Building Corporate Resilience to Supply Chain Disruptions

Leandro Sanz*

April 17, 2026

Abstract

I examine how exposure to disruption risk from hard-to-replace inputs affects corporate resilience investments and firms' learning about that risk. Using a new dataset on 11,000 foreign suppliers to U.S. manufacturers, I show that firms with fewer alternative suppliers hold more inventory and less cash, and have higher leverage. Exploiting natural disasters that disrupt suppliers, I find that firms update their beliefs about disruption risk and make persistent changes to corporate policies in response. Consistent with learning, the response is strongest after first-time shocks. Finally, firms with higher inventory buffers are better protected against performance losses when disruptions occur.

JEL Codes: G31, G32, F23, L23

Keywords: Supply chain risk; corporate liquidity; inventories; capital structure; operational hedging; learning under uncertainty

*University of Notre Dame, Mendoza College of Business, Notre Dame, IN 46556. Email: lsanz@nd.edu. I am deeply grateful for the invaluable guidance and support of my committee: Isil Erel, Rüdiger Fahlenbrach, René Stulz, and Michael Wittry. I also thank Kenneth Ahern (discussant), Taylor Begley, Itzhak Ben-David, Daniel Carvalho, Nuri Ersahin (discussant), Kayla Freeman (discussant), Janet Gao (discussant), Gerard Hoberg, Maria Marchica (discussant), Holger Mueller, Mikael Paaso (discussant), Sergio Salgado, Luke Taylor, Alireza Tahbaz-Salehi, Andrea Vedolin, Michael Weisbach, Yufeng Wu, Lu Zhang, and seminar participants at the 9th Annual Marshall PhD Conference in Finance, the 2023 EFA Doctoral Tutorial, the Nova SBE Finance PhD Final Countdown, the 2023 FMA Doctoral Consortium, the 2023 Rising Scholars Conference held at MIT, the 36th Australasian Finance and Banking conference, the 2024 AFA Annual Meeting, the 2024 Eastern Finance Association Annual Meeting, the 2024 NFA Annual Meeting, the 2025 SFS Cavalcade North America, the Ohio State University, the University of Southern California, the University of Pennsylvania, the University of Kentucky, Pennsylvania State University, Indiana University, Washington University in St. Louis, the University of Notre Dame, and Boston University for valuable comments and suggestions. All errors are my own.

1 Introduction

Tariffs, pandemics, and natural disasters expose firms to operational shocks that financial buffers alone may not fully absorb. The standard corporate finance prescription of maintaining liquidity to mitigate adverse shocks may therefore be less effective when the relevant constraint is not financial. If a shock restricts access to hard-to-substitute inputs, firms may be unable to purchase the inputs needed to sustain production even when they have ample liquidity. Therefore, the effects of such shocks on output and revenue may be large. Cash then becomes a less effective hedge, and firms may be better off relying on operational rather than financial buffers.

In this paper, I study how firms build resilience against disruptions of hard-to-replace inputs and how they learn about that risk. Because production networks can generate fragility that is difficult to observe ex ante (Elliott, Golub, and Leduc, 2022), firms may be uncertain about their exposure to input supply disruptions. Observed shocks can therefore reveal information about that risk and lead firms to revise their policies. I show that firms exposed to inputs with fewer suppliers hold more inventories and less cash, and have higher leverage, consistent with a reallocation from financial to operational buffers in response to input supply disruption risk. I also show that firms update these policies after disruptions to hard-to-replace inputs, and that their responses are strongest when such shocks reveal new information about risk.

I construct a new dataset on the global supply chains of publicly listed U.S. manufacturing firms from 2007 to 2019 using administrative bill of lading (BOL) records collected by U.S. Customs and Border Protection. These data provide shipment-level information on firms' maritime imports, allowing me to observe the specific inputs firms source and the foreign suppliers from which they obtain them. I show, using other data sources, that the maritime relationships captured in the BOL data are informative about firms' broader global supply chains.

Using these data, I develop a new measure of firms' ex ante exposure to disruptions in the supply of inputs, which I call supplier scarcity. Supplier scarcity captures the availability of alternative suppliers in the global input-supplier network. Inputs sourced from suppliers with few alternatives should be more difficult to replace following a disruption and, as a result, should expose firms to greater input supply risk.

I validate this assumption by showing that the supplier scarcity measure indeed identifies exposure to inputs that are difficult to replace after shocks. I refer to such inputs as fragile, and to inputs that can be readily sourced from alternative suppliers as generic. Using over 2,000 flood incidents as natural experiments, I find that shipments from affected suppliers decrease significantly in the weeks following a flood. The ability to reallocate sourcing differs substantially across input types. Firms increase shipments of generic inputs from unaffected suppliers to offset the disruption, but show no significant reallocation for fragile inputs. These results are consistent with supplier scarcity capturing real differences in firms' ability to substitute across suppliers shortly after disruptions.

Motivated by these findings, I develop a model that links firms' exposure to input supply disruptions to corporate financial policy. A firm combines a generic input, available on the spot market, with a fragile input, available only from scarce suppliers. The key friction is that when a disruption cuts off the supply of the fragile input, the firm cannot readily replace it with shipments from other suppliers, and imperfect substitutability prevents it from fully offsetting the loss by using more of the generic input. The evidence above that firms are less able to replace inputs with fewer alternative suppliers after disruptions provides empirical support for this mechanism.

The model implies that when disruptions to fragile inputs are harder to absorb, firms change the composition of their buffers. As disruptions become more costly, inventories of fragile inputs provide more effective insurance against production losses. In contrast, cash becomes less useful because it cannot immediately restore access to the missing input. Firms therefore substitute away from cash and toward inventories, and finance part of that adjustment with debt. These predictions distinguish corporate resilience investment from the standard precautionary savings motive, which would instead predict greater cash holdings if firms could costlessly convert liquidity into inputs when needed. Because firms may be ex ante uncertain about their true exposure to disruptions of fragile inputs, the model also incorporates a learning channel through which observed shocks lead firms to revise their beliefs and adjust their policies over time.

I test the model's predictions in two steps. First, I estimate the ex ante relationship between firms' exposure to disruptions of fragile inputs and corporate policies using cross-sectional variation in the firm-level, import-volume-weighted average scarcity of suppliers across a firm's inputs. Given that supplier scarcity is highly persistent within firms, I interpret this measure as an ex ante proxy of disruption risk and focus on cross-sectional differences across firms within industries rather than within-firm variation over time. Consistent with the model's predictions, I find that a one standard deviation increase in supplier scarcity is associated with 4.6% higher input inventory holdings, 13.0% lower cash holdings, and 5.7% higher book leverage relative to their respective means.

A concern with these estimates is that unobserved features of firms' production technologies may jointly determine supply chain structure and financial policy. For example, firms with more specialized production processes may both rely on inputs with fewer suppliers and hold more inventory for reasons unrelated to disruption risk. To address this concern, I construct an instrument for firm-level supplier scarcity based on changes in the number of globally active suppliers across the firm's input markets, aggregated using the firm's lagged import volume shares. Identification comes from time-series variation in supplier entry and exit in global product markets, which is plausibly exogenous to any given firm's financial policies. The instrument strongly predicts supplier scarcity in the first stage. The two-stage least squares coefficient estimates confirm the baseline results but are larger in magnitude, consistent with attenuation bias from measurement error in the observed scarcity measure.

The exclusion restriction requires that variation in the global supplier pool for a given input affect corporate policies only through firms' exposure to disruption risk. This restriction is plausible because the instrument is driven by aggregate supply-side changes in global product markets rather than by individual firms' choices. I nevertheless consider several potential threats to this restriction. The most direct is that a decline in the supplier pool could increase input prices, leading firms to adjust inventories, cash, and leverage in response to higher costs rather than greater disruption risk. I test this channel using a firm-level input price index constructed from the BOL data and find no significant relationship between the index and the instrument. I also rule out technology cycles and product obsolescence as alternative channels.

Several additional results corroborate the cross-sectional evidence. The relation between supplier scarcity and book leverage is driven primarily by higher accounts payable and greater use of revolving credit facilities, consistent with firms using short-term financing to build inventories. The results are also robust to controlling for import intensity, which addresses the concern that supplier scarcity may proxy for firms with lumpier ordering patterns (Hornok and Koren, 2015; Blum, Claro, Dasgupta, and Horstmann, 2019; Khan and Khederlarian, 2025). In addition, they are virtually the same when I interact supplier scarcity with a measure of the importance of foreign relative to domestic suppliers, suggesting that supplier scarcity captures an inherent feature of the production process that applies to both domestic and foreign sourcing of inputs.

A second prediction of the model is that firms learn about disruption risk from realized shocks and increase their investment in operational resilience over time. To test this prediction, I use supplier floods as exogenous information shocks and examine how firms adjust their corporate policies in response. When a flood unexpectedly disrupts a supplier, managers receive a signal about their exposure to disruption risk and update their beliefs accordingly. I focus on floods affecting suppliers of inputs with few alternatives sourcing options, since these shocks map most closely to the source of risk captured by the model. I then examine changes in imports, inventories, cash holdings, and book leverage over a five-year window around these shocks. Because the median flood lasts just seven days, persistent changes over a five-year window are more naturally interpreted as reflecting updated beliefs about disruption risk than as direct responses to the disruption itself.

Treated firms persistently increase import volume, build up input inventories, draw down cash, and increase book leverage relative to matched control firms in the five years following a flood. These responses level off over time, suggesting that firms converge to a new optimal level of operational resilience as they update their beliefs. Several additional findings support this learning interpretation. First, the response is substantially larger for firms experiencing their first supplier flood than for firms experiencing repeat floods, consistent with first-time shocks generating the largest informational surprise. Second, firms that were least prepared for supply chain disruptions, as measured by abnormally low residual inventory holdings before the shock, show the strongest treatment effects. Third, firms increase the quantity and intensity of supply chain language in their 10-K filings after floods, suggesting that managers actively reassess their exposure to supply chain risk after shocks.

Firms build their post-flood inventories primarily by establishing new supplier relationships, and 82% of these new suppliers are located in a different country than the disrupted supplier. Imports from the disrupted supplier fall modestly and remain persistently below their pre-flood level, suggesting that the supplier allocates scarce post-flood output selectively across customers. As a result, firms seem unable to rely on the disrupted supplier to accommodate the additional sourcing needed to build inventories and instead turn to new suppliers. This interpretation is reinforced by the fact that firms that were more important customers of the disrupted supplier before the flood expand sourcing and inventories less afterward, consistent with suppliers allocating scarce output preferentially to their most valuable trading relationships.

Finally, the existing literature documents that supply chain shocks propagate through production networks and reduce the output of affected firms (Barrot and Sauvagnat, 2016; Boehm, Flaaen, and Pandalai-Nayar, 2019). The findings documented above thus raise a natural question about whether firms can insulate themselves from such shocks. I show that firms with high pre-treatment input inventory buffers experience no significant decrease in operating income, sales, or gross profit relative to undisrupted peers, while firms with lower inventory buffers suffer significant performance losses across all three measures. These results suggest that input inventory buffers serve as a form of operational insurance that preserves production continuity during disruptions.

This paper contributes to three literatures. First, I show that the nature of the risk determines which asset provides effective insurance, extending the precautionary savings literature (Almeida, Campello, and Weisbach, 2004; Bates, Kahle, and Stulz, 2009), which focuses on settings where the binding constraint is financial, and cash is therefore the natural buffer. When inputs lack spot markets, firms switch from cash to inventories, changing the form of precautionary savings and linking capital structure to supply chain structure. Bates, Kahle, and Stulz (2009) and Graham and Leary (2018) document the secular decline in inventories and the increase in cash holdings and attribute them partly to changes in production technology. Kulchania and Thomas (2017) argue that firms with low inventories face higher disruption costs and accumulate cash to draw upon during shocks. My results show that the reverse margin is also important. When inputs with scarce suppliers cannot be sourced on the spot market during disruptions, firms use cash and debt to build inventories in anticipation of shocks.

A related literature examines how trade frictions affect inventory dynamics. Alessandria, Kaboski, and Midrigan (2011) show that importing firms hold larger inventory stocks than domestic firms because international transactions involve longer lead times and higher per-shipment costs. During recessions, these elevated stocks are run down before firms reorder, causing imports to decrease faster and further than domestic production, thereby amplifying trade volatility. Alessandria, Khan, Khederlarian, Mix, and Ruhl (2023) focus on the supply shock of 2020 to 2022 and show that firms with sufficient inventories could continue producing while awaiting delayed shipments. Firms with depleted stocks could not, so low initial inventory levels amplified the aggregate contraction from shipping delays. The motive I identify for holding inventories is different. In these papers, inputs are delayed but will eventually arrive, and inventories bridge the gap between order and delivery. I focus on settings in which inputs with scarce suppliers are difficult to procure from alternative sources during a disruption, regardless of the firm's willingness to pay. In this setting, inventories provide the most direct means of sustaining production during disruptions.

Second, I contribute to the literature on production networks and supply chain risk. Barrot and Sauvagnat (2016) show that supply chain shocks propagate through production networks, and Ersahin, Giannetti, and Huang (2024) and Pankratz and Schiller (2025) document how firms restructure supply chains in response. I focus on the case in which alternative suppliers are scarce, so that firms cannot quickly replace a disrupted supplier. In this setting, firms build inventory buffers by establishing new supplier relationships, a process that unfolds over several years. I show that firms with larger inventory buffers are protected from the performance consequences of subsequent disruptions, suggesting that inventories can dampen the propagation of shocks through production networks when they serve as operational insurance against inaccessible inputs.

Third, this paper contributes to the literature on financial and operational hedging (Allayannis, Ihrig, and Weston, 2001; Hoberg and Moon, 2017; Acharya, Almeida, Amihud, and Liu, 2025). This literature generally studies how firms use operational flexibility to complement financial hedging and how financial constraints shape the tradeoff between the two. Acharya et al. (2025) show that financially constrained firms cut operational hedges to preserve cash. My analysis shows the converse. Even financially unconstrained firms adjust cash and leverage when exposure to input supply risk makes physical buffers more valuable than financial ones. The paper therefore shows that input supply risk can itself be a determinant of financial policy.

The rest of the paper is organized as follows. Section 2 describes the data and introduces the supplier scarcity measure. Section 3 discusses the empirical strategy. Section 4 validates the scarcity measure using flood shocks. Section 5 develops the theoretical framework. Section 6 presents the cross-sectional and instrumental variables results on supplier scarcity and corporate policies. Section 7 tests the learning predictions of the model. Section 8 examines whether input inventory buffers mitigate the performance effects of disruptions. Section 9 concludes.

2 Data

This section describes the sample construction, the global supply chains data, the supplier scarcity measure, and other data used in the empirical analysis.

2.1 Sample construction

The sample begins with a list of U.S. publicly listed manufacturing firms (SIC codes 2000–3999) in Compustat from Q1 2007 to Q4 2019. I restrict the sample to Q4 2019 to avoid confounding effects from supply chain disruptions during the COVID-19 pandemic. Firms are required to have non-missing total assets and sales, and at least three consecutive years of data, resulting in 2,346 manufacturing firms. From this set, the study focuses on 907 firms with global supply chains, defined as firms that import from foreign suppliers. Import data is obtained from Panjiva, which provides granular data on U.S. maritime imports derived from bill of lading (BOL) documents, including product descriptions, supplier identifiers, and logistical details such as ports and shipping vessels.[1] I focus on three measures of import intensity: shipment counts, import weight, and import volume (measured in TEUs, the volume of a standard cargo container). For each import, I consider the supplier to be the ultimate parent of the supplier listed on the BOL.

In addition to data on global supply chains, I obtain financial data from Compustat and Capital IQ, equity data from the Center for Research in Security Prices (CRSP), and data on natural disasters from the Dartmouth Flood Observatory at the University of Colorado.

2.2 BOL data global supply chains

Flaaen, Haberkorn, Lewis, Monken, Pierce, Rhodes, and Yi (2023) compare the BOL data to aggregate data on containerized vessel import value and confidential administrative datasets from the U.S. Census Bureau. They show that the BOL data is highly consistent with the U.S. Census Bureau's data while often providing greater granularity on firms' inputs and suppliers. However, a limitation of the BOL data is that it does not cover non-maritime trade.

Internet Appendix A examines the representativeness of global supply chains constructed from the BOL data by analyzing the share of U.S. imports transported via maritime shipments. I show

[1] Internet Appendix A includes an example of a BOL form.

that, by both weight and value, maritime shipments are the predominant mode of transportation. Although certain products, such as pharmaceuticals, primarily rely on air transport, pharmaceutical manufacturing firms account for only 7% of the sample. Moreover, specialized components heavily used by manufacturing firms, such as electronic equipment and machinery, are well represented in both maritime and air imports, collectively accounting for approximately one-third of the imports among the top ten HS codes for each mode of transport. These statistics suggest that the characteristics of firms' global supply chains derived from BOL data are likely representative of their overall supply chains.

Internet Appendix A also presents several stylized facts about U.S. manufacturing firms' global supply chains. First, imports predominantly consist of raw materials and intermediate goods (collectively referred to as inputs) rather than finished products. Second, there is a significant concentration of suppliers in certain countries. For example, China accounted for 31% of all suppliers in 2019. Third, there is substantial heterogeneity in input-supplier networks. While some inputs are generic and sourced from many suppliers worldwide, many inputs are sourced from only a few suppliers or even a single one.

2.3 Supplier scarcity

I leverage the detailed input-supplier network in the BOL data to develop a new measure of firms' exposure to disruptions in input supply. The measure quantifies, for each firm, the number of alternative suppliers available for its inputs, and thus captures the extent to which firms face binding constraints on their ability to substitute away from disrupted suppliers. Specifically, for each input a firm uses, I count the number of alternative suppliers from which other firms globally source that same input, excluding the firm's current supplier. Figure 2 illustrates this variation in supplier availability, showing a bimodal distribution in the number of global suppliers across inputs. A substantial fraction of inputs have only one or a few suppliers, while many have 20 or more.

This measure offers several advantages. It captures exposure to disruptions arising from different sources, including input specificity, geographic concentration, and market structure, through a single metric. Moreover, rather than inferring constraints from suppliers' or inputs' characteristics, it directly observes the availability of alternative sourcing options. The measure also varies at the firm-input level, capturing substantial heterogeneity in supply chain exposure across firms.

To aggregate to the firm level, I compute the import-volume-weighted average number of alternative suppliers across a firm's inputs, rescaled so that higher values indicate greater scarcity.[2] As discussed in Section 3.1, this firm-level measure is highly persistent within firms, consistent with sticky supply chains and slow-moving production technologies (Antràs and Chor, 2013; Martin, Mejean, and Parenti, 2026).

2.4 Supplier floods

I collect detailed flood data from the Dartmouth Flood Observatory (DFO) at the University of Colorado. The DFO compiles information from various sources, including media reports, governmental records, and satellite imagery. For each flood, the dataset provides the start and end dates, the coordinates of the flood's center, the total area affected in square miles, the number of displaced individuals, and the number of casualties. There are 2,122 floods in the data since 2007.

Figure 3 illustrates the geographic distribution of floods and the locations of affected suppliers, highlighting substantial variation in the occurrence of these natural disasters. The floods are not confined to specific regions or periods, reducing concerns that they might coincide with cyclical trade patterns or be limited to particular areas. Major events in the dataset include the March 2011 tsunami in Japan, triggered by the Tōhoku earthquake, and the Thailand floods that occurred later that year. Both events had significant consequences for U.S. manufacturers reliant on Japanese and Thai suppliers. For example, the tsunami caused severe supply chain disruptions (e.g., Boehm, Flaaen, and Pandalai-Nayar, 2019). Similarly, the Thai floods led to the closure of a key Western Digital factory, with a significant impact on the computer manufacturing industry.[3] Of the 2,122 floods documented in the DFO data, 1,508 impacted the suppliers of the firms in my sample.

2.5 Descriptive statistics

Table 1 presents descriptive statistics for global supply chains (Panel A), accounting characteristics (Panel B), and floods (Panel C) based on a firm-quarter panel of 907 publicly listed U.S. manufacturing firms from Q1 2007 to Q4 2019. During this period, the median firm received 63

[2] The raw measure is the volume-weighted average number of alternative suppliers. I subtract this from the sample maximum and divide by 100, so that the resulting index is increasing in scarcity.

[3] See "Thailand flooding cripples hard-drive suppliers," by Thomas Fuller, November 6, 2011, *The New York Times*.

shipments and 85 TEUs quarterly. However, the distribution of imports is notably skewed, with the average firm receiving 278 shipments and 561 TEUs. These firms typically import a diverse range of products from suppliers located in multiple countries each quarter. For example, the median firm procured 11 distinct products from 11 different suppliers across six countries.

Panel B of Table 1 shows that, from Q1 2007 to Q4 2019, the median total debt and cash-to-asset ratios for sample firms were 22% and 10%, respectively. These firms generally exhibit less financial flexibility than those in the broader Compustat sample, where the median book leverage and cash-to-assets ratios are 15% and 18%, respectively. Moreover, the sample firms tend to be larger, more profitable, and hold more inventory. For example, the median ratio of input inventories to assets is 7% in the sample, compared to zero in the unrestricted dataset.

Panel C of Table 1 provides summary statistics for the flood events. These events affected 16,322 suppliers in 64 countries and impacted 769 firms, or 83% of the firms in the sample at some point. The median flood lasts 7 days, displaces 1,000 individuals, results in 3 casualties, and affects a total area of 23,963.95 square miles.

3 Empirical strategy

This section describes the empirical strategies used to identify the effect of firms' exposure to disruptions in input supply on corporate financial policies. I begin with cross-sectional and instrumental variables estimation of the relationship between supplier scarcity and firm policies, then turn to a stacked difference-in-differences design that exploits supplier floods as exogenous shocks to perceived disruption risk.

3.1 Cross-sectional estimation

To estimate the relationship between supplier scarcity and corporate financial policies, I estimate firm-quarter panel regressions of the form:

$$Y_{i,t} = \alpha_j + \alpha_t + \beta \textit{Supplier scarcity}_{i,t-1} + \gamma \mathbf{X}_{i,t-1} + \epsilon_{i,t} \tag{1}$$

where i indexes firms, j two-digit SIC industries, and t quarters. *Supplier scarcity* is the import-volume-weighted average number of alternative suppliers across the HS codes in firm i's import basket in quarter $t-1$, rescaled so that higher values indicate greater scarcity. $\mathbf{X}$ is a vector of firm-level controls, including the log of assets, the log of sales, market-to-book ratios, R&D expenditures over sales, cash flows over assets, and capital expenditures over net property, plant, and equipment. α_j and α_t denote industry and quarter fixed effects, and standard errors are clustered at the firm and quarter levels. Because *Supplier scarcity* is highly persistent within firms (autocorrelation of 0.7), likely reflecting sticky production technologies and slowly evolving supply chains, I treat it as an ex ante measure of firms' exposure to input supply disruptions and focus on cross-sectional differences across firms within industries.

3.2 Instrumental variables estimation

The OLS cross-sectional estimates may be biased if unobserved characteristics of production technologies simultaneously affect supply chain structure and corporate policies. To address this concern, I construct a shift-share instrument that isolates exogenous variation in supplier scarcity from changes in the pool of available suppliers across a firm's input markets. For each HS code h that firm i imports, I observe the number of distinct active suppliers $N_{h,t}$ in the global market. The firm-level instrument aggregates these product-market supplier counts using firm i's lagged import volume shares $w_{ih,t-1}$ as weights:[4]

$$Z_{i,t} = \sum_h w_{ih,t-1} \left(-\ln N_{h,t}\right) \tag{2}$$

As with the supplier scarcity measure, the negative sign ensures that the instrument is increasing in scarcity.[5] All regressions use $Z_{i,t-1}$, the one-quarter lag of the instrument. The identifying variation comes from changes in $N_{h,t}$, the entry and exit of suppliers in global product markets, which are plausibly exogenous to any individual buyer's corporate policies.

[4] In shift-share designs, the weights are often drawn from a more distant base period to ensure they are predetermined. Because identification in my setting relies on the exogeneity of shifts rather than shares (Borusyak, Hull, and Jaravel, 2022), the lag structure of the import volume weights is less consequential. Nevertheless, Table IA.B.8 shows that the results are robust to constructing the instrument with one-year, two-year, and three-year lagged import volume weights.

[5] Without the negative sign, a decline in the number of suppliers would decrease the instrument, and the first-stage coefficient would be negative. The normalization is without loss of generality.

3.3 Stacked difference-in-differences

I adopt a stacked difference-in-differences approach (Gormley and Matsa, 2011; Cengiz, Dube, Lindner, and Zipperer, 2019; Deshpande and Li, 2019; Baker, Larcker, and Wang, 2022) to validate the supplier scarcity measure and to estimate how firms adjust corporate policies after experiencing disruptions to their input supply. I exploit supplier floods as exogenous shocks in a setting with staggered treatment timing, where the traditional difference-in-differences estimator can be biased (Goodman-Bacon, 2021). The stacked approach constructs separate cohorts for each flood event. Within each cohort, treated firms are those that imported from a flood-affected supplier at any time in the five years preceding the flood.[6] Control firms are those unaffected by the current flood and any prior floods, ensuring that already-treated units never serve as controls.[7]

Validation of the supplier scarcity measure. I first validate the supplier scarcity measure using a short-window specification around flood events. The baseline stacked difference-in-differences regression is:

$$Y_{istj} = \alpha_{isj} + \alpha_{tj} + \beta\, Treated_{isj} \times Post_{tj} + \gamma \mathbf{X}_i \times Post_{tj} + \epsilon_{istj} \tag{3}$$

where i denotes firms, s suppliers, t weeks, and j flood events (cohorts). The model is estimated using a four-week window around each flood event. The dependent variable Y_{istj} is the number of shipments from supplier s to firm i at time t in cohort j. $Treated_{isj}$ equals one for firm-supplier pairs where the supplier was affected by the flood, while $Post_{tj}$ equals one during the four weeks following the flood. $\mathbf{X}_i \times Post_{tj}$ is a vector of fixed firm-level controls measured in the pre-treatment period and interacted with $Post_{tj}$. The specification includes firm×supplier×cohort and calendar week×cohort fixed effects. I estimate these regressions using Poisson pseudo-maximum likelihood (PPML) because the weekly trade data include zero flows (Silva and Tenreyro, 2006). Standard errors are clustered at the firm×supplier and calendar week levels.

To test whether firms can reallocate sourcing to unaffected suppliers after a flood, I also estimate a variant of Equation 3 that replaces the firm×supplier×cohort fixed effects with firm×cohort fixed

[6] The results are robust to other definitions of treatment, including requiring at least one, two, or three years of active importing from the affected supplier.

[7] The results are also robust to alternative staggered difference-in-differences estimators, including those proposed by Borusyak, Jaravel, and Spiess (2024), Sun and Abraham (2021), and de Chaisemartin and D'Haultfœuille (2020).

effects and redefines $Treated_{ij}$ as a firm-level indicator equal to one for firms that had at least one flood-affected supplier in the cohort. I estimate this specification separately on the subsample of unaffected suppliers of fragile inputs and on the subsample of unaffected suppliers of generic inputs. Suppliers are classified as fragile if they provide at least one input for which the total number of global suppliers is below the sample median. In these specifications, standard errors are clustered at the firm and calendar week levels.

I take two additional steps to refine identification. First, I balance treated and control firms using exact matching on two-digit SIC codes and nearest neighbor matching on pre-treatment log assets and aggregate import volume. Second, to account for seasonality, pre-treatment values of the dependent variable are computed by averaging imports from the same calendar weeks as the event weeks, using data from the previous three years.

Corporate policies specification. To examine long-term policy adjustments, I estimate an analogous stacked difference-in-differences specification using a firm-flood-quarter panel and a symmetric five-year window around each flood. The dependent variables are the log of import volume (in TEU), and ratios of input inventories to pre-treatment sales, cash to pre-treatment assets, and total debt to pre-treatment assets. The matching procedure is the same as in the flood validation specification. There are several differences with respect to the supplier scarcity validation specifications. First, the data are quarterly rather than weekly. Second, I estimate by OLS rather than PPML. Third, the pre-treatment values of outcome variables are not seasonally adjusted because these variables do not exhibit the seasonal patterns present in weekly shipment data. Fourth, within each cohort, treated firms are those that imported from the flood-affected supplier of fragile inputs at any time during the five years before the flood. I include firm$\times$flood and calendar quarter$\times$flood fixed effects. Standard errors are clustered at the firm and calendar quarter levels.

4 Validation of the supplier scarcity measure

I estimate the validation specification in Equation 3 using PPML, as described in Section 3.3. Table 2 presents the results for changes in the number of shipments around floods.

Column (1) shows that the number of shipments from flood-affected suppliers decreases by

$-(e^{-0.425} - 1) \times 100 \approx 35\%$ in the four weeks following a flood, relative to the corresponding three-year pre-treatment mean. Columns (2) and (3) examine whether firms can offset this disruption by sourcing from unaffected suppliers. I aggregate shipments from unaffected suppliers at the firm level and split the sample by input type, replacing the firm$\times$supplier$\times$cohort fixed effects with firm$\times$cohort fixed effects. I classify suppliers as fragile if they provide at least one input for which the total number of global suppliers is below the sample median. Column (2) shows that firms do not significantly increase shipments from unaffected suppliers of fragile inputs, consistent with short-run supplier substitution frictions. Column (3) shows that firms successfully reallocate imports of generic inputs, with shipments from unaffected suppliers of those inputs increasing by approximately 78%. These findings suggest that supplier scarcity captures real differences in firms' ability to replace disrupted inputs.

Figure 4 presents the corresponding event study estimates. Panel A shows that shipments from affected suppliers decrease sharply and persistently over the four-week event window, with pre-treatment differences close to zero and statistically insignificant. Panel B plots separate event studies for shipments from unaffected suppliers of fragile and generic inputs. Shipments from unaffected generic suppliers increase after the flood, consistent with successful reallocation. Shipments from unaffected fragile suppliers remain flat, consistent with firms being unable to substitute fragile inputs even from alternative suppliers.

5 Theory and empirical predictions

The evidence in the previous section establishes that supplier scarcity captures real differences in firms' ability to replace disrupted inputs. When a supplier is hit by a flood, firms can reallocate sourcing for generic inputs but not for fragile ones. Motivated by these findings, this section develops a theoretical framework that links firms' exposure to input supply disruptions to corporate financial policies and generates testable predictions.

5.1 Setup

I model supply chain disruption risk as a characteristic of the production process and the inputs it requires. A firm uses two distinct inputs for production: a generic input, denoted by g, and a fragile input, denoted by s. The generic input is available on the spot market, guaranteeing the firm continuous access to alternative suppliers. By contrast, the fragile input is not procurable on the spot market, either because alternative suppliers are scarce or because switching costs make sourcing from new suppliers prohibitive.

The firm operates over two periods. In the first period, t, it decides on production inputs, inventory levels, cash reserves, and debt financing.[8] In the subsequent period, $t+1$, it determines its production inputs and any additional debt financing needs. At the end of period t, the firm realizes interim cash flows from its production decisions. These interim cash flows, combined with the cash saved from period t, help fund the procurement of production inputs in period $t+1$.

The firm's ability to acquire the fragile input in period $t+1$ depends on the parameter $\phi_{s,t+1}$. When $\phi_{s,t+1}=1$, the fragile input's supply chain is disrupted before period $t+1$, forcing the firm to rely exclusively on the inventory of fragile inputs it accumulated in period t. Conversely, when $\phi_{s,t+1}=0$, the supply chain for the fragile input remains functional, allowing the firm to purchase the fragile input on the spot market as needed. Given its production decisions in period $t+1$, the firm generates terminal cash flows, which are distributed to shareholders at the end of that period. Figure 1 summarizes this sequence of events.

5.2 Production technology

The firm produces output using the generic and fragile inputs via a constant elasticity of substitution (CES) production function. Let p denote the output price. The firm's revenue from production in period $\tau \in \{t, t+1\}$ is:

$$F(z_\tau, N_{g,\tau}, N_{s,\tau}) = p \cdot z_\tau \left[\theta N_{g,\tau}^{\rho} + (1-\theta) N_{s,\tau}^{\rho}\right]^{1/\rho} \tag{4}$$

where $\theta \in (0,1)$ governs the relative importance of the generic input, z_τ is a productivity shock whose distribution is known to the firm but whose realization is uncertain at the time decisions are made, and $\rho \in (-\infty, 1)$ determines the elasticity of substitution $\sigma = 1/(1-\rho)$ between the two inputs.[9]

[8] The model focuses on debt as the source of external finance, reflecting the importance of trade credit in financing input purchases and inventory buildup (e.g., Rajan and Zingales, 1995; Barrot, 2016; Yang and Birge, 2018).

[9] I maintain the output price p explicitly rather than normalizing it to one, following Cavallo and Kryvtsov (2023). The results do not depend on this choice.

The CES specification nests several important special cases. When $\rho \to 1$ ($\sigma \to \infty$), the inputs are perfect substitutes and the generic input can fully compensate for the loss of the fragile one, so that supply chain disruption risk is irrelevant. When $\rho \to 0$ ($\sigma \to 1$), the production function reduces to Cobb-Douglas. When $\rho \to -\infty$ ($\sigma \to 0$), the production function approaches Leontief, representing maximum fragility. Any shortfall in the fragile input cannot be compensated by additional generic input. The model thus captures two distinct dimensions of supply chain disruption risk. The no-spot-market assumption determines whether a disruption removes the fragile input from production, corresponding to the empirical finding that inputs with few alternative suppliers cannot be replaced after shocks. The parameter ρ determines how costly that loss is, by modulating the extent to which the firm can compensate with the generic input.

5.3 Firm choices and financing frictions

At the beginning of period t, the firm is endowed with liquid assets W_t and decides on how many units of generic and fragile inputs to purchase $(Q_{g,t}, Q_{s,t})$, how much inventory to hold for each $(i_{g,t}, i_{s,t})$, how much cash to save (C_t), and how much debt to raise (B_t). Input prices are normalized to one. Given input utilization ($N_{g,t} = Q_{g,t} - i_{g,t}$, $N_{s,t} = Q_{s,t} - i_{s,t}$), the firm generates interim cash flows by the end of period t.

In period $t+1$, the firm's decisions are limited to procuring inputs $(Q_{g,t+1}, Q_{s,t+1})$ and issuing debt (B_{t+1}). Because the firm ceases operations at the end of period $t+1$, it does not accumulate savings or hold inventories in that period. If $\phi_{s,t+1} = 1$, the supply chain for the fragile input is disrupted, preventing any spot-market purchases. Fragile input utilization in period $t+1$ is therefore $N_{s,t+1} = (1-\phi_{s,t+1})Q_{s,t+1} + (1-\alpha)i_{s,t}$, where $\alpha \in (0, 1)$ captures inventory carrying costs (e.g., depreciation, storage, obsolescence). Similarly, $N_{g,t+1} = Q_{g,t+1} + (1-\alpha)i_{g,t}$.

The firm finances its period $t+1$ input purchases using interim cash flows, any cash reserves carried over from period t, and new debt raised in period $t+1$. Following Nikolov and Whited (2014), I assume the firm faces quadratic costs of debt $c(B_\tau) = \frac{\gamma}{2}B_\tau^2$, with $\gamma > 0$, which can be motivated by debt overhang (Myers, 1977), moral hazard (Jensen and Meckling, 1976), or adverse selection (Myers and Majluf, 1984).

5.4 Supply chain disruption risk

Recent models of production network formation demonstrate that supply chain fragility can emerge as an equilibrium outcome when firms optimally choose the most productive suppliers. This selection reinforces economies of scale and relationship-specific investments, resulting in a concentrated supplier base. Alternatively, technological advancements that enable firms to incorporate a wider variety of inputs can increase production complexity and thus fragility (Oberfield, 2018; Acemoglu and Azar, 2020; Acemoglu and Tahbaz-Salehi, 2020). Elliott, Golub, and Leduc (2022) formalize this intuition by showing that equilibrium investment in supply relationships can leave networks at a critical threshold where small systemic shocks cause discontinuous drops in production, even though this outcome is always inefficient.

The concept of supply chain disruption risk I adopt is related but different. In Elliott, Golub, and Leduc (2022), fragility is an equilibrium property of aggregate output arising from endogenous network formation. Here, I define supply chain disruption risk as the ex ante probability that shocks disrupt fragile input supply, treating it as a parameter that captures a firm's exposure to disruptions given the characteristics of its inputs and suppliers. This definition reflects the paper's focus on individual firms' policy responses to disruption risk rather than on the network equilibrium that generates it. The disruption indicator $\phi_{s,t+1}$ is a Bernoulli random variable with $\mathbb{P}(\phi_{s,t+1} = 1) = q_s \in (0, 1)$. Disruptions where $\phi_{s,t+1} = 1$ capture situations in which a firm cannot obtain inputs from a supplier, for instance because the supplier is not operating due to a natural disaster. The underlying assumption is that alternative suppliers may not exist, and if they do, high switching costs prevent the firm from sourcing elsewhere. A growing body of empirical evidence motivates these frictions, including Barrot and Sauvagnat (2016), Antràs, Fort, and Tintelnot (2017), Boehm, Flaaen, and Pandalai-Nayar (2019), and Carvalho, Nirei, Saito, and Tahbaz-Salehi (2021). I provide direct evidence that the supplier scarcity measure captures these frictions in Section 4, where I show that firms struggle to substitute suppliers of fragile inputs after flood shocks

5.5 Optimal policies

The firm's expected profit in period $t+1$, conditional on period-t decisions and the disruption outcome $\phi_{s,t+1}$, is:

$$\pi_{t+1} = F(z_{t+1}, N_{g,t+1}, N_{s,t+1}) - Q_{g,t+1} - (1-\phi_{s,t+1})Q_{s,t+1} - \frac{\gamma}{2}B_{t+1}^2 \tag{5}$$

In period t, the firm's liquid assets at the start of period $t+1$ are:

$$W_{t+1} = F(z_t, N_{g,t}, N_{s,t}) + C_t \tag{6}$$

The period-$t+1$ budget constraint requires that input expenditures do not exceed available resources:

$$Q_{g,t+1} + (1-\phi_{s,t+1})Q_{s,t+1} \leq W_{t+1} + B_{t+1} \tag{7}$$

The firm's objective in period t is to choose $(Q_{g,t}, Q_{s,t}, i_{g,t}, i_{s,t}, C_t, B_t)$ to maximize the expected present value of equity to current shareholders. Taking expectations over both the disruption indicator $\phi_{s,t+1} \in \{0,1\}$ (with probability q_s and $1-q_s$, respectively) and the productivity shocks z_t, z_{t+1}, the firm's problem can be written as:[10]

$$\begin{aligned} V_t = \max_{\{Q_{g,t}, Q_{s,t}, i_{g,t}, i_{s,t}, C_t, B_t\}} \; & \mathbb{E}_t\Big[F(z_t, N_{g,t}, N_{s,t})\Big] - Q_{g,t} - Q_{s,t} - C_t - \frac{\gamma}{2}B_t^2 \\ & + \; q_s\, \mathbb{E}_t\Big[\pi_{t+1} \mid \phi_{s,t+1} = 1\Big] + (1-q_s)\, \mathbb{E}_t\Big[\pi_{t+1} \mid \phi_{s,t+1} = 0\Big] \end{aligned} \tag{8}$$

subject to,

$$Q_{g,t} + Q_{s,t} + C_t \leq W_t + B_t \tag{9}$$

$$Q_{g,t} \geq 0,\ Q_{s,t} \geq 0,\ i_{g,t} \geq 0,\ i_{s,t} \geq 0,\ C_t \geq 0,\ B_t \geq 0 \tag{10}$$

where $N_{g,t} = Q_{g,t} - i_{g,t}$ and $N_{s,t} = Q_{s,t} - i_{s,t}$.

The key tradeoff is as follows. When $\phi_{s,t+1} = 1$, the firm cannot purchase fragile inputs in period $t+1$ and must rely solely on the inventory carried over from period t, net of carrying costs. Since

[10] Input inventories are excluded from the liquidity expression W_{t+1} because these inventories are typically tailored to the firm's specific needs and thus are likely to have limited value outside the firm (Shleifer and Vishny, 1991). For expositional clarity, I assume no discounting between periods.

the inputs are imperfect substitutes ($\rho < 1$), additional generic inputs cannot fully compensate for the shortfall in fragile inputs, and the severity of this production loss depends on σ. By contrast, when $\phi_{s,t+1} = 0$, both inputs are available and inventories carried from period t simply augment the firm's production capacity. The firm's period-t choices balance the cost of carrying inventories and holding cash against the insurance value these reserves provide in the disruption state.

Appendix B provides the full derivations. I summarize the key results below.

5.5.1 Disruption risk, input substitutability, and corporate policies

The model identifies two dimensions of supply chain disruption risk that jointly determine optimal corporate policies. The first is the probability that a disruption occurs (q_s), which determines how much weight the firm places on the disruption state when making period-t decisions. The second is the elasticity of substitution between inputs (σ), which determines how costly a disruption is conditional on occurring. Both dimensions push optimal policies in the same direction.

Proposition 1 (Disruption risk and corporate policies)**.** *Suppose that $\rho < 1$ and $\sigma < \infty$. Then, as the disruption probability q_s increases:*

(i) *Optimal inventories of fragile inputs increase: $\partial i^*_{s,t}/\partial q_s > 0$.*

(ii) *Optimal cash holdings decrease: $\partial C^*_t/\partial q_s < 0$.*

(iii) *Optimal debt issuance increases: $\partial B^*_t/\partial q_s > 0$.*

When $q_s = 0$, optimal fragile input inventories converge to zero.

Proof: *See Appendix B.*

Proposition 1 maps directly to the cross-sectional empirical analysis. The supplier scarcity measure captures variation in q_s across firms. Specifically, firms whose inputs are sourced from fewer alternative suppliers face a higher probability that a disruption leaves them unable to procure the input. The proposition predicts that these firms should hold more inventories, less cash, and more leverage.

The second dimension of disruption risk, how costly a disruption is conditional on occurring, depends on the elasticity of substitution between inputs. Proposition 1a formalizes this result.

Proposition 1a (Input substitutability and corporate policies)**.** *Suppose that $q_s \in (0,1)$ and that the CES parameter satisfies $\rho < 1$. Then, as the elasticity of substitution $\sigma = 1/(1-\rho)$ decreases:*

(i) *Optimal inventories of fragile inputs increase: $\partial i^*_{s,t}/\partial\sigma < 0$.*

(ii) *Optimal cash holdings decrease: $\partial C^*_t/\partial\sigma > 0$.*

(iii) *Optimal debt issuance increases: $\partial B^*_t/\partial\sigma < 0$.*

When $\sigma \to \infty$ ($\rho \to 1$), optimal fragile input inventories converge to zero.

Proof: *See Appendix B.*

The intuition is as follows. Consider the firm's decision to carry one additional unit of fragile input inventory into period $t+1$. The marginal benefit of this unit is the expected production gain in the disruption state, which depends on the marginal product of the fragile input when the other input cannot substitute for it. Under the CES technology, the marginal product of the fragile input in the disruption state is:

$$\left.\frac{\partial F}{\partial N_{s,t+1}}\right|_{\phi_{s,t+1}=1} = p \cdot z_{t+1}(1-\theta)\left[\theta\left(\frac{N_{g,t+1}}{N_{s,t+1}}\right)^{\rho} + (1-\theta)\right]^{(1-\rho)/\rho} \tag{11}$$

As σ decreases (i.e., ρ becomes more negative), the marginal product of the fragile input during disruptions increases when the fragile input is scarce relative to generic input, because the CES aggregator penalizes the imbalance between inputs more severely. This raises the insurance value of inventory, increasing $i^*_{s,t}$. In the limit as $\sigma \to 0$, the production function approaches Leontief and the marginal product of the fragile input approaches $p \cdot z_{t+1}$, a finite upper bound that is large relative to the marginal product when inputs are balanced. Conversely, as $\sigma \to \infty$, the marginal product of each input in the disruption state is independent of the input ratio, so inventories provide no additional insurance value.

Because the increase in inventories must be financed, the firm draws down cash reserves and issues additional debt. Cash is useful in both states. In the no-disruption state, it funds purchases of both inputs. In the disruption state, it funds only purchases of generic inputs. As σ decreases, the marginal product of generic inputs in the disruption state declines (because the complementary fragile input is scarce and cannot be compensated for), reducing the value of cash in that state.

Meanwhile, the marginal benefit of fragile input inventory increases. At the margin, the firm optimally reallocates from cash to inventory and issues more debt to finance the gap.

Propositions 1 and 1a establish that the two dimensions of supply chain disruption risk reinforce each other. Firms that face both a high probability of disruption (q_s) and low input substitutability (σ) have the strongest incentive to build operational resilience through inventory accumulation, financed by lower cash and higher leverage. The empirical supplier scarcity measure is most directly related to q_s, since firms with fewer alternative suppliers face a higher probability that a disruption leaves them unable to procure the input. However, the measure may also reflect variation in σ if inputs sourced from concentrated supplier bases tend to be harder to substitute with other inputs in production. Both channels push in the same direction, so that the supplier scarcity measure provides a valid test of the model's predictions regardless of whether it captures variation in disruption probability, input substitutability, or both

5.5.2 Learning about disruption risk

The analysis so far treats the disruption probability q_s as known. In practice, firms are uncertain about the true probability of supply chain disruptions and update their beliefs as they observe shocks. This uncertainty is natural given that severe disruptions to any particular supplier are rare events, and firms may have limited experience from which to form precise estimates of q_s (Gallagher, 2014).

I model learning as Bayesian updating over the disruption probability, following the standard approach in the finance literature on parameter learning (Pástor and Veronesi, 2003). The two-period production model characterized above describes optimal policies for a given belief $\hat{q}_s$. The learning framework developed here governs how beliefs evolve over a longer horizon as the firm accumulates experience with disruptions. Each period, the firm solves the static problem using its current posterior mean and then updates its belief after observing whether a disruption occurred.

The firm has a prior belief $q_s \sim \text{Beta}(a, b)$, where $a > 0$ and $b > 0$ are the shape parameters. The parameters have a natural interpretation: a represents the firm's prior count of disruption events and b the count of non-disruption periods, so that $a + b$ reflects the total prior experience and governs how responsive beliefs are to new observations. The prior mean is $\mathbb{E}[q_s] = a/(a + b)$ and the prior precision is $a + b$. After observing n supplier floods in T periods, the posterior belief is:

$$q_s \mid n, T \ \sim \ \text{Beta}(a+n,\ b+T-n) \tag{12}$$

The conjugate updating rule assumes that disruption events are independent draws, which is a simplification. In the empirical analysis, the use of geographically dispersed flood shocks across different suppliers mitigates concerns about serial correlation in disruption signals. The posterior mean, which the firm uses in place of q_s in its optimization problem, is:

$$\hat{q}_s(n, T) = \frac{a+n}{a+b+T} \tag{13}$$

Because q_s enters the firm's objective linearly as a probability weight over the disruption and no-disruption states, integrating over the posterior distribution of q_s yields the same solution as evaluating the objective at the posterior mean $\hat{q}_s$.

Proposition 2 characterizes how this learning mechanism affects the firm's policy adjustments over time.

Proposition 2 (Learning from supply chain disruptions)**.** *As the firm observes additional supply chain disruptions:*

(i) *The posterior disruption probability $\hat{q}_s$ increases, amplifying the effects in Proposition 1: inventories of fragile inputs increase, cash holdings decrease, and leverage increases.*

(ii) *The marginal effect of each additional disruption on the posterior mean, $\partial \hat{q}_s / \partial n = 1/(a+b+T)$, is decreasing in the total number of observations T. As the firm accumulates experience, the posterior precision $a+b+T$ increases and each additional disruption has a smaller effect on beliefs.*

Proof: *See Appendix B.*

Proposition 2 generates several testable predictions about how firms adjust corporate policies in response to observed disruptions. Each supplier flood increases the firm's perceived disruption probability $\hat{q}_s$, which by Proposition 1 leads to higher inventories, lower cash, and higher leverage. The magnitude of this adjustment depends on the firm's prior precision $a+b$ and the total number of observation periods T. Firms with diffuse priors (low $a+b$) or limited experience (low T) revise their beliefs more aggressively in response to a flood.

The convergence result in part (ii) implies that the marginal impact of each additional flood on corporate policies should decrease as firms accumulate more experience with supply chain disruptions, consistent with the evidence on experience-based learning in household and corporate decisions (Malmendier and Nagel, 2011; Bernile, Bhagwat, and Rau, 2017; Yu, 2025). First-time floods are particularly informative because they change beliefs substantially, whereas subsequent floods are partially anticipated and generate smaller revisions. This prediction distinguishes the learning channel from a simple restocking explanation, in which each flood depletes inventories and firms reorder by the same amount regardless of their history of disruptions.

6 Supplier scarcity, inventories, and corporate financial policy

This section estimates the relationship between supplier scarcity and firms' inventory, cash, and leverage policies. I begin with cross-sectional OLS regressions, and then address endogeneity concerns using a shift-share instrumental variables strategy.

6.1 Baseline estimates

Propositions 1 and 1a predict that firms with greater exposure to input supply disruptions should hold more inventories, less cash, and more debt. I test these predictions using the cross-sectional specification described in Section 3.1, which regresses firm-quarter corporate policy outcomes on lagged supplier scarcity with industry and quarter fixed effects. Table 3 presents OLS and instrumental variables estimates of the effect of supplier scarcity on corporate financial policies. Columns (1) to (3) report OLS estimates. The results suggest that a higher reliance on fragile inputs is associated with higher input inventory holdings, lower cash holdings, and higher book leverage. These results are consistent with the model's predictions and are highly statistically significant. The estimated OLS coefficients imply that a one standard deviation increase in *Supplier scarcity* is associated with 4.6% higher input inventory holdings, 13.0% lower cash holdings, and 5.7% higher book leverage relative to their mean values.

6.1.1 Robustness tests and additional results

Internet Appendix B presents several robustness tests and additional results for the baseline OLS results in Table 3.

Cross-sectional regressions. Because supplier scarcity is highly persistent within firms, the baseline results should also hold in a purely cross-sectional setting. Table IA.B.1 reports regressions with one observation per firm, where all variables are averaged over the sample period. The results are consistent with Table 3.

Debt decomposition. Table IA.B.2 decomposes the leverage effect by debt type. The increase in leverage is primarily driven by higher accounts payable, including trade credit, and greater use of revolving credit facilities, consistent with firms using short-term financing to fund inventory buildup.

Domestic versus foreign suppliers. The BOL data tracks only international shipments. A potential concern is that supplier scarcity may capture risks specific to foreign sourcing. Table IA.B.3 interacts supplier scarcity with the proportion of foreign suppliers derived from Compustat segment data. The coefficients on supplier scarcity are virtually unchanged, and the interaction terms are insignificant, suggesting that the results are not driven by risks specific to foreign sourcing.

Controlling for import intensity. Firms that import more may hold higher inventories due to longer lead times and lumpier ordering patterns (Hornok and Koren, 2015; Blum et al., 2019; Khan and Khederlarian, 2025), not because of disruption risk. Table IA.B.4 adds lagged import intensity as a control. The coefficient on supplier scarcity is virtually unchanged across all three outcomes.

6.1.2 Instrumental variables estimation

The OLS estimates in columns (1) to (3) may be biased if unobserved characteristics of production technologies simultaneously affect supply chain structure and financial policies. For instance, firms using highly specialized production processes may both source from scarce suppliers and hold more inventories for technological reasons unrelated to disruption risk. To address this concern, I use the shift-share instrument described in Section 3.2. The instrument aggregates the negative log number of globally active suppliers across a firm's input markets, weighted by the firm's lagged import volume shares, and is lagged one quarter in all regressions (Equation 2). The identifying variation comes from time-series changes in product-market supplier counts driven by the entry and exit of suppliers in global markets, which are plausibly exogenous to any individual firm's financial policies. This instrument follows the shift-share framework of Bartik (1991). Goldsmith-Pinkham, Sorkin, and Swift (2020) and Borusyak, Hull, and Jaravel (2022) show that valid identification in shift-share designs requires exogeneity of either the shares or the shifts. In my setting, I rely on the exogeneity of shifts, since changes in product-market supplier counts reflect aggregate supply-side conditions rather than individual firm choices.

First stage. Column (4) of Table 3 reports the first-stage regression of *Supplier scarcity* on the instrument. The coefficient on $Z_{i,t-1}$ is positive and highly significant, and shows that a contraction in the pool of available suppliers in a firm's input markets increases the firm's measured supplier scarcity. The Kleibergen-Paap F-statistic is approximately 300, well above conventional thresholds for weak instruments.

Two-stage least squares estimates. Columns (5) to (7) of Table 3 report two-stage least squares (2SLS) estimates of the effect of instrumented supplier scarcity on corporate policies using $Z_{i,t-1}$ as the instrument. The 2SLS estimates confirm the OLS findings. Firms with higher supplier scarcity maintain more input inventories, have lower cash holdings, and have higher leverage. All coefficients are statistically significant. The 2SLS magnitudes are larger than OLS for input inventories and leverage, which could be driven by attenuation bias in the OLS estimates due to measurement error in the observed scarcity measure. These results address the main identification concern. Even when isolating variation in scarcity that comes from exogenous changes in global supplier markets, the estimated effects on corporate policies are large, correctly signed, and statistically significant.[11]

[11] Table IA.B.5 presents reduced form regressions of corporate policy outcomes directly on the instrument. Because reduced form estimates do not depend on the first-stage specification, they provide an independent check that exogenous variation in the global supplier pool moves corporate policies in the directions predicted by the model. The reduced form coefficients are positive for input inventories and leverage and negative for cash, all statistically significant.

Exclusion restriction. The exclusion restriction requires that changes in the global supplier pool affect corporate policies only through their effect on supplier scarcity, not through other channels. I consider three potential threats and provide direct tests for the most plausible one.

The first concern is that supplier exit reflects declining product demand rather than supply-side factors. This channel predicts effects opposite to those documented. Declining demand would lead firms to reduce inventories and leverage, not increase them, so demand contamination would bias estimates against finding the reported effects. The second concern is that a shrinking supplier pool reduces input variety or quality, affecting financial policies through a productivity channel rather than disruption risk. If input quality deterioration were driving the results, one would expect firms to reduce inventories of lower-quality inputs rather than accumulate them, which is the opposite of what the data show.

More broadly, changes in global supplier counts could reflect technology cycles or input obsolescence, either of which could affect corporate policies independently of disruption risk. To assess whether the instrument captures these channels, Table IA.B.6 regresses R&D intensity and year-over-year sales growth on the instrument. If supplier count changes reflected changes in technology or input relevance, they should predict changes in these outcomes. The instrument predicts neither.

The most direct threat is that contractions in the supplier pool drive up input costs, causing firms to adjust policies in response to price changes rather than perceived disruption risk. I test it directly. Table 4 regresses two firm-quarter input price indices on the instrument using the same specification as the first stage. The first index is the equal-weighted mean of within-HS-6 demeaned log unit values (USD per kilogram) across all inputs in a firm's import basket. The second focuses on fragile inputs. In both specifications, the coefficient is statistically insignificant, suggesting that the instrument does not affect corporate policies through an input price channel. A potential concern with this test is that the effect of supplier pool contractions on input prices may manifest as within-product price variation over time, as firms experiencing larger supplier exits pay more for the same inputs. Table IA.B.7 examines this possibility and still finds insignificant coefficients. Finally, although the IV identification strategy relies on the exogeneity of shifts rather than shares, Table IA.B.8 shows that the results are similar when the instrument is constructed with one-year, two-year, and three-year lagged import volume weights.

7 Learning about supply chain disruption risk

The cross-sectional and IV evidence in Sections 6.1 and 6.1.2 establish that firms that rely on production inputs with few alternative suppliers hold more inventories, less cash, and more leverage. However, because severe disruptions to any particular supplier are rare, firms may have limited information about their true exposure to disruption risk. The model captures this through Proposition 2, which predicts that observed shocks cause firms to revise their beliefs about the probability of disruption upward and to increase their investment in operational resilience over time. I test this prediction by exploiting supplier floods as exogenous information shocks. When a flood unexpectedly disrupts a supplier, managers receive a signal about the likelihood of future disruptions and update their beliefs accordingly. I examine whether firms increase imports and inventories, reduce cash holdings, and increase leverage over a five-year window around these shocks. Because the median flood lasts just seven days, persistent changes over this window are more naturally interpreted as reflecting updated beliefs about disruption risk than as direct responses to the disruption itself. A distinctive prediction of the learning channel is that a firm's first flood produces the strongest response, with subsequent floods generating attenuated effects as beliefs converge toward the true disruption probability.

Table 5 presents the long-term treatment effects of flood shocks on corporate policies. Column (1) shows that import volume increases by 0.257 log points for treated firms relative to control firms in the five years following a flood, a 29% increase relative to the pre-treatment mean. Column (2) shows that input inventories increase by 6.6 percentage points relative to pre-treatment sales, a 19% increase relative to the pre-treatment mean. Column (3) shows that cash holdings fall by 3.0 percentage points relative to pre-treatment assets, a 23% decrease relative to the pre-treatment mean. Column (4) shows that book leverage increases by 6.1 percentage points relative to pre-treatment assets, a 23% increase relative to the pre-treatment mean. These long-term changes are consistent with firms revising their beliefs about disruption risk and adjusting their investment in operational resilience accordingly.

Panels A through D of Figure 5 present the event study plots for import volume, input inventories, cash holdings, and book leverage, respectively. In all four panels, pre-treatment differences between treated and control firms are close to zero and statistically insignificant, supporting the parallel trends assumption. Following a flood, Panels A and B show a clear, persistent increase in import volume and input inventories relative to pre-treatment sales. Similarly, Panels C and D show that treated firms gradually decrease their cash holdings and increase their book leverage.[12] The inventory and cash effects appear to level off within the five-year window, suggesting that firms gradually reach an optimal level of operational resilience. This pattern is consistent with the model's prediction that firms balance the marginal benefits of additional inventory against the marginal costs of holding and financing these buffers.

First versus subsequent floods. Proposition 2 predicts that a firm's first flood should trigger the largest policy response, since the firm has the least experience from which to estimate disruption risk and the shock therefore shifts beliefs the most. Subsequent floods should produce attenuated responses as the firm accumulates experience. Table 6 tests this prediction by splitting the sample into first-flood and repeated-flood subsamples.

Panel A reports results for the first-flood subsample. The effects are large and significant across all four outcomes, with import volume increasing by 0.325 log points, input inventories by 6.9 percentage points relative to pre-treatment sales, cash falling by 2.7 percentage points, and leverage increasing by 4.8 percentage points. Panel B reports results for the repeated-flood subsample. The coefficients are uniformly smaller and statistically insignificant across all four outcomes.

This pattern is consistent with Bayesian learning. The first flood shifts beliefs substantially because the firm has limited prior experience, generating large policy adjustments. Subsequent floods, having already moved beliefs toward higher disruption probabilities, provide less new information and produce attenuated responses. This attenuation distinguishes the learning channel from a simple restocking explanation, in which each flood would generate a response of similar magnitude regardless of the firm's history of disruptions.

Under-prepared firms. If the treatment effects reflect learning and adjustment, they should be concentrated among firms that had the most room to adjust their policies. To classify firms by ex ante preparedness, I regress pre-treatment input inventory levels on the same firm character-

[12] Figure IA.B.1 in Internet Appendix B replicates the event study using the alternative staggered difference-in-differences estimators of Borusyak, Jaravel, and Spiess (2024), Sun and Abraham (2021), and de Chaisemartin and D'Haultfœuille (2020). All four estimators produce nearly identical point estimates and confidence intervals.

istics used as controls in the main specification (firm size, sales, market-to-book, R&D intensity, profitability, and capital expenditures) and retain the residual. Firms with below-median residual inventory, those holding less inventory than predicted by their observable characteristics, are classified as under-prepared. Table 7 tests this prediction by interacting the treatment effect with this under-preparedness indicator.

The standalone *Treated* $\times$ *Post* coefficients are insignificant across all four outcomes, suggesting that well-prepared firms barely respond to the shock. The entire adjustment is concentrated among under-prepared firms. The triple interaction (*Treated* $\times$ *Under-prepared* $\times$ *Post*) shows that under-prepared firms increase import volume by an additional 0.451 log points, build an additional 11.7 percentage points of input inventory relative to sales, draw down an additional 9.2 percentage points of cash relative to assets, and increase leverage by an additional 8.4 percentage points. These firms had the most room to adjust their policies toward the new optimum implied by the updated beliefs.

Evidence from 10-K filings. The results above show that firms adjust their financial policies after flood shocks in ways consistent with belief updating. If firms are genuinely reassessing their exposure to supply chain risk, this should also be visible in how they communicate about these risks to investors. Table 8 tests this by examining changes in supply chain language in 10-K filings following flood events. Treated firms significantly increase all four measures of supply chain language after floods, including the frequency of supply chain keywords and the share of action-oriented supply chain sentences. These results suggest that the policy adjustments documented above reflect a deliberate reassessment of supply chain risk.

Sources of post-flood imports. Figure 6 decomposes the firm's shipments into three components based on supplier identity: the supplier directly disrupted by the flood, existing non-disrupted suppliers with whom the firm had a pre-flood trading relationship, and new suppliers that first appear in the firm's import records after the flood. The figure plots each component's share of the firm's total shipments over the event window.

Three patterns emerge. First, the disrupted supplier's share decrease modestly, by approximately 5 percentage points at the time of the flood, and remains at that lower level thereafter. This pattern suggests that the relationship survives the disruption, but that the firm receives less from the disrupted supplier, potentially because scarce output is allocated across customers during the recovery period. Second, existing non-disrupted suppliers do not offset the decrease in imports. Their share instead declines by roughly 18 percentage points over the event window. The entire post-flood reallocation is absorbed by new suppliers, whose share increases from near zero before the shock to approximately 23% of total shipments by year five. This pattern is consistent with firms redirecting finite import and procurement capacity toward sourcing the fragile input, scaling back imports of less critical inputs to accommodate the buildup. Third, the entry of new suppliers unfolds gradually over several years, which is consistent with frictions in replacing sources of fragile inputs.

Table 9 compares observable characteristics of new post-flood suppliers to the disrupted suppliers. New suppliers serve a comparable number of customers (3.7 vs. 4.1) and hold similar HS-code market shares, suggesting that firms are not sourcing from smaller or less established suppliers. However, new suppliers are somewhat more specialized, covering fewer HS-6 codes (15.3 vs. 23.4) and shipping from fewer ports (4.9 vs. 5.9), and 82% are located in a different country than the disrupted supplier.

Supplier allocation and customer importance. The evidence above shows that, on average, firms build inventories of the disrupted input by establishing new supplier relationships rather than by scaling up existing ones. This average response, however, may mask heterogeneity across firms in how the burden of the disruption is shared. When a supplier's capacity is reduced, it must decide how to allocate its residual output across its customer base, and customers need not bear the loss equally. A natural possibility is that suppliers allocate their reduced post-flood capacity preferentially to their most important pre-treatment customers, with whom they have larger and presumably more valuable trading relationships. Under this form of rationing, firms that accounted for a larger share of the disrupted supplier's pre-flood export volume would continue to receive a substantial portion of their pre-flood shipments, while smaller customers would bear a disproportionate share of the capacity loss. Important customers would therefore experience a smaller effective disruption, face less pressure to source from new suppliers, and build less inventory in response. Table 10 tests this by interacting the treatment effect with customer importance, measured as the firm's share of the disrupted supplier's pre-treatment export volume. The interaction is negative and significant for both import volume and input inventories, suggesting that important customers source less aggressively and build less inventory after floods. This pattern is consistent with disrupted suppliers rationing their residual capacity in favor of their most valuable trading relationships, partially insulating those firms from the need to build operational resilience.

Overall, the evidence in this section supports the learning channel formalized in Proposition 2. Supplier floods cause firms to revise their beliefs about disruption risk upward and to increase their investment in operational resilience. The response is strongest for first-time floods, concentrated among firms that had the most room to adjust, and visible in how firms communicate about supply chain risk to investors. Firms build their post-flood inventories primarily by establishing new supplier relationships, with the burden of adjustment falling disproportionately on less important customers of the disrupted supplier. The scale and persistence of these responses is notable given that the median flood lasts just seven days and the disrupted supplier relationship continues at reduced volume, suggesting that the results reflect a sustained reassessment of disruption risk rather than one-time adaptation to a temporary shortfall. These findings are consistent with a deliberate, forward-looking response to supply chain risk.

7.0.1 Robustness tests and additional results

Internet Appendix B presents several robustness tests.

Relationship definition. The baseline treatment definition classifies a firm as treated if it imported from the flood-affected supplier at any point in the five years before the flood. Table IA.B.9 assesses the sensitivity of the results to this definition by requiring treated firms to have at least one, two, or three years of active importing from the affected supplier before the flood. The results are nearly identical across all definitions, showing that the baseline estimates are not sensitive to the choice of relationship threshold.

Alternative staggered DiD estimators. The baseline estimates use a stacked regression approach (Gormley and Matsa, 2011; Cengiz et al., 2019; Deshpande and Li, 2019; Baker, Larcker, and Wang, 2022), which addresses concerns about bias from heterogeneous treatment effects in staggered designs. To verify that the results are not sensitive to the choice of estimator, Figure IA.B.1 replicates the quarterly event study using three alternative estimators that explicitly allow for heterogeneous treatment effects across cohorts and periods. These are the imputation estimator of Borusyak, Jaravel, and Spiess (2024), the interaction-weighted estimator of Sun and Abraham (2021), and the estimator of de Chaisemartin and D'Haultfœuille (2020). All four estimators produce nearly identical point estimates and confidence intervals across the four outcome variables.

Controlling for ordering lumpiness. Firms with scarce suppliers may place fewer, larger orders due to higher fixed ordering costs or longer lead times. If ordering lumpiness correlates with supplier scarcity, the difference-in-differences estimates could partly reflect pre-existing differences in ordering behavior between treated and control firms rather than genuine policy responses to floods. This concern is related to the import intensity test in Table IA.B.4, but the difference-in-differences setting allows it to be addressed at a more granular level using firm-specific ordering behavior around each flood event. Table IA.B.10 adds two pre-treatment measures of ordering regularity as controls. These are the coefficient of variation of quarterly shipments, which captures how lumpy or smooth a firm's ordering pattern is, and the log of mean pre-treatment quarterly shipments, which controls for ordering scale. The treatment effects are virtually unchanged across all four outcomes. While the coefficient of variation has a significant independent effect on some outcomes, suggesting that ordering lumpiness does affect post-flood corporate policies, it does not attenuate or drive the main treatment effect.

8 Do input inventory buffers mitigate the performance effects of supply chain disruptions?

The previous results show that firms build input inventories in response to perceived supply chain disruption risk and finance this buildup by drawing down cash and increasing leverage. A natural next question is whether this reallocation protects firms when disruptions occur. The precautionary savings literature emphasizes cash as the primary buffer against adverse shocks (Almeida, Campello, and Weisbach, 2004; Bates, Kahle, and Stulz, 2009). When the binding constraint is physical availability of inputs, however, the model predicts that input inventories rather than cash should insulate firms from the performance consequences of disruptions. To test this prediction, I interact the treatment effect with an indicator for whether the firm's pre-treatment input inventories are above the sample median. Because floods are transient shocks, any performance effects should be concentrated in the quarters immediately following the disruption. I therefore estimate these regressions using a one-year pre- and post-treatment window. Table 11 reports the results.

Firms with below-median input inventories suffer significant performance losses in the year following a flood. Operating income decreases by 0.7 percentage points of pre-treatment assets (column 1), sales decline by 1.8 percentage points (column 2), and gross profit falls by 1.1 percentage points (column 3). The interaction *Treated* × *Post* × *High buffer* is positive and highly significant across all three outcomes, with magnitudes that fully offset the baseline losses. The net effect for firms with above-median input inventories is not statistically distinguishable from zero, suggesting that pre-treatment inventory buffers mitigate the short-run performance effects of supply chain disruptions. Because both high-buffer and low-buffer firms are exposed to the same flood shock, the fact that performance losses are concentrated among low-buffer firms is consistent with differences in ex ante preparedness rather than exposure to the disruption itself driving the results.

These findings suggest that firms with inventories of fragile inputs in place before a disruption are better able to maintain performance, while firms with insufficient buffers suffer significant losses. More broadly, the operational resilience investment documented in the cross-sectional and difference-in-differences analyses has real consequences for firm performance.

9 Conclusion

When a shock constrains the physical availability of inputs needed for production, maintaining corporate liquidity becomes a less effective buffer. This paper shows that firms respond to this risk by investing in operational resilience. In particular, they reallocate resources from cash to inventories of fragile inputs and increase leverage to finance the additional inventory spending, and they update these policies as they learn about disruption risk from observed shocks.

Using a new transaction-level dataset on the global supply chains of U.S. manufacturing firms, I construct a measure of supplier scarcity that captures the availability of alternative suppliers across the inputs a firm sources. Flood shocks confirm that this measure identifies supply relationships that are difficult to replace in the short run. Firms can reallocate sourcing toward unaffected suppliers for generic inputs after disruptions but show no significant reallocation for fragile inputs.

A model with two inputs and financing frictions formalizes these tradeoffs. When alternative suppliers are scarce, the marginal value of cash falls because the input is physically unavailable during disruptions regardless of the firm's ability to pay. Inventories of fragile inputs, by contrast, ensure operational continuity. The firm therefore reallocates from cash to inventories and issues debt to finance the additional inventory investment. This joint prediction distinguishes investment in operational resilience from standard precautionary savings, which predicts that firms accumulate cash as a buffer against adverse shocks under the assumption that cash can always be converted into the resources they need.

The empirical evidence supports these predictions. Cross-sectional and instrumental variables estimates show that supplier scarcity increases input inventories, decreases cash, and increases leverage. Supplier floods provide an independent test of the model's learning predictions. Firms respond most strongly to their first flood, when the shock to beliefs is largest, and the response is concentrated among firms that were least prepared. Firms build their post-flood inventories primarily by establishing new supplier relationships rather than by expanding existing ones, and the response is weaker for firms that were more important pre-flood customers of the disrupted supplier, consistent with preferential allocation of reduced supplier capacity to the most valuable trading relationships. Firms with higher inventory buffers are protected from the performance losses associated with supply chain disruptions, suggesting that investment in operational resilience pays off.

These findings show that the nature of risk determines which asset provides effective insurance. For risks that financial instruments can hedge, cash is often the optimal buffer. For supply chain disruptions affecting inputs with scarce suppliers, the optimal buffer is the input itself.

References

Acemoglu, D., and P. D. Azar. 2020. Endogenous production networks. *Econometrica* 88:33–82.

Acemoglu, D., and A. Tahbaz-Salehi. 2020. Firms, failures, and fluctuations: The macroeconomics of supply chain disruptions. Working Paper.

Acharya, V., H. Almeida, Y. Amihud, and P. Liu. 2025. Corporate resiliency and the choice between financial and operational hedging. Working Paper.

Alessandria, G., J. P. Kaboski, and V. Midrigan. 2011. US trade and inventory dynamics. *American Economic Review* 101:303–7.

Alessandria, G., S. Y. Khan, A. Khederlarian, C. Mix, and K. J. Ruhl. 2023. The aggregate effects of global and local supply chain disruptions: 2020–2022. *Journal of International Economics* 146:103828–.

Allayannis, G., J. Ihrig, and J. P. Weston. 2001. Exchange-rate hedging: Financial versus operational strategies. *American Economic Review* 91:391–5.

Almeida, H., M. Campello, and M. S. Weisbach. 2004. The cash flow sensitivity of cash. *Journal of Finance* 59:1777–804.

Antràs, P., and D. Chor. 2013. Organizing the global value chain. *Econometrica* 81:2127–204.

Antràs, P., T. C. Fort, and F. Tintelnot. 2017. The margins of global sourcing: Theory and evidence from US firms. *American Economic Review* 107:2514–64.

Antràs, P., and R. W. Staiger. 2012. Offshoring and the role of trade agreements. *American Economic Review* 102:3140–83.

Baker, A. C., D. F. Larcker, and C. C. Y. Wang. 2022. How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics* 144:370–95.

Barrot, J.-N. 2016. Trade credit and industry dynamics: Evidence from trucking firms. *The Journal of Finance* 71:1975–2016.

Barrot, J.-N., and J. Sauvagnat. 2016. Input specificity and the propagation of idiosyncratic shocks in production networks. *The Quarterly Journal of Economics* 131:1543–92.

Bartik, T. J. 1991. *Who benefits from state and local economic development policies?* W.E. Upjohn Institute for Employment Research.

Bates, T. W., K. M. Kahle, and R. M. Stulz. 2009. Why do U.S. firms hold so much more cash than they used to? *The Journal of Finance* 64:1985–2021.

Bernile, G., V. Bhagwat, and P. R. Rau. 2017. What doesn't kill you will only make you more risk-loving: Early-life disasters and CEO behavior. *Journal of Finance* 72:167–206.

Blum, B., S. Claro, K. Dasgupta, and I. Horstmann. 2019. Inventory management, product quality, and cross-country income differences. *American Economic Journal: Macroeconomics* 11:338–88.

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

Borusyak, K., P. Hull, and X. Jaravel. 2022. Quasi-experimental shift-share research designs. *Review of Economic Studies* 89:181–213.

Borusyak, K., X. Jaravel, and J. Spiess. 2024. Revisiting event-study designs: Robust and efficient estimation. *Review of Economic Studies* 91:3253–85.

Carvalho, V. M., M. Nirei, Y. U. Saito, and A. Tahbaz-Salehi. 2021. Supply chain disruptions: Evidence from the great east Japan earthquake. *The Quarterly Journal of Economics* 136:1255–321.

Cavallo, A., and O. Kryvtsov. 2023. What can stockouts tell us about inflation? Evidence from online micro data. *Journal of International Economics* 146:103769–.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer. 2019. The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics* 134:1405–54.

de Chaisemartin, C., and X. D'Haultfœuille. 2020. Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review* 110:2964–96.

Deshpande, M., and Y. Li. 2019. Who is screened out? Application costs and the targeting of disability programs. *American Economics Review* 11:213–48.

Elliott, M., B. Golub, and M. V. Leduc. 2022. Supply network formation and fragility. *American Economic Review* 112:2701–47.

Ersahin, N., M. Giannetti, and R. Huang. 2024. Supply chain risk: Changes in supplier composition and vertical integration. *Journal of International Economics* 147:103854–.

Flaaen, A., F. Haberkorn, L. Lewis, A. Monken, J. Pierce, R. Rhodes, and M. Yi. 2023. Bill of lading data in international trade research with an application to the Covid-19 pandemic. *Review of International Economics* 31:1146–72.

Gallagher, J. 2014. Learning about an infrequent event: Evidence from flood insurance take-up in the United States. *American Economic Journal: Applied Economics* 6:206–33.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift. 2020. Bartik instruments: What, when, why, and how. *American Economic Review* 110:2586–624.

Goodman-Bacon, A. 2021. Differences-in-differences with variation in treatment timing. *Journal of Econometrics* 225:254–77.

Gormley, T. A., and D. A. Matsa. 2011. Growing out of trouble? Corporate responses to liability risk. *The Review of Financial Studies* 24:2781–821.

Graham, J. R., and M. T. Leary. 2018. The evolution of corporate cash. *The Review of Financial Studies* 31:4288–344.

Hassan, T. A., S. Hollander, L. van Lent, and A. Tahoun. 2019. Firm-level political risk: Measurement and effects. *The Quaterly Journal of Economics* 134:2135–202.

Hoberg, G., and S. K. Moon. 2017. Offshore activities and financial vs operational hedging. *Journal of Financial Economics* 125:217–44.

Hornok, C., and M. Koren. 2015. Per-shipment costs and the lumpiness of international trade. *Review of Economics and Statistics* 97:525–30.

Jensen, M. C., and W. H. Meckling. 1976. Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics* 3:305–60.

Khan, S. Y., and A. Khederlarian. 2025. Inventories, input costs, and productivity gains from trade liberalizations. *International Economic Review* 66:175–99.

Kulchania, M., and S. Thomas. 2017. Cash reserves as a hedge against supply-chain risk. *Journal of Financial and Quantitative Analysis* 52:1951–88.

Malmendier, U., and S. Nagel. 2011. Depression babies: Do macroeconomic experiences affect risk taking? *The Quarterly Journal of Economics* 126:373–416.

Martin, J., I. Mejean, and M. Parenti. 2026. Relationship stickiness, international trade, and economic uncertainty. *The Review of Economics and Statistics* 108:179–93.

Myers, S. 1977. Determinants of corporate borrowing. *Journal of Financial Economics* 5:147–75.

Myers, S., and N. S. Majluf. 1984. Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics* 13:187–221.

Nikolov, B., and T. Whited. 2014. Agency conflicts and cash: Estimates from a dynamic model. *The Journal of Finance* 69:1883–921.

Oberfield, E. 2018. A theory of input-output architecture. *Econometrica* 86:559–89.

Pankratz, N. M., and C. M. Schiller. 2025. Climate change and adaptation in global supply-chain networks. *The Review of Financial Studies* 37:1729–77.

Pástor, L., and P. Veronesi. 2003. Stock valuation and learning about profitability. *Journal of Finance* 58:1749–89.

Rajan, R. G., and L. Zingales. 1995. What do we know about capital structure? Some evidence from international data. *The Journal of Finance* 50:1421–60.

Shleifer, A., and R. W. Vishny. 1991. Liquidation values and debt capacity: A market equilibrium approach. *The Journal of Finance* 47:1343–66.

Silva, J. M. C. S., and S. Tenreyro. 2006. The log of gravity. *The Review of Economics and Statistics* 88:641–58.

Sun, L., and S. Abraham. 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics* 225:175–99.

Yang, S. A., and J. R. Birge. 2018. Trade credit, risk sharing, and inventory financing portfolios. *Management Science* 64:3469–970.

Yu, G. Y. 2025. Do bank CEOs learn from banking crises? *Journal of Financial Economics* 166:104009–.

Figure 1: Model timeline

This figure summarizes the timing of the model. In period t, the firm chooses input purchases, inventories, cash reserves, and debt subject to convex financing costs. The CES production technology uses both generic and fragile inputs. Between periods, the fragile input supply chain is disrupted with probability q_s. In the disruption state ($\phi_{s,t+1} = 1$), the firm cannot purchase fragile inputs on the spot market and must rely on inventories carried from period t. In the no-disruption state ($\phi_{s,t+1} = 0$), both inputs are available. The firm then makes its period $t + 1$ production and financing decisions and distributes terminal cash flows to shareholders.

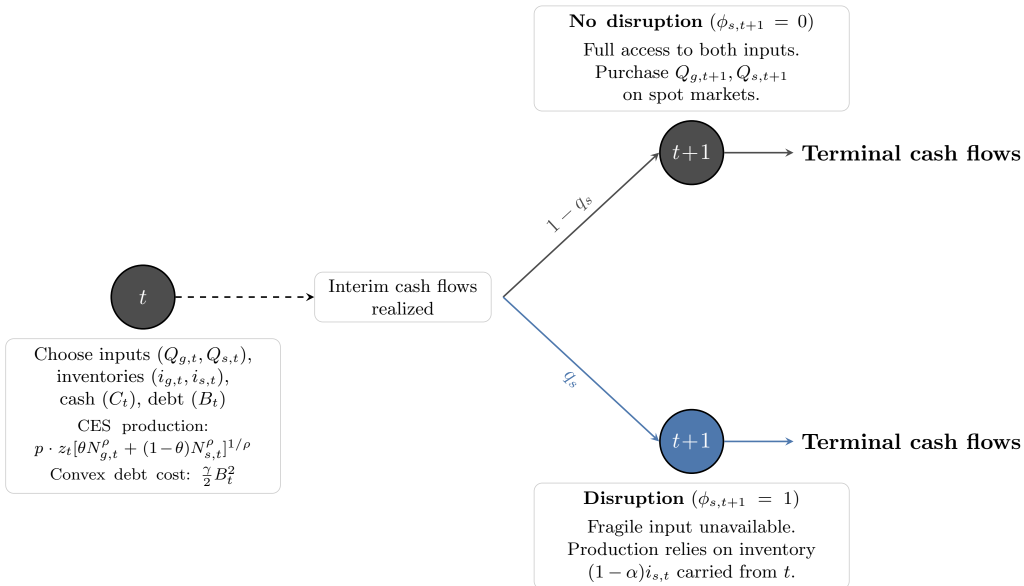

Figure 2: Fraction of inputs by the number of global suppliers

This figure shows the fraction of inputs sourced by firms by the number of suppliers providing those inputs. The sample consists of 907 publicly listed U.S. manufacturing firms with global supply chains from Q1 2007 to Q4 2019.

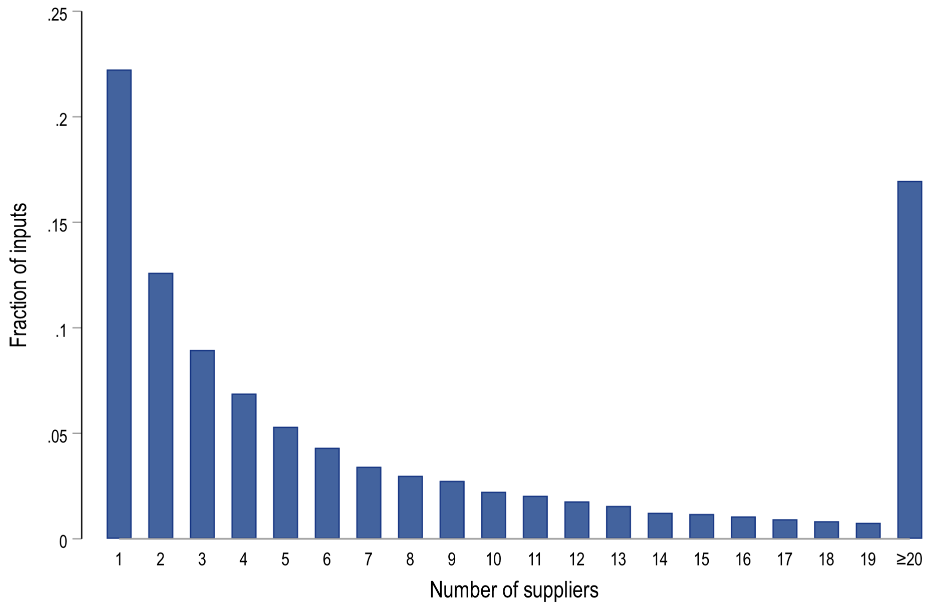

Figure 3: The geographic distribution of suppliers and floods

Panels A and B of this figure show the geographic distribution of foreign suppliers and floods, respectively. The sample consists of 907 publicly listed U.S. manufacturing firms with global supply chains from Q1 2007 to Q4 2019.

Panel A: Geographic distribution of suppliers

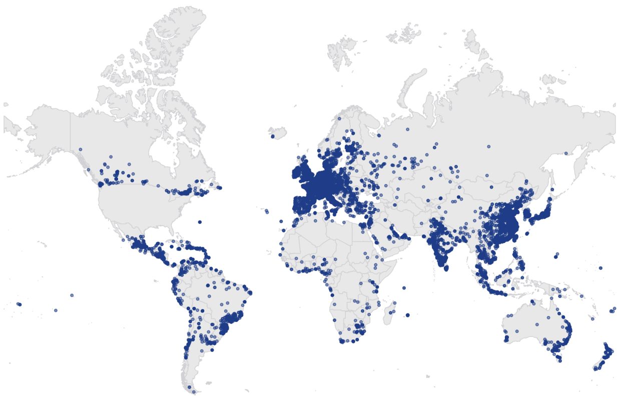

Panel B: Geographic distribution of floods

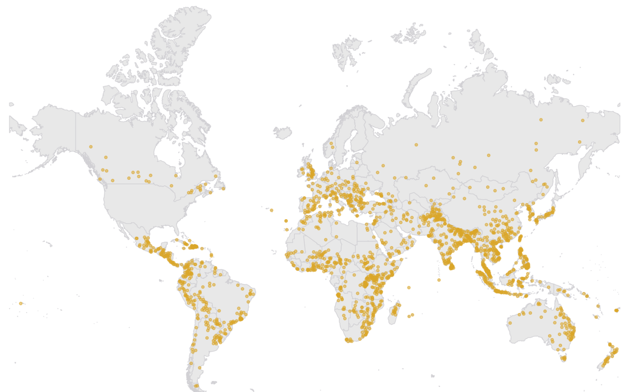

Figure 4: Supplier scarcity validation event study

This figure plots event study coefficients and 90% confidence intervals from Poisson pseudo-maximum likelihood (PPML) regressions of the number of shipments on event-week indicators interacted with treatment status. Panel A estimates the specification from column (1) of Table 2 on all affected firm-supplier pairs. Panel B estimates separate regressions for shipments from unaffected suppliers of fragile inputs and unaffected suppliers of generic inputs, corresponding to columns (2) and (3) of Table 2. Suppliers are classified as fragile if they provide at least one input for which the total number of global suppliers is below the sample median. $Treated$ equals one for firms that had at least one flood-affected supplier at any time in the five years before the flood. Panel A includes firm×supplier×cohort and calendar week×cohort fixed effects, with standard errors clustered at the firm×supplier and calendar week levels. Panel B includes firm×cohort and calendar week×cohort fixed effects, with standard errors clustered at the firm and calendar week levels.

Panel A: All pairs

Panel B: Unaffected pairs

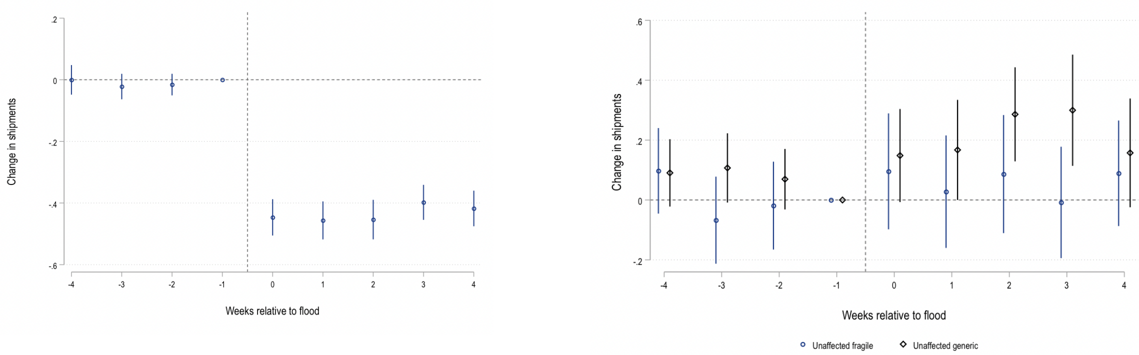

Figure 5: Corporate policies event study

This figure plots the β coefficients and 90% confidence intervals from event study regressions based on a specification analogous to Equation 3: $Y_{itj} = \alpha_{ij} + \alpha_{tj} + \sum_{k \neq -1} \beta_k \mathbb{1}[\text{Event year} = k] \times \text{Treated}_{ij} + \gamma \mathbf{X}_i \times \text{Post}_{tj} + \epsilon_{itj}$, where i denotes firms, t quarters, and j flood events (cohorts). Y_{itj} is *ln(Volume)*, *Input inventories/Sales*, *Cash/Assets*, or *Total debt/Assets*. $Treated_{ij}$ equals one for firms that imported from a flood-affected supplier of fragile inputs at any time in the five years before the flood, and *Event year* is a vector of indicator variables for event years -5 through $+5$. $\mathbf{X}_i \times Post_{tj}$ is a vector of fixed firm-level controls, including the log of assets, log of sales, market-to-book ratios, R&D expenditures over sales, cash flows over assets, and capital expenditures over net property, plant, and equipment measured in the pre-treatment quarter, interacted with $Post_{tj}$, an indicator variable that equals one during the quarters following the start of the flood. α_{ij} and α_{tj} denote firm×cohort and calendar quarter×cohort fixed effects, respectively. Standard errors are clustered at the firm and quarter levels.

Panel A: ln(Volume)

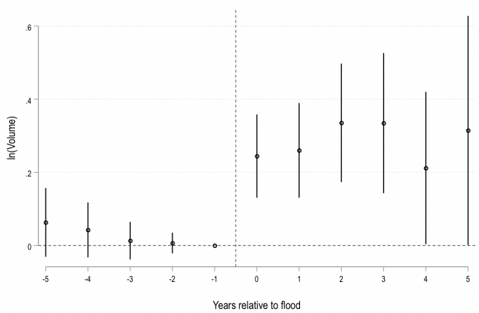

Panel B: Input inventories/Sales

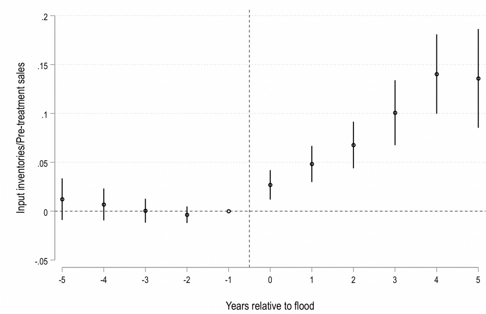

Panel C: Cash/Assets

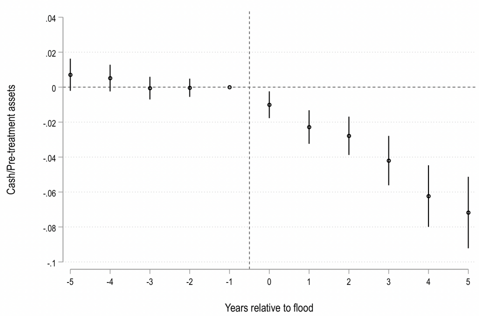

Panel D: Total debt/Assets

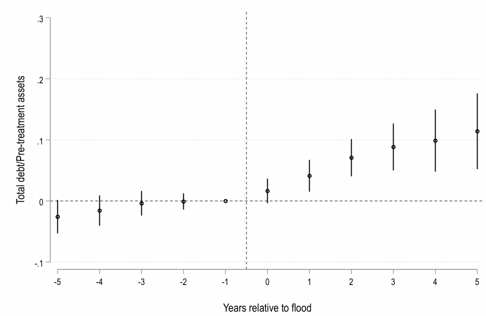

Figure 6: Shipment decomposition by supplier type

This figure decomposes the change in shipments from Figure 5 into three components based on supplier identity: *Disrupted supplier* (the supplier directly hit by the flood), *Existing suppliers* (non-disrupted suppliers with whom the firm had a pre-flood trading relationship), and *New suppliers* (suppliers that first appear in the firm's import records after the flood). The dependent variable is each component's share of the firm's total shipments in a given quarter. Each series plots β coefficients and 90% confidence intervals from event study regressions of the form $Y_{itj} = \alpha_{ij} + \alpha_{tj} + \sum_{k \neq -1} \beta_k \mathbb{1}[\text{Event year} = k] \times \text{Treated}_{ij} + \gamma \mathbf{X}_i \times \text{Post}_{tj} + \epsilon_{itj}$. All specifications include firm$\times$cohort and calendar quarter$\times$cohort fixed effects and the same controls as in Table 5. Standard errors are clustered at the firm and calendar quarter levels.

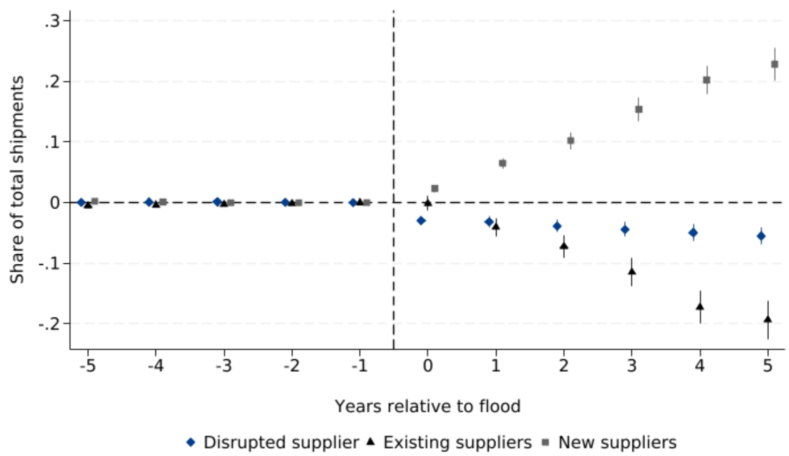

Table 1: Summary Statistics

This table reports summary statistics on global supply chains (Panel A), accounting characteristics (Panel B), and floods (Panel C). The sample in Panels A and B is a firm-quarter panel of 907 publicly listed U.S. manufacturing firms with global supply chains from Q1 2007 to Q4 2019. The sample in Panel C comprises the 1,508 floods that affect firms through their suppliers. Appendix A presents variable definitions.

	Obs	Mean	Std. Dev.	25th Pct.	Median	75th Pct.
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Supply chains						
Import volume (TEU)	31,214	561.27	1,340.52	12.00	85.45	414.62
Number of import countries	31,214	8.72	8.13	3.00	6.00	12.00
Number of products	31,214	22.49	30.23	4.00	11.00	29.00
Number of shipments	31,214	277.89	592.51	13.00	63.00	250.00
Number of suppliers	31,214	24.81	38.48	4.00	11.00	30.00
Panel B: Accounting characteristics						
Assets ($billions)	31,214	5.86	13.32	0.36	1.30	4.38
Sales ($millions)	31,214	1,178.94	2,489.53	97.24	316.63	958.40
Cash/Assets	31,214	0.14	0.15	0.04	0.10	0.20
Cash flow/Assets	31,214	0.03	0.04	0.02	0.03	0.04
Capex/Net PP&E	31,207	0.13	0.11	0.05	0.10	0.17
Input inventories/Sales	31,213	0.38	0.32	0.18	0.30	0.50
M/B	31,214	3.02	5.55	1.34	2.19	3.63
R&D/Sales	31,214	0.12	2.70	0.00	0.01	0.06
Total debt/Assets	31,214	0.23	0.19	0.08	0.22	0.35
Panel C: Floods						
Flood length (days)	1,508	12.53	18.28	4.00	7.00	14.00
Number displaced	1,508	84,851.41	534,770.51	0.00	1,000.00	15,000.00
Number of casualties	1,508	112.51	2,597.52	0.00	3.00	17.00
Area affected (sq miles)	1,508	57,580.99	101,017.61	9,453.93	23,963.95	62,223.51

Table 2: Imports of fragile and generic inputs around supplier floods

This table presents results from stacked difference-in-differences Poisson pseudo-maximum likelihood (PPML) regressions of the number of shipments in a four-week window around flood events. The sample consists of 907 publicly listed U.S. manufacturing firms with global supply chains from Q1 2007 to Q4 2019. Each flood event defines a separate cohort. The data are organized as a firm–supplier–cohort–time panel in column (1) and as a firm–cohort–time panel in columns (2) and (3). The dependent variable in column (1) is the number of shipments from any supplier. In columns (2) and (3), the dependent variable is the aggregate firm-level number of shipments from unaffected suppliers of fragile and generic inputs, respectively. In column (1), *Treated* equals one for firm×supplier pairs where the firm imported from a flood-affected supplier at any time in the five years before the event. In columns (2) and (3), *Treated* equals one for firms that had at least one treated supplier in a cohort. Suppliers are classified as fragile if they provide at least one input for which the total number of global suppliers is below the sample median. Treated and control firms are balanced by exact matching on two-digit SIC codes and nearest matching on the pre-treatment log of assets and aggregate import volume. Column (1) includes firm×supplier×cohort and calendar week×cohort fixed effects. Columns (2) and (3) include firm×cohort and calendar week×cohort fixed effects. Standard errors are clustered at the firm×supplier and calendar week levels in column (1) and at the firm and calendar week levels in columns (2) and (3), and are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively. Appendix A presents variable definitions.

	Number of shipments		
	All pairs	Unaffected pairs	
	All	Fragile	Generic
	(1)	(2)	(3)
Treated × Post	-0.425^{***}	0.058	0.576^{**}
	(0.028)	(0.085)	(0.259)
Controls	✓	✓	✓
Fixed effects			
Firm × Supplier × Cohort	✓		
Firm × Cohort		✓	✓
Calendar week × Cohort	✓	✓	✓
Observations	480,778	32,759	45,321
Pseudo R^2	0.46	0.60	0.98

Table 3: Supplier scarcity and corporate policies

This table reports OLS, first-stage, and two-stage least squares (2SLS) estimates of the effect of supplier scarcity on corporate financial policies. The sample is a firm-quarter panel of 907 publicly listed U.S. manufacturing firms with global supply chains from Q1 2007 to Q4 2019. *Supplier scarcity* is the import-volume-weighted average number of alternative suppliers across the HS codes in firm i's import basket during quarter $t-1$, rescaled so that higher values indicate greater scarcity. *Supplier scarcity* is instrumented with the one-quarter lag of $Z_{i,t} = \sum_h w_{ih,t-1}(-\ln N_{h,t})$, which aggregates the negative log number of globally active suppliers across a firm's input markets using lagged import volume shares as weights. *Input inventories/Sales* is input inventories divided by pre-treatment sales. *Cash/Assets* is cash and short-term investments divided by pre-treatment assets. *Total debt/Assets* is the sum of current and long-term liabilities divided by pre-treatment assets. Columns (1) to (3) report OLS estimates, column (4) reports the first-stage regression of supplier scarcity on the instrument and controls, and columns (5) to (7) report 2SLS estimates. All specifications include industry and calendar quarter fixed effects. Standard errors clustered at the firm and calendar quarter levels are reported in parentheses. The Kleiberger-Paap (KP) F-statistic tests the strength of the first stage. Statistical significance at the 1, 5, and 10 percent levels are denoted by ***, **, *, respectively. Appendix A presents variable definitions.

	OLS			First-stage	Second-stage		
	Input inv./Sales	Cash/Assets	Total debt/Assets	Sup. scarcity	Input inv./Sales	Cash/Assets	Total debt/Assets
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Supplier scarcity$_{t-1}$	0.213***	−0.220***	0.164***		0.841***	−0.214**	0.362**
	(0.076)	(0.039)	(0.047)		(0.223)	(0.098)	(0.149)
$Z_{i,t-1}$				0.020***			
				(0.001)			
Ln(Assets)$_{t-1}$	0.113***	0.035***	0.032**	−0.006	0.115***	0.035***	0.033**
	(0.021)	(0.009)	(0.013)	(0.003)	(0.021)	(0.009)	(0.013)
Ln(Sales)$_{t-1}$	−0.154***	−0.047***	−0.003	0.023***	−0.168***	−0.048***	−0.008
	(0.024)	(0.009)	(0.014)	(0.004)	(0.025)	(0.009)	(0.015)
Market-to-book$_{t-1}$	0.000	0.001***	0.001	−0.000	0.000	0.001***	0.001
	(0.001)	(0.000)	(0.001)	(0.000)	(0.001)	(0.000)	(0.001)
R&D/Sales$_{t-1}$	−0.002	−0.001	0.001	0.001***	−0.003	−0.001	0.001
	(0.002)	(0.001)	(0.002)	(0.000)	(0.002)	(0.001)	(0.002)
Cash flow/Assets$_{t-1}$	−0.707***	0.001	−0.484***	0.030	−0.735***	0.001	−0.492***
	(0.262)	(0.134)	(0.178)	(0.050)	(0.261)	(0.134)	(0.174)
Capex/Net PP&E$_{t-1}$	−0.367***	0.348***	−0.226***	−0.052***	−0.323***	0.349***	−0.212***
	(0.073)	(0.036)	(0.041)	(0.017)	(0.073)	(0.037)	(0.042)
Fixed effects							
Industry	✓	✓	✓	✓	✓	✓	✓
Quarter	✓	✓	✓	✓	✓	✓	✓
KP F-statistic						299.8	
Observations	28,714	28,714	28,714	28,714	28,714	28,714	28,714
Adj. R^2	0.183	0.224	0.212	0.271			

Table 4: Supplier network contraction and input prices

This table tests whether the instrument for supplier scarcity, $Z_{i,t-1}$, affects input prices. Each column regresses an equal-weighted input price index on the instrument. The dependent variable is the simple mean of within-HS-6 demeaned log unit values (USD/kg) across a firm's imported products. Column (1) uses all inputs in a firm's import basket. Column (2) focuses on fragile inputs. Control variables include lagged log assets, log sales, market-to-book, R&D/sales, cash flows/assets, and capex/net PP&E. All specifications include industry and calendar quarter fixed effects. Standard errors clustered at the firm and calendar quarter levels are reported in parentheses. Statistical significance at the 1, 5, and 10 percent levels are denoted by ***, **, *, respectively. Appendix A presents variable definitions.

	All inputs	Fragile inputs
	(1)	(2)
$Z_{i,t-1}$	−0.033	-0.464
	(0.420)	(0.335)
Controls	✓	✓
Fixed effects		
Industry	✓	✓
Quarter	✓	✓
Observations	13,730	13,730
Adj. R^2	0.03	0.01

Table 5: Firm response to fragile input disruptions

This table presents results from stacked difference-in-differences regressions of corporate policies in a five-year window around flood events (cohorts). The sample consists of 907 publicly listed U.S. manufacturing firms with global supply chains from Q1 2007 to Q4 2019. The data form a firm–flood event–quarter panel. *ln(Volume)* is the log of import volume measured in twenty-foot equivalent units (TEU). *Input inventories/Sales* is input inventories divided by pre-treatment sales. *Cash/Assets* is cash and short-term investments divided by pre-treatment assets. *Total debt/Assets* is the sum of current and long-term liabilities divided by pre-treatment assets. *Treated* equals one for firms that imported from a flood-affected supplier of fragile inputs at any time in the five years before the flood. Treated and control firms are balanced by exact matching on two-digit SIC codes and nearest matching on the pre-treatment log of assets and aggregate imports. All specifications include firm×cohort and calendar quarter×cohort fixed effects, and controls (log assets, log sales, market-to-book, R&D/sales, cash flow/assets, and capex/net PP&E) interacted with *Post*. Standard errors clustered at the firm and calendar quarter levels are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels are denoted by ***, **, *, respectively. Appendix A presents variable definitions.

	ln(Volume)	Input inventories/Sales	Cash/Assets	Total debt/Assets
	(1)	(2)	(3)	(4)
Treated × Post	0.257***	0.066***	−0.030***	0.061***
	(0.092)	(0.017)	(0.007)	(0.021)
Controls	✓	✓	✓	✓
Fixed effects				
Firm × Flood	✓	✓	✓	✓
Calendar quarter × Flood	✓	✓	✓	✓
Observations	198,238	205,318	205,318	205,318
Adj. R^2	0.82	0.76	0.74	0.66

Table 6: Learning from the first and subsequent supply chain disruptions

This table presents results from stacked difference-in-differences regressions of corporate policies in a five-year window around flood events (cohorts) estimated on the subsample of firm–flood pairs where the flood is either the first to affect the firm's suppliers of fragile inputs, or a subsequent flood. The sample consists of 907 publicly listed U.S. manufacturing firms with global supply chains from Q1 2007 to Q4 2019. The data form a firm–flood event–quarter panel. *ln(Volume)* is the log of import volume measured in twenty-foot equivalent units (TEU). *Input inventories/Sales* is input inventories divided by pre-treatment sales. *Cash/Assets* is cash and short-term investments divided by pre-treatment assets. *Total debt/Assets* is total debt divided by pre-treatment assets. *Treated* equals one for firms that imported from a flood-affected supplier of fragile inputs at any time in the five years before the flood. Treated and control firms are balanced by exact matching on two-digit SIC codes and nearest matching on the pre-treatment log of assets and aggregate imports. All specifications include firm×cohort and calendar quarter×cohort fixed effects, and controls (log assets, log sales, market-to-book, R&D/sales, cash flow/assets, and capex/net PP&E) interacted with *Post*. Standard errors clustered at the firm and calendar quarter levels are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels are denoted by ***, **, *, respectively.

Panel A: First-flood subsample

	ln(Volume)	Input inventories/Sales	Cash/Assets	Total debt/Assets
	(1)	(2)	(3)	(4)
Treated × Post	0.325***	0.069***	−0.027***	0.048**
	(0.083)	(0.025)	(0.008)	(0.021)
Controls	✓	✓	✓	✓
Fixed effects				
Firm × Flood	✓	✓	✓	✓
Calendar quarter × Flood	✓	✓	✓	✓
Observations	70,078	75,343	75,343	75,343
Adj. R^2	0.67	0.73	0.77	0.67

Panel B: Repeat-flood subsample

	ln(Volume)	Input inventories/Sales	Cash/Assets	Total debt/Assets
	(1)	(2)	(3)	(4)
Treated × Post	0.164	0.040	−0.016	0.030
	(0.154)	(0.029)	(0.012)	(0.031)
Controls	✓	✓	✓	✓
Fixed effects				
Firm × Flood	✓	✓	✓	✓
Calendar quarter × Flood	✓	✓	✓	✓
Observations	128,160	129,975	129,975	129,975
Adj. R^2	0.87	0.78	0.72	0.66

Table 7: Inventory buffer stocks and the tradeoff between supply chain disruption risk and resilience

This table presents results from stacked difference-in-differences regressions of corporate policies in a five-year window around flood events (cohorts), including interactions with a measure of ex ante under-preparedness. The sample consists of 907 publicly listed U.S. manufacturing firms with global supply chains from Q1 2007 to Q4 2019. The data form a firm–flood event–quarter panel. *Under-prepared* firms are those with below-median residual inventory, where residual inventory is the residual from a cross-sectional regression of pre-treatment input inventory levels on the same firm characteristics used as controls in the main specification. *ln(Volume)* is the log of import volume measured in twenty-foot equivalent units (TEU). *Input inventories/Sales* is input inventories divided by pre-treatment sales. *Cash/Assets* is cash and short-term investments divided by pre-treatment assets. *Total debt/Assets* is total debt divided by pre-treatment assets. All specifications include firm×cohort and calendar quarter×cohort fixed effects, and controls (log assets, log sales, market-to-book, R&D/sales, cash flow/assets, and capex/net PP&E) interacted with *Post*. Standard errors clustered at the firm and calendar quarter levels are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels are denoted by ***, **, *, respectively.

	ln(Volume)	Input inventories/Sales	Cash/Assets	Total debt/Assets
	(1)	(2)	(3)	(4)
Treated × Post	0.045	0.009	0.013	0.020
	(0.077)	(0.022)	(0.009)	(0.028)
Treated × Under-prepared × Post	0.451**	0.117***	−0.092***	0.084**
	(0.198)	(0.034)	(0.013)	(0.035)
Under-prepared × Post	−0.578***	−0.141***	0.071***	−0.100***
	(0.184)	(0.028)	(0.012)	(0.022)
Controls	✓	✓	✓	✓
Fixed effects				
Firm × Flood	✓	✓	✓	✓
Calendar quarter × Flood	✓	✓	✓	✓
Observations	196,112	202,266	202,266	202,266
Adj. R^2	0.82	0.75	0.74	0.66

Table 8: Supply chain language in 10-K filings after flood shocks

This table presents results from stacked difference-in-differences regressions of textual measures of supply chain language in 10-K filings in a five-year window around flood events (cohorts). The dependent variables are constructed from the full 10-K filing text. *Ln(1+SC bigrams)* is the log of one plus the share of supply-chain-related bigrams (two-word sequences such as "supply chain," "raw materials," and "shipping delays") in total bigrams, following Hassan, Hollander, van Lent, and Tahoun (2019), scaled by 100. *Ln(1+SC keywords)* is the log of one plus the number of supply chain keyword mentions per 1,000 words. *Ln(1+SC action)* is the log of one plus the number of supply chain action-oriented sentences per 1,000 words. *SC sentence share* is the number of supply chain sentences per 100 words of filing text. All specifications include firm×cohort and calendar quarter×cohort fixed effects. Standard errors clustered at the firm and calendar quarter levels are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels are denoted by ***, **, *, respectively.

	ln(1+SC bigrams)	ln(1+SC keywords)	ln(1+SC action)	SC sentence share
	(1)	(2)	(3)	(4)
Treated × Post	0.055***	0.004***	0.002**	0.017**
	(0.018)	(0.001)	(0.001)	(0.007)
Fixed effects				
Firm × Flood	✓	✓	✓	✓
Calendar quarter × Flood	✓	✓	✓	✓
Observations	197,620	197,620	197,620	197,620
Adj. R^2	0.44	0.64	0.62	0.72

Table 9: Characteristics of new vs. disrupted suppliers

This table compares observable characteristics of new post-flood suppliers to the disrupted suppliers they replace. For each supplier, I compute the number of distinct Compustat importers it serves, its average market share across the HS-6 codes it supplies, the number of origin countries, the number of shipping ports, and the number of HS-6 product codes it covers. *In different country* is the percentage of new suppliers located in a different country than the disrupted supplier. Standard errors are reported in parentheses. Statistical significance at the 1, 5, and 10 percent levels are denoted by ***, **, *, respectively.

	Disrupted supplier	New supplier	Difference
Number of customers served	4.12	3.69	-0.43^{***}
			(0.064)
Average HS-code market share	0.034	0.036	0.001^{*}
			(0.001)
Number of origin countries	1.64	1.33	-0.30^{***}
			(0.018)
Number of shipping ports	5.88	4.94	-0.94^{***}
			(0.048)
Number of HS-6 codes	23.4	15.3	-8.08^{***}
			(0.371)
In different country (%)	—	82.2	
Number of N supplier–flood pairs	29,277	33,742	

Table 10: Supplier allocation and customer importance

This table presents results from stacked difference-in-differences regressions of import volume and input inventories in a five-year window around flood events (cohorts), including interactions with customer importance. The sample consists of 907 publicly listed U.S. manufacturing firms with global supply chains from Q1 2007 to Q4 2019. The data form a firm–flood event–quarter panel. *ln(Volume)* is the log of import volume measured in twenty-foot equivalent units (TEU). *Input inventories/Sales* is input inventories divided by pre-treatment sales. *Customer importance* is the firm's share of the disrupted supplier's pre-treatment export volume. *Treated* equals one for firms that imported from a flood-affected supplier of fragile inputs at any time in the five years before the flood. Treated and control firms are balanced by exact matching on two-digit SIC codes and nearest matching on the pre-treatment log of assets and aggregate imports. All specifications include firm×cohort and calendar quarter×cohort fixed effects, and controls (log assets, log sales, market-to-book, R&D/sales, cash flow/assets, and capex/net PP&E) interacted with *Post*. Standard errors clustered at the firm and calendar quarter levels are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels are denoted by ***, **, *, respectively. Appendix A presents variable definitions.

	ln(Volume)	Input inv./Sales
	(1)	(2)
Treated × Post	0.285***	0.076***
	(0.091)	(0.018)
Treated × Post × Customer importance	−0.026**	−0.009***
	(0.013)	(0.003)
Controls		
Firm × Flood	✓	✓
Quarter × Flood	✓	✓
Observations	198,238	205,318
Adj. R^2	0.82	0.76

Table 11: Input inventory buffers and performance during supply chain disruptions
This table presents results from stacked difference-in-differences regressions of firm performance measures in a one-year window around flood events (cohorts), including interactions with a measure of pre-treatment input inventory buffers. The sample consists of 907 publicly listed U.S. manufacturing firms with global supply chains from Q1 2007 to Q4 2019. The data form a firm–flood event–quarter panel. *High buffer* equals one if the firm's pre-treatment ratio of input inventories to assets is above the sample median. *OI/Assets* is operating income before depreciation divided by pre-treatment assets. *Sales/Assets* is sales divided by pre-treatment assets. *Gross profit/Assets* is sales minus cost of goods sold divided by pre-treatment assets. All specifications include firm×cohort and calendar quarter×cohort fixed effects, and controls (log assets, log sales, market-to-book, R&D/sales, cash flow/assets, and capex/net PP&E) interacted with *Post*. Standard errors clustered at the firm and calendar quarter levels are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels are denoted by ***, **, *, respectively.

	OI/Assets	Sales/Assets	Gross profit/Assets
	(1)	(2)	(3)
Treated × Post	-0.007^{***}	-0.018^{***}	-0.011^{***}
	(0.002)	(0.006)	(0.003)
Treated × High buffer × Post	0.012^{***}	0.030^{***}	0.018^{***}
	(0.003)	(0.010)	(0.004)
High buffer × Post	-0.010^{***}	-0.034^{***}	-0.022^{***}
	(0.002)	(0.007)	(0.003)
Controls	✓	✓	✓
Fixed effects			
Firm × Flood	✓	✓	✓
Quarter × Flood	✓	✓	✓
Observations	134,595	134,595	134,595
Adj. R^2	0.62	0.83	0.84

Appendix A: Variable definitions

Variable	Definition
Panel A: Dependent variables	
Cash/Assets	Cash and marketable securities divided by total assets (or pre-treatment assets in DiD specifications).
Gross profit/Assets	Sales minus cost of goods sold, divided by pre-treatment assets.
Input inv./Sales	Materials and supplies for use in production and in-process goods not ready for sale, divided by quarterly sales (or pre-treatment sales in DiD specifications).
ln(Volume)	The natural log of import volume measured in twenty-foot equivalent units (TEU).
OI/Assets	Operating income before depreciation divided by pre-treatment assets.
Sales/Assets	Quarterly sales divided by pre-treatment assets.
Total debt/Assets	Long-term debt plus debt in current liabilities divided by total assets (or pre-treatment assets in DiD specifications).
Panel B: Key independent variables	
Supplier scarcity	The import-volume-weighted average number of alternative suppliers across HS codes in a firm's import basket, rescaled so that higher values indicate greater scarcity.
Treated	In firm–supplier–cohort panels: equals one for firm–supplier pairs where the firm imported from a flood-affected supplier at any time in the five years before the flood. In firm–cohort panels: equals one for firms that imported from a flood-affected supplier of fragile inputs at any time in the five years before the flood.
Post	Equals one during the periods following the start of a flood event (cohort).
High buffer	Indicator equal to one if the firm's pre-treatment ratio of input inventories to assets is above the sample median.

Variable	Definition
Panel B: Key independent variables (continued)	
Under-prepared	Indicator equal to one if the firm's pre-treatment residual inventory (from a regression of input inventory levels on firm characteristics) is below the sample median.
Panel C: Control variables	
Capex/Net PP&E	Capital expenditures divided by net property, plant, and equipment.
Cash flow/Assets	EBITDA minus interest, taxes, and common dividends, divided by total assets.
Ln(Assets)	The natural log of total assets.
Ln(Sales)	The natural log of quarterly sales.
Market-to-book ratio (M/B)	Market value of equity plus book value of debt, divided by total assets.
R&D/Sales	Research and development expenditures divided by sales.

Appendix B: Derivations and proofs

This appendix derives the first-order conditions for the firm's optimization problem described in Section 5 and provides proofs for Propositions 1 and 2.

B.1 First-order conditions

Recall the CES production function from Equation 4:

$$F(z_\tau, N_{g,\tau}, N_{s,\tau}) = p \cdot z_\tau \left[\theta N_{g,\tau}^{\rho} + (1-\theta) N_{s,\tau}^{\rho}\right]^{1/\rho}$$

The marginal products of the generic and fragile inputs are:

$$\frac{\partial F}{\partial N_{g,\tau}} = p \cdot z_\tau \cdot \theta \cdot N_{g,\tau}^{\rho-1} \left[\theta N_{g,\tau}^{\rho} + (1-\theta) N_{s,\tau}^{\rho}\right]^{(1-\rho)/\rho} \tag{B1}$$

$$\frac{\partial F}{\partial N_{s,\tau}} = p \cdot z_\tau \cdot (1-\theta) \cdot N_{s,\tau}^{\rho-1} \left[\theta N_{g,\tau}^{\rho} + (1-\theta) N_{s,\tau}^{\rho}\right]^{(1-\rho)/\rho} \tag{B2}$$

The firm's expected profit in the two disruption states can be written as follows. When $\phi_{s,t+1} = 0$ (no disruption), input utilization in period $t+1$ is $N^0_{g,t+1} = Q^0_{g,t+1} + (1-\alpha) i_{g,t}$ and $N^0_{s,t+1} = Q^0_{s,t+1} + (1-\alpha) i_{s,t}$. When $\phi_{s,t+1} = 1$ (disruption), $N^1_{g,t+1} = Q^1_{g,t+1} + (1-\alpha) i_{g,t}$ and $N^1_{s,t+1} = (1-\alpha) i_{s,t}$ (no spot-market purchases of the fragile input). The period $t+1$ purchase decisions $Q^j_{g,t+1}$ and $Q^j_{s,t+1}$ are state-contingent because the firm optimizes after observing whether a disruption occurred.

The Lagrangian for the period-t problem in Equation 8 is:

$$\begin{aligned}
\mathcal{L} = \mathbb{E}_t\Big[F(z_t, N_{g,t}, N_{s,t})\Big] &- Q_{g,t} - Q_{s,t} - C_t - \frac{\gamma}{2} B_t^2 \\
&+ q_s\, \mathbb{E}_t\Big[\pi_{t+1} \mid \phi_{s,t+1} = 1\Big] + (1-q_s)\, \mathbb{E}_t\Big[\pi_{t+1} \mid \phi_{s,t+1} = 0\Big] \\
&+ \mu_t\Big[W_t + B_t - Q_{g,t} - Q_{s,t} - C_t\Big] \\
&+ \mu^j_{t+1}\Big[W_{t+1} + B_{t+1} - Q_{g,t+1} - (1-\phi_{s,t+1}) Q_{s,t+1}\Big] \\
&+ \delta_{g,t}\, i_{g,t} + \delta_{s,t}\, i_{s,t} + \kappa_t\, C_t
\end{aligned} \tag{B3}$$

where μ_t and μ^j_{t+1} are the multipliers on the period t and period $t+1$ budget constraints (Equation 9 and Equation 7), respectively, with $j \in \{0, 1\}$ indexing the disruption state, and $\delta_{g,t}$, $\delta_{s,t}$, κ_t are the multipliers on the non-negativity constraints for generic inventory, fragile inventory, and cash, respectively. The period $t+1$ multiplier appears because period t choices affect W_{t+1} through interim cash flows and cash reserves, which in turn affect the tightness of the period $t+1$ budget constraint.

The first-order conditions with respect to the period t choice variables are:

Generic input purchases ($Q_{g,t}$):

$$\mathbb{E}_t\left[\frac{\partial F}{\partial N_{g,t}}\,(1+\mu_{t+1})\right]=1+\mu_t \tag{B4}$$

Fragile input purchases ($Q_{s,t}$):

$$\mathbb{E}_t\left[\frac{\partial F}{\partial N_{s,t}}\,(1+\mu_{t+1})\right]=1+\mu_t \tag{B5}$$

The $(1+\mu_{t+1})$ term reflects the fact that period t production generates interim cash flows F_t that carry over into $W_{t+1}=F_t+C_t$. Each additional unit of input purchased in period t raises F_t and therefore relaxes the period $t+1$ budget constraint, with marginal value μ_{t+1}.

Generic input inventory ($i_{g,t}$):

$$(1-\alpha)\left[q_s\,\mathbb{E}_t\left[\frac{\partial F^1}{\partial N_{g,t+1}}\right]+(1-q_s)\,\mathbb{E}_t\left[\frac{\partial F^0}{\partial N_{g,t+1}}\right]\right]-\mathbb{E}_t\left[\frac{\partial F}{\partial N_{g,t}}\,(1+\mu_{t+1})\right]=-\delta_{g,t} \tag{B6}$$

Fragile input inventory ($i_{s,t}$):

$$(1-\alpha)\left[q_s\,\mathbb{E}_t\left[\frac{\partial F^1}{\partial N_{s,t+1}}\right]+(1-q_s)\,\mathbb{E}_t\left[\frac{\partial F^0}{\partial N_{s,t+1}}\right]\right]-\mathbb{E}_t\left[\frac{\partial F}{\partial N_{s,t}}\,(1+\mu_{t+1})\right]=-\delta_{s,t} \tag{B7}$$

where $F^j\equiv F(z_{t+1},N^j_{g,t+1},N^j_{s,t+1})$ for $j\in\{0,1\}$ denotes period $t+1$ production in state j, and $\mu_{t+1}\equiv q_s\mu^1_{t+1}+(1-q_s)\mu^0_{t+1}$ is the expected shadow value of the period $t+1$ budget constraint. The opportunity cost of diverting inputs to inventory includes the continuation value through W_{t+1}, since reducing $N_{g,t}$ (or $N_{s,t}$) lowers F_t and tightens the period $t+1$ budget.

The marginal product of the generic input depends on $N_{s,t+1}$ through the CES aggregator, and because $N_{s,t+1}$ differs across disruption states (the fragile input is unavailable when $\phi_{s,t+1}=1$), the FOC requires probability-weighting across both states.

Cash holdings (C_t):

$$(1 - q_s)\,\mathbb{E}_t\left[\frac{\partial \pi^0_{t+1}}{\partial W_{t+1}}\right] + q_s\,\mathbb{E}_t\left[\frac{\partial \pi^1_{t+1}}{\partial W_{t+1}}\right] = 1 + \mu_t - \kappa_t \tag{B8}$$

Cash carried to period $t+1$ increases W_{t+1}, which funds additional input purchases and reduces the need for costly debt.

Debt (B_t):

$$\gamma B_t = \mu_t \tag{B9}$$

This condition equates the marginal cost of debt (γB_t) to the shadow value of relaxing the budget constraint (μ_t). Because $\gamma > 0$, the optimal debt level is an interior solution $B_t^* = \mu_t/\gamma$, which increases when the budget constraint is tighter (higher μ_t).

B.2 Proof of Proposition 1

Proposition 1 (Disruption risk and corporate policies)**.** *Suppose that $\rho < 1$ and $\sigma < \infty$. Then, as the disruption probability q_s increases: (i) optimal fragile input inventories increase; (ii) optimal cash holdings decrease; (iii) optimal leverage increases. When $q_s = 0$, optimal fragile input inventories converge to zero.*

Proof. The firm's objective (Equation 8) weights the disruption and no-disruption states by q_s and $1 - q_s$, respectively. As q_s increases, the disruption state receives more weight. In that state, the fragile input is unavailable and its marginal product is high (given $\sigma < \infty$), so that the insurance value of fragile input inventories increases. Conversely, cash is less valuable in the disruption state because it can only fund purchases of the generic input, whose marginal product is depressed by the shortage of its complement. The firm therefore reallocates from cash to inventory and issues additional debt to finance the gap, since the marginal cost of debt (γB_t) remains finite at the interior optimum. When $q_s = 0$, the disruption state receives zero weight and inventories provide no insurance value.

More formally, consider the first-order condition for fragile input inventory (Equation B7). The marginal benefit of inventory is $(1 - \alpha)$ times the probability-weighted expected marginal product across the two states. Increasing q_s puts more weight on the disruption state, where the marginal product of the fragile input is high because $N^1_{s,t+1} = (1-\alpha)i_{s,t}$ is small relative to $N_{g,t+1}$. This increases the marginal benefit of inventory, so that $i^*_{s,t}$ increases. For cash, the first-order condition (Equation B8) shows that the marginal value of cash is a weighted average of its value across states. In the disruption state, cash can only fund generic input purchases, whose marginal product is depressed by the complementarity with the scarce fragile input. Higher q_s shifts weight toward this lower-value state, reducing C^*_t. Debt increases because the tighter budget constraint (more spending on inventory, less internal funding from cash) increases μ_t, and $B^*_t = \mu_t/\gamma$.

$\square$

B.3 Proof of Proposition 1a

Proposition 1a (Input substitutability and corporate policies)**.** *As the elasticity of substitution* $\sigma = 1/(1-\rho)$ *decreases: (i) optimal fragile input inventories increase; (ii) optimal cash holdings decrease; (iii) optimal leverage increases.*

Proof. The proof proceeds in three steps.

Step 1: Inventory of fragile inputs. Consider the first-order condition for $i_{s,t}$ in Equation B7. The marginal benefit of an additional unit of fragile input inventory is $(1-\alpha)$ times the probability-weighted expected marginal product across the two disruption states. In the disruption state ($\phi_{s,t+1} = 1$), the fragile input available is $N^1_{s,t+1} = (1-\alpha)i_{s,t}$, and the marginal product from Equation B2 is:

$$\frac{\partial F^1}{\partial N_{s,t+1}} = p \cdot z_{t+1}(1-\theta)N^{\rho-1}_{s,t+1}\left[\theta(N^1_{g,t+1})^\rho + (1-\theta)(N^1_{s,t+1})^\rho\right]^{(1-\rho)/\rho}$$

As ρ decreases (i.e., σ falls), the CES aggregator becomes more concave in its arguments. In the disruption state, the fragile input available is limited to inventory carried over from period t, so that $N^1_{s,t+1} = (1-\alpha)i_{s,t}$, which is strictly less than the generic input $N^1_{g,t+1}$ at any interior solution where the firm uses more generic than fragile input.[13] Because $N^1_{s,t+1}/N^1_{g,t+1} < 1$, reducing ρ increases the marginal product of the scarce input relative to the abundant one. Formally, $\partial^2 F^1/(\partial N_{s,t+1}\partial\rho) < 0$ when the fragile input is scarce relative to the generic input. This means the marginal product

[13] This condition holds at the optimum whenever $q_s < 1$ and $\alpha > 0$, since the firm does not stockpile enough fragile input inventory to fully replicate the no-disruption input ratio. The condition can be verified by contradiction. If $N^1_{s,t+1} \geq N^1_{g,t+1}$, the firm would be over-investing in inventory relative to the insurance value it provides.

of fragile input inventory in the disruption state increases as substitutability falls, increasing the marginal benefit of holding inventory. By the implicit function theorem applied to the FOC, $\partial i^*_{s,t}/\partial\sigma < 0$.

In the limiting case $\sigma \to \infty$ ($\rho \to 1$), the production function becomes $F = p \cdot z_\tau[\theta N_g + (1-\theta)N_s]$, and the marginal product of the fragile input is a constant $(1-\theta)pz_{t+1}$ regardless of the input ratio. Carrying inventory provides no additional insurance value beyond what cash (used to buy generic inputs) provides, and the carrying cost α makes inventory strictly dominated by cash. Hence $i^*_{s,t} \to 0$.

Step 2: Cash holdings. From the FOC for cash (Equation B8), the marginal benefit of cash is the probability-weighted expected marginal value of additional liquid assets in period $t+1$. Cash is useful in both states. In the no-disruption state, it funds purchases of both inputs. In the disruption state, it funds only purchases of generic inputs. As σ decreases, the marginal product of the generic input in the disruption state falls. This occurs because the complementary fragile input is scarce and, under lower substitutability, the CES aggregator penalizes the imbalance between inputs more severely. The value of cash in the disruption state therefore declines. Meanwhile, the marginal benefit of inventory (which directly alleviates the fragile input shortage) increases, as shown in Step 1. The reallocation from cash to inventory follows from the budget constraint (Equation 9). Since the firm has higher demand for fragile inventory, which provides greater insurance value per dollar than cash at the margin, it optimally reduces cash holdings. Hence $\partial C^*_t/\partial\sigma > 0$.

Step 3: Leverage. From the debt FOC (Equation B9), $B^*_t = \mu_t/\gamma$, so leverage increases if and only if the shadow value of the budget constraint μ_t increases. From the budget constraint (Equation 9), $Q_{g,t} + Q_{s,t} + C_t = W_t + B_t$. As σ decreases, Step 1 shows that $i^*_{s,t}$ increases, requiring higher $Q_{s,t}$ since $Q_{s,t} = N_{s,t} + i_{s,t}$. Step 2 shows that C^*_t decreases, freeing resources. However, the increase in inventory spending exceeds the decrease in cash. To see this, note that as σ falls, the marginal product of the fragile input in the disruption state increases (because $N^1_{s,t+1}/N^1_{g,t+1} < 1$ and lower substitutability penalizes this imbalance more severely), creating strong demand for additional inventory. The marginal value of cash, by contrast, depends on the marginal product of the generic input, which is bounded and less sensitive to σ. Total resource needs therefore increase, tightening the budget constraint and increasing μ_t. Hence $B^*_t = \mu_t/\gamma$ increases, and $\partial B^*_t/\partial\sigma < 0$. $\square$

B.4 Proof of Proposition 2

Proposition 2 (Learning from supply chain disruptions)**.** *As the firm observes additional supply chain disruptions: (i) the posterior disruption probability increases, amplifying the effects of Proposition 1; (ii) the marginal effect of each additional disruption decreases.*

Proof.

Part (i). The posterior mean from Equation 13 is $\hat{q}_s(n, T) = (a + n)/(a + b + T)$. When the firm observes a disruption, both the number of disruptions n and the total number of periods T increase by one. The resulting belief update is:

$$\hat{q}_s(n+1,\, T+1) - \hat{q}_s(n,\, T) = \frac{b + T - n}{(a + b + T + 1)(a + b + T)} > 0$$

which is strictly positive whenever $T > n$ (i.e., whenever disruptions are rarer than non-disruption periods). Since each observed disruption increases $\hat{q}_s$, and the comparative statics in Proposition 1 depend on q_s through the disruption probability entering the firm's objective (Equation 8), the optimal policies change in the direction predicted by Proposition 1. Specifically, higher $\hat{q}_s$ increases the weight on the disruption state, increasing the insurance value of fragile input inventories and reducing the relative value of cash.

Part (ii). Consider the discrete belief update when the firm observes one additional disruption in one additional period. The change in the posterior mean is:

$$\hat{q}_s(n+1,\, T+1) - \hat{q}_s(n,\, T) = \frac{b + T - n}{(a + b + T + 1)(a + b + T)}$$

This expression is strictly positive whenever $T > n$, which holds whenever disruptions are rarer than non-disruption periods. The magnitude of the update is decreasing in T. As the firm accumulates more observations, the denominator grows while the numerator grows at a slower rate, so each additional disruption has a progressively smaller effect on beliefs. This reflects Bayesian convergence. The posterior precision $a + b + T$ increases with experience, and the information content of any single disruption diminishes relative to the accumulated knowledge.

$\square$

Internet Appendix for

Building Corporate Resilience to Supply Chain Disruptions

Leandro Sanz

Mendoza College of Business, University of Notre Dame

IA.A Data

This section provides further details on the global supply chain data, presents a set of stylized facts, and reports additional analyses that validate the quality of the data.

IA.A.1 Bill of lading form example

Figure IA.A.1 shows an example of a BOL form. As illustrated, the form collects detailed information about the vessel transporting the goods and includes key logistical details such as the date of departure and arrival, the ports of lading and unlading, as well as the weight and volume occupied by the imported goods in the vessel. It also captures comprehensive data on the imported products, including product descriptions and details about both the recipient and supplier.

IA.A.2 Stylized facts

In this section, I present several stylized facts about the global supply chains of the sample U.S. manufacturing firms. First, I examine the evolution of aggregate imports and establish that the majority of these imports over the sample period consist of inputs. Next, I analyze the geographic distribution of suppliers and the structure of the network linking inputs to suppliers. Despite the global reach of these networks, imports are predominantly concentrated in a few countries, and the connections between suppliers and inputs are often sparse.

Figure IA.A.2 presents data on the total number of shipments (Panel A) and the total import volume in TEUs (Panel B) for sample manufacturing firms. Since 2007, these firms have received over 500,000 shipments per year via maritime transport, totaling more than 1 million TEUs annually. The figures peaked in 2014 and 2015 when firms received nearly 1 million shipments and 1.8 million TEUs per year. However, during economic downturns, such as the global financial crisis, both the total number of shipments and import volume experienced a marked decline.

The figure also breaks down the total imports of U.S. manufacturing firms into raw materials and intermediate goods (collectively referred to as inputs), and finished goods. Over the sample period, inputs represent 55% of imports, while finished products account for the remaining 45%. These statistics are consistent with the growing importance of fragile intermediate inputs relative to final goods in international trade (Antràs and Staiger, 2012).

Panel A of Figure IA.A.3 illustrates the country distribution of foreign suppliers for sample U.S. manufacturing firms in 2019, highlighting the extensive reach of these global supply chains, which include over 25,000 foreign suppliers across 156 countries. However, the distribution is highly uneven. For example, in 2019, 31% of the foreign suppliers in the BOL data were located in China. Panel B of Figure IA.A.3 further shows that the network structure of inputs and suppliers is highly heterogeneous. While there is a large number of suppliers for some inputs, others are provided by only a few, as indicated by the isolated node pairs on the periphery of the graph.

IA.A.3 Maritime imports vs. other modes of transportation

As discussed in Section 2.2, one limitation of the BOL data is that it only records imports obtained via maritime shipments. In this section, I examine whether the global supply chains constructed from the BOL data are representative of firms' overall global supply chains. To do so, I compare the evolution of aggregate U.S. import value and weight by mode of transportation, including maritime shipments, air shipments, and other modes, using U.S. Census Bureau data.

Panel A of Figure IA.A.4 shows that maritime shipments' import value has consistently outpaced that of other transportation modes from 2003 to 2019. This trend is even more pronounced for import weight. As shown in Panel B, maritime imports far exceed those of any other transportation mode from 2009 to 2019.[14]

Next, I examine how the aggregate value of maritime and air imports differs across HS product codes, focusing on the top ten HS codes ranked by the dollar value of aggregate imports for each transportation mode. Panels A and B of Figure IA.A.5 show the annual percentage of imports attributed to these HS codes from 2009 to 2019 for maritime and air imports, respectively.[15]

The composition of the top ten HS codes for maritime and air imports highlights notable differences in the mode of transportation used across product categories. For example, mineral fuels—which, on average, account for 26% of maritime imports among the top ten HS codes—are absent from the top ten HS codes for air imports. Conversely, pharmaceuticals consistently represent a significant share of air imports each year, yet they are missing from the top ten HS codes for maritime imports.

[14] The time series for import weight by mode of transportation began in 2009. The average weight of air imports during this period is only 0.007 million metric tons.

[15] The decomposition of imports across HS product categories is available starting in 2009.

The fact that pharmaceutical imports are not well represented in maritime shipments is not a significant issue for this analysis, as pharmaceutical manufacturing firms make up only about 7% of the sample. Moreover, specialized components heavily used by other manufacturing firms, such as electronic equipment and machinery, are well represented in both air and maritime imports, collectively accounting for approximately one-third of the imports among the top ten HS codes for each mode. These statistics suggest that analyses based on maritime import data accurately represent the supply chain characteristics relevant to the majority of manufacturing firms in the sample.

Figure IA.A.1: Bill of lading form example

This figure shows an example of a bill of lading (BoL) form.

Form Approved OMB No. 1651-0001

1. Name of Vessel		2. Nationality of Ship	3. IMO No.	4. Voyage No.	Page No. of	U.S. DEPARTMENT OF HOMELAND SECURITY Bureau of Customs and Border Protection **INWARD** **CARGO DECLARATION** 19 CFR 4.7, 4.7a, 4.8, 4.33, 4.34, 4.38, 4.84, 4.85, 4.86, 4.91, 4.93, 4.99	
5. Name of Master	6. Last Foreign Port Before U.S.	7. Port of Discharge	8. Date of Departure from Port of Loading	9. Time of Departure from Port of Loading (Zulu)			
				Answer Col. 14 **OR** Col. 15			
10. Shipper (SH) Consignee (CO) Notify address (NF)	11. Bill of Lading No.	12. Marks & Nos. (MN) Container Nos. (CN) Seal Nos. (SN)	13. No. & Kind of Packages Description of Goods Hazardous Materials (Must Provide UN Code)	14. Gross Wt. (lb. or kg.)	15. Measurement (per HTS)	16. First Port/Place Where Carrier Takes Possession of Cargo	17. Foreign Port Where Cargo is Laden on Board

PAPERWORK REDUCTION ACT NOTICE: This request is in accordance with the Paperwork Reduction Act. We ask for the information in order to carry out the Customs laws of the United States. This form is used by vessel carriers to list all inward cargo on board and for the clearance of all cargo on board with commercial forms. It is mandatory. The estimated average burden associated with this collection of information is 10 minutes per respondent or record keeper depending on individual circumstances. Comments concerning the accuracy of this burden estimate and suggestions for reducing this burden should be directed to Bureau of Customs and Border Protection, Information Services Branch, Washington, DC 20229 and to the Office of Management and Budget, Paperwork Reduction Project (1651-0001), Washington, DC 20503.

CBP Form 1302 (02/03)

Figure IA.A.2: Quarterly import activity

This figure shows the quarterly total number of shipments (Panel A) and import volume in twenty-foot equivalent units (TEUs, Panel B) sourced by sample firms from Q1 2007 to Q4 2019. In both panels, the light portion of the bars represents shipments or import volumes attributed to finished goods, while the dark potions correspond to inputs, which include raw materials and intermediate goods. Imports are classified into inputs or finished goods based on World Bank Harmonized System (HS) code definitions. The sample consists of 907 publicly listed U.S. manufacturing firms with global supply chains.

Panel A: Number of shipments

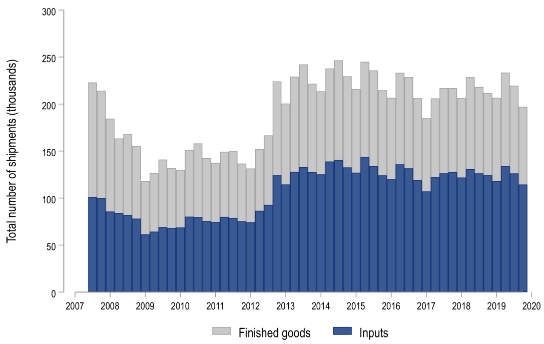

Panel B: Import volume (TEUs)

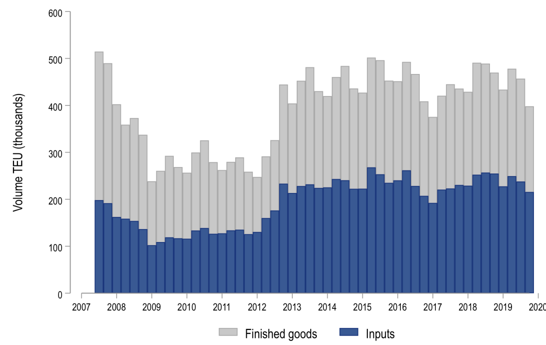

Figure IA.A.3: Country distribution and network structure of global suppliers
Panel A of this figure displays the country distribution of foreign suppliers for the sample firms in 2019. Panel B illustrates the network structure of inputs and their suppliers over the period Q1 2007 to Q4 2019. In Panel B, blue nodes represent unique inputs classified by their Harmonized System (HS) codes, while gray dots denote individual suppliers. The edges between nodes indicate sourcing relationships, illustrating which suppliers provide specific inputs to a node. Larger blue nodes correspond to more generic inputs. The sample consists of 907 publicly listed U.S. manufacturing firms with global supply chains.

Panel A: Country distribution of suppliers

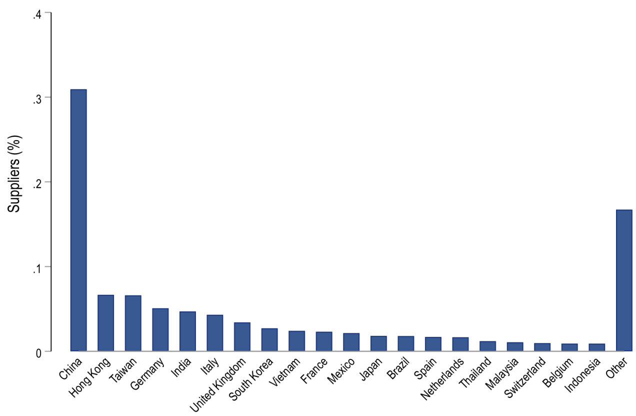

Panel B: Input-supplier network

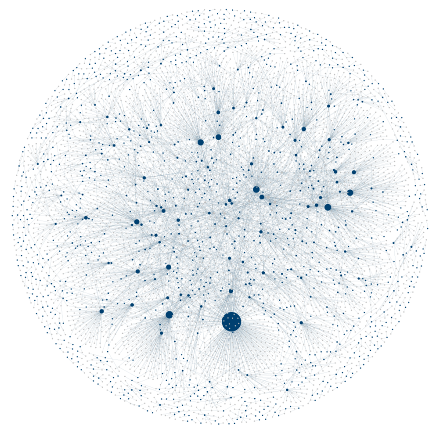

Figure IA.A.4: Evolution of imports by mode of transportation

This figure shows the time series of aggregate U.S. imports by mode of transportation (Water, Air, and Other). The "Other" category includes pipeline, truck, and rail imports. In Panel A, imports are measured in dollar values for the period 2003 to 2019. In Panel B, they are measured in weight for the period 2009 to 2019. These data are from the U.S. Census Bureau.

Panel A: Import value

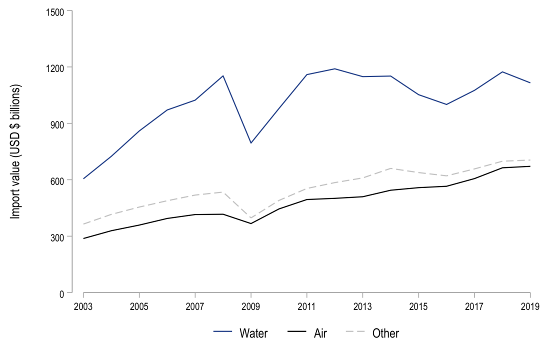

Panel B: Import weight

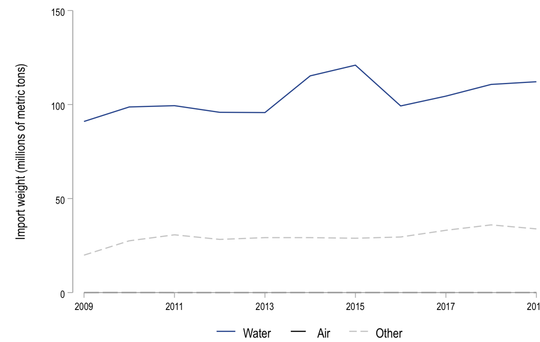

Figure IA.A.5: Product composition of maritime and air imports
This figure shows the composition of maritime (Panel A) and air (Panel B) imports across the ten most common Harmonized System (HS) codes ranked by the dollar value of imports from 2009 to 2019. These data are from the U.S. Census Bureau.

Panel A: Maritime imports

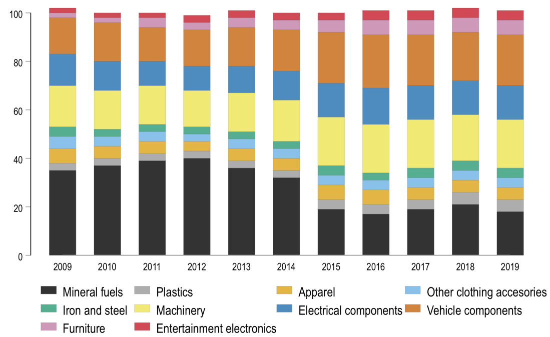

Panel B: Air imports

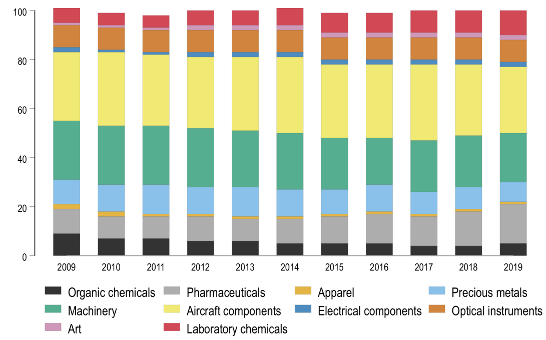

IA.B Additional results and robustness tests

This section presents additional results and robustness tests that are described and referenced in the manuscript.

IA.B.1 Cross-sectional results

As explained in Section 6.1, firm-level supplier scarcity is highly persistent over time, suggesting that firm-level variation in the reliance on fragile inputs is predominantly cross-sectional. Therefore, in this section, I examine the relationship between corporate policies and supply chain disruption risk using purely cross-sectional variation. To do this, I estimate cross-sectional regressions with one observation per firm, where variables are averaged over the sample period. Standard errors in these regressions are robust to heteroskedasticity. The cross-sectional results in Table IA.B.1 are consistent with those in Table 3.

IA.B.2 Debt decomposition

As shown in Section 6, supplier scarcity significantly increases firm book leverage. To explore this further, I decompose the effect of supplier scarcity on leverage by debt type using annual data.[16] Specifically, I examine accounts payable (including trade credit), revolving credit facilities, term loans, commercial paper, and bonds. The findings, reported in Table IA.B.2, suggest that the overall increase in book leverage stems mainly from higher use of accounts payable (trade credit) and revolving credit facilities. Both instruments are tied to short-term operational financing. Trade credit is extended by suppliers as part of input transactions, while revolving credit facilities provide flexible liquidity that can be drawn as inventory needs arise. The absence of significant effects on term loans, commercial paper, or bonds suggests that the increase in leverage reflects the financing of inventory buildup rather than a broader change in capital structure.

IA.B.3 Controlling for the importance of domestic suppliers

A limitation of the BOL data is that it tracks only international shipments. A natural concern is therefore that a measure derived from these data might capture risks specific to foreign sourcing

[16] I use annual rather than quarterly data because debt-type data is better populated at the yearly frequency.

rather than a broader feature of the firm's input requirements. If that were true, supplier scarcity should matter only for firms with predominantly foreign supply chains and should be largely irrelevant for firms with a high proportion of domestic suppliers.

To test this, I use Compustat segment data based on financial disclosures to calculate the proportion of foreign suppliers for each firm. These data list both domestic and foreign suppliers as long as they are material enough to be named in the firm's disclosures. I compute *Foreign suppliers (%)*, the fraction of disclosed suppliers that are foreign, and interact it with the supplier scarcity measure. If foreign sourcing drives the results, this should be reflected in the interaction term. Table IA.B.3 presents the results.

The estimated coefficients on *Foreign suppliers (%)* and the interaction terms are not statistically significant in any specification. Additionally, the estimated coefficients on supplier scarcity are nearly identical to those reported in Table 3. These results suggest that the relationship between supplier scarcity and corporate policies is not driven by risks specific to foreign sourcing.

IA.B.4 Controlling for import intensity

A potential omitted variable in the cross-sectional regressions of Table 3 is the intensity of firms' international sourcing. Firms that import more may hold higher inventory levels due to longer lead times, higher ordering costs, and lumpier shipments, not because of supplier scarcity per se. If the supplier scarcity measure correlates with import intensity, the baseline coefficient may reflect this import-dependence channel rather than a response to fragile input supply. Table IA.B.4 addresses this concern by adding lagged import intensity, measured as total import volume (in TEU) scaled by total assets, as an additional control. In each pair of columns, the odd column reproduces the baseline specification from Table 3 and the even column adds the import intensity control. The coefficient on supplier scarcity is virtually unchanged across all three outcomes, moving from 0.201 to 0.225 for input inventories, from −0.217 to −0.203 for cash, and from 0.164 to 0.159 for leverage. Import intensity itself is statistically insignificant for inventories and leverage, and only marginally significant for cash holdings.

IA.B.5 Reduced form estimates

As an additional check on the IV results in Section 6.1.2, Table IA.B.5 presents reduced form estimates that regress each corporate policy outcome directly on the instrument $Z_{i,t-1}$. Because reduced form estimates do not depend on the first-stage specification, they provide an independent check that the instrument moves corporate policies in the predicted directions. The instrument is positively associated with input inventory holdings and leverage, and negatively associated with cash holdings, confirming that exogenous contractions in the global supplier pool directly predict the corporate policy adjustments identified in Table 3.

IA.B.6 Placebo tests for the exclusion restriction

If the instrument captures technology cycles or input obsolescence rather than disruption risk, it should also predict firm outcomes affected by these channels. Table IA.B.6 regresses R&D intensity and year-over-year sales growth on the instrument with industry and quarter fixed effects. I omit the firm-level controls used in the main specification because several of these variables are closely related to the placebo outcomes. The instrument does not predict either outcome, suggesting that contractions in the global supplier pool do not operate through these alternative channels.

IA.B.7 Supplier network contraction and within-product input prices

The price test in Table 4 examines whether the instrument predicts firm-level input price indices. A more targeted concern is that the effect of a contraction in the supplier pool on input prices may manifest as within-product price variation over time, as firms that experience larger supplier exits pay more for the same inputs. Table IA.B.7 addresses this by constructing the dependent variable as the equal-weighted mean of within-HS-6 demeaned log unit values (USD per kilogram) across a firm's imported products, scaled to percentage points. By demeaning within each HS-6 product, the index removes cross-product price level differences and isolates whether firms facing larger supplier contractions experience increasing prices for the same inputs over time. Column (1) uses all inputs in a firm's import basket and column (2) focuses on fragile inputs. In both specifications, the coefficient on the instrument is statistically insignificant and negative, providing no evidence that supplier network contractions increase within-product input prices. These results reinforce the conclusion from Table 4 that the instrument does not operate through the input cost channel.

IA.B.8 Robustness to instrument weight lags

In the baseline specification, the instrument uses one-quarter lagged import volume shares as weights. In Table IA.B.8, I reconstruct the instrument using one-year, two-year, and three-year lagged import volume weights. The supplier counts $N_{h,t}$ remain contemporaneous, so that the identifying variation is unchanged. The first stage remains strong across all three weight lags. The 2SLS coefficients are similar to those in Table 3.

IA.B.9 Relationship definition robustness

The baseline treatment definition classifies a firm as treated if it imported from the flood-affected supplier at any point in the five years before the flood. Table IA.B.9 assesses the sensitivity of the results to this definition by requiring treated firms to have at least 4 quarters, 8 quarters, or 12 quarters of positive pre-treatment shipments in Panels A, B, and C, respectively. The results are nearly identical across all thresholds, showing that the baseline estimates are not sensitive to the choice of relationship threshold.

IA.B.10 Alternative staggered difference-in-differences estimators

The baseline estimates use a stacked regression approach (Baker, Larcker, and Wang, 2022; Cengiz et al., 2019), which addresses concerns about bias from heterogeneous treatment effects in staggered designs. To verify that the results are not sensitive to the choice of estimator, Figure IA.B.1 replicates the quarterly event study using the imputation estimator of Borusyak, Jaravel, and Spiess (2024), the interaction-weighted estimator of Sun and Abraham (2021), and the estimator of de Chaisemartin and D'Haultfœuille (2020). All four estimators produce nearly identical point estimates and confidence intervals across the four outcome variables.[17]

[17] The Borusyak, Jaravel, and Spiess (2024) and de Chaisemartin and D'Haultfœuille (2020) estimators do not support two-way clustering. For consistency, all three alternative estimators cluster standard errors at the firm level.

IA.B.11 Controlling for ordering lumpiness

Firms with scarce suppliers may place fewer, larger orders due to higher fixed ordering costs or longer lead times. If ordering lumpiness correlates with supplier scarcity, the difference-in-differences estimates could partly reflect pre-existing differences in ordering behavior between treated and control firms rather than genuine policy responses to floods. This concern is related to the import intensity test in Table IA.B.4, but the difference-in-differences setting allows it to be addressed at a more granular level using firm-specific ordering behavior around each flood event. Table IA.B.10 adds two pre-treatment measures of ordering regularity as controls. These are the coefficient of variation of quarterly shipments computed over the pre-flood window, which captures how lumpy or smooth a firm's ordering pattern is, and the log of mean pre-treatment quarterly shipments, which controls for ordering scale. Both measures are fixed at their pre-treatment values and interacted with the post-flood indicator, following the same approach as the other pre-treatment controls.

The treatment effects are virtually unchanged across all four outcomes when controlling for ordering lumpiness. While the coefficient of variation has a significant independent effect on some outcomes, indicating that ordering lumpiness does predict post-flood corporate policies, it does not attenuate or drive the main treatment effect. These results confirm that the documented policy adjustments following supply chain disruptions are not driven by pre-existing differences in ordering patterns between treated and control firms.

Figure IA.B.1: Alternative staggered difference-in-differences estimators

This figure overlays event study coefficients and 90% confidence intervals from four staggered difference-in-differences estimators for each outcome variable. The baseline stacked specification estimates $Y_{itj} = \alpha_{ij} + \alpha_{tj} + \sum_{k\neq-1} \beta_k \mathbb{1}[\text{Event year} = k] \times \text{Treated}_{ij} + \gamma \mathbf{X}_i \times \text{Post}_{tj} + \epsilon_{itj}$, where Event year is a vector of indicator variables for event years -5 through $+5$ and year -1 is the omitted category. The three alternative estimators are the imputation estimator of Borusyak, Jaravel, and Spiess (2024), the interaction-weighted estimator of Sun and Abraham (2021), and the estimator of de Chaisemartin and D'Haultfœuille (2020). All specifications include firm×flood and calendar quarter×flood fixed effects and the same controls as in Table 5. Standard errors are clustered at the firm and calendar quarter levels for the stacked specification and at the firm level for the three alternative estimators.

Panel A: ln(Volume)

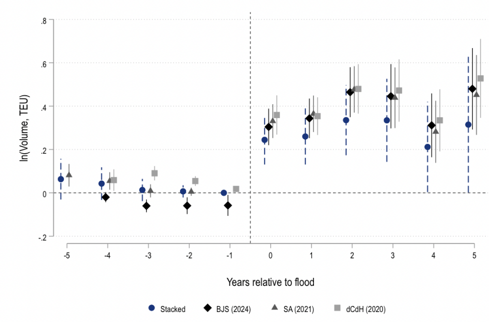

Panel B: Input inv./Sales

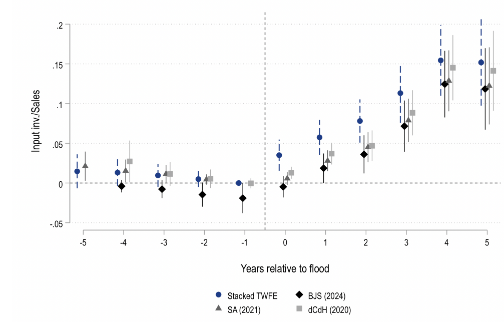

Panel C: Cash/Assets

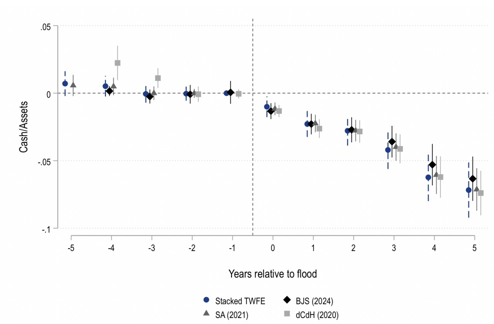

Panel D: Total debt/Assets

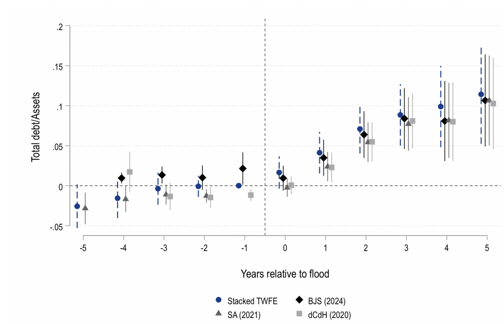

Table IA.B.1: Supplier scarcity and corporate policies, cross-sectional results

This table reports cross-sectional regressions of corporate policies on supplier scarcity. The sample consists of 907 publicly listed U.S. manufacturing firms with global supply chains from Q1 2007 to Q4 2019, with all variables averaged within each firm over the sample period. *Supplier scarcity* is the import-volume-weighted average number of alternative suppliers across the HS codes in firm i's import basket, rescaled so that higher values indicate greater scarcity. *Input inventories/Sales* is input inventories divided by pre-treatment sales. *Cash/Assets* is cash and short-term investments divided by pre-treatment assets. *Total debt/Assets* is the sum of current and long-term liabilities divided by pre-treatment assets. All specifications include industry fixed effects. Robust standard errors are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels are denoted by ***, **, *, respectively. Appendix A presents variable definitions.

	Input inventories/Sales	Cash/Assets	Total debt/Assets
	(1)	(2)	(3)
Supplier scarcity$_{t-1}$	0.332^{**}	-0.410^{***}	0.349^{***}
	(0.135)	(0.077)	(0.090)
Controls	✓	✓	✓
Fixed effects			
Industry	✓	✓	✓
Observations	907	907	907
Adj. R^2	0.17	0.28	0.13

Table IA.B.2: Debt decomposition

This table decomposes the leverage effect from Table 3 by debt type. The sample is a firm-year panel of 907 publicly listed U.S. manufacturing firms with global supply chains from 2007 to 2019. *Supplier scarcity* is the import-volume-weighted average number of alternative suppliers across the HS codes in firm i's import basket during year $t-1$, rescaled so that higher values indicate greater scarcity. Accounts payable is scaled by sales; all other debt instruments are scaled by total assets. All specifications include industry and calendar year fixed effects. Standard errors clustered at the firm and calendar year levels are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels are denoted by ***, **, *, respectively. Appendix A presents variable definitions.

	Accounts payable	Revolving credit	Term loans	Commercial paper	Bonds
	(1)	(2)	(3)	(4)	(5)
Supplier scarcity$_{t-1}$	0.639**	0.019***	0.011	-0.001	-0.062
	(0.270)	(0.005)	(0.011)	(0.001)	(0.064)
Controls	✓	✓	✓	✓	✓
Fixed effects					
Industry	✓	✓	✓	✓	✓
Year	✓	✓	✓	✓	✓
Observations	7,702	7,702	7,702	7,702	7,702
Adj. R^2	0.33	0.06	0.05	0.08	0.01

Table IA.B.3: Controlling for the importance of domestic suppliers

This table tests whether the baseline results from Table 3 are driven by risks specific to foreign sourcing. The sample is a firm-quarter panel of 907 publicly listed U.S. manufacturing firms with global supply chains from Q1 2007 to Q4 2019. *Supplier scarcity* is the import-volume-weighted average number of alternative suppliers across the HS codes in firm i's import basket during quarter $t-1$, rescaled so that higher values indicate greater scarcity. *Foreign suppliers (%)* is the proportion of a firm's suppliers that are foreign, computed from Compustat segment data. *Input inventories/Sales* is input inventories divided by pre-treatment sales. *Cash/Assets* is cash and short-term investments divided by pre-treatment assets. *Total debt/Assets* is the sum of current and long-term liabilities divided by pre-treatment assets. All specifications include industry and calendar quarter fixed effects and the same controls as in Table 3. Standard errors clustered at the firm and calendar quarter levels are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels are denoted by ***, **, *, respectively. Appendix A presents variable definitions.

	Input inventories/Sales	Cash/Assets	Total debt/Assets
	(1)	(2)	(3)
Supplier scarcity$_{t-1}$	0.192**	-0.214^{***}	0.163***
	(0.076)	(0.039)	(0.047)
Foreign suppliers (%)	-0.329	0.066	-0.017
	(0.228)	(0.088)	(0.082)
Supplier scarcity$_{t-1}$ × Foreign suppliers (%)	0.904	-0.090	-0.123
	(0.667)	(0.265)	(0.276)
Controls	✓	✓	✓
Fixed effects			
Industry	✓	✓	✓
Quarter	✓	✓	✓
Observations	28,785	28,785	28,785
Adj. R^2	0.18	0.23	0.21

Table IA.B.4: Controlling for import intensity

This table tests whether the baseline results from Table 3 are robust to controlling for lagged import intensity. The sample is a firm-quarter panel of 907 publicly listed U.S. manufacturing firms with global supply chains from Q1 2007 to Q4 2019. *Import intensity* is total import volume (in TEU) scaled by total assets. For each outcome, the odd columns reproduce the baseline specification from Table 3 and the even columns add the import intensity control. *Supplier scarcity* is the import-volume-weighted average number of alternative suppliers across the HS codes in firm i's import basket during quarter $t-1$, rescaled so that higher values indicate greater scarcity. *Input inventories/Sales* is input inventories divided by pre-treatment sales. *Cash/Assets* is cash and short-term investments divided by pre-treatment assets. *Total debt/Assets* is the sum of current and long-term liabilities divided by pre-treatment assets. All specifications include industry and calendar quarter fixed effects and the same controls as in Table 3. Standard errors clustered at the firm and calendar quarter levels are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels are denoted by ***, **, *, respectively. Appendix A presents variable definitions.

	Input inventories/Sales		Cash/Assets		Total debt/Assets	
	(1)	(2)	(3)	(4)	(5)	(6)
Supplier scarcity$_{t-1}$	0.201**	0.225***	−0.217***	−0.203***	0.164***	0.159***
	(0.076)	(0.079)	(0.039)	(0.039)	(0.046)	(0.046)
Import intensity$_{t-1}$		−0.010		−0.006*		0.002
		(0.007)		(0.003)		(0.005)
Controls	✓	✓	✓	✓	✓	✓
Fixed effects						
Industry	✓	✓	✓	✓	✓	✓
Quarter	✓	✓	✓	✓	✓	✓
Observations	28,785	28,785	28,785	28,785	28,785	28,785
Adj. R^2	0.18	0.18	0.22	0.23	0.21	0.21

Table IA.B.5: Reduced form estimates

This table reports reduced form regressions of corporate financial policies directly on $Z_{i,t-1}$. The sample is a firm-quarter panel of 907 publicly listed U.S. manufacturing firms with global supply chains from Q1 2007 to Q4 2019. *Input inventories/Sales* is input inventories divided by pre-treatment sales. *Cash/Assets* is cash and short-term investments divided by pre-treatment assets. *Total debt/Assets* is the sum of current and long-term liabilities divided by pre-treatment assets. All specifications include industry and calendar quarter fixed effects and the same controls as in Table 3. Standard errors clustered at the firm and calendar quarter levels are reported in parentheses. Statistical significance at the 1, 5, and 10 percent levels are denoted by ***, **, *, respectively. Appendix A presents variable definitions.

	Input inventories/Sales	Cash/Assets	Total debt/Assets
	(1)	(2)	(3)
$Z_{i,t-1}$	0.017^{***}	-0.004^{**}	0.007^{**}
	(0.004)	(0.002)	(0.003)
Controls	✓	✓	✓
Fixed effects			
Industry	✓	✓	✓
Quarter	✓	✓	✓
Observations	28,714	28,714	28,714
Adj. R^2	0.19	0.21	0.21

Table IA.B.6: Placebo tests for the exclusion restriction

This table regresses R&D intensity and sales growth on the instrument $Z_{i,t-1}$. *R&D/Sales* is research and development expenditures divided by sales. *Sales growth* is year-over-year quarterly sales growth. All specifications include industry and calendar quarter fixed effects. No other controls are included. Standard errors clustered at the firm and calendar quarter levels are reported in parentheses. Statistical significance at the 1, 5, and 10 percent levels are denoted by ***, **, *, respectively. Appendix A presents variable definitions.

	R&D/Sales	Sales growth
	(1)	(2)
$Z_{i,t-1}$	−0.012	−0.002
	(0.012)	(0.002)
Fixed effects		
Industry	✓	✓
Quarter	✓	✓
Observations	29,718	26,435
Adj. R^2	0.00	0.10

Table IA.B.7: Supplier network contraction and within-product input prices

This table tests whether the instrument for supplier scarcity, $Z_{i,t-1}$, affects within-product input prices. The dependent variable is the equal-weighted mean of within-HS-6 demeaned log unit values (USD/kg) across a firm's imported products, scaled to percentage points. By demeaning within each HS-6 product, the index isolates whether firms experiencing larger supplier exits pay more for the same inputs over time. Column (1) uses all inputs in a firm's import basket. Column (2) focuses on fragile inputs (HS-6 codes with four or fewer global suppliers). All specifications include the same controls, industry fixed effects, and calendar quarter fixed effects as in Table 3. Standard errors clustered at the firm and calendar quarter levels are reported in parentheses. Statistical significance at the 1, 5, and 10 percent levels are denoted by ***, **, *, respectively. Appendix A presents variable definitions.

	All inputs	Fragile inputs
	(1)	(2)
$Z_{i,t-1}$	−0.033	-0.464
	(0.420)	(0.335)
Controls	✓	✓
Industry FE	✓	✓
Quarter FE	✓	✓
Observations	13,730	13,730
Adj. R^2	0.03	0.01

Table IA.B.8: Robustness to instrument weight lags

This table reports first-stage and 2SLS estimates using instruments constructed with one-year, two-year, and three-year lagged import volume weights. The supplier counts $N_{h,t}$ remain contemporaneous, so the identifying variation is unchanged. Panel A reports the first stage, in which the dependent variable is *Supplier scarcity*. *Input inventories/Sales* is input inventories divided by pre-treatment sales. *Cash/Assets* is cash and short-term investments divided by pre-treatment assets. *Total debt/Assets* is the sum of current and long-term liabilities divided by pre-treatment assets. Panels B, C, and D report 2SLS estimates using one-year, two-year, and three-year lagged weights, respectively. All specifications include the same controls as in Table 3 and industry and calendar quarter fixed effects. Standard errors clustered at the firm and calendar quarter levels are reported in parentheses. The Kleibergen-Paap (KP) F-statistic tests the strength of the first stage. Statistical significance at the 1, 5, and 10 percent levels are denoted by ***, **, *, respectively. Appendix A presents variable definitions.

Panel A: First stage			
	1-year lag	2-year lag	3-year lag
	(1)	(2)	(3)
$Z_{i,t-1}^{1\text{yr}}$	0.018^{***}		
	(0.001)		
$Z_{i,t-1}^{2\text{yr}}$		0.018^{***}	
		(0.001)	
$Z_{i,t-1}^{3\text{yr}}$			0.019^{***}
			(0.001)
Controls	✓	✓	✓
Fixed effects			
Industry	✓	✓	✓
Quarter	✓	✓	✓
Observations	28,714	28,714	28,714
Adj. R^2	0.262	0.259	0.267

Panel B: 2SLS with 1-year lagged weights			
	Input inventories/Sales	Cash/Assets	Total debt/Assets
	(1)	(2)	(3)
Supplier scarcity$_{t-1}$	0.863^{***}	-0.208^{**}	0.362^{**}
	(0.232)	(0.102)	(0.155)
Controls	✓	✓	✓
Fixed effects			
Industry	✓	✓	✓
Quarter	✓	✓	✓
KP F-statistic	259.6	259.6	259.6
Observations	28,714	28,714	28,714

Panel C: 2SLS with 2-year lagged weights			
	Input inventories/Sales	Cash/Assets	Total debt/Assets
	(1)	(2)	(3)
Supplier scarcity$_{t-1}$	0.858***	-0.195^{*}	0.381**
	(0.236)	(0.101)	(0.156)
Controls	✓	✓	✓
Fixed effects			
Industry	✓	✓	✓
Quarter	✓	✓	✓
KP F-statistic	227.1	227.1	227.1
Observations	28,714	28,714	28,714

Panel D: 2SLS with 3-year lagged weights			
	Input inventories/Sales	Cash/Assets	Total debt/Assets
	(1)	(2)	(3)
Supplier scarcity$_{t-1}$	0.800***	-0.208^{**}	0.361**
	(0.226)	(0.098)	(0.152)
Controls	✓	✓	✓
Fixed effects			
Industry	✓	✓	✓
Quarter	✓	✓	✓
KP F-statistic	191.0	191.0	191.0
Observations	28,714	28,714	28,714

Table IA.B.9: Robustness to alternative relationship definitions

This table tests whether the baseline results from Table 5 are sensitive to the definition of the treatment group. The sample consists of 907 publicly listed U.S. manufacturing firms with global supply chains from Q1 2007 to Q4 2019. The data form a firm–flood event–quarter panel. *ln(Volume)* is the log of import volume measured in twenty-foot equivalent units (TEU). *Input inventories/Sales* is input inventories divided by pre-treatment sales. *Cash/Assets* is cash and short-term investments divided by pre-treatment assets. *Total debt/Assets* is the sum of current and long-term liabilities divided by pre-treatment assets. The baseline results define a firm as treated if it imported from the flood-affected supplier of fragile inputs at any time in the five years before the flood. Panel A restricts treated firms to those with at least 4 quarters of positive pre-treatment shipments, Panel B requires at least 8 quarters, and Panel C requires at least 12 quarters. Treated and control firms are balanced by exact matching on two-digit SIC codes and nearest matching on the pre-treatment log of assets and aggregate imports. All specifications include firm×cohort and calendar quarter×cohort fixed effects and the same controls as in Table 5. Standard errors clustered at the firm and calendar quarter levels are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels are denoted by ***, **, *, respectively. Appendix A presents variable definitions.

Panel A: Treated = $\geq$ 4 quarters of pre-treatment shipments

	ln(Volume)	Input inventories/Sales	Cash/Assets	Total debt/Assets
	(1)	(2)	(3)	(4)
Treated × Post	0.256***	0.066***	−0.030***	0.060***
	(0.092)	(0.017)	(0.007)	(0.021)
Controls	✓	✓	✓	✓
Firm × Flood	✓	✓	✓	✓
Calendar quarter × Flood	✓	✓	✓	✓
Observations	196,900	203,868	203,868	203,868
Adj. R^2	0.82	0.75	0.74	0.66

Panel B: Treated = $\geq$ 8 quarters of pre-treatment shipments

	ln(Volume)	Input inventories/Sales	Cash/Assets	Total debt/Assets
	(1)	(2)	(3)	(4)
Treated × Post	0.261***	0.069***	−0.029***	0.063***
	(0.093)	(0.017)	(0.007)	(0.021)
Controls	✓	✓	✓	✓
Firm × Flood	✓	✓	✓	✓
Calendar quarter × Flood	✓	✓	✓	✓
Observations	190,031	196,901	196,901	196,901
Adj. R^2	0.82	0.75	0.74	0.66

Panel C: Treated = $\geq$ 12 quarters of pre-treatment shipments

	ln(Volume)	Input inventories/Sales	Cash/Assets	Total debt/Assets
	(1)	(2)	(3)	(4)
Treated × Post	0.252**	0.076***	−0.028***	0.067***
	(0.095)	(0.017)	(0.007)	(0.022)
Controls	✓	✓	✓	✓
Firm × Flood	✓	✓	✓	✓
Calendar quarter × Flood	✓	✓	✓	✓
Observations	172,627	179,337	179,337	179,337
Adj. R^2	0.81	0.75	0.74	0.65

Table IA.B.10: Controlling for ordering lumpiness

This table tests whether the baseline results from Table 5 are robust to controlling for pre-treatment ordering patterns. The sample consists of 907 publicly listed U.S. manufacturing firms with global supply chains from Q1 2007 to Q4 2019. The data form a firm–flood event–quarter panel. *ln(Volume)* is the log of import volume measured in twenty-foot equivalent units (TEU). *Input inventories/Sales* is input inventories divided by pre-treatment sales. *Cash/Assets* is cash and short-term investments divided by pre-treatment assets. *Total debt/Assets* is the sum of current and long-term liabilities divided by pre-treatment assets. The specification adds two controls to the results in Table 5: the coefficient of variation (CV) of raw quarterly shipments computed over the pre-treatment period, which captures ordering lumpiness, and the log of mean pre-treatment quarterly shipments, which captures ordering scale. Both controls are interacted with $Post_{tj}$. *Treated* equals one for firms that imported from a flood-affected supplier of fragile inputs at any time in the five years before the flood. Treated and control firms are balanced by exact matching on two-digit SIC codes and nearest matching on the pre-treatment log of assets and aggregate imports. All specifications include firm×cohort and calendar quarter×cohort fixed effects and the same controls as in Table 5. Standard errors clustered at the firm and calendar quarter levels are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels are denoted by ***, **, *, respectively. Appendix A presents variable definitions.

	ln(Volume)	Input inventories/Sales	Cash/Assets	Total debt/Assets
	(1)	(2)	(3)	(4)
Treated × Post	0.362***	0.060***	−0.034***	0.063***
	(0.133)	(0.018)	(0.008)	(0.021)
CV of quarterly shipments × Post	−0.756	−0.264**	0.115***	−0.103
	(0.576)	(0.102)	(0.035)	(0.095)
Ln(mean pre-treatment shipments) × Post	−0.097*	−0.003	0.006**	−0.004
	(0.054)	(0.007)	(0.003)	(0.010)
Controls	✓	✓	✓	✓
Firm × Flood	✓	✓	✓	✓
Calendar quarter × Flood	✓	✓	✓	✓
Observations	196,112	202,266	202,266	202,266
Adj. R^2	0.82	0.75	0.74	0.66