures and reports the measurement result. Section 4 sets out the event-study and panel designs and the placebo architecture. Section 5 follows the shock from the plant to the input bill to capital allocation. Section 6 treats heat. Section 7 asks why the shock stops, Section 8 states what the design cannot see, and Section 9 concludes.

2 Related Literature

Three literatures meet in this paper: the transmission of physical climate risk to firms, the propagation of shocks through production networks, and the measurement of capital misallocation. Our contribution sits at a point where the first two have not yet reached the third.

The closest study to ours is Pankratz and Schiller (2024), who match supplier locations to daily weather and show that realised heat and flood exposure depresses supplier performance and that customers respond by terminating relationships—by seven percent when realised exposure exceeds prior expectations. Their design establishes both that physical risk reaches firms and that firms adapt. We take the second finding as our starting point and ask what remains after the adaptation: if customers reallocate away from stressed suppliers, does anything survive to distort the allocation of capital in the customers' own economy? Our answer qualifies the premise as well as answering the question. The extensive-margin regressions of Section 7.2 reproduce their termination result at the industry-country level; but adding Korean exports to the same partners as a comparison shows the decline is not specific to imports, which is what a relationship termination would require. At this level of aggregation the adaptation their firm-level design identifies is not visible, and what looks like it is better read as contraction in the partner economy.

Blaum et al. (2026) study the same adaptation margin with an entirely different measure of risk. Using US transaction-level customs data and wave conditions along 331,000 maritime routes, they isolate the component of shipping time induced by weather and show that importers exposed to it diversify across more routes and suppliers while reducing import values. Their risk is what happens to goods *in transit*; ours is what happens at the *place of production*. The two are complementary and, notably, imply opposite adjustments: transit risk leads firms to add suppliers, while the production-location risk we measure is associated with fewer firms trading with the affected country.

At the aggregate end, the cross-border spillovers of physical risk are modelled in general equilibrium by International Monetary Fund (2021), and Cruz and Rossi-Hansberg (2024) embed warming in a spatial growth model. Both deliver magnitudes for whole economies but are silent on allocation across firms within an industry, which is the object of the misallocation literature and of this paper.

A larger literature measures how physical climate risk reaches firms directly, at the place where the firm itself operates, and it is worth being explicit about the outcomes it uses, because ours is different. Zhang et al. (2018) estimate temperature effects on output and factor reallocation across half a million Chinese manufacturing plants, and Somanathan et al. (2021) document losses in labour productivity and labour supply in Indian manufacturing, following Graff Zivin and Neidell (2014) on the allocation of time. Addoum et al. (2020) look at establishment sales for US public firms and find no significant effect. Tarsia (2025) estimates output, value added, capital, employment and firm exit for European firms and finds the damage concentrated among less productive ones, while Ponticelli et al. (2023) report effects on local industry concentration. The outcomes are, in short, levels: sales, output, productivity, employment, investment, survival. Because that is the convention and ours is not, we report every one of these outcomes that our microdata can construct—twenty-four in all, including shipments, value added, employment, capital, investment, labour productivity and concentration—against every hazard, in Section Appendix C. None of them moves.

That natural disasters propagate along supply links is well established. Barrot and Sauvagnat (2016) identify shocks from 28 major US disasters and find substantial output losses for customers of affected suppliers, concentrated where inputs are specific. Boehm et al. (2019) trace the 2011 Tōhoku earthquake through Japanese parent-affiliate links and document near-proportional transmission to US affiliates in the short run. Both use exposure measured at the level of the individual relationship, which is the resolution their questions require and the resolution we lack.

This matters for interpreting our results, and we take the point seriously enough to test it. Sections 5.1 and 5.2 measure exposure at plant coordinates for firms whose overseas subsidiaries we can locate, which is the resolution class of Pankratz and Schiller (2024), and find that what reaches the parent is a write-down on the consolidated balance sheet and nothing further. Aggregation alone therefore does not account for what we find.

Our outcome follows Hsieh and Klenow (2009), who infer distortions from the dispersion of marginal revenue products across establishments. The interpretation of that dispersion has since been substantially qualified. Bils et al. (2021) show that measurement error inflates estimated reallocation gains—by their estimates, overstating them by a factor of two or more—and Asker et al. (2014) show that adjustment costs generate dispersion even absent distortions, so that a high level of sd(log *MRPK*) need not indicate inefficiency. David and Venkateswaran (2019) decompose dispersion into informational and distortionary components and find both matter. The underlying premise, that policy and institutional frictions misallocate inputs across heterogeneous producers, is Restuccia and Rogerson (2008).

One study shares our outcome variable exactly. Liu and Xu (2025) regress the dispersion of log MRPK at the region-sector level on temperature bins for firms in thirty-two countries, and find that hot days raise it: capital cannot be reallocated within the year, so heterogeneity in firms' heat sensitivity shows up as dispersion in capital returns. Their exposure is domestic, measured where the firm operates. Ours is imported, measured where its suppliers operate. The contrast is the sharper statement of what this paper finds: the same dispersion measure that responds to local temperature does not respond to temperature realised in supplier economies. The null we report is therefore not a property of an insensitive outcome variable. It is a property of the channel.

These critiques bear on levels rather than responses. Since our specification includes industry fixed effects and asks whether dispersion *moves* with exposure, a constant measurement-error or adjustment-cost component is differenced out. We are careful throughout to make claims about the response and never about the level, and Section 4 states this as a maintained restriction rather than an assumption we can test.

A parallel literature asks how distortions aggregate. Baqaee and Farhi (2020) and Baqaee and Farhi (2019) characterise the aggregate consequences of microeconomic distortions in production networks, and Bigio and La'O (2020) show that financial frictions propagate along input-output links. Our exposure measure is built on the same value-added accounting, following Los et al. (2015), but we use it to construct treatment rather than to compute aggregate effects.

Korean work on climate and financial stability has largely been prospective and scenario-based. Hwang (2022) conduct physical risk stress testing of the Korean banking system, Chae et al. (2024) review the policy questions facing the Bank of Korea, and Chun and Kang (2024) and Kim (2024) assess climate exposures of Korean firms. These studies establish that Korean supervisors regard imported physical risk as material. None measures it using observed international sourcing structure, which is the gap the *FPR* measure of Section 3.1 is built to fill.

Against these literatures the paper's contribution is threefold. First, a measure of imported physical risk that any country with input-output and trade statistics can construct, together with the concordance work that makes it operational, and a plant-level counterpart built from disclosure data. Second, a null result with stated bounds: rather than reporting that no effect was found, we report what magnitude of effect each hazard's design would have detected, hazard by hazard, and defend those bounds with a geophysical placebo that separates a climate finding from a generic disaster-exposure finding. Third, evidence on what absorbs the shock: a sourcing structure diversified across eighteen to twenty-three countries and inventories concentrated where it is not, together with the finding that import-specific reallocation after a shock is absent to within two percent. Absorption here is structural rather than behavioural, and establishing that is what makes the null informative rather than merely negative.

3 Data and Measurement

Measuring a country's exposure to *foreign* physical climate risk requires answering two questions: which foreign locations does a domestic producer depend on, and how much climate stress did those locations experience? The existing literature has generally answered both at the firm level. Firm-to-firm supply data of the kind used by Barrot and Sauvagnat (2016) and Pankratz and Schiller (2024) do so directly, but such data exist for few countries and are rarely public.

This section builds an alternative that uses only public sources, at the cost of coarser resolution. It combines exposure weights from an international input-output system, physical risk from disaster and reanalysis archives, and a concordance that makes the two commensurable with Korean industrial statistics. Section 3.7 then describes a second, finer-resolution measure built from geocoded overseas plants, which the ownership-channel analysis of Section 5.1 uses.

3.1 The FPR Measure

Let $\omega_{s,c,t}$ denote the share of country c's value added embodied in one unit of output of Korean industry s in year t, and let $PhysRisk_{c,t}$ be a standardised measure of physical climate stress realised in country c in year t. Foreign physical risk exposure is their inner product over foreign origins,

$$FPR_{s,t} \;=\; \sum_{c \neq KOR} \omega_{s,c,t-1}\, PhysRisk_{c,t}. \tag{1}$$

Weights are lagged one year relative to the risk realisation. This is not a timing refinement but an identification requirement: contemporaneous weights would respond to the disaster itself—a country struck by a flood supplies less—so that $FPR_{s,t}$ would embed the very adjustment whose consequences we are trying to measure. With lagged weights, variation in $FPR_{s,t}$ for a given industry comes from weather realised in a supplier composition fixed before that weather was known.

Because $\sum_c \omega_{s,c,t} = 1$ including Korea itself, the weights entering (1) total the foreign value-added share $1 - \omega_{s,KOR,t-1}$, which varies across industries and over time. We renormalise them to sum to one, so that $FPR_{s,t}$ is the weighted average physical risk of an industry's foreign suppliers and is not mechanically larger for industries that happen to source more abroad. The level of foreign dependence enters the specification separately, as the GVC participation measure. The unnormalised form, in which exposure scales with foreign dependence, is reported in the robustness analysis of Section 5.4. Renormalisation also means the measure is defined only over origins for which risk data exist, which matters for one industry and is taken up in Section 3.6.

3.2 Value-Added Exposure Weights

The weights come from the Asian Development Bank Multi-Regional Input–Output tables (Asian Development Bank, 2025), covering 63 economies and 35 sectors for 2000 and annually for 2007–2024. Write A for the global matrix of intermediate input coefficients and v for the row vector of value-added shares in gross output. The value added of each origin embodied in each unit of final output is then given by the rows of $\hat{v}(I - A)^{-1}$ (Los et al., 2015). Taking the columns that correspond to Korean industries, and aggregating over the sectors of each origin economy, gives $\omega_{s,c,t}$.

One implementation detail governs whether the decomposition is correct. The value-added shares v can be read directly from the tables' value added row, or derived from the input coefficients as $v = \mathbf{1} - \mathbf{1}'A$. These are equal in an internally consistent table, but in the published tables they are not. Taking the value added row at face value produces columns of $\hat{v}(I - A)^{-1}$ that sum to 0.98 rather than 1—a two percent leakage that would be silently absorbed into the weights. Defining v residually from A enforces the accounting identity and yields column sums equal to one to eight decimal places in all 19 years. We adopt the residual definition throughout. One sector (c35) has zero Korean gross output and is dropped; its column is identically zero by construction.

Table 1 reports the resulting foreign value-added share of Korean manufacturing output, which ranges from 28 to 41 percent with a peak in 2011–2012 and a trough in 2020. This is the quantity that makes Korea an informative case: a manufacturing sector sourcing roughly a third of its embodied value added abroad, in an economy whose own territory sees comparatively little climate disruption.

Where that foreign value added originates matters as much as how much of it there is, since a hazard can only propagate through the origins an industry actually depends on. Figure 1 maps the weights aggregated over manufacturing sectors for the most recent five years. Exposure is concentrated but not dominated: China contributes the largest single share at 19 percent of foreign value added over 2020–2024, the United States 14 percent and Japan 7 percent, with the remainder spread across some sixty economies. Averaged over the full estimation period the corresponding shares are 16, 11 and 9 percent, China having gained and Japan lost ground across the sample. This distribution is what the absorption argument of Section 7 ultimately rests on, and Section 5.3 converts it into a magnitude: no single origin holds a large enough share for a disaster there to move an industry's input bill appreciably.

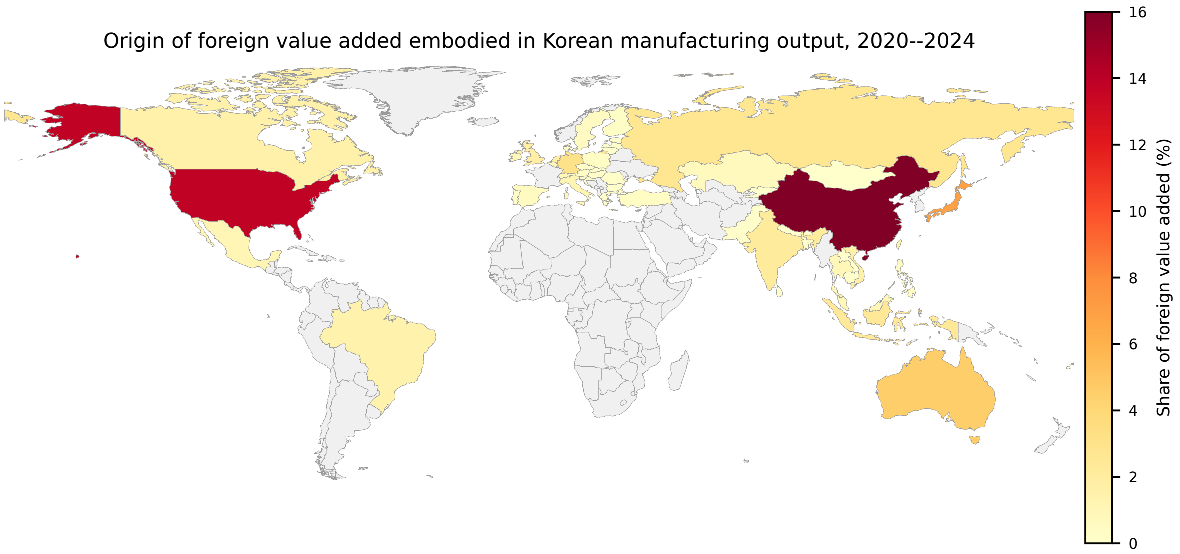

Figure 1: Origin of foreign value added embodied in Korean manufacturing
Notes. Each economy is shaded by its share of the foreign value added embodied in Korean manufacturing gross output, averaged over 2020–2024 and across sectors c3–c16, computed from the Leontief decomposition described in this section. Shares sum to one hundred across all foreign origins. Economies in grey fall inside the MRIO rest-of-world aggregate and cannot be attributed individually.

3.3 Physical Risk Indicators

Which hazards should enter? The literature offers two templates. One pools events into a single disaster indicator: Barrot and Sauvagnat (2016) identify shocks from 28 major US disasters spanning hurricanes, blizzards, floods, wildfires and earthquakes, without distinguishing among them. The other selects hazards with a clear transmission mechanism: Pankratz and Schiller (2024), the study closest to ours in design, measure supplier exposure to *heat and floods* and to nothing else. Supervisory frameworks are more expansive still, with the NGFS acute-hazard set covering heat, drought, flood, tropical cyclone, wildfire and landslide.

We follow the second template but widen it to the NGFS acute categories that country-year data can support. Pooling hazards would make it impossible to say which mechanism operates; reporting only two would leave the reader unable to judge whether a result

Table 1: Foreign value-added share of Korean manufacturing output

Year	Share	Year	Share	Year	Share
2000	0.305	2009	0.352	2017	0.365
2007	0.325	2010	0.367	2018	0.313
2008	0.397	2011	0.411	2019	0.312
		2012	0.407	2020	0.282
		2013	0.383	2021	0.336
		2014	0.367	2022	0.380
		2015	0.379	2023	0.343
		2016	0.358	2024	0.327

Notes. $1 - \omega_{s,KOR,t}$ averaged over manufacturing sectors (ADB MRIO sectors c3–c16), weighted by gross output. Author's calculations from Asian Development Bank (2025). Column sums of $\hat{v}(I - A)^{-1}$ equal 1 to eight decimal places in every year; 2001–2006 are not published.

is hazard-specific or generic. Table 7 accordingly reports heat, flood, drought, storm, wildfire and landslide, together with two composite measures.

Two data systems supply these, with opposite weaknesses: disaster reports record damage but under-report it where reporting capacity is weak, while reanalysis records weather uniformly but is silent on whether anything was damaged. We use both and report each.

Weather measures come from the World Bank Climate Change Knowledge Portal (World Bank, 2025), which aggregates ERA5 reanalysis to country-year frequency for 1950–2023. Heat enters as days above 35°C and days above 40°C, flood as maximum five-day precipitation and days with precipitation above 20 mm, drought as consecutive dry days. For heat we also construct the annual maximum temperature, which contrasts instructively with the threshold measures: cumulative exposure above a threshold and the temperature of the single hottest day are different constructs, and Section 5.4 finds they behave differently. Because levels differ across climates in ways that say nothing about disruption, all weather measures enter as anomalies relative to each country's own 1961–1990 baseline, standardised across the panel.

Storms require separate treatment, because neither data source above measures wind. High winds damage plant through a different physical mechanism than precipitation, and the damage is steeply non-linear in wind speed—the dissipated power of a cyclone scales with the cube of its maximum wind. A precipitation index cannot represent this, and an event count treats a category-one storm and a category-five storm identically. The CCKP collection contains no wind variable.

We therefore build country-level storm exposure from the IBTrACS best tracks used in Section 3.7, assigning each track point to countries whose boundary lies within 200 km, since a cyclone damages a wide swath without making landfall. For each country-year this yields storm counts at tropical-storm and typhoon strength, maximum wind speed, accumulated cyclone energy ($\sum V^2$), and the power dissipation index ($\sum V^3$), the last being the closest available proxy for wind damage. The resulting climatology is a check on the assignment: mean annual storm counts of 11.4 for Japan, 9.3 for China, 7.8 for the Philippines and 4.2 for Taiwan reproduce observed tropical cyclone frequencies. Having five measures that differ only in how heavily they weight intensity also provides a diagnostic, used in Section 5.4, that a single index would not.

Disaster measures come from EM-DAT (Delforge et al., 2025), covering 223 countries over 1990–2025. We use event counts by type (flood, storm, drought, extreme temperature, wildfire, wet mass movement) and two composites: total climate-related event counts and the affected population as a share of population. Reported monetary damage is available for only 36 percent of events, against 77 percent for affected population, and the missingness concentrates in low-income countries and smaller events—exactly the cells whose exposure matters here—so damage enters only as a secondary check.

EM-DAT also records geophysical disasters, which we retain deliberately. Earthquakes are unrelated to climate but disrupt production through the same physical mechanism as storms and floods—damaged plant, severed transport—which makes exposure to them a placebo for the climate measures. If our exposure construction were picking up something generic about supplier countries rather than climate itself, earthquake exposure should behave like climate exposure. Section 5.4 reports this test.

Vulnerability indicators come from ND-GAIN (Notre Dame Global Adaptation Initiative, 2026), which we use to construct an alternative weighting in which the same physical event counts more where adaptive capacity is lower. We use the time-varying components (vulnerability, readiness, sensitivity, capacity) and not the exposure sub-index, which is a projection from climate models and is constant over the sample period.

3.4 Concordance

The three data systems classify countries and industries differently, and reconciling them is a substantial part of the measurement contribution. Korea's Trade by Enterprise Characteristics statistics, which we use for validation and for the analysis in Section 7, report partner countries in a bespoke scheme that mixes individual countries with regional aggregates and reports both a parent aggregate and its members. Summing the published cells therefore double-counts.

We construct a mutually exclusive and exhaustive partition of 67 cells: 60 individual countries, 6 regional residuals defined by subtracting enumerated members from their parent aggregate, and one global residual. Of these, 48 map one-to-one onto MRIO economies; the remainder map onto the MRIO rest-of-world aggregate. Building the partition surfaced one inconsistency worth recording: the published Eastern Europe aggregate exceeds the sum of its stated members by 0.36 to 0.75 percent, because several member countries are also counted within the EU-28 aggregate. We exclude that cell from the partition rather than allocate the overlap by assumption.

Industries are bridged across three classifications—the seven TEC manufacturing groups, the 35 MRIO sectors, and two-digit KSIC—at the finest level that admits a clean correspondence. Because KSIC was revised during the sample (8th revision through 2006, 9th revision and later thereafter), the bridge maps each vintage's divisions directly to MRIO sectors rather than chaining through a KSIC-to-KSIC concordance.

3.5 Dependent Variables

Capital misallocation is measured by the dispersion of marginal revenue products following Hsieh and Klenow (2009). For establishment i in industry s,

$$MRPK_i = \alpha_s \frac{VA_i}{K_i}, \qquad TFPR_i = \frac{VA_i}{K_i^{\alpha_s} L_i^{1-\alpha_s}}, \tag{2}$$

where K_i is the average of beginning- and end-of-year tangible fixed assets, L_i is the annual wage bill, and the capital share $\alpha_s = 1 - \sum_i W_i / \sum_i VA_i$ is computed within each industry-year and bounded to $[0.1, 0.9]$. The outcome variables are sd(log $MRPK$) and sd(log $TFPR$) within each industry-year cell, each winsorised at one percent, together with the implied Hsieh and Klenow (2009) reallocation gain.

The source is the Mining and Manufacturing Survey, a census of establishments with ten or more employees conducted by Statistics Korea, for 2002–2024. Working with the microdata required resolving column definitions that change almost every year; two problems were consequential enough to state. First, tangible fixed assets and inventories are both reported as beginning- and end-of-year balances, and in the 2011 and 2012 files the column headings do not distinguish them—the relevant totals are labelled only by the number of components they sum. We identify the capital block by the accounting identity end = begin + additions − disposals − depreciation, which holds for 59 of 60 sampled establishments in the block we select and fails systematically in the alternative. Second, two years are dropped. The 2010 census file contains no tangible asset columns at all. In 2015 the missing rate for tangible assets is 37 percent against 3–4 percent in adjacent years, and the gaps are concentrated entirely in two establishment-type codes rather than scattered; retaining the year on its available subsample would therefore induce selection. The estimation sample comprises 1.11 million establishment-year observations across 21 survey years, aggregating to 484 KSIC-division-year cells and, after mapping to MRIO sectors, to the 210 industry-year observations (14 sectors × 15 years) used in the main regressions.

Table 2 reports summary statistics for the estimation panel.

Table 2: Descriptive statistics, industry-year panel

	N	Mean	S.D.	Min	Max
Outcomes					
sd(log $MRPK$)	210	1.549	0.303	0.768	2.296
sd(log TFPR)	210	1.158	0.178	0.744	1.622
Reallocation gain	210	1.056	0.502	0.319	3.094
Exposure					
FPR (disaster count)	210	0.217	0.395	−0.466	1.362
FPR (drought)	196	0.431	0.380	−0.284	1.497
FPR (flood)	196	0.053	0.373	−0.860	0.791
FPR (affected population)	210	−0.461	0.597	−2.645	0.531
Controls					
GVC backward participation	210	0.339	0.111	0.187	0.765
Establishments per cell	210	3,892	3,232	33	11,380

Notes. MRIO manufacturing sectors c3–c16 × 2008–2024 excluding 2010 and 2015. FPR measures are constructed as in equation (1) from standardised country-year risk indicators; the drought and flood measures lose one MRIO sector-year block where the underlying CCKP series ends in 2023.

3.6 Validation

The exposure weights are model-based, and the natural check is whether they reproduce what Korean firms actually import. The TEC statistics report import values by industry and partner country from customs records, so for each industry-year we can compare the ranking of origins implied by ω against the ranking implied by observed import shares. The two are conceptually distinct—value added embodied in output attributes to the country where value was created, whereas customs data record the country of last shipment—so agreement should be strong but not perfect, and rank correlation is the appropriate statistic.

Table 3 reports the comparison, restricted to the partition cells with individual MRIO correspondence and renormalised within that set. Rank correlations range from 0.75 to 0.92 across industries and are stable over time (0.80–0.88 by year), and the top five origins agree on average in 3.8 of 5 positions. Coverage of the individually matched cells is 83–84 percent of import value in six of seven industries.

Coke, petroleum and chemicals is the exception, at 44 percent. The reason is structural rather than a defect of the concordance: Korea's crude oil imports originate overwhelmingly in Gulf producers that the MRIO treats within its rest-of-world aggregate and thus cannot be matched individually. The rank correlation for that industry remains high (0.87) among the origins that are matched, but its FPR series describes a minority of its import value, and we treat it separately in the robustness analysis.

Table 3: Validation of exposure weights against customs data

Industry group	Coverage of imports	Rank corr. ρ	Top-5 agreement
Food, beverages, tobacco	84.2%	0.918	4.0/5
Textiles, leather	92.6%	0.806	2.8/5
Wood, paper, printing	83.5%	0.847	3.8/5
Coke, petroleum, chemicals	44.3%	0.872	4.0/5
Non-metallic minerals, metals	83.6%	0.896	4.0/5
Electrical, electronic, machinery	98.0%	0.821	3.9/5
Transport equipment, other	97.2%	0.746	4.0/5

Notes. Averages over 2010–2024. Comparison is between $\omega_{s,c,t}$ from the MRIO decomposition and import shares from Statistics Korea's Trade by Enterprise Characteristics, restricted to the 48 partition cells with one-to-one MRIO correspondence and renormalised within that set. Coverage is the share of the industry's import value falling in those cells. Spearman rank correlation across origins.

3.7 Overseas Plant Coordinates

The measure described so far is a country-by-industry annual aggregate, and its resolution is its principal limitation. For the ownership channel studied in Section 5.1 we can do considerably better, because Korean listed firms disclose the address of each consolidated subsidiary.

From the annual reports of 2,240 listed manufacturers we extract the standardised subsidiary table, which reports each subsidiary's name, address, principal business, and total assets. Filtering to subsidiaries whose stated business indicates production and whose address lies outside Korea yields 2,269 records with 1,859 distinct addresses.

Geocoding these addresses is harder than it first appears, and the difficulty is instructive about what disclosure data can support. Naive submission of the raw strings to a gazetteer resolves 48 percent to city level or better. The failures are not random: Korean filings render Chinese and Vietnamese place names in Korean transliteration or Sino-Korean readings, and industrial-park names that are locally canonical are absent from open gazetteers. Three corrections raise the resolution rate to 92 percent on a 200-address development sample: a transliteration dictionary for Chinese, Vietnamese and other Asian place names; sliding-window queries over address tokens, which recover cases where the city name sits in the middle of the string rather than at its end; and a mapping of recurring industrial-park names to their host cities. Applied to all 1,859 addresses the pipeline resolves 73 percent to city level or better. The gap to the development sample reflects park names encountered only once and addresses that disclose nothing below the province.

Two classes of *apparent* successes had to be removed, and they matter more than the failures because they are silent. Ninety-six records carry romanised Korean addresses that the industry filter did not catch; the gazetteer resolves them correctly, but to Korea, so they do not belong in a measure of foreign exposure. A further 96 resolve to a country that contradicts a country name stated in the address itself—an address reading only `U.S.A` matching a street of that name in the Philippines, a Vietnamese address matching a building in Canada. We drop both classes. The final panel comprises 1,399 plants with city-level or better coordinates, in 55 countries, belonging to 474 parent firms.

Physical risk at these coordinates is measured with tropical cyclone best tracks from IBTrACS (Knapp et al., 2010), using the 2,775 storms reaching tropical-storm strength ($\geq$ 34 knots) over 1995–2025; a plant-year is exposed if a track point of such a storm passes within a given radius. Tropical cyclones are the hazard for which coordinates buy the most, since exposure within a country varies sharply with coastal proximity. The resulting geography is itself a check on the coordinates: exposure rates of 79 percent in the Philippines and 78 percent in Japan, against 34 percent in China, 0.3 percent in near-equatorial Indonesia, and zero in inland Europe, reproduce tropical cyclone climatology without any tuning.

Cyclone tracks are one hazard; the others require gridded daily weather at each site. We evaluate ERA5-Land temperature (0.1°) and ERA5 precipitation and ten-metre wind (0.25°) at every plant, sales office and exited plant for 2015–2024 through the Open-Meteo archive, which serves the native reanalysis grids as point queries. To confirm that the point service reproduces the grids, we downloaded the ERA5-Land daily statistics for the Chinese sites directly from the Copernicus Climate Data Store and compared the two series day by day: the correlation is 1.000 and the root-mean-square difference 0.03°C, which is rounding. From the daily series we build, for each site-year, the ETCCDI-style indicators used throughout—days above 35°C and 40°C, the annual maximum and minimum, maximum five-day precipitation, days above 20 mm, consecutive dry days, total precipitation, and maximum wind and gust—so that plant-level and country-level hazards share definitions and differ only in where they are evaluated. The same series underlie the sales-office placebo of Section 4.3, which uses 1,844 offices of the same parents, and the exit analysis of Appendix F.2, which adds 176 plants that had left the filings by 2024 and were geocoded from their 2022 addresses.

3.8 National Aggregates versus Plant Coordinates

Having both a country-year and a plant-coordinate measure of the same hazards makes it possible to ask how much the country-year convention misstates what firms actually face. Table 4 compares, for the twelve countries hosting the most Korean plants, the hazard averaged over plant sites with the national average published by the World Bank Climate Change Knowledge Portal from the same ERA5 reanalysis.

The two disagree systematically, and in opposite directions for heat and for rain. National averages overstate extreme heat wherever Korean plants cluster in the temperate

Table 4: What plants experience versus what national averages report

		Days > 35°C		Rx5day (mm)		CDD (days)		
Country	Plants	Plant	National	Plant	National	Plant	National	Heat ratio
China	434	5.8	7.0	145	59	30	47	1.2
United States	228	9.9	14.6	94	45	30	30	1.5
Vietnam	194	16.5	14.2	169	135	29	27	0.9
India	78	74.7	81.1	173	132	66	67	1.1
Indonesia	46	2.7	2.3	141	118	19	12	0.9
Mexico	39	31.4	38.4	87	68	50	62	1.2
Germany	26	0.6	0.6	56	50	19	17	1.1
Poland	25	0.1	0.4	64	50	18	18	3.5
Japan	24	0.8	0.8	180	144	17	13	1.1
Malaysia	19	0.0	4.2	148	124	9	8	—
Brazil	19	7.7	32.5	116	90	30	55	4.2
Thailand	15	14.3	51.4	124	113	31	41	3.6

Notes. Annual values 2015–2024, averaged over years. "Plant" is the mean over Korean-owned manufacturing plants in the country, each evaluated at its coordinates from ERA5-Land (0.1°, temperature) and ERA5 (0.25°, precipitation). "National" is the country-year value of the Climate Change Knowledge Portal, also from ERA5. Heat ratio is national over plant for days above 35°C.

or coastal parts of a hot country: the ratio is 3.6 for Thailand, where plants sit around Bangkok and the eastern seaboard rather than the interior, and 4.2 for Brazil. They understate extreme rainfall wherever plants sit on the wet coast of a country with a dry interior: China's national five-day maximum is 59 millimetres against 145 at its plants, and the United States' 45 against 94. The discrepancy is not noise; it is the geography of where production is. Any exposure measure built from country-year data, including the value-added-weighted measure of Section 3.1, carries it. For the trade channel the bias works against finding an effect on heat and towards finding one on flood; the results of Section 5 should be read with that in mind, and the plant-level results are free of it.

4 Empirical Strategy

4.1 The Ownership Channel: Events at the Principal Plant

A firm with several overseas plants is exposed to a disaster at any one of them only in proportion to that plant's weight, so an asset-weighted average of hazards across plants dilutes the shock before it is measured. We therefore define exposure as an event at the firm's *principal* plant, the overseas manufacturing subsidiary with the largest disclosed assets, and treat it as a binary indicator:

$$Hit^h_{i,t} = \mathbf{1}\Big[\text{the principal plant of firm } i \text{ experiences a severe event of type } h \text{ in year } t\Big]. \tag{3}$$

Severe events are defined on absolute physical thresholds: a tropical cyclone passing within 100 km with maximum sustained winds of at least 64 knots; at least 45 days with maximum temperature above 35°C; five-day precipitation of at least 300 mm; and at least 90 consecutive dry days. The thresholds were fixed before estimation by requiring an annual event frequency between two and ten percent, so that events are rare enough to be severe and common enough to be estimable; they yield 5.2, 4.9, 3.6 and 3.9 percent of firm-years respectively, with 59 to 173 distinct firms affected. Relative thresholds (a plant's own historical percentile) are not used: with ten years of data a 95th percentile is mechanically the top year, so every hazard would have exactly ten percent frequency by construction. Typhoons and floods coincide in 49 percent of firm-years, almost all of them in 2024 when Typhoon Yagi crossed northern Vietnam and southern China, so we also define a compound indicator.

The outcome is the difference between what the consolidated statements report and what the separate statements report for the same firm and year,

$$D_{i,t} = y_{i,t}^{\mathrm{CFS}} - y_{i,t}^{\mathrm{OFS}}, \tag{4}$$

for y equal to the change in log total assets, return on assets, capital expenditure over lagged assets, the change in log revenue, and the tangible-asset share. Consolidated statements include overseas subsidiaries and separate statements exclude them, so the difference is the net contribution of the overseas operations and parent-side variation cancels. We report the two components as well.

The specification is an event study with two-way fixed effects,

$$D_{i,t} = \sum_{k=-1}^{2} \beta_k \, Hit_{i,t-k}^{h} + \delta_i + \delta_{c(i)\times t} + \varepsilon_{i,t}, \tag{5}$$

where δ_i is a firm effect and $\delta_{c(i)\times t}$ is an effect for the principal plant's host country interacted with year. The latter absorbs everything that happens to a country in a year—exchange rates, demand, policy—so identification comes only from the contrast between plants that were hit and plants in the same country that were not. Because events are transient rather than absorbing, no reference year is omitted and β_{-1} is a direct test of pre-trends. A static version with $k = 0$ only uses the full 2015–2024 window; the dynamic version uses 2017–2023. Standard errors are clustered by firm (474 clusters) and p-values come from a wild cluster bootstrap with the null imposed, 9,999 draws. We report the minimum detectable effect for every coefficient and apply the Benjamini–Hochberg procedure across the full set of tests.

Two further specifications discipline the interpretation. Replacing the principal-plant indicator with an indicator for an event at *any* of the firm's plants tests whether a response is specific to the principal plant or reflects a network-wide reaction; and estimating equation (5) on the firm's principal *sales office* instead of its principal plant provides the placebo described in Section 4.3.

4.2 The Trade Channel: Value-Added-Weighted Exposure

The main specification relates the dispersion of marginal revenue products within an industry to that industry's exposure to foreign physical risk,

$$\sigma_{s,t} \;=\; \beta\, FPR_{s,t} \;+\; \delta_s \;+\; \delta_t \;+\; \varepsilon_{s,t}, \tag{6}$$

where $\sigma_{s,t}$ is sd($\log MRPK$) in industry s and year t, δ_s and δ_t are industry and year fixed effects, and $FPR_{s,t}$ is standardised so that β is the response to a one standard deviation increase in exposure. Industry fixed effects absorb permanent differences in the level of measured dispersion across industries, which is essential given that the level of sd($\log MRPK$) reflects technology and measurement as much as misallocation (Bils et al., 2021). Year fixed effects absorb the global business cycle and any worldwide climate shock common to all suppliers. Identification therefore comes from an industry's exposure moving differently from other industries' in the same year, which happens because industries source from different countries and those countries experience different weather.

The panel has 210 observations—14 MRIO manufacturing sectors over 15 years—and clustering is by industry, giving 14 clusters.

Fourteen clusters is few, and cluster-robust standard errors are downward-biased in that range, so conventional tests over-reject (Cameron et al., 2008). We therefore report, alongside asymptotic p-values, the wild cluster restricted bootstrap with the null imposed and Rademacher weights. Because $2^{14} = 16{,}384$ sign vectors is a tractable number, we enumerate the full set rather than sampling from it, which makes the bootstrap p-value exact rather than simulation-based. The size-bin specification below has $G = 7$ and is enumerated in full as well ($2^7 = 128$).

This choice is not cosmetic. As Section 5.4 reports, the one coefficient in the entire analysis that is significant under asymptotic inference is not significant under the bootstrap.

Equation (6) has a weakness that the fixed effects create. Once year effects are removed, the residual variation in $FPR_{s,t}$ is small, because the physical risk realised in a given year is correlated across the supplier countries that different Korean industries draw on. We quantify this in Section 5.4; it is the binding constraint on precision.

To obtain variation that survives a more demanding set of controls, we build exposure at the industry-by-size-bin level. Establishments of different sizes within the same industry source from different countries, and Korea's Trade by Enterprise Characteristics statistics report import values by industry and by firm size separately. We recover the joint distribution by iterative proportional fitting of the two margins, which delivers $FPR_{s,k,t}$ for size bin k, and estimate

$$\sigma_{s,k,t} \;=\; \beta\, FPR_{s,k,t} \;+\; \delta_{s\times t} \;+\; \delta_k \;+\; \varepsilon_{s,k,t}. \tag{7}$$

Industry-by-year fixed effects here absorb everything common to an industry in a year, including any global shock and any industry-specific trend. Identification comes only from differences across size bins within the same industry and year. The cost is that $FPR_{s,k,t}$ is an approximation—the raking procedure assumes no three-way interaction among industry, size, and origin—and that the panel covers 2016–2024 with seven TEC industry groups, so $G = 7$ clusters.

One further matter belongs here, since it governs how the estimates should be read. A paper reporting mostly nulls invites the suspicion that something was withheld, and a paper reporting one significant coefficient among many invites the opposite suspicion. We address both by reporting the full grid rather than a selection: every hazard measure constructed in Section 3.3 is estimated against every outcome, and Table 7 shows all of them. Nothing was estimated and set aside.

With nineteen hazard measures, some significant coefficient is expected by chance, so we report multiple-testing corrections alongside the bootstrap p-values and treat a result as informative only if it survives them. We also apply two additional filters that a chance finding is unlikely to pass: consistency across the three outcome variables, and the geophysical placebo of Section 3.3. A coefficient that is significant on one outcome, fails the correction, or is matched by an equally significant earthquake coefficient is not evidence of anything.

4.3 Placebo Architecture

A null result on climate exposure is open to three objections that are usually inseparable: that the design would not detect any disaster effect, climate or otherwise; that whatever moves the outcome is a parent-side shock correlated with where plants happen to be; and that overseas presence of any kind, not production, is what matters. The three comparisons in Table 5 hold one of these fixed at a time.

Earthquakes damage plants and infrastructure through the same physical mechanism as storms and floods but are unrelated to climate; if climate hazards showed an effect and earthquakes did not, the effect would be about climate, and if neither showed one, the null would be about disaster exposure in general. Consolidated and separate statements are filed by the same firm for the same year; an effect of an overseas event must appear in the former and not the latter. Sales offices are located in the same countries and hit by the same weather as plants but produce nothing; an effect of production disruption must appear at plants and not at sales offices.

Table 5: Three placebo comparisons and what each removes

Comparison	What it holds fixed	Confound it removes	Result
Earthquakes vs. climate hazards	Physical destruction of productive capacity	Any effect of disaster exposure as such	Null at equal precision; climate hazards are not different from earthquakes
Consolidated vs. separate statements	Same firm, same year	Parent-side shocks correlated with where plants are located	Damage signal appears in consolidated, not separate
Manufacturing plants vs. sales offices	Same firm, same country, same hazard	Any "overseas presence" effect unrelated to production	Null at sales offices; typhoon at equal power

Notes. The earthquake comparison is estimated on the trade channel (Section 5.4); the other two on the ownership channel (Sections 5.1–5.2).

5 Results: How Far the Shock Travels

We follow the shock from the point of impact outward. The ownership channel comes first, because it is the channel with the finer resolution and the more direct link from damage to accounts; the trade channel follows.

5.1 Damage Reaches the Consolidated Balance Sheet

Table 6 reports the static coefficient β_0 from equation (5) for each hazard and each outcome, on the consolidated-minus-separate difference and on its two components. Figure 2 shows the dynamic coefficients for the three hazards with enough events to estimate them.

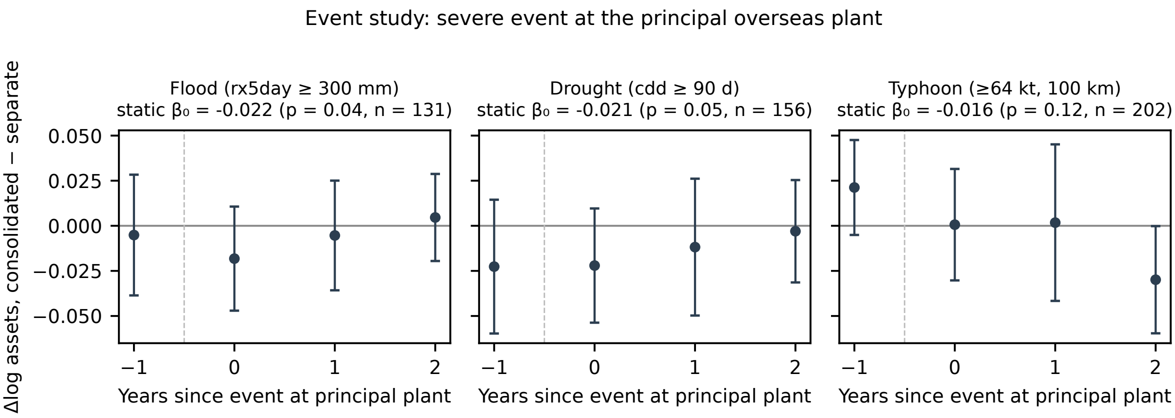

Figure 2: Dynamic event study: change in log assets, consolidated minus separate
Notes. Coefficients β_k from equation (5) with 95 percent confidence intervals from firm-clustered standard errors. $k = -1$ is the year before the event. The dashed line separates pre- from post-event years.

One pattern stands out. After a severe flood or drought at the principal plant, consolidated asset growth is about two percentage points lower than separate-statement growth

Table 6: Severe events at the principal overseas plant and parent-firm outcomes

Event at principal plant	Δ log assets	ROA	Capex / lagged assets	Events
Consolidated − separate				
Flood	−0.0215** (0.04)	−0.0151* (0.07)	−0.0270* (0.08)	119
Drought	−0.0210** (0.05)	−0.0053 (0.29)	+0.0009 (0.83)	152
Typhoon	−0.0155 (0.12)	−0.0019 (0.89)	−0.0188 (0.13)	183
Heat	+0.0005 (0.97)	+0.0083 (0.88)	+0.0003 (0.98)	192
Typhoon ∩ flood	−0.0398* (0.09)	−0.0389* (0.07)	−0.0579 (0.22)	64
Consolidated				
Flood	−0.0195 (0.33)	−0.0097 (0.43)	+0.0114 (0.94)	123
Drought	−0.0453* (0.06)	+0.0020 (0.84)	+0.0043 (0.62)	152
Typhoon	−0.0154 (0.43)	−0.0071 (0.63)	+0.0136 (0.57)	186
Heat	−0.0198 (0.40)	+0.0082 (0.49)	−0.0167 (0.13)	192
Typhoon ∩ flood	−0.0146 (0.64)	−0.0139 (0.54)	+0.0369 (0.90)	67
Separate				
Flood	+0.0169 (0.46)	+0.0040 (0.79)	+0.0368 (0.65)	123
Drought	+0.0010 (0.95)	+0.0134 (0.28)	+0.0032 (0.66)	156
Typhoon	+0.0025 (0.90)	−0.0058 (0.69)	+0.0303 (0.23)	190
Heat	−0.0293 (0.21)	+0.0018 (0.92)	−0.0157 (0.40)	193
Typhoon ∩ flood	+0.0258 (0.45)	+0.0204 (0.52)	+0.0892 (0.79)	66

Notes. Static β_0 from equation (5) with firm and host-country-by-year fixed effects; wild cluster bootstrap p-values in parentheses (474 firm clusters, 9,999 draws). Outcomes are the consolidated-minus-separate difference and its components. N = 3,500–3,975 firm-years. ** $p < 0.05$, * $p < 0.10$. Of 375 coefficients, 24 reach five percent against 18.8 expected under the null; none survives Benjamini–Hochberg at $q = 0.05$.

(−0.022, $p = 0.043$; −0.021, $p = 0.047$). The sign of the components is what an overseas effect requires: consolidated growth falls and separate growth does not, so the difference is not a parent-side shock. It is specific to the principal plant—defining exposure as an event at any of the firm's plants, with two to three times as many events, gives coefficients indistinguishable from zero—and it is not stronger in single-plant firms, where dilution would predict it to be, but in firms with several plants. The typhoon and compound coefficients have the same sign and are less precise; heat shows nothing.

We report this as a trace rather than a finding. Neither coefficient survives correction for the 375 tests in the table, the drought pre-trend coefficient is as negative as the event coefficient (Figure 2), and the compound indicator has a significant pre-trend in revenue. Only the flood pattern is clean. What the next two outcomes establish is what kind of trace it is.

5.2 It Stops Before Profits, Investment and Survival

If a two-percent decline in consolidated asset growth reflected lost production, it should show in profitability; if it reflected a decision to invest less at the damaged site, it should show in capital expenditure; if it reflected abandonment, it should show in plant exit. It shows in none of them.

Return on assets does not move. Pooling all four hazards into a single severe-event indicator gives a consolidated-minus-separate coefficient of -0.0005 ($p = 0.96$); by hazard the largest estimate is the flood coefficient of -0.015 ($p = 0.07$), which appears with the same magnitude when the event is at a sales office rather than a plant and is therefore not a production effect. Capital expenditure over lagged assets—the investment decision itself rather than the book value that accumulates from it—is a precisely estimated zero for drought ($+0.001$, $p = 0.83$, minimum detectable effect 0.18σ) and null for every other hazard in every dynamic lead and lag. Plant survival is unchanged: among 1,189 plants observed in the 2022 filings with exposure, those hit by a severe event in 2022 had disappeared from the 2024 filings at 12.8 percent against about 15 percent for the rest, a difference of +1.8 percentage points ($p = 0.65$), though this design can only exclude effects that would roughly double the exit rate. Plant-level asset growth between the two filings, for the 920 plants with assets reported in both, is lower after an event by a statistically insignificant margin under parent fixed effects.

Figure 3 places the plant estimates beside the sales-office placebo. At sales offices the asset-growth coefficients are zero or positive, the ROA coefficients match the plant ones, and capital expenditure is null throughout. The placebo is underpowered for flood and drought—its minimum detectable effects are two to three times the plant coefficient—but at equal power for typhoons, where it is null.

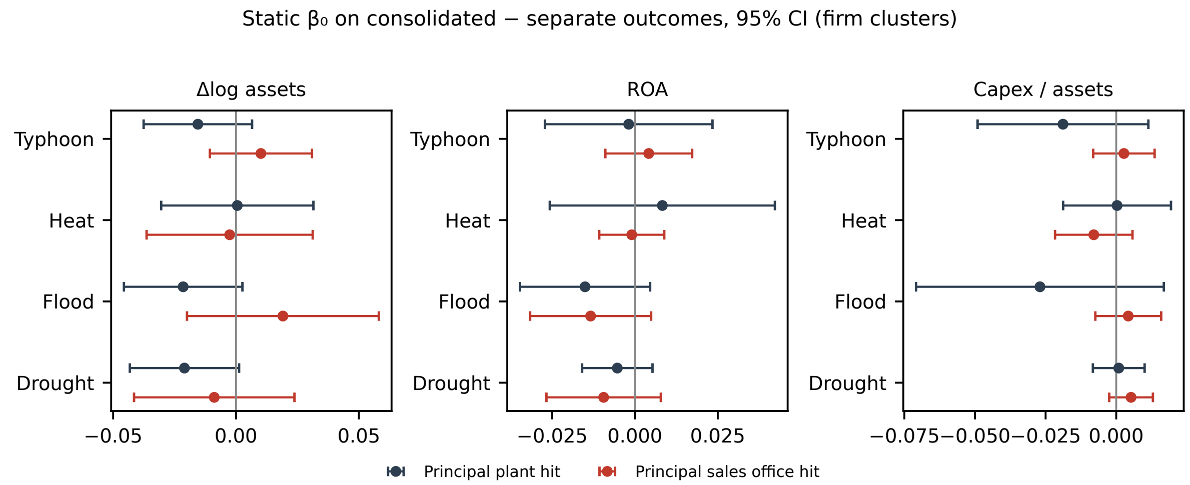

Figure 3: Principal plant versus principal sales office: static β_0 on consolidated-minus-separate outcomes

Notes. Static coefficients from equation (5) with 95 percent confidence intervals. Plant estimates use 1,399 geocoded manufacturing plants; sales-office estimates use 1,844 geocoded sales offices of the same firms with the same thresholds.

Taken together, the trace on consolidated assets is a write-down: damage that the accounts record and that the firm neither loses money over, responds to with investment, nor walks away from. This is where the ownership channel stops.

5.3 The Trade Channel Reaches the Input Bill as a Fifth of a Percent

The two preceding subsections supply the ingredients for a calculation that the qualitative claim—sourcing is diversified, so shocks are diluted—leaves unanswered: is the dilution quantitatively sufficient? Every term below is estimated in this paper rather than assumed.

Foreign value added accounts for 32.8 percent of Korean manufacturing gross output (Section 3.2). Within that foreign component, the largest single origin averages 23.5 percent across industries (Section 7.1). And a one standard deviation increase in a partner country's drought exposure reduces Korean imports from it by 2.8 percent, the intensive-margin estimate of Table 15. Whether that decline reflects supply disruption or contraction in the partner economy does not matter here: either way it is the reduction in inputs that Korean industry actually experiences. The three magnitudes multiply. A drought of that size in an industry's *largest* supplier moves its input bill by

$$0.328 \times 0.235 \times 0.0284 \;=\; 0.22 \text{ percent.}$$

Suppose instead, counterfactually, that every foreign supplier were struck simultaneously with a one standard deviation drought. Dropping the 23.5 percent concentration term gives $0.328 \times 0.0284 = 0.93$ percent. Even that upper bound leaves the input bill essentially unmoved.

A second dilution operates on top of the first, and follows from what the outcome variable measures. Dispersion of log $MRPK$ within an industry is invariant to any shock that hits its establishments equally: a common proportional change shifts every log $MRPK$ by the same constant and leaves the standard deviation unchanged. For exposure to move our outcome, establishments within an industry must differ in how much they source from the affected country. The portion of the 0.22 percent that reaches the dispersion measure is therefore smaller still, by an amount we cannot observe because firm-level sourcing is not in the data.

This is what the null means. The effect implied by the observed sourcing structure lies an order of magnitude below what a design at this level of aggregation could detect. Resilience here is arithmetic: an industry drawing on roughly twenty origins, none dominant, cannot be moved appreciably by weather in any one of them.

It does not follow that the estimation was redundant, and the reason matters for how the contribution should be read. The calculation above is a first-order accounting, and a substantial literature holds that first-order accounting understates what happens in production networks. Baqaee and Farhi (2019) and Baqaee and Farhi (2020) show that microeconomic distortions aggregate through input-output linkages in ways that simple share weighting misses. Barrot and Sauvagnat (2016) and Boehm et al. (2019) find that geographically confined disasters impose losses on customers well beyond the affected supplier's share of their inputs. Had such amplification operated here, a 0.22 percent input shock could have produced a detectable allocative effect. The estimates say it did not. The contribution of the null is therefore not to confirm the arithmetic but to bound the wedge between the arithmetic and the outcome: whatever amplification operates in this setting is small enough that direct exposure shares remain an adequate guide to consequences.

It bears emphasis that this is not because the sample lacks severe events. It contains the largest climate disasters of the period in Korea's principal suppliers—India in 2015 with 347 million people affected, China in 2010 with 180 million, the Philippines in 2013 and 2024—and the exposure measure registers them. What the arithmetic says is that even disasters of that magnitude, diluted across a portfolio of this shape, arrive as a fraction of a percent of the input bill.

5.4 It Does Not Reach the Allocation of Capital

Table 7 reports estimates of equation (6) for every hazard measure. Flood, drought, wildfire, landslide, the storm intensity measures and the composites are all statistically indistinguishable from zero, most with negative point estimates—the opposite sign to what a disruption channel predicts.

Three of the nineteen coefficients are significant, which is more than chance would deliver but not by a wide margin, and both exceptions dissolve on inspection. The two heat threshold measures enter positively: a one standard deviation increase in exposure to days above 35°C is associated with a 0.069 increase in sd(log $MRPK$), about a quarter of the outcome's sample standard deviation, and the 40°C measure gives 0.053. Unusually, their bootstrap p-values are *smaller* than the asymptotic ones. Section 6 shows that the association is confined to a single supplier country and runs opposite to supplier climate elsewhere. The raw storm count is also significant, with a negative sign, and Section 5.6 shows why it cannot be read as evidence of wind damage. Neither survives as a claim about hazards.

5.5 What the Nulls Rule Out

Table 8 reports minimum detectable effects: the smallest coefficient the design would reject the null against, at five percent significance and eighty percent power, expressed in standard deviations of the outcome.

The headline point is that precision varies by a factor of seven across hazards, so a single summary figure would misrepresent the evidence. The storm intensity measures are the most precise in the paper: an effect above 0.12σ for the power dissipation index,

Table 7: Foreign physical risk and capital misallocation, by hazard

Hazard measure	β	s.e.	t	Asym. p	Boot. p	N
Heat						
Days $> 35°$C	+0.0688	(0.0267)	+2.58	0.023	**0.007**	196
Days $> 40°$C	+0.0530	(0.0202)	+2.62	0.021	**0.003**	196
Annual maximum temperature	+0.0579	(0.0295)	+1.97	0.071	0.120	196
Extreme temperature events (EM-DAT)	−0.0257	(0.0251)	−1.02	0.325	—	210
Flood						
Max. 5-day precipitation	+0.0267	(0.0492)	+0.54	0.596	0.865	196
Days with precip. > 20 mm	+0.0081	(0.0439)	+0.18	0.857	—	196
Flood events (EM-DAT)	−0.0125	(0.0257)	−0.48	0.636	—	210
Drought						
Consecutive dry days	−0.0115	(0.0121)	−0.95	0.361	0.488	196
Drought events (EM-DAT)	−0.0157	(0.0102)	−1.55	0.146	—	210
Storm (IBTrACS, by intensity weighting)						
Storm count (≥ 34 kt)	−0.0421	(0.0108)	−3.88	0.002	0.019	210
Typhoon count (≥ 64 kt)	−0.0209	(0.0108)	−1.94	0.075	0.145	210
Maximum wind speed	−0.0090	(0.0068)	−1.32	0.210	0.320	210
Cyclone energy ($\sum V^2$)	−0.0152	(0.0123)	−1.24	0.238	0.479	210
Dissipation index ($\sum V^3$)	−0.0121	(0.0116)	−1.04	0.318	0.553	210
Storm events (EM-DAT)	−0.0004	(0.0249)	−0.02	0.987	—	210
Other hazards (EM-DAT event counts)						
Wildfire	−0.0429	(0.0253)	−1.70	0.113	—	210
Landslide (wet mass mvmt.)	−0.0131	(0.0109)	−1.20	0.251	—	210
Composite						
All climate disaster counts	−0.0129	(0.0174)	−0.74	0.471	0.453	210
Disaster-affected population	−0.0364	(0.0361)	−1.01	0.332	0.375	210
Disaster count × vulnerability	−0.0109	(0.0170)	−0.64	0.533	0.528	210

Notes. Dependent variable is sd(log $MRPK$). Estimates of equation (6); all specifications include industry and year fixed effects. FPR is standardised, so coefficients are responses to a one standard deviation change in exposure. Standard errors clustered by industry (14 clusters). Bootstrap p-values from the wild cluster restricted bootstrap with the null imposed, enumerating all 16,384 Rademacher sign vectors; computed for the measures discussed in the text. Weather measures are anomalies relative to the country's 1961–1990 baseline and end in 2023, which costs one year relative to the EM-DAT measures.

or above 0.07σ for maximum wind speed, would have been detected. Drought is nearly as precise at 0.12σ. In economic terms, these bound the response of sd(log $MRPK$) to a one standard deviation increase in exposure at between one and two percent of its sample mean.

At the other extreme, the precipitation-based flood measures detect only effects above 0.44 to 0.49σ, or roughly ten percent of the dispersion—a weak constraint, and the one place where the null as stated would be uninformative. Section Appendix D addresses it directly.

Table 8: Minimum detectable effects by hazard

Hazard measure	MDE (σ of outcome)	Verdict
Maximum wind speed	0.07	precise null
Drought events (EM-DAT)	0.10	precise null
Landslide events	0.11	precise null
Storm count, typhoon count	0.11	see §5.6
Drought (consecutive dry days)	0.12	precise null
Power dissipation index, ACE	0.12	precise null
Disaster count (composite)	0.17	precise null
Heat, days > 40°C	0.20	**effect found**
Storm, wildfire events (EM-DAT)	0.25	moderate
Extreme temperature events	0.25	moderate
Flood events (EM-DAT)	0.26	moderate
Heat, days > 35°C	0.27	**effect found**
Annual maximum temperature	0.29	moderate
Disaster-affected population	0.36	weak
Flood, days precip. > 20 mm	0.44	weak
Flood, max. 5-day precipitation	0.49	weak

Notes. MDE = $(t_{0.975,G-1} + t_{0.80,G-1}) \times \text{s.e.}(\hat{\beta})/\text{sd}(y)$, computed from equation (6) with sd(log $MRPK$) as the outcome and $G = 14$ industry clusters. Ordered by precision.

5.6 Storms: An Intensity Gradient

The storm block of Table 7 contains the third significant coefficient and, we argue, the paper's most instructive non-finding. Reading down that block, the measures are ordered by how heavily they weight intensity—a raw count treats all storms alike, a typhoon count applies a threshold, maximum wind takes the peak, and ACE and PDI weight by V^2 and V^3. The bootstrap p-values rise monotonically along that ordering: 0.019, 0.145, 0.320, 0.479, 0.553.

If wind damage were transmitting to Korean capital allocation, this gradient would run the other way. Damage scales with dissipated power, so the $\sum V^3$ measure should carry the signal and the unweighted count should be the noisiest proxy for it. We observe the reverse: the measure with no intensity information is significant, and the measure that best represents damage is the weakest in the table. Whatever produces the storm-count coefficient, it is not the mechanism by which storms destroy productive capacity.

The coefficient is also negative—more storms, less misallocation—which no disruption mechanism predicts, and it does not survive correction for the nineteen hazard measures now reported (Bonferroni threshold 0.0026 against $p = 0.019$). It is, however, stable across all three outcomes ($p = 0.002$, 0.007, 0.003), so we report it as an unexplained regularity rather than dismissing it as noise. We tested one candidate explanation, that storm frequency proxies sourcing from East Asia, and found the correlation too weak to account for it (0.10). We build nothing on this coefficient.

What the storm measures do deliver is precision. Their minimum detectable effects, 0.07 to 0.12σ, are the tightest in the paper. The reader who suspects that a strengthening cyclone could damage a supplier without producing a flood can be answered directly: that channel is measured here, at the intensity where damage occurs, and an effect above roughly a tenth of a standard deviation is ruled out.

5.7 The Ladder

Table 9 collects the estimates in the order the shock would travel, with the bound each null carries and the placebo that qualifies it.

The shock is visible at the first rung and at the sixth, in each case at a magnitude the firm can absorb, and it is absent at every rung that involves a decision.

6 Extreme Heat: Divided Evidence

6.1 The Trade-Channel Coefficient and Its Limits

The heat coefficients are the largest in the table and the only ones with the sign a disruption channel predicts, so they deserve scrutiny rather than either dismissal or acceptance. We apply the filters set out in Section 4.2—multiple-testing correction, consistency across outcomes, and the geophysical placebo—and then two further tests that the filters do not anticipate. The estimate passes the first three and fails the last two.

Consider first the multiple-testing correction. What it delivers depends on what counts as the family of tests, and we report both defensible answers rather than the more favourable one.

Treating each of the nineteen measures as a separate test gives a Bonferroni threshold of 0.0026, which the 40°C measure ($p = 0.003$) narrowly fails. The Benjamini-Hochberg procedure at a five percent false discovery rate rejects nothing: the ordered p-values 0.003, 0.007 and 0.019 all exceed their thresholds 0.0026, 0.0053 and 0.0079.

That calculation, however, treats the five storm measures as five independent tests

Table 9: How far the shock travels: estimates, bounds and placebos by stage

Stage	Estimate	Bound (95%)	Placebo	Verdict
Ownership channel: severe event at the principal overseas plant				
1. Consolidated asset growth	flood -0.0215^{**}, drought -0.0210^{**}	MDE 0.31–0.33σ	consolidated $< 0 <$ separate; absent at sales offices	Trace of write-downs; not FDR-robust
2. Profitability (ROA)	drought -0.0053	MDE 0.26σ	same at sales offices	Null
3. Investment (capex)	drought $+0.0009$	MDE 0.18σ	null at sales offices	Precise null: write-down, not disinvestment
4. Plant survival	any event +1.8 pp	MDE 11 pp	—	Null; only large effects excluded
5. Plant-level assets (2022–24)	any event -0.18	MDE 0.66σ	parent FE	Null; underpowered
Trade channel: physical risk in supplier economies				
6. Input bill	0.22% of inputs (largest supplier, 1σ drought); 0.93% if all suppliers	arithmetic	—	Diluted before arrival
7. Re-sourcing	import/export ratio -0.0008 ($p = 0.961$)	MDE 0.017	trade falls symmetrically	Absorption is structural, not behavioural
8. Capital allocation, sd(log MRPK)	null across 24 outcomes × 13 hazards	storms, drought, landslide $\leq$2%; flood $\leq$5% of dispersion	earthquakes null	Bounded null

Notes. Stages 1–5 are the ownership channel (event at the principal overseas plant, consolidated minus separate); stages 6–8 the trade channel (value-added-weighted supplier hazards). Bounds are minimum detectable effects in standard deviations of the outcome, or the percentage of dispersion excluded, as indicated. ** $p < 0.05$.

when they are five operationalisations of one hazard, and likewise for the three heat measures; the tests within a hazard are strongly correlated, so the correction is conservative by construction. Defining the family at the level of hazards—heat, flood, drought, storm, wildfire, landslide, and the composites—gives seven tests, a Bonferroni threshold of 0.0071, and both heat measures clear it, as they clear the corresponding Benjamini-Hochberg thresholds.

The honest summary is that the heat result is significant under a hazard-level correction and marginal-to-failing under a measure-level one. We therefore do not rest the claim on the correction, and turn to two filters that do not depend on how the family is drawn.

A second filter is consistency across outcomes. Table B3 shows that the threshold heat measures enter positively on all three outcome variables, with p-values between 0.021 and 0.105. Flood and drought are null on all three. A coefficient that appeared by chance on one outcome would not be expected to reappear on the other two with the same sign.

The third filter is the geophysical placebo. Table 10 reports the test motivated in Section 3.3. Exposure to earthquakes, constructed identically to the climate measures, is unrelated to capital misallocation: $\beta = 0.017$ with a bootstrap p-value of 0.642. This is not a power failure. The earthquake specification detects effects down to 0.28σ, essentially the same precision as the 35°C measure at 0.27σ; it has the power to find an effect of the size heat produces, and finds nothing. Whatever drives the heat coefficient is therefore not a generic property of supplier countries that experience and report natural disasters.

Table 10: Geophysical placebo

Exposure measure	Climate-related?	β	Cluster s.e.	Boot. p	MDE
Earthquake events	No	+0.0168	(0.0281)	0.642	0.28σ
All geophysical events	No	+0.0206	(0.0330)	0.707	0.33σ
Days > 35°C	Yes	+0.0688	(0.0267)	**0.007**	0.27σ
Days > 40°C	Yes	+0.0530	(0.0202)	**0.003**	0.20σ
Max. 5-day precipitation	Yes	+0.0267	(0.0492)	0.865	0.49σ
Consecutive dry days	Yes	−0.0115	(0.0121)	0.488	0.12σ

Notes. Geophysical exposure is built from EM-DAT earthquake, volcanic activity and dry mass movement counts using exactly the construction of equation (1) and the same weights. Dependent variable sd($\log MRPK$); industry and year fixed effects; bootstrap enumerates all 16,384 sign vectors. MDE is the minimum effect detectable at five percent significance and eighty percent power, in units of the outcome's standard deviation.

Two features of the heat result sit awkwardly and we report them rather than smooth them over. First, the EM-DAT extreme-temperature event count is not significant and carries the opposite sign (-0.026, $p = 0.325$). Declared heat-wave disasters and cumulative days above a temperature threshold are different constructs, and the former depends on national declaration practice, but the divergence means the heat result rests on the reanalysis measures alone. Second, within the reanalysis measures the two threshold counts are significant while the annual maximum temperature is not ($p_{\text{boot}} = 0.120$). A reading consistent with both facts is that what matters is accumulated operating time under extreme heat rather than peak intensity, which is also the construct Pankratz and Schiller (2024) find to matter for supplier performance. We offer this as an interpretation of the pattern, not as a tested hypothesis.

A more consequential limitation emerges when the exposure measure is decomposed by origin. Splitting *FPR* into the contribution of China—Korea's largest supplier, accounting for 16 percent of the foreign value added embodied in its manufacturing—and the contribution of all other countries, then entering both in the same regression, gives

$$\hat{\beta}_{\text{China}} = 0.0859 \;\; (p_{\text{boot}} = 0.039), \qquad \hat{\beta}_{\text{rest}} = 0.0063 \;\; (p_{\text{boot}} = 0.766).$$

The effect is entirely Chinese in origin. This is not a matter of statistical power: the non-China component is estimated *more* precisely than the China component (minimum detectable effect 0.22σ against 0.33σ), and its confidence interval, $[-0.036, 0.049]$, excludes an effect of the magnitude found for China. Heat in Korea's other supplier countries does not move domestic capital allocation, and we can say so with reasonable precision. Dropping supplier countries one at a time confirms the asymmetry: removing China costs significance ($p = 0.628$), while removing the United States, Germany, Japan, Indonesia, the United Kingdom or Malaysia leaves the estimate significant throughout.

Concentration in one country raises an identification problem that year fixed effects cannot solve. If the coefficient is driven by Chinese hot years, anything else particular to those years in China—power rationing, pandemic closures, trade disputes—would produce the same pattern. We therefore apply the origin decomposition to the other hazards, including the geophysical placebo. Within China, earthquake exposure is unrelated to misallocation ($\hat{\beta} = -0.013$, $p = 0.632$), as are drought ($p = 0.910$), flood ($p = 0.139$) and storms ($p = 0.203$). Whatever the Chinese component is picking up, it is specific to heat rather than to Chinese years in general.

Is heat then the mechanism? The within-China evidence suggests it is heat rather than Chinese years in general. But if heat damages production, the effect should be strongest where suppliers are hottest. We test this by splitting foreign exposure three ways: China; other suppliers that are genuinely hot, defined by a long-run average of at least five days above 35°C and comprising some 21 percent of foreign value added; and the remainder, which are temperate. Table 11 reports the result, and it does not support a heat mechanism.

The coefficients run opposite to supplier heat. India, Australia and Thailand experience an order of magnitude more extreme-heat days than China, and their combined coefficient is negative and insignificant, at a precision better than China's own (0.24σ against 0.33σ), so this is not a power failure. Meanwhile the temperate suppliers—

Table 11: Heat exposure decomposed by supplier climate

Component	Days > 35°C (long-run mean)	β	Boot. p	MDE
China	5.8	+0.0711	**0.042**	0.33σ
Other hot suppliers	high	−0.0373	0.185	0.24σ
Temperate suppliers	≈ 0	+0.0361	0.098	0.17σ

Notes. All three components enter the same regression, built from a common denominator so that they sum to total exposure. Dependent variable sd(log *MRPK*); days > 35°C measure is the exposure variable throughout. Hot suppliers include India (78.7 days), Australia (80.2), Thailand (43.6), Brazil (22.1), the United States (12.5) and Vietnam (10.0); temperate suppliers include Japan (0.6), Germany (0.3), the United Kingdom and Taiwan (both ≈ 0). Classification uses long-run levels and is fixed before estimation.

Japan, Germany, the United Kingdom, Taiwan, countries that essentially never exceed 35°C—carry a positive coefficient significant at ten percent, which no account of heat damaging production can rationalise. A heat mechanism predicts the ordering hot > China > temperate; the data give China > temperate > hot.

Two facts therefore hold simultaneously: within China the association is specific to heat, and across countries it bears no relation to how much heat suppliers actually experience. We cannot reconcile them. One possibility is that China's response to heat is idiosyncratic rather than its heat being severe—its power system is unusually exposed, through hydropower dependence and industrial rationing, as the 2021–2022 episodes illustrated—but we have no way to test that here. We therefore report the association as confined to one supplier country and unreconciled with supplier climate, which is one of the two sides of the divided evidence set out in Section 6.2.

6.2 Robustness of the Heat Result

Table 12 subjects the heat estimates to four alternative samples and definitions, with drought and flood shown alongside as controls that should stay null throughout.

Consider first what survives. Excluding the poorly-measured petroleum industry leaves both estimates significant. Excluding the pandemic years strengthens them, which rules out the possibility that the result reflects COVID-era disruption coinciding with hot years. A leave-one-industry-out exercise is more demanding still: dropping each of the fourteen industries in turn, the 35°C coefficient ranges from 0.048 to 0.091 and its bootstrap p-value from 0.003 to 0.029, remaining below five percent in all fourteen cases; the 40°C coefficient behaves likewise ($p = 0.001$ to 0.034). No single industry drives the finding.

Three things break the estimate, and all deserve to be stated plainly rather than buried.

Table 12: Robustness of the heat estimates

Specification	Days > 35°C β	Days > 35°C Boot. p	Days > 40°C β	Days > 40°C Boot. p
Baseline	+0.0688	**0.007**	+0.0530	**0.003**
Excluding coke and petroleum	+0.0482	**0.014**	+0.0703	**0.026**
Excluding 2020–2021	+0.1002	**0.001**	+0.0815	**0.025**
Unnormalised *FPR*	+0.0075	0.831	+0.0221	0.587
Unnormalised *FPR*, controlling foreign share	−0.0320	0.802	−0.0096	0.779
Renormalised *FPR*, controlling foreign share	+0.0568	**0.003**	+0.0469	**0.008**
Hazard in days, not standardised	−0.0209	0.712	−0.0288	0.707
Non-China component only	+0.0063	0.766	−0.0478	0.322
China component only	+0.0859	**0.039**	+0.0557	0.050
Controls: same variants, drought and flood				
Drought, all variants	$p = 0.48$–0.91 throughout			
Flood, all variants	$p = 0.14$–0.90 throughout			

Notes. Estimates of equation (6) with sd(log *MRPK*) as outcome. Bootstrap p-values enumerate all 2^G Rademacher sign vectors. $N = 196$ except excluding coke and petroleum (182) and excluding 2020–2021 (168).

The first is the origin decomposition reported in Section 6: the estimate is Chinese in origin, and the non-China component is a precisely estimated zero rather than an imprecise one. We repeat it in Table 12 because it belongs with the other robustness evidence, and because it is the sharpest limitation on what the finding supports.

The second is normalisation. The estimate does not survive replacing the renormalised exposure measure with the unnormalised one ($p = 0.831$ and 0.587). Since FPR^{unnorm} is the product of average supplier risk and the foreign value-added share, an obvious conjecture is that the unnormalised specification suffers from omitting that share, and that controlling for it would restore the estimate. It does not: with the foreign share entered separately, the unnormalised coefficient is -0.032 ($p = 0.802$), while the renormalised coefficient is unaffected ($+0.057$, $p = 0.003$). The conjecture is therefore rejected, and we are left without an account of why the two constructions differ so sharply. What can be said is descriptive: after fixed effects the unnormalised measure retains less than half the identifying variation of the renormalised one (residual standard deviation 0.103 against 0.242) and the two correlate at only 0.50. Section 3.1 gives the theoretical case for renormalisation—it is the standard shift-share object, with shares summing to one—but a reader who regards total rather than average exposure as the quantity of interest will read the evidence differently, and we cannot settle that with these data.

The third is the standardisation of the hazard itself, and it is related to the second. Our climate measures enter as anomalies scaled by each country's own 1961–1990 standard deviation, which is standard practice but has an unusual consequence for a threshold count like days above 35°C. That baseline standard deviation ranges from 0.08 days in Germany and 0.17 in Japan to 8.6 in India and 10.8 in Australia—two orders of magnitude. One additional hot day is therefore a twelve standard deviation event in Germany and a tenth of one in India. The z-measure consequently places overwhelming weight on low-variance suppliers, and China, with a baseline standard deviation of 0.80 days, is among them.

The natural alternative, if damage scales with the physical quantity of heat rather than with its statistical unusualness, is the deviation measured in days. Under that construction the coefficient becomes -0.021 ($p = 0.712$) for the 35°C measure and -0.029 ($p = 0.707$) for 40°C. Both vanish and change sign. The two constructions are each defensible ex ante and give opposite answers, and we cannot adjudicate between them.

Taken together, the evidence on heat points in both directions, and we set out each side rather than choosing between them.

Four checks support reading the coefficient as a real association. It is robust to sample composition: no single industry, and not the pandemic, accounts for it. It passes both the geophysical placebo and its within-China counterpart. It is present under five separate constructions of the dispersion measure—the standard deviation and variance of log MRPK, the interdecile range, the reallocation gain, and the dispersion of log TFPR—which an artefact of a particular summary statistic would not survive (Section Appendix C). And it appears on the capital margin alone: the dispersion of log MRPL, which measures the same distortion for labour, is flat against both heat measures (-0.009, $p = 0.840$; -0.011, $p = 0.833$). That asymmetry is what a mechanism operating through slow capital adjustment would predict, and it is the pattern Liu and Xu (2025) document for domestic temperature exposure.

Three checks cut the other way. The association holds for one supplier country, under one of two defensible normalisations of exposure, and under one of two defensible standardisations of the hazard. Estimated in days rather than standardised anomalies the coefficient reverses sign and loses significance, whereas the drought and flood nulls are entirely insensitive to the same choices: drought gives -0.006 ($p = 0.514$) against -0.012 ($p = 0.488$), and flood $+0.030$ ($p = 0.603$) against $+0.027$ ($p = 0.865$). The measurement sensitivity is specific to heat, and heat is the only hazard whose measure is a threshold count with near-degenerate baselines. Nor is the coefficient stable across thresholds once industry-specific trends are allowed: the 35°C estimate strengthens, while the 40°C estimate disappears altogether.

We therefore do not claim that imported heat propagates, and we do not claim that it does not. What the design can establish is the bound: the estimate is small enough that even at face value it would not overturn the aggregate picture, and it is the only hazard for which our checks disagree among themselves. The other hazards are not in this position; for them the nulls hold under every measurement convention we can construct.

Settling heat requires data this design does not have, and the plant panel of Section 3.7 is the first step towards it. Two features of the country-year measure produce the ambiguity: a threshold count aggregated over a whole country cannot distinguish a supplier whose plants sit in the affected region from one whose plants do not, and the baseline variance used to standardise it differs across suppliers by two orders of magnitude. Daily temperature evaluated at plant coordinates removes both, which is what the next subsection does for the plants Korean firms own; there the test is weak only because those plants rarely see days above 40°C, not because of the measure. The remaining constraint is the panel itself: fourteen industry clusters over fifteen years is thin for a hazard whose effect, if present, is a few percent of dispersion, and firm-level sourcing data would raise the number of clusters by orders of magnitude. We regard imported heat as the open question this paper hands on, with the measurement problem solved and the sample problem not.

6.3 Heat at Plant Level

The country-year measure cannot separate a mechanism from a measurement convention, and the plant-level data were assembled in part to try. At plant level heat is measured in absolute days above a threshold, so no standardisation by a country's baseline variance is involved, and the host-country-by-year fixed effect removes any national-level story. The result is nothing: 233 firm-years with at least 45 days above 35°C at the principal plant, and no outcome moves in any specification (Table 6). This weighs against the mechanism reading of the trade-channel coefficient, but not decisively, because Korean plants sit where extreme heat is rare—the 95th percentile of days above 40°C across principal plants is zero, and the threshold had to be lowered to 35°C for events to exist at all. A null on a hazard that the sample barely contains is a weak null.

7 Why the Shock Stops

A null is informative only with an account of why the effect is absent. This section provides one, and arrives at it by ruling out the more obvious alternative.

The obvious account is behavioural: firms notice that a supplier country has been struck and source elsewhere, so the shock is neutralised before it reaches domestic production. Section 7.2 tests this and rejects it. Trade with climate-stressed countries does fall, but Korean exports to those countries fall by the same amount as Korean imports from them, which is the signature of a contraction in the partner economy rather than of re-sourcing. Import-specific reallocation is absent, and absent at a precision that would have detected effects above two percent.

What remains is a structural account, and the section builds it in three steps. Section 7.1 establishes the structure: Korean manufacturers source from eighteen to twenty-three countries on average, with no origin dominant. Section 7.2 establishes how much trade actually responds to a hazard, which is the quantity the structure operates on. Section 5.3 then combines the two and computes how large a shock survives the dilution—a fifth of a percent of the input bill for a severe drought in an industry's largest supplier. That calculation is the answer to the question in the section title. Sections 7.3 and 7.4 draw out what the account does and does not explain, and add inventories as a second structural buffer.

7.1 Sourcing Diversification

Table 13 reports the distribution of Korean manufacturing imports across firms grouped by the number of countries they trade with. The data come from the Trade by Enterprise Characteristics (TEC) statistics, which link the business register to customs records and therefore describe actual importers rather than modelled flows.

The concentration of sourcing is low. In the most recent year, firms trading with ten or more partner countries account for roughly 75 percent of manufacturing imports, while firms dependent on a single partner account for about 6 percent. The value-weighted average number of partner countries ranges from 18 to 23 in six of the seven industry groups. For a firm sourcing from twenty countries, a disaster in any one of them affects on the order of five percent of its inputs—an exposure unlikely to move the dispersion of marginal revenue products of capital within its industry.

One industry departs from this pattern. Textiles and leather has a value-weighted average of 8.5 partner countries, and only a third of its imports come from firms trading with ten or more partners. If concentrated sourcing were sufficient for propagation, this is where it should appear. We return to this in Section 7.4.

Table 13: Sourcing diversification by industry

Industry group	Avg. partner countries (value-wtd.)	Import share of firms with ≥ 10 partners	Import share of single-partner firms
Coke, petroleum, chemicals	23.1	93.8%	0.9%
Electrical, electronic, machinery	21.7	89.5%	2.0%
Transport equipment, other	21.2	87.5%	2.5%
Non-metallic minerals, metals	20.5	84.0%	3.5%
Food, beverages, tobacco	18.1	81.2%	2.4%
Wood, paper, printing	13.1	55.5%	11.5%
Textiles, leather	8.5	33.1%	20.2%

Notes. Author's calculations from Statistics Korea, Trade by Enterprise Characteristics (table `DT_1TEC_P316`), imports, most recent year. Partner-country counts are reported in brackets (1, 2, 3–5, 6–9, 10–14, 15–19, 20+); averages use bracket midpoints weighted by import value.

Diversification is also stable over time. The value-weighted average number of partner countries declines only modestly, from 19.4 in 2010 to 18.0 in 2025, and the import share of firms with ten or more partners falls from 81 to 75 percent. The absence of propagation is therefore not attributable to a recent shift in sourcing structure; the buffer was in place throughout the sample.

7.2 The Extensive Margin of Sourcing

Diversification alone describes a static buffer. We now ask whether firms also respond to realised physical risk by reallocating across suppliers. Pankratz and Schiller (2024) show, using firm-to-firm supply relationships, that customers are more likely to terminate a supplier relationship when the supplier's realised climate exposure exceeds prior expectations. The TEC data report the *number of importing firms* by industry and partner country, which permits an aggregate analogue of that test.

The answer, once the appropriate comparison is made, is that they do not. Trade with climate-stressed countries does fall, but it falls in both directions and by the same amount, which is what a contraction in the partner economy produces rather than what re-sourcing by Korean firms would produce. We set out the count regression first, because it is the direct analogue of the firm-level literature, and then the comparison that overturns its natural reading.

We estimate

$$\log N_{s,c,t} \;=\; \beta\, \mathit{PhysRisk}_{c,t} \;+\; \delta_{s\times t} \;+\; \delta_{s\times c} \;+\; \varepsilon_{s,c,t}, \tag{8}$$

where $N_{s,c,t}$ is the number of Korean firms in industry s importing from country c in year t. Industry-by-year fixed effects absorb global shocks and any industry-specific trend; industry-by-country fixed effects absorb the persistent component of each trading relationship. Identifying variation therefore comes from movements within an industry–country pair over time. Standard errors are clustered by country, and p-values are computed both asymptotically and by the wild cluster bootstrap of Cameron et al. (2008) with 9,999 Rademacher draws.

Results appear in Table 14. Every measure except the storm index enters negatively. Drought is the most precisely estimated: a one standard deviation increase in a partner country's consecutive dry days is associated with a 1.6 percent decline in the number of Korean firms importing from that country, with a bootstrap p-value of 0.001.

Heat also enters negatively and significantly at the 35°C threshold, but at 0.67 percent the response is only 42 percent as large as the drought response. This comparison is central to the argument of Section 7.3, and we return to it there. Two qualifications belong here. The two confidence intervals overlap and the difference between the coefficients is not itself significant ($p = 0.13$), so the contrast is one of magnitude rather than of one effect being present and the other absent. The heat sample is also smaller (4,485 against 5,472 observations, 49 against 60 clusters), because the threshold measure covers fewer countries.

Flood, storms and disaster counts produce no detectable extensive-margin response, and the affected-population measure is marginal.

Table 14: Physical risk and the number of importing firms

	Dependent variable: $\log N_{s,c,t}$					
Physical risk measure	β	s.e.	Asym. p	Boot. p	N	MDE
Drought (consecutive dry days)	−0.0161	(0.0052)	0.003	**0.001**	5,472	0.009
Heat (days > 35°C)	−0.0067	(0.0035)	0.064	**0.013**	4,485	0.006
Heat (days > 40°C)	−0.0122	(0.0075)	0.115	0.121	2,558	0.013
Disaster-affected population	−0.0127	(0.0074)	0.094	0.089	5,474	0.013
Disaster count	−0.0112	(0.0080)	0.168	0.153	5,880	0.014
Flood (max. 5-day precipitation)	−0.0034	(0.0072)	0.639	0.630	5,472	0.013
Storm (power dissipation index)	+0.0030	(0.0041)	0.474	0.459	3,421	0.008

Notes. Estimates of equation (8). All regressions include industry-by-year and industry-by-country fixed effects. Risk measures are standardised, so coefficients are semi-elasticities with respect to a one standard deviation change. Standard errors are clustered by country (52–60 clusters). Bootstrap p-values follow Cameron et al. (2008) and impose the null, using 9,999 Rademacher draws. Physical risk measures are anomalies relative to the 1961–1990 country baseline; disaster measures are constructed from EM-DAT (Delforge et al., 2025) and climate measures from ERA5 aggregates (World Bank, 2025).

Taken at face value, this looks like the aggregate counterpart of Pankratz and Schiller (2024): fewer firms trade with countries that experience climate stress, with the smaller magnitude that aggregation would predict. The specification is also considerably better identified than the capital-allocation regressions of Section 5.4—the sample is an order of magnitude larger and the number of clusters four times as great—so the nulls for the other hazards are informative rather than merely imprecise, with minimum detectable effects under two percent throughout.

The natural reading is nonetheless wrong, and the test that shows this is simple. If Korean firms were re-sourcing away from climate-stressed suppliers, the decline should be specific to imports. If instead the partner economy is contracting, both directions of trade should fall together: lower output reduces its capacity to export, and lower income reduces its demand for imports. The TEC statistics report Korean exports to the same partners, so the two explanations can be separated directly.

Table 15 reports the comparison. Drought reduces Korean imports from an affected country by 2.8 percent per standard deviation of exposure—and Korean exports to it by 2.8 percent as well. The two coefficients are indistinguishable. Applying the test in its sharpest form, by using the log ratio of imports to exports as the dependent variable, gives a coefficient of −0.0008 with a bootstrap p-value of 0.96 and a minimum detectable effect of 0.017. Import-specific reallocation larger than about two percent is ruled out.

The same holds for every other hazard.

Table 15: Is the trade decline specific to imports?

	Coefficient (bootstrap p)			
Hazard	Imports	Exports	Import/export ratio	MDE (ratio)
Drought	−0.0284 (0.027)	−0.0282 (0.010)	−0.0008 (0.961)	0.017
Heat > 35°C	−0.0064 (0.605)	−0.0072 (0.495)	+0.0010 (0.909)	0.011
Flood	+0.0006 (0.976)	+0.0002 (0.992)	−0.0010 (0.953)	0.021
Storm (PDI)	+0.0088 (0.490)	−0.0064 (0.546)	—	—
Disaster count	+0.0026 (0.898)	−0.0019 (0.865)	+0.0057 (0.795)	0.025

Notes. Estimates of equation (8) with the dependent variable replaced by log Korean import value from the partner, log Korean export value to the partner, and their log ratio. Industry-by- year and industry-by-country fixed effects throughout; country clusters (49–60); wild cluster bootstrap with 9,999 draws. N = 3,533–6,230.

Two consequences follow. The first is interpretive: the decline in the number of importing firms is not evidence that Korean manufacturers reallocate away from climate risk. It is consistent with, and better explained by, a contraction in the affected economy that reduces trade in both directions. We accordingly do not treat the extensive margin as an adaptation channel, and Section 7.3 revises the account of absorption in light of this.

The second is that the absence of reallocation is itself measured precisely. The design would have detected import-specific reallocation above two percent and finds none, for any hazard. Firms that source from eighteen to twenty-three countries and hold substantial inventories may simply have no need to re-source in response to a single supplier country's weather, and the evidence is consistent with their not doing so.

One hazard deserves a separate note. Heat reduces the number of importing firms (−0.67 percent, $p = 0.013$) without reducing import value (−0.64 percent, $p = 0.605$) and without moving exports. Fewer firms buying the same total implies consolidation among larger importers rather than reallocation across countries. The effect is small and we do not build on it, but it is the one pattern in the table that is specific to the import side.

7.3 Why Absorption Is Structural, Not Behavioural

The three preceding subsections point to a particular account of the null, and it is not the one we expected to find. Exposure is diluted before the shock arrives, by a sourcing structure that spreads inputs across eighteen to twenty-three countries and reduces even a severe drought in the largest supplier to a fifth of a percent of the input bill. It is not further attenuated afterwards: import-specific reallocation is absent, and absent at a precision that would have detected effects above two percent. Absorption in this setting is a property of how sourcing is organised, not of how firms respond.

That distinction matters for what the paper claims. A behavioural account would say that Korean firms notice climate stress and act on it. The evidence does not support that, and the natural reading of the count regression—that firms exit relationships with stressed suppliers—dissolves once exports are used as a comparison. A structural account says instead that a firm sourcing from twenty countries has little reason to act, precisely because the arithmetic of Section 5.3 applies at the level of the firm no less than at the level of the industry: a disaster in any one origin moves a small share of its input bill, and inventories cover the interruption. On this reading the absence of a response is not a puzzle alongside the absence of an effect; it is the same fact seen twice.

The structural account also explains why the hazards differ in the way they do. Floods, storms, landslides and wildfires are spatially confined: a flood severe enough to halt production in one river basin leaves suppliers a few hundred kilometres away untouched, so a diversified importer barely notices. Extreme heat is not confined in this way. A summer that pushes one supplier country past 35°C for an unusual number of days typically does the same across the region, so the alternatives available to a Korean importer are drawn from the same overheated distribution. Diversification across twenty countries provides little protection against a hazard correlated across all twenty. This is consistent with heat being the only hazard that leaves any trace in Table 7, and with the finding that threshold measures matter while peak temperature does not: a single very hot day is a local, transient event, whereas a season above a threshold is a sustained regional condition.

Two limits on this account should be stated. It is an interpretation that organises the pattern of results, not a tested mechanism; a direct test would measure the spatial correlation of each hazard across a firm's actual supplier set and show that the propagation coefficient rises with it, which country-year exposure data cannot support. And it does not explain the geography of the heat coefficient. Limited substitutability is a claim about hazards, so it predicts that heat should leave a trace wherever suppliers are hot; Section 6 shows the association is instead absent among suppliers far hotter than China. Whatever accounts for that, this account does not.

7.4 Remaining Explanations

Two further candidates merit brief discussion.

The first is inventories. Firms holding larger stocks relative to sales can smooth input disruptions without altering production. We compute, from the same Mining and Manufacturing Survey underlying the capital-allocation measures, the value-added-weighted ratio of average inventory (finished goods, semi-finished goods, work in progress, raw materials, and supplies) to annual shipment value, by industry group.

Table 16 shows that Textiles and leather—the one industry that departed from the diversification pattern in Section 7.1—also holds the largest inventory buffer, at 23.7

percent of shipment value in 2024, roughly twice the 10–17 percent held by other industries. The pattern holds over the full sample as well as in the most recent year. This is consistent with inventories substituting for diversification: the industry with the least diversified supplier base compensates with the largest stockpile, which would explain why concentrated sourcing in textiles does not appear to propagate into capital misallocation despite lacking the buffer described in Section 7.1. We do not have a direct test that inventories are the mechanism connecting the two facts, and offer this as a plausible complementary explanation rather than a demonstrated one.

Table 16: Inventory-to-shipment ratio by industry

Industry group	Full sample (2007–2024,[a] VA-wtd.)	2024 (VA-wtd.)
Textiles, leather	0.197	0.237
Non-metallic minerals, metals	0.124	0.158
Electrical, electronic, machinery	0.108	0.165
Coke, petroleum, chemicals	0.103	0.131
Food, beverages, tobacco	0.089	0.098
Wood, paper, printing	0.082	0.097
Transport equipment, other	0.079	0.106

Notes. Author's calculations from Statistics Korea, Mining and Manufacturing Survey (establishments with 10+ employees). Inventory is the average of beginning- and end-of-year totals across five categories; the denominator is annual shipment value. Ratios are value-added-weighted averages of establishment-level ratios within each industry-year cell, then averaged across years for the full-sample column. Industry groups match Table 13.
[a] Excludes 2011–2012, when the survey's repeated-column layout for that year could not be resolved with the same rule used elsewhere, and 2002–2006, which use a different inventory field layout not yet mapped.

The second is aggregation and measurement resolution. The studies that do detect propagation measure exposure at a far finer resolution than our imported-input measure. Barrot and Sauvagnat (2016) observe firm-to-firm supply contracts; Pankratz and Schiller (2024) observe plant coordinates matched to daily weather; Boehm et al. (2019) observe parent–affiliate links. The FPR measure, by contrast, is a country-by-industry annual aggregate, and because Korea's principal suppliers are large, internally diversified economies, a national disaster index for such a country is itself an average over many locations, most of them unaffected. Sections 5.1 and 5.2 close part of this gap: for the ownership channel we measure exposure at plant coordinates—the same resolution class as Pankratz and Schiller (2024)—and still find no response of parent outcomes at annual frequency. The resolution critique therefore retains its force for the imported-input channel but cannot, on its own, explain away the firm-level null. What survives of it is the frequency dimension: all of our outcomes are annual, and disruptions absorbed within a quarter would not register in either channel.

8 What the Design Cannot See

The bounds in Table 9 are the paper's claim, and four features of the design determine how far they reach. Outcomes are annual, so a disruption that depresses one quarter and is recovered within the year leaves no trace; the trace on consolidated assets is therefore a lower bound on transient damage and the null on profits is a statement about the fiscal year. The sample is ten years and contains between 53 and 212 events per hazard at the principal plant; the drought and compound pre-trends in Figure 2 are the kind of noise that sample size permits. The firms are listed manufacturers with disclosed subsidiaries, the segment of Korean industry most likely to hold insurance, inventories and alternative sites; the nulls do not extend to smaller firms. And the plant-level heat test is weak by construction, as Section 6.3 explains.

What the design also does not observe is the mechanism of absorption on the ownership side. Insurance recoveries, production shifted to other sites, and inventories are all consistent with a write-down that leaves profits untouched, and the data do not distinguish among them. On the trade side the mechanism is visible: sourcing is spread across eighteen to twenty-three countries and the shock is diluted before it arrives.

9 Conclusion

Physical climate shocks realised at Korean manufacturers' overseas suppliers and plants reach their accounts and stop there. A severe flood or drought at a firm's principal plant leaves a trace of about two percent on consolidated asset growth, consistent with write-downs and absent from the parent's separate statements; profitability, capital expenditure and plant survival do not respond, and the nulls are bounded. A severe drought in an industry's largest supplier reaches the input bill as a fifth of a percent, trade with the stressed partner falls symmetrically in both directions, and the allocation of capital across domestic establishments is unmoved to within two to five percent of its dispersion. Earthquakes are null at the same precision as climate hazards, consolidated accounts move where separate accounts do not, and sales offices show nothing where plants show a trace.

The measurement result stands on its own. National-aggregate exposure overstates the heat that plants face by a factor of three to four in the hottest supplier economies and understates their flood exposure by half, because production sits on coasts and in river basins and the national average does not. Supervisory frameworks that assign foreign physical risk from country-year data are assigning the wrong hazards to the wrong firms, and the direction of the error differs by hazard.

What the results do not say is that foreign climate risk is harmless. They say that for the segment of Korean industry that discloses its overseas plants, and over a decade that included the largest disasters of the period in its principal supplier economies, the shock is absorbed before it reaches any decision that would show in annual accounts. The margin on which it is absorbed—sourcing structure on the trade side, and whatever combination of insurance, inventories and alternative sites on the ownership side—is itself costly to maintain, and that cost is the part of imported physical risk this design does not price.

References

Addoum, J. M., Ng, D. T., and Ortiz-Bobea, A. (2020). Temperature shocks and establishment sales. *Review of Financial Studies*, 33(3):1331–1366.

Asian Development Bank (2025). Multiregional input-output tables, 2025 edition. Asian Development Bank.

Asker, J., Collard-Wexler, A., and De Loecker, J. (2014). Dynamic inputs and resource (mis)allocation. *Journal of Political Economy*, 122(5):1013–1063.

Baqaee, D. R. and Farhi, E. (2019). The macroeconomic impact of microeconomic shocks: Beyond Hulten's theorem. *Econometrica*, 87(4):1155–1203.

Baqaee, D. R. and Farhi, E. (2020). Productivity and misallocation in general equilibrium. *Quarterly Journal of Economics*, 135(1):105–163.

Barrot, J.-N. and Sauvagnat, J. (2016). Input specificity and the propagation of idiosyncratic shocks in production networks. *Quarterly Journal of Economics*, 131(3):1543–1592.

Bigio, S. and La'O, J. (2020). Distortions in production networks. *Quarterly Journal of Economics*, 135(4):2187–2253.

Bils, M., Klenow, P. J., and Ruane, C. (2021). Misallocation or mismeasurement? *Journal of Monetary Economics*, 124(S):S39–S56.

Blaum, J., Esposito, F., and Heise, S. (2026). Input sourcing under climate risk: Evidence from U.S. manufacturing firms. *Review of Economic Studies*. Forthcoming.

Boehm, C. E., Flaaen, A., and Pandalai-Nayar, N. (2019). Input linkages and the transmission of shocks: Firm-level evidence from the 2011 tōhoku earthquake. *Review of Economics and Statistics*, 101(1):60–75.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3):414–427.

Chae, H.-Y., Kwon, S., Kim, J., Joo, D., and Han, S. (2024). Climate change and the role of the Bank of Korea. *Economic Analysis (Bank of Korea)*, 30(4). In Korean.

Chun, H. M. and Kang, S. H. (2024). The impact of climate risk on financial stability. *Korea International Accounting Review*, 116:23–52. In Korean.

Cruz, J.-L. and Rossi-Hansberg, E. (2024). The economic geography of global warming. *Review of Economic Studies*, 91(2):899–939.

David, J. M. and Venkateswaran, V. (2019). The sources of capital misallocation. *American Economic Review*, 109(7):2531–2567.

Delforge, D., Wathelet, V., Below, R., Lanfredi Sofia, C., Tonnelier, M., van Loenhout, J. A. F., and Speybroeck, N. (2025). EM-DAT: The emergency events database. *International Journal of Disaster Risk Reduction*.

Graff Zivin, J. and Neidell, M. (2014). Temperature and the allocation of time: Implications for climate change. *Journal of Labor Economics*, 32(1):1–26.

Hsieh, C.-T. and Klenow, P. J. (2009). Misallocation and manufacturing TFP in China and India. *Quarterly Journal of Economics*, 124(4):1403–1448.

Hwang, J. (2022). Physical risk stress testing on Korean banks' corporate lending. *Journal of Korean Economic Studies*, 40(3):89–125. In Korean.

International Monetary Fund (2021). We are all in the same boat: Cross-border spillovers of climate risk through international trade and supply chain. IMF Working Paper 2021/013, International Monetary Fund.

Kim, S. (2024). Climate change-induced physical risks' impact on Korean commercial banks and property insurance companies in the long run. *Atmosphere (Korean Meteorological Society)*, 34(2):107–121. In Korean.

Knapp, K. R., Kruk, M. C., Levinson, D. H., Diamond, H. J., and Neumann, C. J. (2010). The international best track archive for climate stewardship (IBTrACS): Unifying tropical cyclone data. *Bulletin of the American Meteorological Society*, 91(3):363–376.

Liu, T. and Xu, Z. (2025). The (mis)allocation channel of climate change: Evidence from global firm-level microdata. Working paper, Cornell University.

Los, B., Timmer, M. P., and de Vries, G. J. (2015). How global are global value chains? a new approach to measure international fragmentation. *Journal of Regional Science*, 55(1):66–92.

Notre Dame Global Adaptation Initiative (2026). ND-GAIN country index. University of Notre Dame.

Pankratz, N. M. C. and Schiller, C. (2024). Climate change and adaptation in global supply-chain networks. *Review of Financial Studies*, 37(6):1729–1777.

Ponticelli, J., Xu, Q., and Zeume, S. (2023). Temperature and local industry concentration. Working Paper 31533, National Bureau of Economic Research.

Restuccia, D. and Rogerson, R. (2008). Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic Dynamics*, 11(4):707–720.

Rosvold, E. L. and Buhaug, H. (2021). GDIS, a global dataset of geocoded disaster locations. *Scientific Data*, 8:61.

Somanathan, E., Somanathan, R., Sudarshan, A., and Tewari, M. (2021). The impact of temperature on productivity and labor supply: Evidence from Indian manufacturing. *Journal of Political Economy*, 129(6):1797–1827.

Tarsia, R. (2025). Heterogeneous effects of weather shocks on firm economic performance. Working Paper 414, Grantham Research Institute on Climate Change and the Environment, London School of Economics.

World Bank (2025). Climate change knowledge portal. World Bank Group. Country-level aggregates of ERA5 and CRU reanalysis.

Zhang, P., Deschênes, O., Meng, K., and Zhang, J. (2018). Temperature effects on productivity and factor reallocation: Evidence from a half million Chinese manufacturing plants. *Journal of Environmental Economics and Management*, 88:1–17.

Appendix A Trade Channel: Additional Results

This appendix does two things. Sections A.1 and A.2 present graphical summaries of the exposure measures and the estimates, which are easier to read than the tables in the text. Sections A.3 and A.4 report analyses that were conducted but are not part of the confirmatory set defined in Section 4.2, so that the reader can see the full extent of what was estimated.

A.1 The Geography of Exposure

Figure A1 maps the plant-level exposure measure: the 1,399 overseas manufacturing plants of Korean listed firms located in Section 3.7, coloured by tropical cyclone exposure. The concentration in coastal East and Southeast Asia is what makes the ownership channel worth measuring separately, and the geography of storm exposure across the sample—dense in the Philippines, Vietnam and southern China, absent in inland Europe—is the informal check on the geocoding reported in Section 3.7.

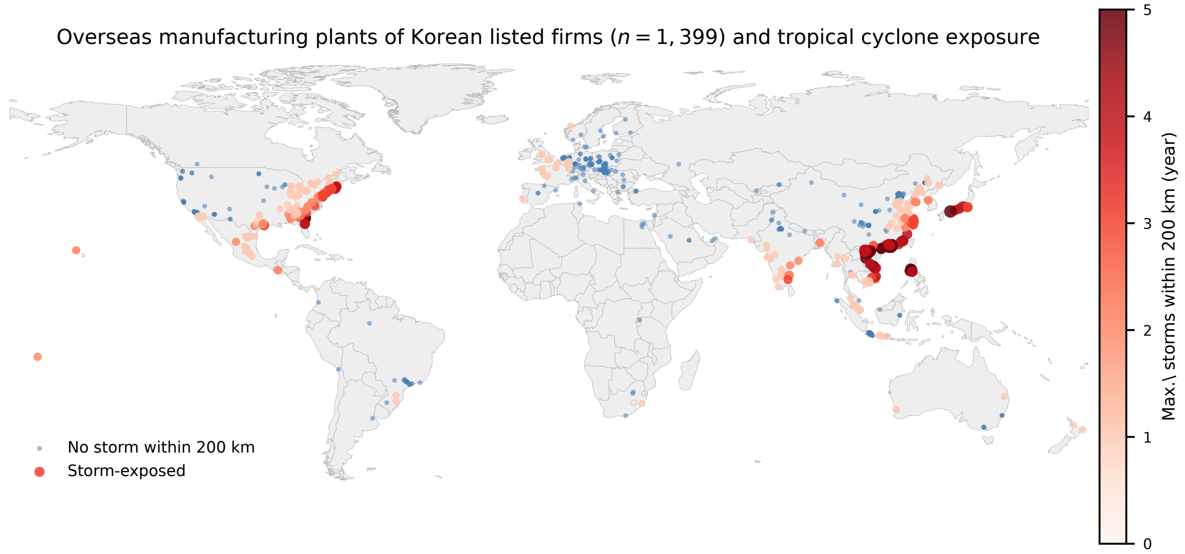

Figure A1: Overseas plants of Korean listed firms and cyclone exposure
Notes. Each point is one overseas manufacturing subsidiary of a Korean listed firm, geocoded to city level or better from disclosure filings. Red points experienced at least one tropical cyclone of tropical-storm strength passing within 200 km in some sample year, with colour intensity and size increasing in the maximum annual storm count; blue points never did. Storm tracks are from IBTrACS (Knapp et al., 2010), 1995–2025.

Figure A2 shows extreme heat in the principal supplier economies, in levels and as anomalies against the 1961–1990 baseline. The two panels make visible a point that matters for interpreting the heat result of Section 6: China experiences far fewer days above 35°C than India, Thailand or Vietnam in levels, so an estimate driven by Chinese exposure is not picking up the suppliers where extreme heat is most severe.

A.2 Estimates at a Glance

Figure A3 plots every coefficient from Table 7 together with its confidence intervals, including the geophysical placebo. The figure conveys what the table cannot at a glance: the estimates cluster tightly around zero, the intervals for the storm and drought measures are narrow enough to exclude economically meaningful effects, and the flood intervals are visibly wider—which is the imprecision that Section Appendix D addresses. The two heat coefficients are the only ones whose 95 percent intervals exclude zero, and the earthquake placebo sits on zero with an interval comparable in width to theirs.

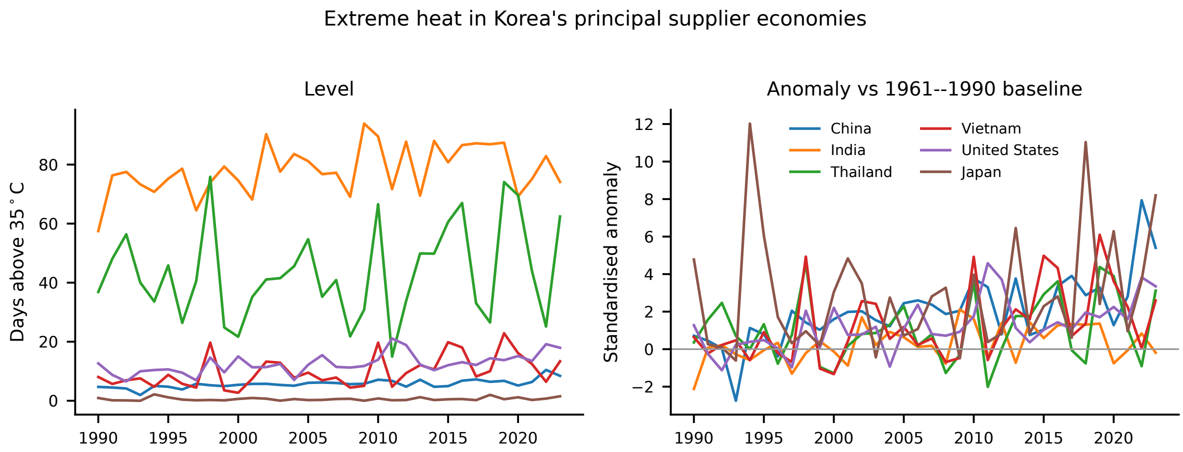

Figure A2: Extreme heat in Korea's principal supplier economies
Notes. Left panel, annual days with maximum temperature above 35°C. Right panel, the same series expressed as standardised anomalies relative to each country's own 1961–1990 mean and standard deviation, which is the transformation used to construct *FPR*. Source: ERA5 reanalysis aggregated by World Bank (2025).

A.3 Diversification Interaction

Section 7.1 argues that diversified sourcing dilutes exposure. The natural test is whether propagation is stronger in industries whose sourcing is concentrated, which we implement by interacting *FPR* with an indicator for below-median partner-country counts. Table A1 reports the interaction coefficients.

Table A1: Exposure interacted with low sourcing diversification

Physical risk measure	$\beta_{FPR\times\text{low div.}}$	Bootstrap p	Direction
Drought	+0.0155	0.030	as predicted
Disaster count	−0.0274	0.010	against prediction
Flood	−0.0130	0.044	against prediction

Notes. Coefficients on the interaction between standardised *FPR* and an indicator for industries below the median number of partner countries, added to equation (6). Bootstrap p-values from 9,999 Rademacher draws. Dependent variable sd(log *MRPK*).

The results do not support a clean interpretation. All three interactions are significant but their signs are inconsistent, and reporting only the drought coefficient would be selective. We do not treat any of the three as evidence.

The likely reason is that the low-diversification group is effectively an industry group rather than a sourcing category. Splitting seven industries at the median separates textiles and leather (8.5 partner countries) and wood, paper and printing (13.1) from the remainder (18–23), so the interaction largely reproduces an industry dummy and picks up whatever else is idiosyncratic to those industries' time series. A test of the diversification mechanism requires within-industry variation in sourcing concentration, which the size-

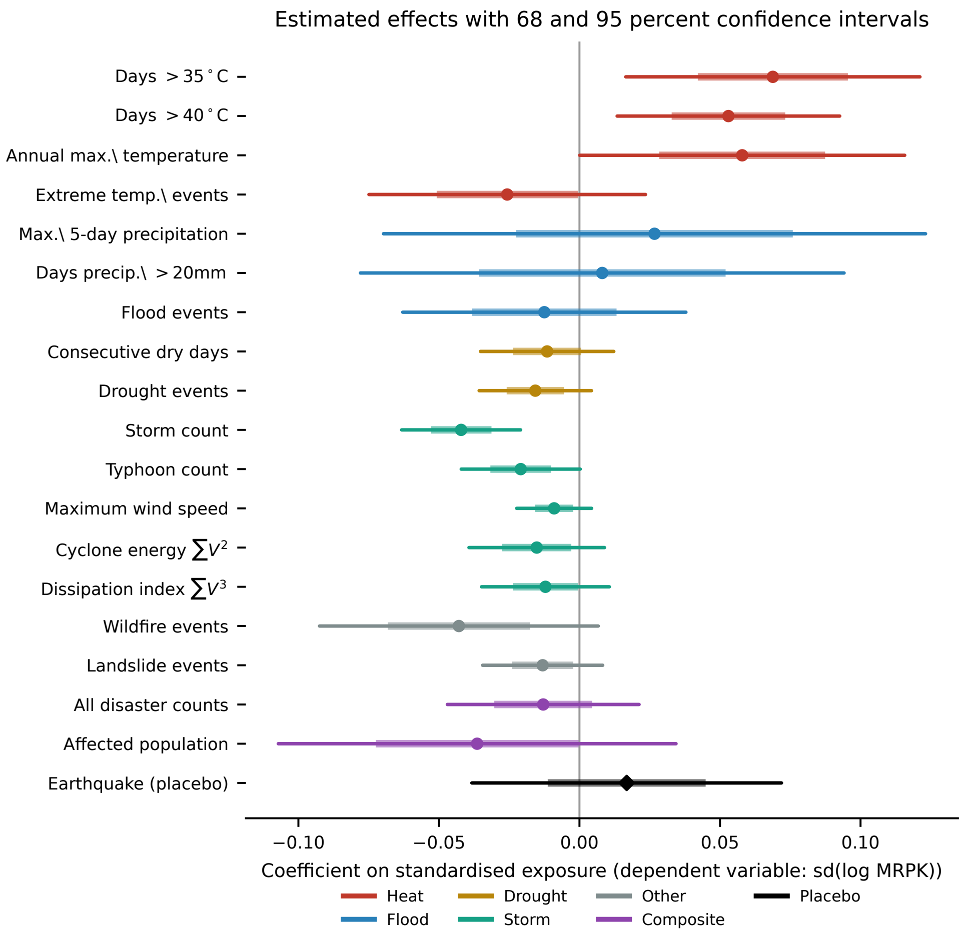

Figure A3: Estimated effects of foreign physical risk, by hazard measure
Notes. Point estimates from equation (6) with sd(log $MRPK$) as the dependent variable and standardised exposure, so that coefficients are comparable across hazards. Thick bars are 68 percent and thin bars 95 percent confidence intervals, both computed from industry-clustered standard errors. The earthquake measure, plotted as a diamond, is the geophysical placebo of Section 6.

bin specification of equation (7) provides in principle; there, the sourcing concentration measure is insignificant (Table B2).

We also tried the continuous version, interacting *FPR* with the effective number of suppliers (1/HHI) and, separately, with the inventory-to-shipment ratio, on the fourteen-sector panel. None of the six interactions is significant, and the two that come closest carry the wrong sign ($+0.010$, $p = 0.075$ for drought interacted with inventories; $+0.007$, $p = 0.063$ for drought interacted with diversification).

We do not read this as evidence against buffering, because the test cannot deliver such evidence. An interaction asks whether a buffer *attenuates* an effect, and the main effects here are null throughout; there is no attenuation to detect. Establishing that diversification or inventories buffer propagation would require a setting in which propagation is observed among the unbuffered, which this sample does not provide. The quantitative argument of Section 5.3 is the alternative route we take instead, and it does not depend on an interaction being estimable.

A.4 Heat and Baseline Climate

Section 5.4 shows that the heat coefficient disappears when exposure is measured in days rather than standardised anomalies. One economically motivated reading of that sensitivity is adaptation: an additional hot day should matter more in a country where such days are rare and production is not built for them, and standardising by each country's own variability is one way of encoding exactly that. On this reading the standardised measure is the right one and the level measure is wrong.

We tested the implication directly, splitting foreign exposure by whether a supplier's baseline exceeds five days above 35°C per year and measuring the hazard in days. For the 35°C measure the ordering is as adaptation predicts: suppliers with cool baselines give $+0.096$ ($p = 0.038$) and those with hot baselines -0.028 ($p = 0.211$). Against this, the 40°C measure shows nothing in either group (-0.001 and -0.032), and within the cool-baseline group the significance is again carried by China alone, whose own coefficient ($+0.079$, $p = 0.042$) is larger than that of suppliers with still cooler baselines ($+0.027$, $p = 0.127$).

The adaptation reading therefore organises part of the pattern and not the rest. We report it because it is the most plausible economic account of a result we otherwise cannot explain, and we do not promote it to the main text because it fails at the stricter threshold and does not remove the country-specificity.

Appendix B Further Specifications

Table B2 reports equation (7), in which industry-by-year fixed effects absorb all common shocks and identification rests on differences across size bins within an industry-year. The coefficients are indistinguishable from zero and uniformly negative across the three weather measures. The sourcing concentration measure, included because concentrated sourcing is the condition under which propagation is most plausible, is likewise insignificant.

The heat coefficient in this table requires comment, since it is negative and insignificant while Table 7 reports a positive and significant one. There is no contradiction, because the two specifications estimate different objects. Industry-by-year fixed effects absorb, by construction, everything that moves an industry's establishments together within a year—including an industry's exposure to heat abroad, which is common to all size bins that source from the same countries. Equation (7) can only detect an effect that operates *differentially across size bins*, and is uninformative about one that operates at the industry-year level. Its role here is to provide a demanding test for the hazards where Table 7 finds nothing, not to adjudicate the heat result.

Table B2: Size-bin specification with industry-by-year fixed effects

Exposure measure	β	Cluster s.e.	t	Asym. p	Boot. p	N
Flood	−0.0069	(0.0662)	−0.10	0.920	0.953	208
Drought	−0.0113	(0.0542)	−0.21	0.842	0.672	208
Heat (days > 35°C)	−0.0127	(0.0150)	−0.85	0.430	0.594	208
Sourcing concentration (HHI)	+0.0702	(0.0649)	+1.08	0.321	0.500	234

Notes. Dependent variable is sd(log $MRPK$). Estimates of equation (7), 2016–2024, seven TEC industry groups × four size bins. All specifications include industry-by-year and size-bin fixed effects. Standard errors clustered by industry (7 clusters); bootstrap enumerates all 128 sign vectors. $R^2 =$ 0.80–0.81. Exposure by size bin is recovered from TEC marginals by iterative proportional fitting.

The heat coefficient passed the filters of Section 6. One coefficient fails them, and the contrast is what justifies the procedure.

Among the interaction specifications, the interaction between flood exposure and GVC backward participation is significant under asymptotic inference. Adding main effects and the interaction to equation (6) gives

$$\hat{\beta}_{FPR\times GVC} = 0.0152 \ \ (0.0053), \qquad t = 2.86, \qquad p_{\text{asym}} = 0.013,$$

suggesting that flood exposure raises dispersion more in GVC-intensive industries. Under the wild cluster bootstrap on the same specification, enumerating all 16,384 sign vectors, the p-value is **0.200**.

The bootstrap moves this coefficient in the direction opposite to heat: there, imposing the null and enumerating sign vectors took the p-value from 0.023 to 0.007; here it takes it from 0.013 to 0.200. With fourteen clusters the bootstrap is the procedure to believe in both cases, and it delivers opposite verdicts on two coefficients that asymptotic inference would have treated alike.

We report the failed one rather than omit it because it shows what the inference procedure is doing. An analyst reporting asymptotic p-values alone would have concluded that imported flood risk propagates in globally integrated industries—and, having found flood significant, might never have looked at heat. We do not treat the interaction as suggestive evidence, and no part of the argument that follows relies on it.

Table B3 re-estimates equation (6) with two alternative dependent variables: the dispersion of total factor productivity revenue, sd(log $TFPR$), which reflects distortions to both capital and labour, and the Hsieh and Klenow (2009) reallocation gain, the counterfactual output increase from equalising marginal products.

The division seen in Table 7 reappears. Both threshold heat measures enter positively on both alternative outcomes, significantly or marginally so; the annual maximum temperature, flood and drought are null throughout. That heat's coefficient on sd(log $TFPR$) is smaller than on sd(log $MRPK$) is consistent with the capital margin being the one primarily affected.

Table B3: Alternative dependent variables

	sd(log $TFPR$)		Reallocation gain	
Hazard measure	β	p	β	p
Days > 35°C	+0.0380 (0.0183)	0.059	+0.1253 (0.0646)	0.074
Days > 40°C	+0.0377 (0.0216)	0.105	+0.1595 (0.0736)	0.049
Annual maximum temperature	+0.0188 (0.0214)	0.395	+0.0941 (0.0925)	0.328
Flood (max. 5-day precip.)	+0.0027 (0.0348)	0.938	−0.0646 (0.0797)	0.433
Drought (consecutive dry days)	−0.0015 (0.0079)	0.856	+0.0290 (0.0290)	0.335
Disaster count	−0.0162 (0.0179)	0.384	−0.0559 (0.0490)	0.275
Disaster-affected population	−0.0251 (0.0441)	0.579	−0.1313 (0.1504)	0.398

Notes. Estimates of equation (6) with the dependent variable replaced. Industry and year fixed effects throughout; standard errors in parentheses, clustered by industry; p-values asymptotic. $N = 196$–210.

Taken together, the results describe a null across most of the grid. Imported flood, drought, storm, wildfire and landslide risk leave the allocation of capital across Korean manufacturing establishments undisturbed, on every outcome and under both specifications, and those nulls are insensitive to how exposure is normalised and to whether the hazard is measured in standardised anomalies or in physical units.

Heat is the one hazard on which the evidence does not settle. The coefficient is positive and significant, and it holds across several constructions of the dispersion measure. It is also confined to Chinese suppliers, absent among suppliers that are far hotter, and it survives neither an alternative normalisation of exposure nor an alternative standardisation of the hazard—measurement choices to which the drought and flood nulls are wholly indifferent. We report the checks on both sides in Section efsec:power-heat without adjudicating between them: what the data support is that the evidence on heat is divided, not that heat propagates or that it does not.

Appendix C The Full Outcome Grid

The outcome used so far is a dispersion measure, and the literature reviewed in Section 2 uses levels. That difference invites an obvious objection: perhaps imported physical risk does move output, sales or employment, and we have simply not looked. We therefore construct every outcome the establishment microdata supports and estimate equation (6) on all of them.

Real quantities come first, because the survey reports values in current prices and year fixed effects absorb only the common component of inflation, not industry-specific price movements that vary in exactly the dimension our treatment does. Two of the outcomes need no deflator at all: the production and shipment indices published by Statistics Korea are real quantity indices published by KSIC division, and we aggregate them to our sectors. For the monetary variables we construct an implicit deflator as nominal shipments divided by the real shipment index and apply it to value added, capital and the wage bill. It behaves as it should, rising by seventy-five percent over the sample in wood and paper and by five percent in electronics.

The remaining outcomes need no deflation for a different reason. Employment and establishment counts are physical units. The dispersion and concentration measures are algebraically invariant to any industry-year deflator, which shifts $\log MRPK$ by a constant within the cell and therefore leaves its standard deviation, interdecile range, skewness and value-added shares unchanged—a property of the Hsieh-Klenow outcome worth stating, since it means the paper's main results carry no price-measurement risk at all.

Beyond the three dispersion measures already reported, the panel adds the variance of $\log MRPK$, which is the exact object of Liu and Xu (2025); the interdecile range and skewness of $\log MRPK$; and $\mathrm{sd}(\log MRPL)$, which measures the same distortion on the labour rather than the capital margin. Real output corresponds to Zhang et al. (2018) and to the sales variable of Addoum et al. (2020); employment, capital and the establishment count to Tarsia (2025); labour productivity to Somanathan et al. (2021); and the Herfindahl index and top-decile share to Ponticelli et al. (2023).

Table C4 reports, for each outcome, the hazard that fits it best—the smallest bootstrap p-value among the thirteen—together with how many of the thirteen reach five percent. Reporting the minimum is deliberately generous to the alternative: it is the largest association the data offer for that outcome, selected after the fact.

Twenty-five of the 351 regressions reach five percent, against 17.6 expected under the null, and neither the Benjamini-Hochberg procedure at $q = 0.05$ nor a Bonferroni correction leaves a single survivor. The outcomes the literature uses behave no differently from the dispersion measures we began with. Real output is the clearest case: the production and shipment indices, which require no deflator and are measured independently of our data, are unmoved by every hazard, with smallest p-values of 0.175 and 0.187.

Whether that pattern is informative depends on how large an effect the design could have detected, which the next subsection establishes with minimum detectable effects; Section 7 then asks why so little propagates.

Appendix D Improving the Weakest Null

Flood is the one hazard for which Table 8 licenses no conclusion: at 0.49σ, an effect raising the dispersion of marginal products by nearly ten percent would have gone undetected. Since floods are the hazard most often assumed to disrupt supply chains, leaving the matter there would be unsatisfying, and the imprecision turns out to be remediable.

Its source is mechanical. Flood exposure is more collinear with the fixed effects than drought exposure (93 percent of its variance absorbed against 88), because heavy-precipitation anomalies are more synchronised across supplier countries within a year. More importantly, its residuals are strongly correlated within industry clusters, inflating the clustered standard error relative to the homoskedastic one by a factor of 2.7 against 0.9 for drought—and essentially all of that inflation comes from one industry. Excluding coke and refined petroleum, whose exposure measure covers the smallest share of actual sourcing (Section 3.6), cuts the flood standard error from 0.049 to 0.021.

Table D5 pursues this together with alternative measurements of the same hazard.

Two routes roughly halve the minimum detectable effect, and they are independent of one another. Restricting to industries whose exposure is well measured gives 0.25σ; replacing the precipitation proxy with counts of realised flood disasters, on the full sample, gives 0.26σ. That two different corrections converge on the same figure is more reassuring than either alone. Every specification remains statistically indistinguishable from zero, and six of the eight carry negative point estimates.

The flood null can therefore be stated with roughly the same confidence as the others: effects larger than about five percent of the dispersion of marginal products are ruled out, rather than the ten percent that Table 8 alone would suggest. We retain the conservative full-sample precipitation figure in Table 8 and report the improvement here.

One conjecture failed and is worth recording, since it bears on how the measures relate. We expected that averaging the three flood measures would cancel their independent measurement errors and improve precision. It does the opposite: the composite is the *least* precise specification in the table (0.53σ). Precipitation intensity, precipitation frequency and disaster counts evidently capture different things rather than the same thing with

Table C4: Every outcome the microdata supports, against every hazard

Outcome	N	Best-fitting hazard of thirteen			
		Hazard	β	Boot. p	$p < 0.05$
Real output and deflated levels					
log production index	196	Heat, days > 35°C	−0.0691	0.175	0/13
log shipment index	196	Heat, days > 35°C	−0.0691	0.187	0/13
log real value added	192	Heat, days > 35°C	−0.1432	0.135	0/13
log real shipments	192	Heat, days > 35°C	−0.0743	0.160	0/13
log real capital	192	Drought, consec. dry days	−0.0265	0.234	0/13
log real wage bill	206	Wildfires (count)	−0.0999	0.208	0/13
log real labour productivity	206	All disasters (count)	−0.0845	**0.021**	5/13
Δ log production index	182	Landslides (count)	−0.0211	0.058	0/13
Δ log shipment index	182	Landslides (count)	−0.0207	**0.042**	2/13
Δ log real value added	160	Flood, days > 20mm	+0.0610	**0.002**	2/13
Δ log real shipments	160	Flood, days > 20mm	+0.0230	**0.005**	2/13
Δ log real capital	174	Floods (count)	−0.0704	**0.048**	1/13
Physical quantities					
log employment	196	Heat, annual max. temp.	−0.1726	**0.006**	1/13
log establishments	196	Heat, annual max. temp.	−0.1134	**0.049**	1/13
Δ log employment	168	Drought, consec. dry days	−0.0070	0.283	0/13
Δ log establishments	182	Landslides (count)	−0.0281	**0.009**	1/13
Net investment rate	210	Droughts (count)	+0.0285	0.133	0/13
log capital intensity	210	Floods (count)	−0.1127	**0.025**	1/13
Misallocation					
sd(log *MRPK*)	196	Heat, days > 40°C	+0.0469	**0.015**	2/13
var(log *MRPK*)	196	Heat, days > 40°C	+0.1976	**0.016**	2/13
sd(log *MRPL*)	210	Floods (count)	−0.0065	0.079	0/13
sd(log *TFPR*)	196	Heat, days > 40°C	+0.0339	0.056	0/13
Reallocation gain	196	Heat, days > 40°C	+0.1501	**0.030**	1/13
p90–p10 of log *MRPK*	196	Heat, days > 40°C	+0.1924	**0.039**	3/13
Skewness of log *MRPK*	210	Landslides (count)	+0.0571	**0.027**	1/13
Concentration					
HHI of value added	210	Wildfires (count)	+0.0247	0.223	0/13
Top-10% VA share	210	Droughts (count)	−0.0085	0.258	0/13

Notes. Each cell of the grid is an estimate of equation (6) with the dependent variable and hazard replaced, industry and year fixed effects throughout, standard errors clustered by industry (14 clusters), and bootstrap p-values from full enumeration of all 2^{14} Rademacher sign vectors with the null imposed. The table reports the hazard giving the smallest bootstrap p-value for each outcome, which is a selected maximum rather than a pre-specified test, and the final column counts how many of the thirteen hazards reach five percent. Of the 351 regressions shown, 25 do, against 17.6 expected under the null; none survives Benjamini-Hochberg at $q = 0.05$ or a Bonferroni correction. Production and shipment indices are from Statistics Korea (KOSIS DT_1F02001, 2020 = 100), aggregated from KSIC divisions to our sectors by value-added weights; monetary variables are deflated by nominal shipments over the real shipment index. The current-price counterparts of the deflated rows give the same conclusion and are reported in the replication file. Outcomes are built from the Mining and Manufacturing Survey as described in Section 3.5; the dispersion measures reproduce the series used in Table 7 with a correlation of 0.98.

Table D5: Alternative flood measures and samples

Measure	Sample	β	Boot. p	MDE	Implied dispersion
Max. 5-day precipitation	All industries	$+0.0267$	0.865	0.49σ	9.6%
	Excl. petroleum	-0.0306	0.215	$\mathbf{0.25}\boldsymbol{\sigma}$	4.2%
Days precip. > 20 mm	All industries	$+0.0081$	0.787	0.44σ	8.6%
	Excl. petroleum	-0.0302	0.380	0.35σ	5.8%
Flood events (EM-DAT)	All industries	-0.0125	0.666	$\mathbf{0.26}\boldsymbol{\sigma}$	5.0%
	Excl. petroleum	-0.0381	0.124	0.29σ	4.8%
Composite index	All industries	$+0.0091$	0.866	0.53σ	10.3%
	Excl. petroleum	-0.0465	0.185	0.38σ	6.3%

Notes. Dependent variable sd(log $MRPK$). "Implied dispersion" converts the MDE into the percentage change in sd(log $MRPK$) at its sample mean, so that the columns read as the smallest effect the design would detect. The composite index averages the standardised values of the three component measures. Petroleum exclusion is motivated by the coverage diagnostic of Section 3.6, which is fixed before estimation.

noise, so averaging them adds variance instead of removing it.

The deeper limitation is not the choice of index but spatial aggregation. A country-average five-day precipitation maximum says little about whether a large country's industrial areas flooded, since most of its territory contains no factories. The principled remedy is production-weighted flood exposure, built from satellite flood extents such as the Global Flood Database, Dartmouth Flood Observatory event polygons, or the geocoded disaster locations of Rosvold and Buhaug (2021), weighted by gridded production. That is beyond the data assembled here for the imported-input channel, which is irreducibly country-level. For the ownership channel it is already done: exposure there is evaluated at plant coordinates (Section 3.7), the flood measure is five-day precipitation at the plant itself, and Section 5.1 reports the trace it leaves.

Appendix E Ownership Channel: First-Pass Design

The trade channel measures exposure through imported inputs, aggregated by country and by industry, which leaves two objections open. First, aggregation may simply dilute the signal: a national disaster index for a large, internally diversified supplier country averages over many locations, most of them unaffected. Second, imports are not the only channel through which foreign physical risk reaches Korean balance sheets. Korean manufacturers own and operate production plants abroad, and damage to an owned plant reaches the parent directly rather than through the prices of arm's-length trade. The first design we ran on the plant panel addressed both at once by constructing plant-level exposure from geocoded subsidiary addresses and testing whether it moves parent-firm outcomes.

E.1 Firm-Level Exposure

Section 3.7 described the construction of the plant panel: 1,399 overseas manufacturing plants in 55 countries, belonging to 474 Korean listed parents, with city-level or better coordinates, and tropical cyclone tracks from IBTrACS (Knapp et al., 2010) evaluated at those coordinates. A plant-year's exposure is the number of distinct storms of tropical-storm strength or above passing within 200 km. Firm-level exposure, $FPR^{FDI}_{i,t}$, is the average of plant exposures weighted by disclosed subsidiary assets, so that a storm affecting a firm's principal overseas plant counts for more than one affecting a minor one.

E.2 A Built-In Placebo: Consolidated versus Separate Statements

Korean firms file two sets of financial statements: consolidated statements (CFS), which incorporate overseas subsidiaries, and separate statements (OFS), which cover the parent alone. If typhoon strikes on overseas plants damage operations, the damage must appear in consolidated outcomes and should be absent from separate ones. Estimating every specification on both sets therefore provides a placebo test within the same firms and years. We collect both statements for the 474 parents with locatable overseas plants over 2015–2024 and estimate

$$y_{i,t} \;=\; \beta\, FPR^{FDI}_{i,t} \;+\; \delta_i \;+\; \delta_t \;+\; \varepsilon_{i,t}, \tag{9}$$

with firm and year fixed effects and standard errors clustered by firm, for three outcomes: asset growth, revenue growth, and the tangible-asset share of total assets (an investment-intensity measure).

Table E6 reports the results. Every coefficient is small and statistically indistinguishable from zero, in the consolidated and separate statements alike; the smallest p-value in the table is 0.20. The samples are not small—roughly 450 firms and 3,500 firm-year observations per specification—so the nulls are not a power artefact of an exotic subsample.

The placebo structure never comes into play: there is no consolidated effect for the separate statements to fail to match. What the design delivers instead is corroboration of the aggregate null at a much finer resolution, for the hazard it covers. Imported storm risk does not move capital allocation in Table 7, and storm risk striking a firm's own plants—measured at plant coordinates rather than country aggregates—does not move its balance sheet either. The resolution critique of country-year exposure measures therefore cannot explain the storm null.

Two features of this first pass explain why the main text uses a different design. Its exposure is continuous and asset-weighted, so a storm at a minor plant and a storm at the principal plant enter in proportion to assets, and a firm-year with several weak storms can

Table E6: Typhoon exposure of overseas plants and parent-firm outcomes

Outcome	Statements	β	Cluster s.e.	p	N
Δ log(assets)	Consolidated	+0.0019	(0.0093)	0.834	3,497
	Separate	+0.0023	(0.0093)	0.806	3,566
Δ log(revenue)	Consolidated	−0.0034	(0.0150)	0.823	3,494
	Separate	−0.0067	(0.0138)	0.628	3,552
Tangible assets / total assets	Consolidated	−0.0042	(0.0033)	0.201	3,484
	Separate	−0.0022	(0.0033)	0.507	3,548

Notes. Estimates of equation (9), 2015–2024. $FPR^{FDI}_{i,t}$ is the asset-weighted number of tropical cyclones (winds $\geq$ 34 kt) passing within 200 km of firm i's overseas manufacturing plants in year t. All regressions include firm and year fixed effects; standard errors are clustered by firm (450–455 clusters). Consolidated statements include overseas subsidiaries; separate statements cover the parent alone. Financial statements and subsidiary tables are from the DART electronic disclosure system; storm tracks from IBTrACS (Knapp et al., 2010).

score higher than one with a single destructive strike; the event indicators of Section 4.1 replace it with a threshold at the principal plant, where a hit is a hit. Its outcomes are the consolidated and separate statements taken one at a time, so the placebo is read across two regressions; the difference outcome of equation (4) puts the comparison inside one regression and removes the parent's own domestic shocks with it. And it covers tropical cyclones only, the hazard for which public track data existed before the daily reanalysis at plant coordinates was assembled; the main text adds heat, flood, drought and compound events at the same sites. The first-pass null is consistent with what the sharper design finds for typhoons, which is nothing, and it is reported here because it was run before the design was fixed and its null is part of the record.

Appendix F Ownership Channel: Further Estimates

Sections 5.1 and 5.2 report the event study at the principal plant and the outcomes on which it finds nothing. This appendix adds the alternative exposure definitions and firm splits that the heterogeneity discussion draws on, and the plant-level responses—asset growth and exit between the 2022 and 2024 filings—that bound what happens at the plant itself.

F.1 Exposure Definitions and Firm Splits

Table F7 re-estimates the static coefficient of equation (5) on the difference outcome under alternative definitions of exposure and across fixed firm splits. Replacing the principal-plant indicator with an indicator for an event at any of the firm's plants removes the signal, which is why the main text reads the trace as specific to the principal plant rather than as a portfolio effect. Across splits, the flood and drought coefficients are larger for multi-plant firms than for single-plant firms, the opposite of what dilution across plants would produce, and this is the comparison Section 5.1 relies on to reject the dilution reading.

Table F7: Exposure definitions and firm splits: change in log assets, consolidated minus separate

Hazard	Exposure / split	β	Boot. p	MDE	Events
Flood $\rightarrow \Delta \log$ *assets, consolidated* $-$ *separate*					
	Principal plant hit	-0.0215^{**}	0.043	0.33	131
	Any plant hit	-0.0089	0.232	0.22	287
	Principal sales office hit	$+0.0191$	0.304	0.54	53
	All firms	-0.0215^{**}	0.048	0.33	131
	Single-plant firms	-0.0261	0.138	0.54	72
	Multi-plant firms	-0.0195	0.137	0.41	59
	Single-country firms	-0.0253	0.133	0.48	81
	Multi-country firms	-0.0222^{*}	0.097	0.45	50
Drought $\rightarrow \Delta \log$ *assets, consolidated* $-$ *separate*					
	Principal plant hit	-0.0210^{**}	0.047	0.31	156
	Any plant hit	-0.0024	0.728	0.21	414
	Principal sales office hit	-0.0089	0.557	0.45	170
	All firms	-0.0210^{**}	0.046	0.31	156
	Single-plant firms	-0.0058	0.721	0.47	78
	Multi-plant firms	-0.0301^{**}	0.035	0.42	78
	Single-country firms	-0.0131	0.384	0.41	85
	Multi-country firms	-0.0263^{*}	0.073	0.46	71
Typhoon $\rightarrow \Delta \log$ *assets, consolidated* $-$ *separate*					
	Principal plant hit	-0.0155	0.121	0.30	202
	Any plant hit	$+0.0031$	0.655	0.21	413
	Principal sales office hit	$+0.0101$	0.276	0.29	108
	All firms	-0.0155	0.123	0.30	202
	Single-plant firms	-0.0122	0.405	0.43	123
	Multi-plant firms	-0.0209	0.132	0.44	79
	Single-country firms	-0.0126	0.379	0.40	131
	Multi-country firms	-0.0219^{*}	0.093	0.43	71
Any event $\rightarrow \Delta \log$ *assets, consolidated* $-$ *separate*					
	All firms	-0.0159^{**}	0.017	0.20	548
	Single-plant firms	-0.0127	0.200	0.29	278
	Multi-plant firms	-0.0201^{**}	0.029	0.28	270
	Single-country firms	-0.0153	0.102	0.27	300
	Multi-country firms	-0.0187^{**}	0.036	0.29	248

Notes. Static β_0 from equation (5). "Any plant hit" replaces the principal-plant indicator with an indicator for an event at any of the firm's plants. Firm splits are fixed at the firm level. The split by principal-plant asset share is omitted because its median is one: for more than half of firms the principal plant holds all disclosed overseas plant assets.

F.2 Plant-Level Assets and Exit

The consolidated trace could reflect a write-down at the plant that was hit or something else on the consolidated balance sheet. The subsidiary tables report each plant's assets in every filing, so the question can be asked of the plant directly, at the cost of a two-year window and a cross-section. Table F8 reports asset growth between the 2022 and 2024 filings for plants hit in 2022, and disappearance from the 2024 filings. Plant assets fall by eighteen percent after any event ($p = 0.088$), with no single hazard significant and minimum detectable effects of 0.66 to 1.15σ; exit is 1.8 percentage points higher for hit plants ($p = 0.65$), against a baseline disappearance rate of about fifteen percent for plants with exposure, and the design would detect only a doubling of that rate. For comparison, raw disappearance rates between filings are 12.9 percent for manufacturing plants and 11.7 percent for sales offices, both upper bounds inflated by residual formatting artefacts in the filings, and the similarity of the two is a first indication that physical risk does not drive closure over this period. These are bounds on the plant-level mechanism rather than estimates of it; the direction is consistent with a write-down at the plant, and the sample is what the disclosure cycle allows.

Table F8: Plant-level responses between the 2022 and 2024 filings

Outcome	Hazard	FE	β	Boot. p	MDE	Hit plants
Δ log plant assets 2022–24	Any event	country	-0.132^{*}	0.071	0.36	193
Δ log plant assets 2022–24	Any event	country+parent	-0.183^{*}	0.088	0.66	193
Δ log plant assets 2022–24	Typhoon	country	-0.055	0.474	0.39	75
Δ log plant assets 2022–24	Typhoon	country+parent	-0.120	0.238	0.66	75
Δ log plant assets 2022–24	Heat	country	-0.141	0.290	0.64	101
Δ log plant assets 2022–24	Heat	country+parent	-0.139	0.477	1.15	101
Δ log plant assets 2022–24	Flood	country	$+0.079$	0.541	0.61	20
Δ log plant assets 2022–24	Flood	country+parent	-0.202	0.159	0.81	20
Δ log plant assets 2022–24	Drought	country	-0.195	0.155	0.66	31
Δ log plant assets 2022–24	Drought	country+parent	-0.154	0.333	0.99	31
Exit by 2024 (pp)	Any event	country	$+1.8$	0.648	11	180
Exit by 2024 (pp)	Any event	country+parent	-0.3	0.963	17	180
Exit by 2024 (pp)	Typhoon	country	-4.8	0.380	15	68
Exit by 2024 (pp)	Typhoon	country+parent	-7.2	0.418	27	68
Exit by 2024 (pp)	Heat	country	$+8.3$	0.250	20	94
Exit by 2024 (pp)	Heat	country+parent	$+6.4$	0.426	27	94
Exit by 2024 (pp)	Drought	country	$+12.3$	0.147	23	32
Exit by 2024 (pp)	Drought	country+parent	$+13.5$	0.180	30	32

Notes. Cross-sections of plants listed in the 2022 filings. Asset growth is for the 920 plants with assets reported in both filings after converting each filing's stated unit to won and trimming $|\Delta \log| > 2$; exit is disappearance from the 2024 filings for 1,189 plants with exposure, including 176 exited plants geocoded from their 2022 addresses. Events are in 2022. Standard errors clustered by parent; MDE in standard deviations of the outcome (percentage points for exit).