

Input Sourcing under Climate Risk: Evidence from U.S. Manufacturing Firms*

Joaquin Blaum
Boston University

Federico Esposito
Fordham University

Sebastian Heise
NY Fed

January 2026

Abstract

We study the effect of risk on how firms organize their supply chains. We use transaction-level data on U.S. manufacturing imports to construct a novel measure of input sourcing risk based on the historical volatility of ocean shipping times. Our measure isolates the unexpected component of shipping times that is induced by weather conditions along more than 331,000 maritime routes. We first document that unexpected shipping delays significantly reduce importers' sales, profits, and employment. We then show that firms actively diversify weather risk by using more routes and foreign suppliers, although their import values decline. To rationalize these findings, we introduce shipping time risk into a general equilibrium model of importing with firm heterogeneity. Our quantitative analysis predicts substantial costs for the U.S. economy associated with supply chain risk.

*Blaum: jblaum@bu.edu. Esposito: fesposito4@fordham.edu. Heise: sebastian.heise@ny.frb.org. We thank our discussants Davin Chor, Simona Giglioli and Shafaat Yar Khan, as well as Laura Alfaro, David Atkin, Teresa Fort, Matthew Gudgeon, Fadi Hassan, Georg Schaur, Peter Schott and participants at numerous seminars and conferences for their valuable comments. The views and opinions expressed in this work do not necessarily represent the views of the Federal Reserve Bank of New York or the Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 1883 (CBDRB-FY25-P1883-R12596 and CBDRB-FY26-P1883-R12776).


1 Introduction

The past decades have seen a dramatic transformation in the international organization of production, with intermediate inputs accounting for over two-thirds of global trade and complex global value chains spanning multiple countries (Johnson and Noguera (2012), Antràs and Chor (2022)). For many firms, the timely delivery of inputs is a crucial element of their production process (Hummels and Schaur (2013)). However, the increased reliance on imported intermediates has exposed firms to a host of supply chain risks that can adversely impact the timeliness of their inputs. Salient recent examples include the increased frequency of extreme weather events associated with climate change or the strain on port infrastructure that followed the Covid pandemic (e.g., Brancaccio et al. (2024)). How do these and other supply chain risks impact firms' import behavior? Do firms adapt their supply chains to hedge against the delay risk stemming from these shocks? Answering these questions is challenging because of the inherent difficulty in developing credible measures of firm-level risk.

We shed light on these questions by focusing on a specific but important source of risk: weather. We start by establishing that weather conditions have a significant effect on the ocean shipping times of U.S. manufacturing imports. To do so, we rely on transaction-level import data on ocean shipments provided by the U.S. Census Bureau as well as detailed data on oceanic wave conditions along more than 331,000 maritime routes. We exploit this relationship to measure the component of shipping times that is induced by weather, which we interpret as unexpected by U.S. importers given the unpredictability of high frequency ocean conditions.

Armed with this measure, we establish two key empirical results. We first show that unexpected shipping delays induced by weather shocks have large and disruptive effects on U.S. importers' production levels and profit margins. We then build a measure of risk based on the volatility of the weather-induced shipping times. As Figure 1 shows, the standard deviation of wave height has increased in many locations over the past decade. Our second finding is that firms systematically respond to this type of weather risk along different margins of adjustment. More exposed firms choose to rely on more routes and foreign suppliers, and lower both their imports and the concentration of expenditure across routes and suppliers. We next incorporate risky shipping times into a general equilibrium model of firm-level importing, and calibrate the model to match salient features of the data. We use our framework to quantify the impact of two scenarios of heightened risk: climate change and port congestion. Overall, we find that these shocks trigger an important risk diversification response by importers, but nevertheless reduce U.S. real income.

Figure 1: Change in Standard Deviation of Wave Height, 2011-2023



Source: WaveWatch III Global Wave Model, University of Hawaii. Notes: We compute the standard deviation of average daily wave height across all days of each year at each coordinate in the oceans and then average across years in 2011-2013 and in 2021-2023. The figure shows the percentage change at each grid point between these two periods.

The cornerstone of our analysis is the U.S. Census Bureau’s Longitudinal Firm Trade Transactions Database (LFTTD), which provides transaction-level data recording the identity of the U.S. importer and its foreign supplier, as well as information about the product, quantity, and value transacted for the universe of U.S. imports. Importantly, the data record the delivery time between the foreign port of exit and the U.S. port of entry and, for ocean shipments, the vessel identity. Since the customs data do not contain details on each vessel’s journey across the ocean, we develop an algorithm that uses the vessel name, foreign port stops, and U.S. port of entry to determine the intermediate stops each vessel made on its way to the U.S. We then construct the shipment route by finding the shortest maritime route for each trip segment of the vessel’s journey using data from Eurostat’s SeaRoute program.¹ Finally, we combine this data with information on the oceanic wave conditions along each shipment’s route, which we measure using hourly data on wave height and wave direction at the 0.5 degree level from the National Oceanic and Atmospheric Administration (NOAA) for the years 2011-2019.

A central feature of our analysis is the measurement of the shipping times that are unexpected by importers. To this end, we extract the components of shipping times that are induced by weather conditions—specifically, the realized wave height and direction observed

¹Ganapati et al. (2024) show that vessels on average follow the optimal maritime routes very closely. Using AIS tracking data, we confirm that the major routes we construct are close to the actual routes followed by vessels.

along the route of each shipment. We also use a rich set of fixed effects and controls to remove other predictable components that might influence how weather conditions affect shipping time, such as the route, the vessel, the month, and the shipping charges. To interpret the variation in the weather-induced shipping times as unexpected, our identifying assumption is that the *realized* wave conditions along the entire maritime route are not anticipated by the importers when they place their orders, beyond seasonal patterns that are picked up by route-month fixed effects. We view this assumption as plausibly satisfied in the data. On the one hand, most maritime shipments to the U.S. involve multi-week ocean crossings, and import orders are typically placed many weeks before production is finalized and goods are shipped (see [Deloitte \(2024\)](#)). On the other hand, weather forecasts are reasonably accurate up to 7 days into the future, and predict only general patterns beyond 2 weeks—with ocean wave height being particularly hard to predict given the chaotic nature of ocean dynamics ([Alley et al. \(2019\)](#), [Zhang et al. \(2022\)](#), and [Mishra et al. \(2022\)](#)).

We use our measure to first analyze the effects of shipping delays induced by weather shocks on firm performance. We identify for each year the shipments that were extremely delayed, which we define as having a weather-induced delivery time greater than the 95th percentile of its distribution for a given route. We estimate panel regressions for the years 2011-2019 and document that U.S. importers with a higher share of delayed inputs in total costs experienced significant declines in sales, profits and employment. A one standard deviation increase in the share of input costs that are weather-delayed reduces firms' sales by 5.3%, profits by 3.5% and employment by 0.9% in the same year. These large negative effects highlight the substantial impact of supply chain disruptions on firms' production and suggest that firms are typically unable to fully hedge their supply chain risk through insurance or financial instruments.

We next study whether U.S. importers adjust their sourcing strategy and import demand ex-ante to reduce the potential impact of weather shocks. To this end, we build a measure of risk based on the volatility of weather-induced shipping times. In particular, we measure the riskiness of each foreign supplier-route-product combination as the standard deviation of the weather-induced shipping times over 3-year rolling windows. We construct a shift-share measure of exposure to risk for each importer as a weighted average of the risk of its suppliers and routes over the previous 3 years, using pre-determined import shares as weights. We then estimate panel regressions of firms' sourcing behavior on risk exposure at the importer-product-year level and include a rich set of fixed effects and controls. Our results indicate that U.S. importers diversify weather-induced risk along the extensive and intensive margins. Going from the 25th to the 75th percentile of the shipping risk distribution increases the number of routes and the number of foreign suppliers used by 9.0% and 5.9%,

respectively. Moreover, it reduces the total value imported by 4.3%. Thus, importers with ex-ante riskier supply chains spread their input expenditures among more routes and foreign suppliers, and import less overall. Taken together, these results highlight the detrimental impact of supply chain uncertainty on international trade.

To rationalize these findings, we incorporate shipping risk into a standard model of importing with firm heterogeneity, along the lines of [Blaum et al. \(2018\)](#), [Gopinath and Neiman \(2014\)](#), and [Halpern et al. \(2015\)](#). Firms can source their inputs domestically or from foreign suppliers. We follow [Hummels and Schaur \(2013\)](#) in their treatment of timeliness by assuming that input qualities are reduced when inputs take longer to arrive, for example due to spoilage, absence of key inputs, etc. The key departure from the literature is that firms are uncertain about shipping times when placing input orders. While firms are risk-neutral, the presence of market power with elastic demand and the imperfect substitutability between inputs introduce curvature in revenues, making expected revenues fall with more volatile input qualities. Firms can diversify their shipping time risk by sourcing from multiple foreign suppliers, or equivalently by using multiple routes, although this strategy is limited by per-supplier fixed costs. We provide conditions under which firms increase their number of foreign suppliers and reduce their import values after a mean-preserving spread to supplier qualities.

We consider a calibrated version of the model to assess whether the theory can come to terms with the empirical evidence. Firms are heterogeneous both in their productivity and in the shipping time risk they face. The calibration targets our estimate of the effect of shipping time risk on the extensive margin of importing to capture the role of risk, and requires the model to match the negative association between sales and average shipping times observed in the data to discipline the role of supplier timeliness. To quantify the aggregate effects of risk, we also target the joint distribution of firm sales and risk observed in the data and, in particular, the fact that larger importers are typically matched with safer foreign suppliers. The calibrated model replicates well the key moments of shipping time risk and import demand. We can therefore use the model as a laboratory to evaluate the impacts of any scenario involving changes in shipping time risk.

We assess the impact of two risk-related scenarios that have recently received significant attention, namely, climate change and port congestion. The volatility of ocean wave height has increased on average by 0.34% per year between 2011-2023, consistent with work in oceanography suggesting an increasing likelihood of extreme wave heights ([Young et al. \(2011\)](#)). We evaluate the effects of an increase in the volatility of ocean wave heights that continues along this trend over the next 50 years in our model. In a second exercise, we

consider the greater variability of waiting times at ports associated with the rise in port congestion that took place in the post Covid period of 2021-2022. In both exercises, we find a strong risk diversification response along the extensive margin together with a substantial fall in imports, as firms reduce their risk exposure by substituting towards domestic inputs. This shift increases production costs and prices, reducing U.S. real income by 0.5-1.1%.

Related Literature. Our paper contributes to several strands of the literature. First, it relates to work that investigates the importance of timeliness in international trade, both in theory and in the data ([Evans and Harrigan \(2005\)](#) and [Hummels and Schaur \(2013\)](#)). While these seminal papers focus on the role of the level of shipping times, we study the effect of the variance of shipping times, controlling for their mean. Our empirical results show that uncertainty around shipping times has an additional negative effect on import demand. We propose a theory of the firm that incorporates this mechanism in a way that is both tractable and amenable to quantitative analysis.

Second, we contribute to a broader literature that analyzes the impact of uncertainty on firms. Most of the international trade literature on this topic has focused on exports and FDI (e.g., [Ramondo et al. \(2013\)](#), [Fillat and Garetto \(2015\)](#), [Esposito \(2022\)](#), [Baley et al. \(2020\)](#), and [De Sousa et al. \(2020\)](#)). In contrast, we analyze risk on the input side and how it affects firms' sourcing decisions. Only a few papers have studied the effects of sourcing uncertainty on international trade (e.g., [Gervais \(2018\)](#), [Grossman et al. \(2023\)](#), and [Handley et al. \(2024\)](#)). Our contribution to this literature is to develop a novel and plausibly exogenous measure of firm-level shipping time risk using weather shocks, which we use to study the causal impact of risk on importers' performance in the United States. We combine weather data with comprehensive firm-level administrative data and show that importers actively adjust the intensive and extensive margins of importing in response to weather risk.² Complementary to our work are [Balboni et al. \(2023\)](#) and [Castro-Vincenzi et al. \(2024\)](#), who study how firms diversify their sourcing locations in Pakistan and India, respectively. In contrast to our focus on maritime shipping risk and international trade, these works focus on flood risk and on domestic trade.

Third, we contribute to work that studies the effects of supply chain disruptions on firms ([Carvalho et al. \(2021\)](#), [Barrot and Sauvagnat \(2016\)](#), [Boehm et al. \(2019\)](#), [Khanna et al. \(2022\)](#), [Alessandria et al. \(2023\)](#), [Lafrogne-Joussier et al. \(2023\)](#)).³ Relative to this literature,

²The diversification mechanism we highlight is complementary to firms' use of inventories, documented by [Alessandria et al. \(2023\)](#) and [Carreras-Valle \(2024\)](#).

³Also related is work that studies the effects of climate shocks on firms, e.g., [Pankratz and Schiller \(2024\)](#) and [Dunbar et al. \(2023\)](#). More broadly, we contribute to a large literature that studies firm-to-firm

we provide a new way to identify granular supply shocks using readily available weather data, rather than large, aggregate shocks—such as the Japanese earthquake or the Covid lockdowns. Our measure therefore lends itself to a wide range of applications that require exogenous shocks to firms.

Finally, we contribute to the literature that incorporates input trade into quantitative models with firm heterogeneity—e.g., [Gopinath and Neiman \(2014\)](#), [Halpern et al. \(2015\)](#), [Antràs et al. \(2017\)](#), and [Blaum et al. \(2018\)](#). Existing models of importing typically abstract from supplier risk considerations. Our contribution is to extend a sourcing model to allow for risk and to quantify its impact in general equilibrium. Our calibrated model can serve as a laboratory to estimate the impact of any type of sourcing risk on U.S. importers.

The remainder of the paper proceeds as follows. Section 2 describes our data, measurement of shipping times, and the construction of the weather shocks, while Section 3 discusses our empirical results. Section 4 presents the model, which we calibrate to perform our quantitative analysis in Section 5. Section 6 concludes.

2 Measuring Shipping Times, Routes and Weather Conditions

In this section, we describe how we measure shipping times, routes, and weather conditions for U.S.-bound import transactions and how we construct our measure of weather-induced shipping times. Using detailed U.S. Census transaction-level import data, we first develop a novel methodology to infer the shipping routes followed by vessel-borne shipments headed to the U.S. between 2011 and 2019, yielding more than 331,000 distinct routes. We then use highly granular data on ocean wave conditions to measure the realized weather conditions along each vessel’s journey. These will form the basis of our measure of shipping time risk, which we use in the empirical analysis in Section 3.

2.1 U.S. Census Data

Our empirical analysis relies on the Longitudinal Firm Trade Transactions Database (LFTTD) provided by the U.S. Census Bureau. This dataset comprises the entire universe of international trade transactions made by U.S. firms. We focus on all the import transactions during the period 2011-2019. Each transaction is associated with an identifier of the U.S. importer, the HS-10 product code traded, the mode of transportation (vessel, air, etc.), as well as the

relationships, see e.g., [Dhyne et al. \(2021\)](#), [Bernard et al. \(2019\)](#), [Esposito and Hassan \(2023\)](#), and [Heise \(2024\)](#). This literature does not typically focus on uncertainty.

value, weight, and quantity shipped. The data also report an identifier of the foreign seller and an indicator of whether the transaction is between related parties.⁴ We calculate prices as the value of the shipment divided by the quantity shipped and keep both related party and arm’s-length transactions.

The LFTTD contains several additional variables that are critical to construct our measure of risk. First, each customs record reports the export and the import dates (after customs are cleared), which allows us to construct shipping times. Second, for seaborne imports, we also observe the foreign port of departure, the U.S. port of arrival, and the vessel name. We use this information to construct shipping routes, as explained below. For non-seaborne imports we only know the shipping company’s name instead of the vessel name, and we only know the country of departure instead of the departure port. Since in some cases an import transaction spans multiple customs records, we collapse the data to the supplier (x) - product (h) - foreign country (c) - origin port (p_e) - destination port (p_i) - foreign export date (t_e) - import date (t_i) - vessel (v) - importer (f) - related party status (a) level.⁵ We call such an observation a *transaction*. We describe in detail the data cleaning process in Appendix A.1.

We merge the LFTTD data with the Longitudinal Business Database (LBD), which reports the annual employment and the main industry of each U.S. firm. Given our focus on supply chains, we restrict our analysis to firms that operate in the manufacturing sector, as their imports are most likely intermediates into production. We also obtain firms’ total sales, cost of materials, and employee compensation from the Census of Manufactures (CMF) in census years (2012, 2017) and from the Annual Survey of Manufacturers (ASM) for non-census years. We construct profits as sales minus cost of materials and payroll.

Table 1 reports summary statistics of our dataset. The first column considers all manufacturing imports over the period 2011-2019. The second column includes only seaborne transactions, which we use to construct our measure of shipping risk. Our dataset covers about 5.5 trillion dollars of imports (in 2009 dollars), of which about 40 percent are seaborne. For these vessel-based transactions, we observe 250 U.S. ports, 1,600 foreign ports, and 145,000 unique vessels—crucial pieces of information for constructing shipping routes, which we turn to next.

⁴The foreign seller is identified by a Manufacturer ID (MID), which is an alphanumeric code that combines information on the seller’s country, name, street address, and city. We follow [Kamal et al. \(2015\)](#) and [Kamal and Monarch \(2018\)](#) in combining sellers with the same country and name into one. We use the concordance by [Pierce and Schott \(2012\)](#) to transform the HS-10 codes into time-consistent product codes. Note that we do not observe domestic suppliers, only foreign ones. See Appendix A.1 for more details.

⁵For non-vessel transactions, the origin port is replaced with the origin country and the vessel is blank.

Table 1: U.S. Import Transaction Summary Statistics

	All	Seaborne
Total Imports	5,470	2,090
Unique Importers (f)	114,000	55,000
Unique Exporters (x)	485,000	225,000
Number of Transactions (millions)	58.6	18.4
Number of U.S. Ports of Entry (p_i)		250
Number of Foreign Ports (p_e)		1,600
Number of Origin-Destination Port Pairs		28,000
Unique Vessels (v)		145,000

Source: LFTTD and authors' calculations. The table summarizes U.S. imports from 2011 to 2019. Values in the first row are reported in billions of 2009 dollars.

2.2 Construction of Shipping Times and Routes

Shipping Times For all shipments, irrespective of their mode of transportation, we calculate the shipping time as the difference, in days, between the import date in the U.S. and the export date from the foreign country. We report statistics of the distribution of shipping times in Table 2. Vessel shipments take on average 19 days to arrive to the U.S., which is substantially more than for all other modes of transportation. Air and truck shipments arrive in the U.S. on average within the same day, while train shipments arrive on average in 5 days. Importantly, there is a high degree of dispersion in vessel shipping times.⁶

Routes and Journeys For seaborne shipments, we develop an algorithm to construct ocean shipping routes and vessels' journeys between ports from the information on the port of entry and port of origin. We assign each transaction in the customs data to a *trip*, defined as a journey of a vessel that begins with the loading of cargo at a foreign port and ends, possibly after some intermediate stops, with the unloading of cargo at a U.S. port. As a starting point, we sort all transactions involving a given vessel by their foreign departure date. We then take all the vessel's transactions and assign them to a single trip ("Trip 1"). Next, we find the earliest arrival date of the vessel in the U.S. for this trip. If there exists any transaction of the same vessel with an export departure date abroad that is later than this earliest arrival date in the U.S., we assign these transactions to a new trip ("Trip 2"). We continue splitting trips into sub-trips until no further splits are possible.⁷ We then use the

⁶Of course, predictable factors such as the origin country or the time of year affect vessel shipping times to the U.S. We show additional statistics on shipping times and their determinants in Appendix A.2.

⁷In some instances the departure or arrival date may be misreported, for example because a vessel reports a departure abroad at the same time as unloading cargo at a U.S. port. In Appendix A.1 we explain how we

Table 2: Shipping Times by Mode of Transportation

	Avg. Time	Std. Time	P5	P25	P50	P75	P95	Total Value
Vessel	19.1	19.9	6.1	12.0	15.4	25.0	36.9	2,090
Train	5.4	7.3	0	0	0	10.5	20.9	745
Truck	0.1	0.5	0	0	0	0	0	1,140
Airplane	0.6	1.0	0	0	0	1.0	2.3	656

Source: LFTTD. The table summarizes the distribution of shipping time and value across different regions and modes of transportation. Values are reported in billions of 2009 dollars.

dates of import and export to construct the sequence of ports visited by each shipment, e.g. Le Havre - Southampton - New York - Newport News. We refer to a leg of the trip between two ports as *route segment*.

We determine the path taken by vessels across the ocean on any route segment using Eurostat’s SeaRoute program. This program computes the shortest maritime paths using the network of global shipping lanes and observed vessel movements from satellite data. We obtain a single path for each route segment, which we assign to any vessel going through it. We refer to the entire ordered set of route segments as a *granular route*, and call the associated origin-destination port pair a *route*. Our sample includes around 14,000 route segments, 20,500 routes, and 331,000 granular routes. We show that, for a selected sample, these routes closely follow actual vessel movements reported by AIS data (Appendix A.3).⁸ The upper panel of Figure 2 illustrates the route segments in our data.

For shipments arriving by modes of transportation other than vessel, we observe only the country of departure rather than the departure port. We therefore approximate the shipping route as a foreign country of origin and a U.S. entry point (e.g., airport or border crossing) pair, and end up with 13,500 distinct non-seaborne routes. Since we cannot compute the weather conditions for these transactions, we will assume that their shipping risk is zero for the empirical analysis in Section 3.


Our measures of shipping times and routes are the building blocks of our empirical analysis. The key advantage of relying on the U.S. Census transaction-level data to obtain these measures is its comprehensive nature and extreme detail, which allows for a systematic analysis of the role of shipping risk in the U.S. economy. However, there are a number of limitations due to the nature of the available data. First, we do not have information

identify such cases and how we refine our algorithm to redefine the trips.

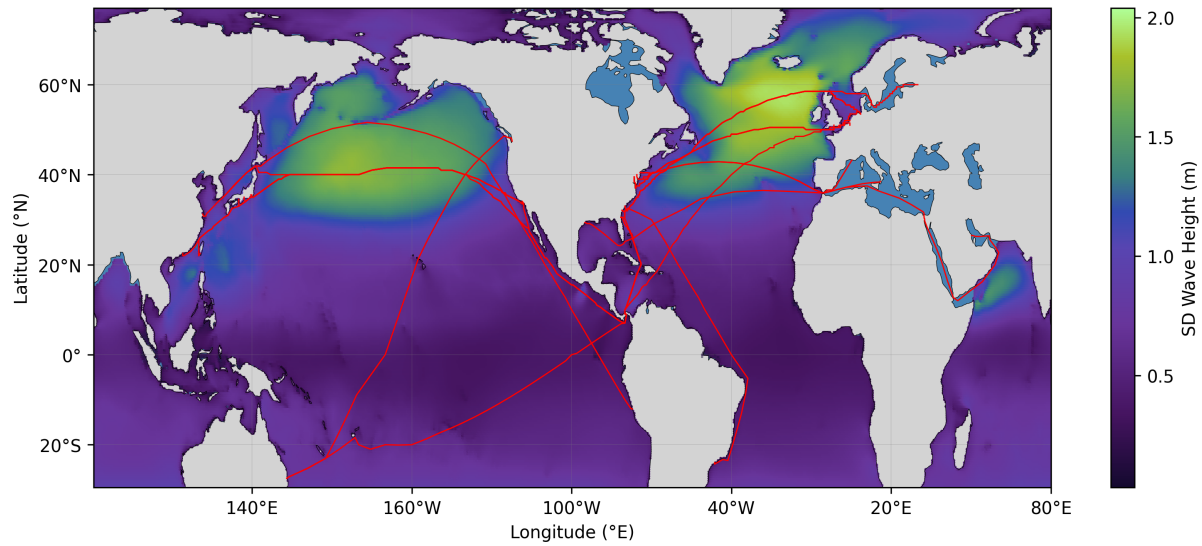
⁸This evidence is consistent with [Ganapati et al. \(2024\)](#), who show that vessels typically follow the minimum-distance routes fairly closely. In addition, two-thirds of world trade in manufacturing travels on container ships, which typically follow fixed itineraries (i.e. the so-called “bus system”, see [Branaccio et al. \(2020\)](#) and [Heiland et al. \(2025\)](#)), which are likely captured by Eurostat’s SeaRoute program.

Figure 2: Weather Conditions and Routes

(a) Network of Route Segments



(b) Volatility of Wave Height



Notes: The upper panel shows the network of shipping routes constructed in our data. The bottom panel shows the standard deviation of wave height across all days from 2011-2019 and some selected shipping routes.

on the voyage from the manufacturer’s production facilities to the foreign port, nor on the journey from the U.S. port of entry to the importer’s plant. Hence we do not capture the risks associated with those parts of the trip. However, typically goods spend several weeks on the vessel to the U.S., and thus this part of the journey is likely a large fraction of the total travel time. Second, because the data report only the foreign ports where goods are loaded onto a vessel bound for the U.S., we do not know whether goods are reloaded from one vessel to another, i.e., “trans-shipped” (Ganapati et al., 2024). Thus, we only observe the journey of the last vessel to the U.S., implying that we underestimate the overall shipping delay risk faced by importers.

2.3 Construction of Weather Conditions

To obtain exogenous variation in shipping times we rely on information on oceanic weather conditions, which we obtain from the WaveWatch III model maintained by the University of Hawaii based on NOAA data. These data report the height and direction (in degrees) of significant waves at hourly or three-hourly frequency for geo-locations at 0.5 degree distances in the oceans from 2011 onward.⁹ There is an extensive literature showing that oceanic wind conditions and waves affect navigation speed (e.g., Filtz et al. 2015 and Viellechner and Spinler, 2020) and increase accident risk (Heij and Knapp, 2015).¹⁰ To reduce computational requirements, we aggregate the hourly information to compute the daily average of significant wave height and direction for each geo-location in the oceans. If information is unavailable at a geo-location, we take a simple average of the weather in the surrounding grid points.¹¹

We combine the weather information with the route segments for seaborne shipments constructed earlier for each coordinate and day. Since the effect of waves on vessel speed depends on the direction of travel, we compute for each route coordinate a “relative wave direction”. This variable is computed by taking the absolute difference between the direction of the waves and the estimated direction of vessels at that point. We estimate the direction of any vessel going through a route segment using the latitude and longitude of subsequent coordinates of the segment’s path. A greater relative direction means that the waves are less aligned with vessels’ likely course.¹² For each route segment, we compute the average

⁹Significant waves are the waves that a trained observer would see when looking at the ocean. Significant wave height is the average height of the highest third of the waves.

¹⁰Ocean currents are also important determinants of navigation speed, but they can be perfectly predicted, and therefore are absorbed by the route-month fixed effects we use in our methodology.

¹¹Note that while our dataset reports weather conditions for locations in the oceans, it does not include information for interior bodies of water, such as the Great Lakes, the Mediterranean sea, and the Baltic Sea.

¹²For example, a wave direction of 75 degrees for a vessel traveling at direction 90 degrees corresponds to a wave direction of $\text{abs}(90 - 75) = 15$ degrees. When this absolute difference exceeds 180, we subtract it

weather (i.e., wave height and relative direction) for each day by averaging across all segment coordinates.

In the final step, we merge the route segments and weather information with the trade transaction data. For each day a vessel spends on a segment, we merge in the corresponding segment-level average weather. We then take an average across the day-level weather measures for each transaction. Our final dataset thus contains, for each transaction, an average wave height and an average relative wave direction along the entire granular route.

To illustrate the source of exogenous variation, the blue and green shading in the bottom panel of Figure 2 report the standard deviation of the average daily wave height for each grid point in our data. The red lines indicate various shipping routes used by U.S. importers. There are significant differences across locations. Routes across the Atlantic and Pacific have higher wave height volatility than routes along the coast of South America. Importantly, there is variation even across routes that are relatively close to each other. Vessels traversing the Northern Atlantic Ocean on their way to the East Coast face a significantly higher standard deviation of wave height than vessels that travel further South. We provide some summary statistics on wave height in Appendix A.4.

2.4 Measuring Unexpected Shipping Times

A central goal of our methodology is to measure shipping delays, that is, instances where goods arrive later than expected. However, we do not observe the shipping times expected by the importers, only the realized ones. We therefore focus on the variation in shipping time that is induced by plausibly unanticipated weather conditions. To this end, we regress the shipping time on the weather conditions realized along a shipment’s maritime route, as well as a rich set of fixed effects and observables that capture factors which are likely anticipated by importers. We treat the predicted effects from the weather terms in this regression as the weather-induced unexpected shipping times.

Our measurement of unexpected shipping times relies on the assumption that, at the time of placing orders, importers do not fully anticipate the weather conditions along the entire maritime route beyond the usual seasonal patterns, which are picked up by the route-month fixed effects in the regression. We believe that this is a reasonable assumption as import orders are placed typically many weeks before production finalizes and goods are shipped (see [Deloitte, 2024](#)). Moreover, most shipments to the U.S. require multi-week ocean crossings

from 360 to get the minimum distance. For instance, if a vessel travels North and the waves go West, the relative direction would be $360 - \text{abs}(0 - 270) = 90$ degrees.

where weather conditions cannot be perfectly predicted, even by shipping companies relying on sophisticated weather forecasting technology.¹³ Importers may also be uncertain about the exact shipping date of their orders.

We next describe a simple statistical model that relates shipping time to realized weather conditions along the route and other observables. We then estimate the model to extract the weather component of shipping times.

A Statistical Model of Shipping Times and Weather Conditions. Consider a buyer f that orders a shipment s of product h from seller x in time period t . The shipment arrives to the U.S. via route r . The time it takes for the shipment to arrive in the U.S., T_s , is a stochastic variable determined by:

$$\ln(T_s) = \alpha_f + \alpha_h + \alpha_x + \alpha_{rt} + \alpha_\nu + \alpha_a + \eta \ln(C_s) + \rho \ln(W_s) + t_{weather,s} + \varepsilon_s, \quad (1)$$

where ν denotes the vessel used, a is an indicator for whether the importer and exporter are related parties, W_s and C_s are the shipment’s weight and shipping charges (freight costs plus insurance), respectively, and $t_{weather,s}$ is the component of the shipping time explained by weather conditions affecting shipment s .

The α terms capture deterministic components of shipping times that might be known to the buyer at the time of ordering. For instance, α_h may capture that some products are harder to ship or take longer to clear at customs, and α_x may reflect the ability of a supplier to arrange logistics with shipping companies. α_{rt} captures route characteristics in a given month-year t (e.g., April 2015), such as the average time it takes to unload a shipment and clear customs in a port. Weather conditions that are anticipated at the time of placing import orders are also captured by α_{rt} .

Shipping times are also determined by stochastic weather conditions which are realized after import orders are placed. Following the literature on ocean shipping, we represent weather conditions with the height and relative direction of the significant waves a vessel encounters on its route (Filtz et al., 2015):

$$t_{weather,s} = \beta_1 \text{Height}_s + \beta_2 \text{Direction}_s + \beta_3 \text{Height}_s \cdot \text{Direction}_s, \quad (2)$$

¹³Forecasts are reasonably accurate only until around 7 days into the future, and only general weather trends can be predicted beyond 2 weeks (see Alley et al., 2019 and Ritchie, 2024). Zhang et al. (2022) and Mishra et al. (2022) argue that despite advancements in machine learning and predictive modeling, accurately forecasting ocean wave height remains a difficult problem due to the chaotic and non-linear nature of ocean dynamics.

where Height_s is the average wave height along the shipment’s route in meters and Direction_s is the average wave direction in degrees relative to the vessel’s direction of travel.

Lastly, the term ε_s includes any random shocks that affect T_s beyond the deterministic components and the weather conditions.

Extracting the Weather Component of Shipping Times. We extract the component of shipping times explained by stochastic weather conditions by regressing the shipping times on the set of fixed effects and observables specified in expressions (1)-(2):

$$\begin{aligned} \ln(T_s) = & \alpha_f + \alpha_h + \alpha_x + \alpha_{rt} + \alpha_\nu + \alpha_a + \eta \ln(C_s) + \rho \ln(W_s), \\ & + \beta_1 \text{Height}_s + \beta_2 \text{Direction}_s + \beta_3 \text{Height}_s \cdot \text{Direction}_s + \varepsilon_s. \end{aligned} \quad (3)$$

The values of Height_s and Direction_s are unique to each vessel’s journey, reflecting the specific weather conditions encountered by the vessel during the days it spent on its shipping route. Since the regression includes route-month fixed effects, the β coefficients are estimated from deviations from the average weather on the route in the given month, thus capturing weather shocks. We measure the component of shipping times explained by unanticipated weather conditions by

$$\hat{t}_{weather,s} = \hat{\beta}_1 \text{Height}_s + \hat{\beta}_2 \text{Direction}_s + \hat{\beta}_3 \text{Height}_s \cdot \text{Direction}_s. \quad (4)$$

Table 3 presents the regression coefficient estimates on the weather terms from specification (3) (the other regressors are omitted for brevity). The weather condition variables are gradually introduced, with the first column including only wave height, the second column adding the relative wave direction, and the third column adding the interaction term. Across all specifications, higher waves reduce shipping times. According to the estimates in the final column, a one standard deviation increase in wave height (1.5m) reduces shipping time by about 8 log-points. Shipping times are also reduced when waves are against the direction of travel: waves that are opposite to the vessel’s direction of travel (180 degrees) reduce the shipping time by 2 log points. While the small positive effect of wave height and direction on vessel speed is possibly surprising, we find similar results when we run these regressions with satellite data (Appendix A.3). These data do not rely on an imputation of routes and report information on vessel speed and direction at exact vessel locations in the ocean, indicating that our results are not driven by our imputation methodology. The results are also in line with earlier findings that have shown a positive effect of wave height on speed (Filtz et al., 2015), and could be consistent with vessels increasing cruising speed when passing through

areas with bad weather.

Table 3: Effect of Weather on Shipping Times

Dep. Var:	$\ln(T_s)$		
Wave Height ^s	-0.051*** (0.000)	-0.051*** (0.000)	-0.054*** (0.000)
Direction ^s		-0.007*** (0.000)	-0.012*** (0.000)
Wave Height ^s × Direction ^s			0.004*** (0.000)
Fixed Effects	Y	Y	Y
Controls	Y	Y	Y
R-Squared	0.789	0.789	0.789
Observations	9,892,000	9,892,000	9,892,000

Notes: The table reports the coefficients on the weather terms from estimating specification (3). Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. The variable wave height is expressed in meters, while the relative direction is expressed in hundreds of degrees. Importer, HS10 product, exporter, route-time, vessel, and related party fixed effects, as well as log charges and log weight, are included in the regression but not reported in the table for brevity. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

3 Empirical Analysis

Armed with our measure of unexpected shipping times at the transaction-level, $\hat{t}_{weather,s}$, we investigate how U.S. importers cope with weather-induced shipping delays and the associated shipping time risk. Our analysis revolves around two questions. First, what are the effects of shipping delays on U.S. manufacturing importers' performance? Second, do importers adjust their sourcing strategy and import demand in response to shipping time risk? We start by documenting that shipping delays, defined as extremely long shipping times induced by weather conditions, have negative and significant consequences for importers' performance. We then build a measure of shipping time risk based on the weather-induced unexpected shipping times, and establish that U.S. manufacturing importers adjust their sourcing strategy and import demand in response to shipping time risk.

3.1 Shipping Delays and Importers' Performance

We define a weather-induced delay as a case where a transaction's weather-induced unexpected shipping time, $\hat{t}_{weather,s}$, is above the 95th percentile of the shipping time distribution within the corresponding product-route. For each importer, we then compute the share of these weather-delayed inputs in total input costs in a year as:

$$FracDelayed_{ft}^{weather} = \frac{\sum_s \mathbb{D}_{ft}^{s,weather} \cdot \text{Imp Value}_{ft}^s}{\text{Total Input Costs}_{ft}}, \quad (5)$$

where $\mathbb{D}_{ft}^{s,weather}$ is an indicator that shipment s to importer f was weather-delayed in year t , Imp Value_{ft}^s is the import value of such shipment, and $\text{Total Input Costs}_{ft}$ are defined as labor plus materials costs, including domestically sourced inputs and imports by all modes of transportation. In this way, $FracDelayed_{ft}^{weather}$ measures the share of an importer’s input expenditures that are subject to extremely long shipping times in a given year. We then estimate the following panel regression for the years 2011-2019:

$$\ln(Y_{ft}^o) = \alpha + \beta_1 FracDelayed_{ft}^{weather} + \gamma_f + \delta_t + \epsilon_{ft}, \quad (6)$$

where Y_{ft}^o is either the sales, operating profits (sales minus materials and labor costs), or number of employees, and γ_f and δ_t are firm and year fixed effects, respectively. Our identifying assumption is that, conditional on firm fixed effects, the fraction of input costs that is subject to extreme shipping delays due to weather is orthogonal to any unobservable characteristics that may affect an importer’s post-delay performance. The construction of weather shocks discussed in the previous section is designed to satisfy this assumption.

Table 4 reports the results, with standard errors clustered at the firm-level. Shipping delays significantly disrupt production levels and profit margins. Increasing the fraction of delayed shipments by one standard deviation (2.02 percentage points, which is almost a seven-fold increase from the average fraction delayed, 0.30%) is associated with a drop in sales of 5.3%, a fall in profits of 3.5%, and in employment of 0.9%. Therefore, extreme unexpected delays have a large and significant impact on U.S. importers. Such large effects of shipping delays are consistent with evidence that the shipping risk is borne primarily by the buyer (see [Herghelegiu and Monastyrenko, 2020](#) and [Eurosender, 2023](#)) and that insurance for supply disruption events is limited and expensive ([Heckmann, 2016](#)).¹⁴

In Appendix B.1, we compute an alternative measure of weather shocks and re-run the regressions. Instead of averaging over the weather conditions of the entire route, we predict where on the route the vessel is on each day by assuming that vessels travel at constant speed, and use local weather conditions around this estimated location to re-construct our

¹⁴Note that most importers are relatively small. Our estimates of the impacts of shipping delays on sales are in line with the effects of other supply chain disruptions found in a recent literature. [Carvalho et al. \(2021\)](#) estimate an elasticity of sales of -3.6% following a shock hitting a domestic supplier. [Barrot and Sauvagnat \(2016\)](#) find that when one of their suppliers is hit by a major natural disaster, firms experience an average drop of 2 to 3 percentage points in sales growth. [Khanna et al. \(2022\)](#) find that firms with one standard deviation higher supplier risk (which they define as the exposure of suppliers to different lockdown policies across India) decreased their output by up to 2.7% after the lockdowns.

Table 4: Effect of Extreme Delays on Firms' Outcomes

	(1)	(2)	(3)
Dependent Variable (in logs):	Sales	Profits	Employees
Frac Delayed	-2.617*** (0.303)	-1.728*** (0.352)	-0.452*** (0.163)
Importer FE	Y	Y	Y
Year FE	Y	Y	Y
R-Squared	0.97	0.90	0.98
Observations	78,000	78,000	78,000

Notes: The table reports the coefficients on the fraction of imports delayed, $FracDelayed_{ft}$, from specification (6). Fixed effects for the firm and the year are included. The number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the firm level. Mean of $FracDelayed_{ft}^{weather}$ is 0.003 and standard deviation is 0.020. R^2 is the overall fit inclusive of the fixed effects. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

measure of weather-induced shipping delays. While we find smaller effects of wave height and wave direction on shipping times under this alternative specification than in Table 3, the relationship between delays and firm outcomes is similar to our baseline in Table 4: a one standard deviation increase in delayed shipments (2.04 percentage points) is associated with a drop in sales of 5.4%, a decline in profits of 3.8%, and a fall in employment of 0.9%.

We also report, in Appendix B.1, the results of the regressions estimated using a broader measure of shipping delays, which we construct as the residual from running equation (3) without the weather terms. This measure captures all types of delays, including those unrelated to weather, such as port congestion or strikes. However, it requires the stronger identification assumption that these residuals do not include any determinants of shipping times that are anticipated by importers. While this assumption is not as clearly satisfied as for weather shocks, we find qualitatively similar, though larger effects than in the baseline: a one standard deviation increase in the fraction of delayed shipments under the broader measure (3.1 percentage points) is associated with a drop in sales of 8.3%, a fall in profits of 4.7%, and in the number of employees of 1.6%.

3.2 Shipping Time Risk and Import Demand

Having shown that shipping delays have large adverse impacts on U.S. importers, we now investigate whether firms take actions to actively diversify this source of risk. To motivate this analysis, we show in Table 5 that firms rely on multiple routes to source the same HS-10 product within the same year. We include all modes of transportation to capture diversification across modes. The average firm uses 2.3 routes per product-year across all modes of transportation. The large standard deviation compared to the mean indicates that

Table 5: Summary Statistics on Foreign Sourcing

	Mean	St. Dev.	P50	P95
Number of Routes	2.32	4.64	1	6.78
Number of Suppliers	1.90	4.17	1	4.68
HHI across Routes	0.84	0.25	1	1
HHI across Suppliers	0.89	0.21	1	1

Source: LFTTD and authors' calculations. Table reports the mean and standard deviation across importer-product-year tuples in our 2011-2019 sample period. The Herfindahl-Hirschman index (HHI) across routes is defined by computing the share of each route in import value at the importer-HS10 product-year level and taking the sum of squared shares. The HHI across suppliers is computed analogously using the share of each supplier in import value at the importer-HS10 product-year level.

there is substantial heterogeneity across buyers. While the median importer uses only one route for a given product, firms at the 95th percentile use nearly 7. Similar patterns arise when looking at the number of foreign suppliers instead of routes. Additionally, import values are highly concentrated among routes and suppliers, with substantial heterogeneity in the degree of concentration across importers.

Is importers' use of multiple routes and suppliers related to hedging against shipping-time risk? To explore this question, we compute a measure of exposure to shipping time risk based on the volatility of the weather-induced shipping times experienced by importers. We then document how changes in exposure to such risk both across importers and over time affect multi-route and supplier sourcing as well as other patterns of import demand.

Measuring Exposure To Risk. We start by computing the standard deviation of the weather-induced shipping times, $\hat{t}_{weather,s}$, over three-year rolling windows for each supplier-product-route-year (x, h, r, t) combination. We denote this standard deviation by $\widehat{StdTime}_{xhrt-3,t-1}$. We assume that non-vessel transactions are riskless and set $\widehat{StdTime}_{xhrt-3,t-1} = 0$ for such (x, h, r, t) cells. Cells with fewer than 10 transactions are dropped. We aggregate this risk measure at the importer-product-year level by taking a weighted average over the importer's suppliers and routes over the previous three years:

$$\widehat{StdTime}_{fht-3,t-1} \equiv \sum_{x,r} \omega_{fxhr,t-3,t-1} \widehat{StdTime}_{xhrt-3,t-1}, \quad (7)$$

where the weights $\omega_{fxhr,t-3,t-1}$ are firm f 's import shares of product h from each supplier-route over the years $t - 3$ to $t - 1$. Our measure is akin to a shift-share exposure measure, as in [Bartik \(1991\)](#), where the supplier-route-product level standard deviations are the ‘‘shift’’, and the import shares are the pre-determined ‘‘shares.’’

Econometric Specification. Armed with our measure of risk, we estimate the following panel specification:

$$\ln(Y_{fht}) = \alpha + \beta_1 \ln(\widehat{StdTime}_{fht-3,t-1}) + \beta_2 X_{fht} + \gamma_f + \mu_h + \delta_t + \epsilon_{fht}, \quad (8)$$

where Y_{fht} is an import demand outcome of importer f in year t for product h . This regression analyzes whether the risk faced by the importer in the previous three years ($t - 3$ to $t - 1$) affects its sourcing patterns in the current year t . Our shift-share measure of risk helps alleviate concerns of reverse causality, that is, the endogeneity of the risk measure through the importers' choice of routes and suppliers. Given the well-known stickiness in buyer-supplier relationships (e.g., [Martin et al. \(2026\)](#), [Heise \(2024\)](#)), our panel specification with firm fixed effects exploits variation over time in importers' exposure to risk driven by within-route changes in the volatility of weather shocks.

We consider the following dimensions of import sourcing as dependent variable: (i) the number of routes, (ii) the number of foreign suppliers, (iii) the concentration of value imported across routes as measured by the Herfindhal–Hirschman index (HHI), (iv) the HHI of imports across suppliers, and (v) the total value imported. X_{fht} is a vector of controls, and γ_f , μ_h and δ_t are importer, product, and year fixed effects, respectively. In our baseline specification, we omit firm-product pairs with only one foreign supplier in year t since the purchase volume may be too small to make diversification viable or the product may be too specialized. We show below that our results are robust to including such firms.

Our controls X_{fht} include the importer's unexpected shipping times, $\hat{t}_{weather,s}$, averaged over the previous three years using the same weights as in (7). This variable accounts for the direct negative effect of shipping times on import demand, as documented in [Hummels and Schaur \(2013\)](#). The average shipping time also controls for the fact that suppliers located in countries further away may mechanically have more volatile shipping times purely because they have more scope for delays. We also include the average unit value paid by the importer for product h in year t . This variable controls for the fact that riskier suppliers may sell cheaper inputs, confounding the relationship between risk and import demand we aim to estimate. Finally, we control for importers' size (proxied by the total imports of product h over the previous 3 years) and for suppliers' size (proxied by the total exports of the product over the previous 3 years), since larger importers or exporters have more shipments, which could mechanically increase their risk.¹⁵

¹⁵For example, exporters shipping greater volumes in a given period may need to send more shipments with more vessels, which may introduce a predictable correlation between risk and exporter size. This effect would not be captured by the exporter fixed effects in specification (3). The reasoning is similar for importers.

Table 6: Shipping Time Risk and Import Demand

	(1)	(2)	(3)	(4)	(5)
Dep. Var.:	Number of Routes	Number of Suppliers	HHI over Routes	HHI over Suppliers	Value Imported
Std Time	0.177*** (0.009)	0.116*** (0.008)	-0.089*** (0.003)	-0.072*** (0.003)	-0.085*** (0.011)
Importer FE	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y
R-squared	0.72	0.67	0.46	0.41	0.88
Observations	129,000	129,000	129,000	129,000	129,000

Notes: The table shows the coefficients on the log of the standard deviation of the weather-induced shipping times, $\widehat{StdTime}_{fht-3,t-1}$, from specification (8) for different dependent variables: the log of the number of routes (col. 1), the log of the number of foreign suppliers (col. 2), log HHI across routes (col. 3), log HHI across suppliers (col. 4), and the log total value imported (col. 5). The number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the firm level. Regression includes controls for three-year average of weather-induced shipping times, average unit value paid, the importer's total imports over the previous three years, and the suppliers' total exports to the U.S. over the previous three years. For brevity, the table does not report the coefficients of these regressors. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

Results. Table 6 presents the findings for the sample period 2011-2019. Standard errors are clustered at the firm level. Column 1 documents a positive and significant relationship between the number of routes used and shipping risk. An increase in risk from the 25th to the 75th percentile of the weather risk distribution (51 log points) increases the number of shipping routes used by 9.0%. Column 2 shows that there is also a positive relationship between the number of foreign suppliers and risk. An increase from the 25th to the 75th percentile of the risk distribution increases the number of suppliers used by 5.9%. The larger coefficient on risk in column 1 than in column 2 suggests that importers diversify weather risk primarily by adding routes rather than suppliers, as weather risk is route-specific and suppliers may share the same routes.

We next look at the relationship between shipping time risk and the concentration of import value across an importer's routes and suppliers. Column 3 reports a negative and significant coefficient between shipping risk and the HHI over routes, suggesting that importers with riskier routes feature a more diversified pattern of expenditure across their routes. Column 4 shows that this effect is similar when we look at the concentration across suppliers. Lastly, in column 5, we find a negative and statistically significant relationship between shipping time risk and total import value. Quantitatively, going from the 25th to the 75th percentile of the risk distribution decreases the route HHI by 4.5%, the supplier HHI by 3.7%, and total imports by 4.3%.

Taking stock, our empirical analysis shows that importers with riskier supply chains rely on more routes and foreign suppliers to source their inputs, and reduce the concentration of their input expenditures. We interpret these results as evidence of risk diversification behavior, operating at both the extensive and intensive margins. In addition, we find that the net effect of these different margins of adjustment is a significant reduction in total imports. Importantly, this negative effect of risk on import demand is estimated controlling for the effect of average shipping times, which is the focus of [Hummels and Schaur \(2013\)](#). We incorporate these channels into a model with risky shipping times in Section 4.

3.2.1 Selection Bias, Robustness, and Diversification via Air Shipping

Our empirical results show that firms that are more exposed to shipping time risk feature a more diversified structure of import demand. The measure of risk exposure, however, takes the firm’s set of suppliers and routes as given, raising the concern of selection bias. We now discuss various forms in which this selection could affect our results. Consider first the case where importers differ in their risk aversion. To the extent that more risk averse importers feature safer suppliers/routes and also more suppliers/routes, this selection works against our empirical findings. That is, it produces a negative relationship between shipping time risk and the number of suppliers or routes. Consider next the role of firm size. In the presence of fixed costs to adding suppliers and routes, larger firms would feature more suppliers/routes. If in addition larger firms feature riskier suppliers and routes, this selection could produce relationships as the ones documented in the previous section. We address this issue by including firm fixed effects and controlling for past imports. Moreover, in our sample we instead find a negative and significant correlation (-0.26) between the size of the firm (proxied with log sales) and our risk measure.

We next perform a number of robustness exercises with our weather risk measure and report their results in Appendix B.2. First, we show that we obtain similar results using an alternative measure of weather risk, which we construct by predicting where the vessel is on the route on each day assuming that the vessel travels at constant speed, and using local weather conditions around the vessel’s likely location. Going from the 25th to the 75th percentile of the risk distribution for this measure (61 log points) increases the number of shipping routes used by 9.9%, the number of suppliers by 6.4%, and reduces total imports by 5.7%.

Second, we report the results using the broader risk measure that is constructed by taking the standard deviation of the residual from running equation (3) without the weather terms. The results are similar both qualitatively and quantitatively: an increase from the 25th to the

75th percentile of the risk distribution (91 log points) increases the number of routes used by 8.4% and the number of suppliers by 5.2%. Moreover, it decreases the route HHI, supplier HHI, and total imports by 4.3%, 3.2%, and 5.6%, respectively.

Third, we show that our results also hold when we include firm-time fixed effects, which control for firm-level shocks that may affect production choices, instead of firm and time fixed effects separately. Fourth, we also obtain similar findings when we include in the sample firms sourcing from a single supplier, and when we focus only on the risk of the importer’s main supplier rather than a weighted average across suppliers. Fifth, our main findings are unchanged when we control for importers’ inventories, which have been recently shown to be an important margin of adjustment to sourcing risk (see [Alessandria et al. \(2023\)](#), [Carreras-Valle \(2024\)](#)).

Lastly, we analyze whether U.S. firms use different modes of transportation to diversify shipping risk. We focus on air shipments, as over half of all importer-product-year combinations are sourced by both vessel and plane. To do so, we construct a dummy variable that is equal to one if a firm has obtained imports by air in year t , and estimate a variant of our main specification, equation (8),

$$d_{fht} = \alpha + \beta_1 \ln(\widehat{StdTime}_{fht-3,t-1}) + \beta_2 X_{fht} + \gamma_f + \mu_h + \delta_t + \epsilon_{fht}, \quad (9)$$

where d_{fht} is a dummy that is equal to one if firm f uses air shipments for HS10 h in year t , and $\widehat{StdTime}_{fht-3,t-1}$ is the same weather-based risk measure as before. The controls X_{fht} are identical to the ones used before. Table 7 documents that higher shipping risk is associated with a higher likelihood of using air shipments. An increase in risk from the 25th to the 75th percentile of the weather risk distribution increases the likelihood of using air shipments by 1.6%. While the effect is small since firms use air shipments likely for many reasons, for example for fast delivery of seasonal goods, our findings suggest that firms use air transportation to hedge ocean shipping risk.¹⁶

Taken together, our empirical findings establish that U.S. importers systematically react to shipping risk along different margins of adjustment. Importers with riskier suppliers or routes feature i) more suppliers and routes, ii) less concentrated import expenditures, iii) lower imports, and iv) use multiple modes of transportation.

¹⁶This result is in line with the findings in [Hummels and Schaur \(2010\)](#), who show that firms use air shipping to smooth demand volatility on international markets.

Table 7: Shipping Time Risk and Import Demand with Air Shipments

Dep. Var.:	Air Shipments
Std Time	0.032*** (0.003)
Avg Time	1.108*** (0.081)
Importer FE	Y
Product FE	Y
Year FE	Y
Controls	Y
R-Squared	0.54
Observations	129,000

Notes: The table shows the coefficient on the standard deviation of the weather-induced shipping times, $\widehat{StdTime}_{fht-3,t-1}$, from specification (9). The dependent variable is a dummy that is equal to one if firm f uses air shipments for HS10 h in year t . The number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the firm level. Regression includes controls for three-year average of weather-induced shipping times, average unit value paid, the importer’s total imports over the previous three years, and the suppliers’ total exports to the U.S. over the previous three years. For brevity, the table does not report the coefficients of these regressors. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

4 A Model of Input Sourcing with Shipping Risk

To rationalize the empirical evidence on shipping time risk and to quantify its aggregate implications, we lay out a theoretical framework that builds on standard models of importing with firm heterogeneity, as in [Halpern et al. \(2015\)](#), [Blaum et al. \(2018\)](#), and [Gopinath and Neiman \(2014\)](#). As in these models, imported intermediate inputs lower firms’ marginal production costs due to production complementarities and differences in qualities and prices, but importing requires the payment of fixed costs. Our key departure from this literature is that inputs’ shipping times are a component of input quality, thus affecting production levels in the spirit of [Hummels and Schaur \(2013\)](#), and that such shipping times are stochastic. As a result, firms have incentives to increase the number of foreign suppliers to mitigate the impact of shipping time risk on expected revenues.

Section 4.1 outlines the environment of the model, while Section 4.2 characterizes the firm’s problem. Section 4.3 provides theoretical results that describe the impact of risk on import demand. Section 4.4 closes the model in equilibrium. We calibrate the model and perform counterfactuals in Section 5.

4.1 Environment

We consider a small open economy populated by a fixed mass of risk-neutral firms that produce differentiated manufacturing varieties which are sold locally. Firms produce by combining labor, a domestic input, and a foreign input according to the following nested structure:

$$y_f = \varphi_f l^{1-\gamma} \left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + \left(\sum_{i=1}^N \alpha_i x_i \right)^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\gamma \frac{\varepsilon}{\varepsilon-1}}, \quad (10)$$

where f denotes a firm, φ_f is firm efficiency, $\gamma \in (0, 1)$ and $\varepsilon > 1$. Labor l is combined with an intermediate input bundle using a Cobb-Douglas aggregator. The intermediate bundle, in turn, is a CES aggregator of a domestic input with quantity x_D and a foreign input that is sourced from N suppliers, with quantity x_i and quality α_i for supplier i . As is standard in the literature, the extensive margin of trade is limited by fixed costs. In particular, each foreign supplier entails the payment of a fixed cost F in units of domestic labor.¹⁷

The firm chooses its suppliers from a pool of unlimited foreign suppliers to whom it is exogenously matched. A central element of our theory is that supplier shipping times are stochastic and unknown to firms at the time of placing orders. Furthermore, motivated by our empirical evidence, shipping delays adversely impact production by reducing the quality of inputs α_i , similarly to [Hummels and Schaur \(2013\)](#). For tractability, we assume that the suppliers of any given firm are *ex-ante* identical but may differ *ex-post* in their realized shipping times—implying that the extensive margin of trade can be summarized by the number of suppliers.

We parametrize the relationship between input quality and shipping time as:

$$\alpha_i = \begin{cases} \bar{\alpha}_i & \text{if } d_i \leq \mathbb{E}[d_i] \\ e^{-\tau \cdot d_i} & \text{if } d_i > \mathbb{E}[d_i], \end{cases} \quad (11)$$

where d_i are the number of days it takes to ship the input of supplier i to firm f , $\bar{\alpha}_i = e^{-\tau \cdot \mathbb{E}[d_i]}$, and $\mathbb{E}[d_i]$ is the expected shipping time. This formulation implies that, if an input arrives earlier than or just as expected, it has a constant level of quality $\bar{\alpha}_i$. Instead, if an input arrives *later* than expected, quality falls with shipping time with a semi-elasticity given by τ .¹⁸ For each importer, the shipping days d_i are i.i.d. and drawn from a CDF given by $G_f(\cdot)$,

¹⁷This production structure is standard in the literature—it corresponds to the one in [Gopinath and Neiman \(2014\)](#) or [Blaum et al. \(2018\)](#) when foreign inputs are perfect substitutes. We abstract from love-of-variety effects for the foreign inputs to focus on the extensive margin of importing as a channel of risk diversification.

¹⁸The specification in equation (11) rules out risk loving behavior by imposing that early input arrivals do

which is known to the firm at the time of placing input orders. Additionally, this distribution is importer-specific to allow for differences across importers in the riskiness of their foreign suppliers. This feature enables the model in our quantitative exercises to account for firms' differential exposure to risk observed in the data.¹⁹

Firms are price takers in input markets and thus can source any quantity of the domestic and foreign inputs and labor at prices p_D , p_M , and w , respectively. We assume that foreign input prices p_M , which incorporate any variable trade costs, are exogenously given and common across all suppliers. In output markets, firms compete under monopolistic competition.

The firms are owned by a representative consumer who is endowed with L units of labor and consumes the locally produced manufacturing goods with preferences given by:

$$C = \left(\int c_f^{\frac{\sigma-1}{\sigma}} df \right)^{\frac{\sigma}{\sigma-1}}, \quad (12)$$

where $\sigma > 1$ and c_f denotes final consumption of the good produced by firm f .²⁰ To allow for input-output linkages across firms, we assume a structure of roundabout production by which firms use the output of all other domestic firms as inputs. In particular, the domestic bundle x_D is produced by a perfectly competitive intermediary sector that uses the same CES aggregator as in (12).²¹

4.2 Firm's Problem under Risk

We next describe the firm's problem of choosing domestic and foreign input quantities and the number of suppliers in the presence of risk. The total sales of firm f , which include demand from both consumers and other firms, are:

$$R_f = y_f^{\frac{\sigma-1}{\sigma}} P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma}, \quad (13)$$

not increase production. This property will be useful in the quantitative exercises of Section 5 to come to terms with the empirical evidence of Section 3.

¹⁹In the quantitative exercises of Section 5, we design the exogenous assignment of firms to suppliers to replicate the negative correlation between firm size and supplier riskiness observed in the data. A micro-foundation where firms search for suppliers, and finding safer suppliers is more costly, would deliver this pattern.

²⁰For simplicity, we abstract from exporting, importing of final goods, and consumption of non-tradable goods. Incorporating these elements, as in [Blaum \(2024\)](#), is feasible but would complicate the analysis without giving additional insights.

²¹The assumption that the CES aggregators for the domestic input and consumer utility coincide is made for tractability. Under this assumption, we do not need to treat sales to the consumer and to the domestic input producer separately in the firm's problem. See also [Blaum et al. \(2018\)](#); [Gopinath and Neiman \(2014\)](#); [Adão et al. \(2025\)](#) for a similar assumption.

where y_f is given by equation (10), P is the price index associated with (12) and S denotes total domestic spending (combining demand by firms and consumers). Both P and S are endogenous variables determined in general equilibrium.

Firms choose the quantities of domestic and foreign inputs, as well as the number of foreign suppliers, before the realization of uncertainty. The quantity of labor is instead chosen after uncertainty is realized. This assumption simplifies the numerical solution to the firm's problem in the quantitative exercises below. Since foreign suppliers are ex-ante symmetric, in the first stage the firm sources the same quantity from all suppliers, i.e., $x_i = x$ for all i . After maximizing out labor, the firm's problem before the realization of uncertainty is given by:

$$\max_{x_D, x, N} \chi_f \mathbb{E} \left[\left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + \left(\sum_{i=1}^N \alpha_i(d_i) \right)^{\frac{\varepsilon-1}{\varepsilon}} x^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\psi} \right] - p_D x_D - N p_M x - w N F, \quad (14)$$

where χ_f and ψ depend on firm efficiency, general equilibrium variables, and parameters, and $\alpha(d_i)$ is given by (11).²² Note that the expectation operator is taken over the possible realizations of d_i and thus depends on the distribution of shipping times $G_f(\cdot)$.

In choosing the number of foreign suppliers, firms trade off the diversification of shipping risk against the payment of the fixed costs. Similarly, the choice of the quantity of the imported input x is limited not only by its price but also by the associated shipping risk. Before turning to the definition of the equilibrium, we illustrate how risk affects the firm's production choices in a simplified environment.

4.3 The Workings of Risk

Can the theory outlined so far come to terms with the evidence documented in Section 3 on the effect of shipping risk on import demand? To answer this question, we now study the effects of increased risk in supplier input quality on import demand. For tractability, we consider in this section a version of the model without the domestic input and we abstract from general equilibrium forces. In the quantitative exercises of Section 5, we allow for such effects and include the domestic input in the production. All derivations of this section are contained in Appendix C.2.

After maximizing out labor and the foreign inputs, the firm's problem in the simplified

²²See Section C.1 of the Appendix for the definitions of χ_f and ψ and a derivation of (14).

version of the model becomes:

$$\max_N \underbrace{\tilde{\chi}_f \left(\mathbb{E} \left[\bar{\alpha}^{\tilde{\psi}} \right] \right)^{\frac{1}{1-\tilde{\psi}}}}_{=\tilde{R}} - wNF, \quad (15)$$

where $\bar{\alpha} \equiv \frac{1}{N} \sum_{i=1}^N \alpha_i$ is the average supplier quality, $\tilde{\psi} \equiv \frac{\gamma(\sigma-1)}{1+\gamma(\sigma-1)} \in (0, 1)$, $\tilde{\chi}_f$ is a constant that depends on firm efficiency, general equilibrium objects (held fixed in this section) and parameters, and \tilde{R} is expected revenues net of labor and foreign input variable costs. Expression (15) makes it clear that volatility in the average supplier quality lowers expected revenues as $\tilde{\psi} < 1$. Relying on a second order approximation of $\bar{\alpha}^{\tilde{\psi}}$ around $\mathbb{E}[\alpha]$, the firm problem can be written as:

$$\max_N \tilde{\chi}_f \left((\mathbb{E}[\alpha])^{\tilde{\psi}} - \tilde{\psi} \frac{(1-\tilde{\psi})}{2} (\mathbb{E}[\alpha])^{\tilde{\psi}-2} \frac{1}{N} \mathbb{V}[\alpha] \right)^{\frac{1}{1-\tilde{\psi}}} - wNF, \quad (16)$$

where $\mathbb{E}[\alpha]$ and $\mathbb{V}[\alpha]$ are the mean and variance of the supplier-level quality. This expression highlights the role of the mean and the variance of supplier quality, as well as of the number of suppliers, in shaping expected revenues. In particular, dispersion in input qualities reduces expected revenues. By increasing the number of suppliers N , the firm lowers the variance of the *average* supplier quality $\mathbb{V}[\bar{\alpha}] = \mathbb{V}[\alpha]/N$, thus reducing the effective amount of risk faced and mitigating its effects on expected revenues. The following result formalizes the effect of increased supplier risk for the case where N is continuous.

Proposition 1. (*Effect of Risk on Inputs*) Let N^* be the optimal number of suppliers and $\epsilon_{N^*, \mathbb{V}[\alpha]} = \frac{\partial N^*}{\partial \mathbb{V}[\alpha]} \frac{\mathbb{V}[\alpha]}{N^*}$ be the elasticity of N^* with respect to the variance of supplier-level quality $\mathbb{V}[\alpha]$. Then $\epsilon_{N^*, \mathbb{V}[\alpha]} > 0$ if and only if

$$\frac{1}{N^*} \frac{\mathbb{V}[\alpha]}{(\mathbb{E}[\alpha])^2} < \frac{2}{\tilde{\psi}}. \quad (17)$$

That is, under condition (17), a higher $\mathbb{V}[\alpha]$ leads to an increase in N^* . Furthermore, $\epsilon_{N^*, \mathbb{V}[\alpha]} < 1$ regardless of whether (17) holds. It follows that a higher $\mathbb{V}[\alpha]$ always reduces the import value.

Proof. See Section C.2 of the Appendix. □

The first part of Proposition 1 states that a mean preserving spread in supplier-level quality increases the number of suppliers, if condition (17) holds. There are two opposite

forces at work. When the variance of quality, $\mathbb{V}[\alpha]$, is higher, holding constant the expected value $\mathbb{E}[\alpha]$, an increase in the number of suppliers leads to a larger reduction in the variance of average supplier quality, $\mathbb{V}[\alpha]/N$, which is what matters for expected profits (16). This force increases the returns to adding suppliers. At the same time, a higher variance of supplier-level quality reduces expected revenue, leading to a lower return to adding suppliers. When supplier risk is a small part of expected revenue, as ensured by condition (17), the negative level effect on expected revenues is dominated by the stronger reduction in the variance of average quality. As a result, a mean preserving spread in supplier quality increases the number of suppliers.

In the second part of the proposition, we turn our attention to import values. To understand this result, note that the import value at the optimal number of suppliers, $N^*x^*p_M$, is proportional to expected revenue and thus is a decreasing function of the variance of average supplier quality $\mathbb{V}[\alpha]/N^*$ —see (16).²³ An increase in the variance of supplier-level quality therefore lowers import value if N^* either decreases or it increases less than proportionally with $\mathbb{V}[\alpha]$. The proposition establishes that $\epsilon_{N^*, \mathbb{V}[\alpha]} < 1$ and hence that import value necessarily falls with more volatility in supplier quality.²⁴

In connecting these results to our findings of Section 3, it should be noted that in the empirical analysis we measure the volatility of suppliers' shipping times, not of input qualities, which are unobservable. In our theory, a mean-preserving spread to the distribution of shipping days affects both the variance and the mean of input qualities. In particular, given the non-linear mapping between days and qualities, the expected input qualities can increase or decrease depending on parameters. In turn, the effect of a given change in expected quality on the returns to adding suppliers also depends on parameters (see Section C.4 of the Appendix for a formal treatment). Ultimately, whether the theory can come to terms with our empirical findings on shipping time volatility and import demand documented above is a quantitative matter which we tackle in Section 5. There, we consider a calibrated version of the model with a domestic input, general equilibrium, and firm heterogeneity. Before turning to this analysis, we next show how the model is closed in equilibrium.

²³As shown in Appendix C.2, total import value is given by

$$Nxp_M \approx (p_M)^{-\frac{\tilde{\psi}}{1-\tilde{\psi}}} \left(\tilde{\psi}\chi_f \left((\mathbb{E}[\alpha])^{\tilde{\psi}} - \tilde{\psi} \frac{(1-\tilde{\psi})}{2} (\mathbb{E}[\alpha])^{\tilde{\psi}-2} \frac{1}{N} \mathbb{V}[\alpha] \right) \right)^{\frac{1}{1-\tilde{\psi}}}.$$

²⁴In Appendix C.3, we also investigate how the firms' response to risk, both in terms of the number of suppliers and import value, depends on productivity.

4.4 Equilibrium

Thus far we have studied the effects of increased shipping time risk for a firm when aggregate domestic spending and the price index are kept fixed. In going forward, we allow this firm-level risk to affect these equilibrium variables. We abstract from aggregate risk by assuming that there is a continuum of firms of unit mass within each type f . We consider an equilibrium where firms maximize profits, the consumer maximizes utility, and goods markets clear. We do not impose labor market clearing and therefore allow the manufacturing sector to run a deficit or a surplus with the rest of the economy or the world. An equilibrium can be fully characterized by the aggregate domestic spending S and the price index P associated to consumer utility. Note that the price of the domestic input bundle is given by $p_D = P$ as the domestic input aggregator is identical to consumer preferences. We normalize the wage to unity. We now describe how S and P are determined.

Consumer expenditure is given by:

$$PC = wL + \Pi, \tag{18}$$

where $\Pi \equiv \int \pi_f df$ are aggregate profits and π_f are expected profits of firm type f . Because there is a unit mass of firms of each type, π_f is also the aggregate profits of type f . Given the roundabout structure by which firms use locally-produced manufacturing products as inputs, aggregate domestic spending satisfies:

$$S = PC + p_D \int x_{Df} df, \tag{19}$$

where x_{Df} is the demand for the domestic input of firm f . Standard calculations imply that:

$$P = \left(\int p_f^{1-\sigma} df \right)^{\frac{1}{1-\sigma}}, \tag{20}$$

where p_f is the price set by firm type f . An equilibrium is attained whenever (18), (19), and (20) are satisfied and firms choose their inputs optimally—see Appendix C.5 for additional details.

5 Counterfactual Analysis: Climate Change and Infrastructure Risk

The empirical evidence in Section 3 shows that firms increase their number of suppliers and reduce their import values when faced with greater shipping time risk, patterns that the theory developed in the previous section can qualitatively generate. We now discipline the model’s parameters with key moments of the data—notably, our reduced-form estimates of the effect of shipping risk on firms’ sourcing decisions—to establish that the model can also quantitatively come to terms with our empirical evidence (Section 5.1). We then employ the model to quantify the aggregate consequences of climate change and strains on port infrastructure (Section 5.2). These exercises highlight how our quantitative model can be used as a laboratory to assess the impact of any risk event affecting ocean shipping on U.S. imports and welfare.

5.1 Calibration

Our calibration strategy requires the model to replicate various moments related to suppliers’ shipping time risk. To assess whether the model can be consistent with the empirical evidence of Section 3, we target our estimate of the effect of shipping time risk on the extensive margin of importing. We do not target the effect of risk on import values, and thus leave this moment for model evaluation. Additionally, we require that the model matches the sensitivity of sales to shipping times. Finally, given our focus on aggregate effects in the counterfactuals, we target the joint distribution of firm size and supplier risk across importers. We next describe how we parameterize this distribution.

Parametrization of Firm Heterogeneity and Shipping Days To generate cross-sectional variation in importer size and exposure to supplier risk, we allow for firm types to be heterogeneous in two dimensions: efficiency φ_f and the standard deviation of shipping days of their suppliers σ_{df} . Firm efficiency φ_f is drawn from a log-normal distribution with standard deviation σ_φ (we normalize average log efficiency). The economy is populated by two types of suppliers, low and high risk: $\sigma_{df} \in \{\sigma_{dL}, \sigma_{dH}\}$. To allow for firms of different sizes to differ in their risk exposure, we let the risk type be correlated with firm efficiency. Specifically, we assign each firm a draw of a latent risk variable σ_{df}^l generated according to:

$$\log(\sigma_{df}^l) = \rho_{\varphi\sigma} \log(\varphi_f) + v_f, \tag{21}$$

where $\rho_{\varphi\sigma}$ is a parameter that controls the correlation between the latent risk variable and efficiency, and v_f is a standard normal random variable. We then assign the high (low) risk type to firms with latent risk above (below) the median.

For any foreign supplier matched to an importer of type f , the distribution of shipping days d_i is log normal with standard deviation σ_{df} and mean μ_{df} . Shipping days are i.i.d. across a firm’s suppliers.

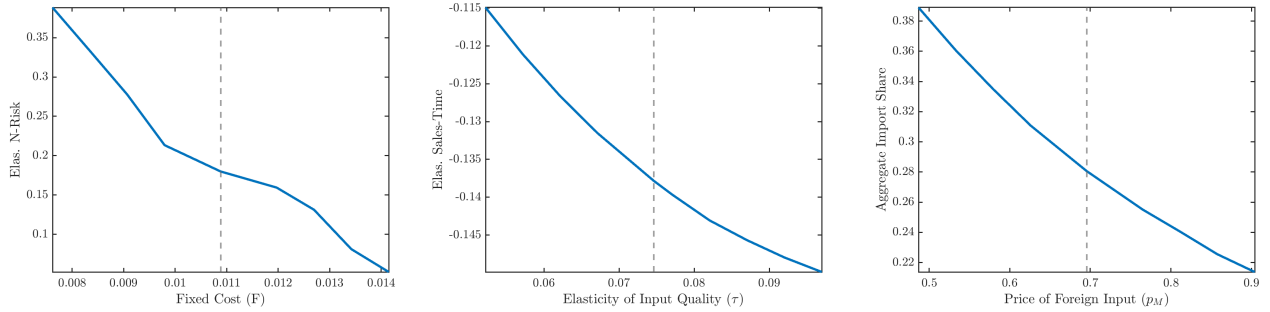
Parameters, Moments, and Identification We measure risk for each type, σ_{dL} and σ_{dH} , using the standard deviation of the residualized log shipping times over three-year rolling windows. We compute the average of this measure within the groups of firms above and below the median risk. To isolate the role of supplier risk, we set expected log shipping days μ_{df} for each type to match a common μ average shipping time of 19.1 days, which is the average across all seaborne shipments reported in the LFTTD (Table 2).

The fixed cost of adding foreign suppliers F , the semi-elasticity of input quality to shipping time τ , the price of imported inputs p_M , the dispersion in firm efficiency σ_φ , and the risk-efficiency correlation parameter $\rho_{\varphi\sigma}$ are chosen to match the following moments of the data: (i) the elasticity of the number of routes with respect to shipping risk, (ii) the elasticity of sales to shipping times, (iii) the aggregate import share, (iv) the coefficient of variation of log sales, and (v) the correlation between log sales and our risk measure. We use the number of routes to measure the extensive margin of importing in (i) since that is the firms’ main margin of diversification associated with our weather-based identification strategy. All moments are measured from the sample of U.S. manufacturing importers used in the empirical analysis of Section 3.

While each moment is affected by all parameters in equilibrium, intuitively, a higher cost of adding suppliers F makes diversification through the extensive margin of trade more costly, and thus controls the elasticity of the number of suppliers with respect to risk (Figure 3, left panel). The parameter τ affects the degree to which qualities, and thus revenues, fall with longer shipping times (Figure 3, center panel). Importantly, this negative association between shipping times and sales predicted by the model is verified in the data. A panel regression of log sales on the average weather-induced shipping time of the firm’s imported inputs yields a negative and statistically significant coefficient.²⁵ By affecting the relative price of imported inputs, p_M controls firms’ expenditure on foreign inputs and thus the aggregate import share (Figure 3, right panel). Finally, σ_φ controls the dispersion in firm size and $\rho_{\varphi\sigma}$ regulates the

²⁵The regression includes firm and year fixed effects and yields an estimated coefficient on the average weather-induced shipping time of -0.27 with a standard error of 0.07 clustered at the firm level.

Figure 3: Identification of Parameters



Notes: Each graph plots a specific moment as a function of the relevant parameter, holding constant the other parameters at their calibrated values. The elasticity of the number of suppliers (N) with respect to risk is the estimated coefficient of a cross-sectional regression of log optimal N on the log of the standard deviation of weather-induced shipping days, controlling for firm efficiency. The elasticity of sales with respect to shipping time is the coefficient of a cross-sectional regression of log sales on the log of shipping time.

correlation between firm size and risk.

Lastly, we set the elasticity of substitution between domestic and foreign inputs to $\varepsilon = 2.38$ and the demand elasticity to $\sigma = 3.83$ as in [Blaum et al. \(2018\)](#). We set the output elasticity with respect to materials to $\gamma = 0.58$ to match the average material expenditure share, defined as the ratio of material costs to material plus labor costs, in our sample.²⁶

Calibration Results We report the calibrated parameter values in Table 8. The model is able to closely match the targeted moments. The average fixed costs paid are \$121,300 dollars, which are in line with the literature ([Fieler et al. \(2018\)](#); [Antràs et al. \(2017\)](#)).²⁷ An important feature of our calibration is that the model roughly matches the elasticity of imports with respect to risk estimated in Section 3 (Table 6), despite not targeting this moment directly (see bottom of Table 8). This feature, together with the close match of the elasticity of the number of routes with respect to risk, implies that our quantitative model is able to come to terms with the key empirical findings of Section 3.

²⁶These parameter values are standard in the literature. Estimates of the demand elasticity σ typically range between 3 and 6 ([De Loecker \(2011\)](#)), with [Antràs et al., 2017](#) estimating $\sigma = 3.85$ for US manufacturing firms. These authors also estimate an elasticity of substitution of $\varepsilon = 2.8$. [Halpern et al. \(2015\)](#) find a value of $\varepsilon = 4$ using Hungarian data—which is the same value used by [Gopinath and Neiman \(2014\)](#). Finally, [Blaum et al. \(2018\)](#) find an estimate of $\gamma = 0.61$.

²⁷We back out the average fixed costs from the ratio of average sales to average fixed costs and the assumption that average sales are the same as in the data (\$79 million in our sample).

Table 8: Calibrated Parameters and Targeted Moments

Parameter			Moment	Model	Data
Fixed Cost per Supplier	F	0.01	Elast. of N w.r.t. Risk	0.18	0.18
Elasticity of Input Quality	τ	0.07	Sales Elast. w.r.t. Ship. Time	-0.14	-0.27
Foreign Input Price	p_M	0.70	Aggregate Import Share	0.28	0.28
Std. Dev. Log Efficiency	σ_φ	0.16	Coef. of Variation Log Sales	0.21	0.22
Corr. High Risk and Efficiency	$\rho_{\varphi\sigma}$	-1.30	Corr. Log Sales and Risk	-0.22	-0.26
Std. Dev. of Shipping Times (High)	σ_{dH}	0.28	Avg. $\widehat{StdTime}$ above median	0.28	0.28
Std. Dev. of Shipping Times (Low)	σ_{dL}	0.09	Avg. $\widehat{StdTime}$ below median	0.09	0.09
Expected Log Shipping Days (High)	μ_{dH}	2.91	Average Shipping Days	19.1	19.1
Expected Log Shipping Days (Low)	μ_{dL}	2.95	Average Shipping Days	19.1	19.1
<i>Not calibrated:</i>			<i>Not targeted:</i>		
Elasticity Domestic-Foreign Inputs	ε	2.38	Elast. of Imports w.r.t. Risk	-0.16	-0.09
Demand Elasticity	σ	3.83			
Output Elast. w.r.t. Materials	γ	0.58			

Notes: The elasticities are coefficient estimates of cross-sectional regressions of $\log N$, \log sales, and \log import value on \log of the standard deviation of shipping days (controlling for firm efficiency), or \log shipping times. The aggregate import share is the fraction of material expenditure accounted by foreign inputs. $\widehat{StdTime}$ is the measure of risk defined in Section 3.2. The moments in the data are measured in the 2011-2019 period. Sources: U.S. Census Bureau and authors' calculations.

5.2 Counterfactual Analysis

Armed with the calibrated model, we assess the impacts of two scenarios of heightened risk: climate change and port congestion. We also study the consequences of a complete removal of shipping time risk.

Climate Change. We simulate an increase in weather volatility due to climate change over the next 50 years under the assumption that future weather conditions will continue to follow their historical trend. We use the matched dataset of shipping routes and weather conditions from our empirical analysis, but extend the weather data to 2023 to capture a trend over a longer period. We use daily data to compute the standard deviation of wave height across days for each year and location, and then compute the annual growth rate of these standard deviations between 2011 and 2023 (see Figure 1 in the Introduction). The average annual growth rate across all locations is 0.34%. This pattern is consistent with the findings of Young et al. (2011) for 1985-2008, and with work suggesting an increasing

likelihood of extreme wave heights (e.g., [Shi et al. \(2024\)](#)).²⁸ Compounding this growth over 50 years, we obtain a long-run growth rate in the standard deviation of wave height of 18.5%. Assuming that the log-linear model of weather conditions of Section 2.4 is stable over time, this implies an equivalent growth rate in the standard deviation of log shipping times, which we feed into the model.²⁹

Port Congestion. In the second exercise, we evaluate the economic effects of the greater variability of waiting times at ports due to the rise of port congestion that occurred globally in the aftermath of the Covid pandemic. We capture congestion using the Average Congestion Time (ACT) measure developed by [Bai et al. \(2024\)](#), which reflects the average number of hours a container ship waits at port before docking at the berth for the top-50 container ports worldwide. Starting in late 2020 and until the end of 2022, both the level and the volatility of waiting times increased as strains on global supply chains intensified. In particular, using monthly data between January of 2017 and June of 2024, we find that the standard deviation of the ACT was 52% higher in the January 2021-December 2022 period than in the rest of the sample. We feed this shock into our model as a change in the standard deviation of shipping times of 52% for all firms.

Results. Table 9 summarizes the aggregate effects of the counterfactuals on the U.S. manufacturing sector. The two scenarios feature qualitatively similar results, although the port congestion simulation implies a stronger impact due to the larger magnitude of the shock. The average number of suppliers grows by 18.6% (31.9%) in the climate (port congestion) counterfactual, as firms diversify the increased risk of international shipping delays. The increase in N is larger for firms with high-risk suppliers, which on average raise N by 42% (71.5%).

Total manufacturing imports fall by 1.3% (2.8%), corresponding to a decline in U.S. imports of about \$15.4 (\$33.1) billion dollars relative to 2019.³⁰ The decline in imports is stronger for high risk firms, which on average contract their imports by about -4% (-8.8%). As importers substitute risky foreign inputs with domestic ones, production costs and prices rise, with the price index growing by 0.6% (1.4%). Despite an increase in total domestic

²⁸Work in the oceanography literature typically does not report changes in the standard deviation of wave height and instead focuses on the mean and the 99th percentile. [Young et al. \(2011\)](#) find no annual change in the average wave height, an annual increase of 0.25% of the 90th percentile, and an annual increase of 0.50% of the 99th percentile, suggesting a moderate annual increase in the dispersion of wave height.

²⁹In all counterfactuals, we hold the average α_i constant across risk types, effectively feeding into the model a mean preserving spread to supplier qualities.

³⁰The U.S. imported \$1.17 trillion in 2019, excluding consumption and capital goods, based on import data from the U.S. Census Bureau.

Table 9: Counterfactuals

Growth in (%)	Climate Change	Port Congestion	Removing Risk
Average N	18.58	31.86	-11.06
Total Import Value	-1.32	-2.83	5.21
Import Share	-1.70	-2.53	7.29
Price Index	0.63	1.45	-2.93
Total Domestic Spending	0.42	0.87	-1.96
Real Income	-0.49	-1.13	2.27

Notes: The table reports aggregate statistics associated with (i) an increase in the standard deviation of log shipping times of 18.5% (first column), (ii) an increase in the standard deviation of shipping times of 52% (second column), and (iii) the elimination of risk (third column).

spending stemming from the shift towards domestic input production, U.S. real income is reduced by 0.5% (1.1%).³¹

Removing Shipping Time Risk. We conclude this section by quantifying the effects of a complete removal of shipping time risk (Table 9, column 3). When risk is removed, there is a 11.1% reduction in the average number of foreign suppliers used by importers. This happens because, without risk, foreign suppliers are as safe as the domestic ones and firms save on the fixed costs of foreign sourcing. The removal of shipping uncertainty also implies an increase of 5.2% in aggregate imports, or \$61 billion relative to 2019, and an increase of 7.3% in the aggregate import share. Total domestic spending falls as there is lower demand for the domestically produced input. Overall, removing shipping risk lowers production costs and prices, resulting in a 2.9% reduction in the price index and a 2.3% increase in U.S. real income.

6 Conclusions

In this paper, we combine U.S. Census shipment-level data with information on ocean wave conditions along shipments' maritime routes to construct a novel measure of weather-based supply chain risk. We document substantial negative effects of shipping delays on firms' sales, profits, and employment, and study how exposure to shipping time risk correlates with the pattern of import demand of U.S. manufacturing firms at the intensive and extensive margins. Our results show that U.S. importers that are more exposed to shipping time

³¹In Appendix C.6, we examine the counterfactual effects on sales, profits, and employment and compare them to a back-of-the-envelope calculation based on the reduced-form coefficients from Table 4.

volatility subsequently feature lower imports, a larger number of routes and suppliers, and a lower concentration of expenditure across routes and suppliers, which indicates that firms actively diversify this source of risk. To rationalize this evidence, we introduce risky delivery times into a quantitative model of firm-level importing. We show that increases in shipping risk associated with climate change and port congestion can have significant impacts on supply chains and welfare.

Our findings carry relevance at a time of increasing climate and geopolitical risk, and contribute to a rapidly growing literature discussing the implications of increasing fragmentation, re-shoring, and supply chain diversification. Our results suggest that there may be limits to firms' willingness to concentrate their sourcing too strongly on any one country or region if it comes at the expense of higher delivery risk. Shedding more light on the dynamics of supplier selection and on how firms adjust their supplier portfolio as they grow remain important questions for further research.

References

- Adão, R., Arkolakis, C., Esposito, F., 2025. General equilibrium indirect effects in space: Theory and measurement. NBER Working Paper No. 25544.
- Alessandria, G., Khan, S.Y., Kheralarian, A., Mix, C., Ruhl, K.J., 2023. The aggregate effects of global and local supply chain disruptions: 2020-2022. *Journal of International Economics* 146, 103788.
- Alley, R.B., Emanuel, K.A., Zhang, F., 2019. Advances in weather prediction. *Science* 363, 342--344.
- Antràs, P., Chor, D., 2022. Global value chains. *Handbook of International Economics* 5, 297--376.
- Antràs, P., Fort, T.C., Tintelnot, F., 2017. The Margins of Global Sourcing: Theory and Evidence from US Firms. *American Economic Review* 107, 2514--2564. doi:[10.1257/aer.20141685](https://doi.org/10.1257/aer.20141685).
- Bai, X., Fernández-Villaverde, J., Li, Y., Zanetti, F., 2024. The causal effects of global supply chain disruptions on macroeconomic outcomes: Evidence and theory. NBER Working Paper No. 32098.
- Balboni, C., Boehm, J., Waseem, M., 2023. Firm adaptation in production networks: Evidence from extreme weather events in pakistan Mimeo.
- Baley, I., Veldkamp, L., Waugh, M., 2020. Can global uncertainty promote international trade? *Journal of International Economics* 126, 103347.
- Barrot, J.N., Sauvagnat, J., 2016. Input specificity and the propagation of idiosyncratic shocks in production networks. *The Quarterly Journal of Economics* 131, 1543--1592.
- Bartik, T.J., 1991. Who benefits from state and local economic development policies? WE Upjohn Institute for Employment Research.
- Bernard, A.B., Moxnes, A., Saito, Y.U., 2019. Production networks, geography, and firm performance. *Journal of Political Economy* 127, 639--688.
- Blaum, J., 2024. Global firms in large devaluations. *The Quarterly Journal of Economics* 139, 2427--2474.
- Blaum, J., Lelarge, C., Peters, M., 2018. The Gains from Input Trade with Heterogeneous Importers. *American Economic Journal: Macroeconomics* 10, 77--127.
- Boehm, C.E., Flaaen, A., Pandalai-Nayar, N., 2019. Input linkages and the transmission of shocks: firm-level evidence from the 2011 tōhoku earthquake. *Review of Economics and Statistics* 101, 60--75.
- Brancaccio, G., Kalouptsi, M., Papageorgiou, T., 2020. Geography, transportation, and endogenous trade costs. *Econometrica* 88, 657--691.
- Brancaccio, G., Kalouptsi, M., Papageorgiou, T., 2024. Investment in infrastructure and trade: The case of ports. NBER Working Paper No. 32503.
- Carreras-Valle, M.J., 2024. Increasing inventories: The role of delivery times. Mimeo.
- Carvalho, V.M., Nirei, M., Saito, Y.U., Tahbaz-Salehi, A., 2021. Supply Chain Disruptions: Evidence from the Great East Japan Earthquake. *The Quarterly Journal of Economics* 136, 1255--1321.
- Castro-Vincenzi, J., Khanna, G., Morales, N., Pandalai-Nayar, N., 2024. Weathering the storm: Supply chains and climate risk. NBER Working Paper No. 32218.
- De Loecker, J., 2011. Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity. *Econometrica* 79, 1407--1451. URL: <http://onlinelibrary.wiley.com/doi/10.3982/ECTA7617/abstract>.
- De Sousa, J., Disdier, A.C., Gaigné, C., 2020. Export decision under risk. *European Economic Review* 121,

103342.

- Deloitte, 2024. Restructuring the supply base: Prioritizing a resilient, yet efficient supply chain. URL: <https://www2.deloitte.com/us/en/insights/industry/manufacturing/global-supply-chain-resilience-amid-disruptions.html>.
- Dhyne, E., Kikkawa, A.K., Mogstad, M., Tintelnot, F., 2021. Trade and domestic production networks. *The Review of Economic Studies* 88, 643--668.
- Dunbar, G.R., Steingress, W., Tomlin, B., 2023. Climate variability and international trade. Bank of Canada Staff Working Paper 23-8.
- Esposito, F., 2022. Demand risk and diversification through international trade. *Journal of International Economics* 135, 103562.
- Esposito, F., Hassan, F., 2023. Import competition, trade credit, and financial frictions in general equilibrium. Mimeo.
- Eurosender, 2023. Incoterms for sea freight and ocean transport. URL: <https://www.eurosender.com/en/incoterms/sea-freight>.
- Evans, C.L., Harrigan, J., 2005. Distance, time, and specialization: Lean retailing in general equilibrium. *American Economic Review* 95, 292--313.
- Fieler, A.C., Eslava, M., Xu, D.Y., 2018. Trade, Quality Upgrading, and Input Linkages: Theory and Evidence from Colombia. *American Economic Review* 108, 109--146.
- Fillat, J., Garetto, S., 2015. Risk, return and multinational production. *The Quarterly Journal Economics* 130, 2027--2073.
- Filtz, E., De La Cerda, E.S., Weber, M., Zirkovits, D., 2015. Factors affecting ocean-going cargo ship speed and arrival time, in: *International Conference on Advanced Information Systems Engineering*, Springer. pp. 305--316.
- Fort, T.C., Klimek, S.D., 2018. The Effects of Industry Classification Changes on US Employment Composition. U.S. Census Bureau Center for Economic Studies Working Paper 18-28.
- Ganapati, S., Wong, W.F., Ziv, O., 2024. Entrepot: Hubs, scale, and trade costs. *American Economic Journal: Macroeconomics* 16, 239--278.
- Gervais, A., 2018. Uncertainty, risk aversion and international trade. *Journal of International Economics* 115, 145--158.
- Gopinath, G., Neiman, B., 2014. Trade Adjustment and Productivity in Large Crises. *The American Economic Review* 104, 793--831.
- Grossman, G.M., Helpman, E., Lhuillier, H., 2023. Supply chain resilience: Should policy promote international diversification or reshoring? *Journal of Political Economy* 131, 3462--3496.
- Halpern, L., Koren, M., Szeidl, A., 2015. Imported Inputs and Productivity. *American Economic Review* 105, 3660--3703.
- Handley, K., Limão, N., Ludema, R.D., Yu, Z., 2024. Firm input choice under trade policy uncertainty. *Journal of International Economics* 150, 103909.
- Heckmann, I., 2016. Supply chain risk analytics: Towards supply chain risk analytics: Fundamentals, simulation, optimization. *Springer Fachmedien Wiesbaden* , 131--162.
- Heij, C., Knapp, S., 2015. Effects of wind strength and wave height on ship incident risk: regional trends and seasonality. *Transportation Research Part D: Transport and Environment* 37, 29--39.
- Heiland, I., Moxnes, A., Ulltveit-Moe, K.H., Zi, Y., 2025. Trade from space: Shipping networks and the global implications of local shocks. *Review of Economics and Statistics* , 1--45.

- Heise, S., 2024. Firm-to-firm relationships and the pass-through of shocks: Theory and evidence. *Review of Economics and Statistics* , 1--45.
- Herghelegiu, C., Monastyrenko, E., 2020. Risk- and cost-sharing in firm-to-firm trade. DEM Discussion Paper Series 20-24, Department of Economics at the University of Luxembourg.
- Hummels, D.L., Schaur, G., 2010. Hedging price volatility using fast transport. *Journal of International Economics* 82, 15--25.
- Hummels, D.L., Schaur, G., 2013. Time as a Trade Barrier. *American Economic Review* 103, 2935--2959.
- Johnson, R.C., Noguera, G., 2012. Accounting for intermediates: Production sharing and trade in value added. *Journal of International Economics* 86, 224--236.
- Kamal, F., Krizan, C.J., Monarch, R., 2015. Identifying Foreign Suppliers in US Merchandise Import Transactions. U.S. Census Working Paper No. CES-15-11 .
- Kamal, F., Monarch, R., 2018. Identifying foreign suppliers in u.s. import data. *Review of International Economics* 26, 117--139.
- Khanna, G., Morales, N., Pandalai-Nayar, N., 2022. Supply chain resilience: Evidence from indian firms. NBER Working Paper No. 30689.
- Lafrogne-Joussier, R., Martin, J., Mejean, I., 2023. Supply shocks in supply chains: Evidence from the early lockdown in china. *IMF economic review* 71, 170--215.
- Martin, J., Mejean, I., Parenti, M., 2026. Relationship stickiness, international trade, and economic uncertainty. *Review of Economics and Statistics* 108, 179--193.
- Mishra, J.P., Singh, K., Chaudhary, H., 2022. Intelligent ocean wave height prediction system using light gbm model. *International Journal of Systematic Innovation* 7, 64--77.
- Pankratz, N.M., Schiller, C.M., 2024. Climate change and adaptation in global supply-chain networks. *The Review of Financial Studies* 37, 1729--1777.
- Pierce, J.R., Schott, P.K., 2012. Concording u.s. harmonized system categories over time. *Journal of Official Statistics* 28, 53--68.
- Ramondo, N., Rappoport, V., Ruhl, K.J., 2013. The proximity-concentration tradeoff under uncertainty. *Review of Economic Studies* 80, 1582--1621.
- Ritchie, H., 2024. Weather forecasts have become much more accurate; we now need to make them available to everyone. *Our World in Data* <https://ourworldindata.org/weather-forecasts>.
- Shi, J., Feng, X., Toumi, R., Zhang, C., Hodges, K.I., Tao, A., Zhang, W., Zheng, J., 2024. Global increase in tropical cyclone ocean surface waves. *nature communications* 15, 174.
- Viellechner, A., Spinler, S., 2020. Novel data analytics meets conventional container shipping: predicting delays by comparing various machine learning algorithms. *Proceedings of the 53rd Hawaii International Conference on System Sciences* .
- Young, I., Zieger, S., Babanin, A.V., 2011. Global trends in wind speed and wave height. *Science* 332, 451--455.
- Zhang, S., Jiang, H., Yin, Y., 2022. Characteristics of ocean waves in the southern ocean and seasonal prediction of difficulty for the vessels crossing the westerlies. *Atmosphere* 13, 916.

Online Appendix (Not for Publication)

A Data Construction and Summary Statistics

A.1 Data Construction

In this section, we describe the steps taken to clean the U.S. Census Bureau data.

We prepare the Longitudinal Firm Trade Transactions Database (LFTTD) by first dropping all transactions with an invalid date, zero or negative transaction value, missing vessel name, and cases with a missing importer and exporter ID, as well as transactions that are likely to be incorrectly recorded as indicated by a blooper ID. Second, we drop warehousing transactions and observations where the foreign exporter is recorded as being in the U.S. Third, we use the concordance by [Pierce and Schott \(2012\)](#) to generate time consistent 10-digit Harmonized System (HS-10) codes, and calculate prices as unit values by dividing the value of the shipment by the quantity shipped. Fourth, we translate the nominal shipment values into real values in 2009 prices using the U.S. GDP deflator.

Since the manufacturer ID (MID) differs across establishments of the same firm in different locations and since logistics are likely arranged at the firm-level, we consider MIDs with the same name and country component but with a different street address or city component to belong to the same exporter. Specifically, we replace the MID with a shortened identifier that contains only the country ISO code and the name portion of the ID.³² This approach follows earlier work by [Kamal et al. \(2015\)](#) and [Kamal and Monarch \(2018\)](#). [Kamal et al. \(2015\)](#) compare the number of MIDs in the Census data to the number of foreign exporters for 43 countries from the World Bank’s Exporter Dynamics Database (EDD), which is based on foreign national government statistics and private company data. They show that the number of MIDs in the Census data matches well with the number of sellers in the EDD when the street address or the city component are omitted. [Kamal and Monarch \(2018\)](#) provide further support that the MID is a good identifier of foreign exporters.

The LFTTD also contains an indicator for whether a transaction is conducted between related parties. Based on Section 402(e) of the Tariff Act of 1930, a related-party trade is an import transaction between parties with “any person directly or indirectly, owning, controlling, or holding power to vote, [at least] 6 percent of the outstanding voting stock or shares of any organization.” To correct for missing or incorrect related-party flags, we

³²While different establishments may have different efficiencies or distances to the port, we only observe port-to-port shipping times.

classify an importer-exporter pair as related if it had a related-party flag for any transaction in the given year. Our final dataset includes all related-party trade, and we include the type of relationship (arms' length or related party) as one of the dimensions in the fixed effects in specification (3).

We combine the import data with the Longitudinal Business Database (LBD) by merging both datasets at the firm-year level. We collapse the plant-level observations in the LBD to that level and construct each firm's main industry in each year as the 6-digit NAICS code associated with the highest employment. We use the time-consistent industry codes constructed by [Fort and Klimek \(2018\)](#).

We also merge into our dataset firm-year level observations from the Census of Manufactures (CMF) and from the Annual Survey of Manufacturing (ASM). The CMF captures the universe of all manufacturing plants in census years (2012 and 2017). The ASM captures a representative sample of manufacturing plants in the remaining years. We drop observations that are not included in official census tabulations due to poor quality, and obtain plant-level information on sales, value added, wages and labor costs (payroll), and beginning-of-year total inventories. We set inventory observations that are imputed to missing, and collapse the data to the firm-year level by summing over the plants belonging to the same firm. We obtain materials costs as sales minus value added, and compute profits as sales minus cost of materials minus payroll.

Trips construction We provide some further details on how we construct vessels' trips. As described in the main text, we sort all transactions involving a given vessel by their foreign departure date. We then take all the vessel's transactions and split them into trips using the arrival date in the U.S. and the export departure date abroad. Specifically, for each trip we find the earliest arrival date of the vessel in the U.S., and assign transactions with a later export departure date abroad to a new trip. In some cases, however, the foreign departure or U.S. arrival date is possibly misreported. For example, if transactions 1-5 depart abroad on June 22 and arrive in the U.S. on July 5, transaction 6 departs on June 23 and purportedly arrives on June 24, and transactions 7-10 depart on June 25 and arrive on July 5, then the procedure described above would assign transactions 1-6 to one trip and transactions 7-10 to another, even though almost all shipments arrive on the same day in the U.S. It seems likely that the arrival date for transaction 6 is misreported. We therefore re-combine some of the previously separated trips. For each of the trips assigned in the first step, we compare the latest importation date in the U.S. to the earliest departure date abroad of the *next* trip. If the earliest departure date abroad of the next trip is before the latest importation date of the

earlier trip, then the two trips must have been part of the same journey and we recombine these trips into one. We again iterate through this procedure until no more trips can be combined. Our resulting final dataset contains trips with completely non-overlapping foreign departure and importation dates for each vessel.

A.2 The Determinants of Shipping Times

In this section we examine the factors affecting vessel-borne shipping times that motivate the inclusion of some of the fixed effects in specification (3).

Table A.1 presents the average shipping times and their standard deviation for vessel-based shipments by origin. Shipments from Latin America and Canada tend to arrive fastest in the U.S., while shipments from Oceania and Africa take the longest. There is a large standard deviation of shipping times for all source countries.

Table A.1: Vessel-Based Shipping Times by Origin

	(1)	(2)	(3)
	Avg. Time	Std. Time	Total Value (\$Bill.)
Canada	9.96	22.22	27
Latin America	7.63	27.08	140
South America	21.78	18.65	111
Europe	16.10	18.06	585
Asia	20.37	20.12	1,130
Oceania	29.13	15.97	25
Africa	30.22	23.69	54
Other	25.16	23.73	19

Source: LFTTD. Table summarizes the distribution of shipping time and value across different origin regions for vessel-borne shipments. Values are reported in billions of 2009 dollars.

We next investigate the role of the shipping route more formally by regressing each transaction’s log shipping time separately on fixed effects for the U.S. port of entry (p_i), foreign port of departure (p_e), and the port combination. These regressions yield an R^2 of 0.19, 0.33, and 0.62, respectively, indicating that the port of origin - port of arrival pair explains nearly two thirds of the variation. Replacing the port pair with pair-by-month fixed effects raises the R^2 to 0.69.

We finally analyze the role played by the season of the year, related party status, shipping weight, and charges by running regressions of log shipping time for vessel-based shipments on

Table A.2: Factors Affecting Vessel-Based Shipping Times

Dep. Var.: Log Shipping Times	(1)	(2)	(3)	(4)	(5)
Q2	-0.036*** (0.000)				
Q3	-0.037*** (0.000)				
Q4	-0.024*** (0.000)				
Related-Party		0.006*** (0.000)			
Log Shipment Weight			0.008*** (0.000)		0.007*** (0.000)
Log Shipping Charges				0.005*** (0.000)	0.001*** (0.000)
R^2	0.616	0.616	0.616	0.616	0.617
Port Pair FE	Y	Y	Y	Y	Y
Observations (thousands)	18,260	18,260	18,260	18,260	18,260

Notes: The unit of observation is an importer (f) - exporter (x) - HS10 (h) - vessel (v) - foreign country (c) - origin port (p_e) - destination port (p_i) - foreign export date (t_e) - importation date (t_i) combination. Rows 1, 2 and 3 represent quarter fixed effects in Column (1). Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the country level.

these characteristics. We include fixed effects for the port pair in all regressions. We present the results in Table A.2. In column 1, we test whether seasonality affects shipping times by adding dummies for each quarter of the year. Shipping times are nearly 4% shorter in the summer quarters of the northern hemisphere, highlighting the potential role of weather in affecting shipping routes. In column 2, we find that related party transactions have slightly longer shipping times relative to arms-length transactions. The next columns find a positive relationship between shipping time and shipment weight, and also a positive one between shipping charges and delivery times.

A.3 Analyzing Vessel Movements with AIS Data

In this section, we provide some further analysis of our constructed routes and of the effect of weather conditions on shipping times using satellite Automatic Identification System (AIS) data.

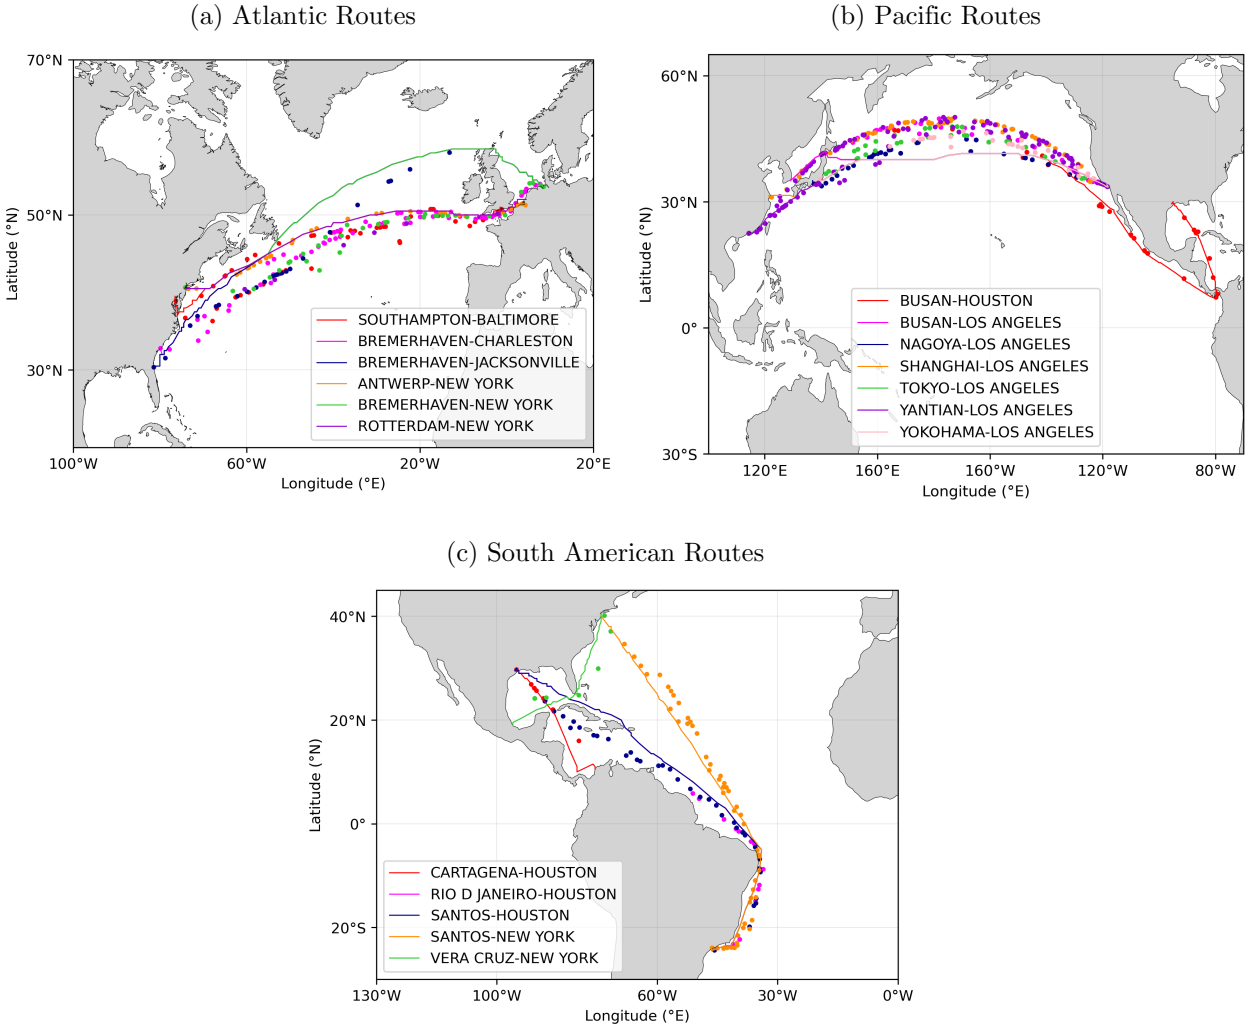
Vessel movements vs. routes As discussed in the main text, we construct the route taken by vessels between ports across the ocean using Eurostat’s SeaRoute program. We compare a subset of these constructed routes to actual vessel movements using AIS data from MarineTraffic, a provider of ship tracking and maritime analytics services.³³ These

³³See <https://www.marinetraffic.com>.

data report the precise location of vessels on the oceans based on transceiver signals of ships. We downloaded detailed geolocations with time stamps for 74 vessels traveling on 19 routes between July 10 and July 16, 2023 and between August 21 and August 28, 2023. For each vessel, we obtained origin and destination port as well as speed, direction, and weather at different locations (latitudes and longitudes) with detailed time stamps along the route. We obtained on average 121 different observations for each vessel along its route, with a range of between 13 to 373 data points. While we have historical data for each vessel, allowing us to observe each vessel from its departure port, the limited time of our data access did not permit us to observe each vessel until its arrival at the destination. On average, we observe 83% of a vessel's full voyage from origin to destination (median: 90%).

Figure A.1a plots some routes across the Atlantic from the SeaRoute program against the observed locations of vessels traveling on these routes. The vessel locations are reasonably close to the routes, though not perfect. In particular, the SeaRoute program suggests that vessels traveling between Bremerhaven and New York mostly follow a route to the North of the United Kingdom, while the vessels we observe making the journey between these ports followed a route to the South of Britain. In contrast, some vessels traveling to Jacksonville followed the Northern route. Figure A.1b plots the routes across the Pacific. Here vessels are closer to the routes near the end points, but follow a more Northern trajectory in the middle. Finally, Figure A.1c plots vessel movements on South American routes. Overall, the analysis shows that vessels tend to broadly follow the routes from the SeaRoute program, but that there is substantial variation. This variation will introduce measurement error into our weather variable that will bias our results towards zero.

Figure A.1: Vessel Movements vs Routes



Source: MarineTraffic and authors' calculations. Notes: The figure shows the locations of vessels traveling between selected ports against the routes from the SeaRoute program used for the analysis.

Vessel Speed and Weather Conditions We examine the effect of weather conditions on vessel speed in the AIS data. MarineTraffic provides for each vessel at each recorded location the course and speed, as well as the wind speed, wind angle, wave height, and wave direction. We can therefore run similar regressions in these data as in the Census data to see whether we find similar effects. Since we observe the vessel speed at each location, we use this variable rather than the overall shipping time to analyze the contemporaneous effect of weather conditions on speed. Specifically, we estimate:

$$\ln(\text{Speed}_{ijt}) = \beta_1 \text{Height}_{ijt} + \beta_2 \text{Direction}_{ijt} + \beta_3 \text{Height}_{ijt} \cdot \text{Direction}_{ijt} + \gamma_i + \epsilon_{ijt},$$

where i indexes the vessel, j the location, and t is the time stamp. Here, Speed_{ijt} is the speed of the vessel at location j and time t , Height_{ijt} is the height of the waves, Direction_{ijt} is the wave direction relative to the direction of travel (where zero indicates that the waves are in the direction of travel), and γ_i are vessel fixed effects.³⁴ The first column of Table A.3 presents the results. As in the Census data, higher waves increase the vessel speed: a one standard deviation increase in significant wave height from the mean increases vessel speed (hence reduces shipping time) by about 4 log points. Also consistent with the Census data, a greater wave angle relative to the direction of travel has a positive effect on speed. When the waves are against the direction of travel (180 degrees), vessel speed is about 13 log points higher, reducing the shipping time.

In the second column of Table A.3 we run a similar regression, but use wind speed and wind direction instead of wave height and direction. Related work on shipping times such as [Filtz et al. \(2015\)](#) also finds a strong relationship between vessel speed and wind speed and direction, which we do not observe in the WaveWatch III data. Wind speed should be positively correlated with wave height, and hence we might expect similar results also with respect to this variable. As expected, we find that vessels are faster when wind speed is higher and when the wind is in the opposite direction of the vessel's course. A one standard deviation increase in wind speed (6 knots) raises vessel speed by 1.4 log points. Wind against the direction of travel (180 degrees) increases speed by 4 log points. Overall, the results are therefore consistent with the relationships between wave height and shipping times that we observe with the Census data.

³⁴We observe each vessel only on one route and so these are effectively vessel-route fixed effects.

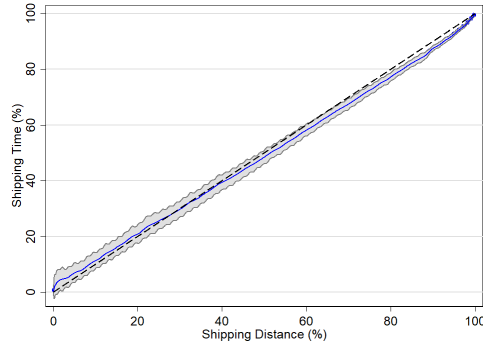
Table A.3: Effect of Weather on Vessel Speed

Dep. Var:	Vessel speed	Vessel speed
Wave Height	0.0584*** (0.0122)	
Direction	0.0007*** (0.0002)	
Wave Height \times Direction	-0.0002** (0.0001)	
Wind Speed		0.0024** (0.0011)
Wind Direction		0.0003** (0.0001)
Wind Speed \times Wind Direction		-0.0000 (0.0000)
Vessel FE	Y	Y
Observations	8,902	8,842

Source: MarineTraffic. Notes: First column shows regression of log vessel speed on wave height and relative wave direction. Second column shows regression of log vessel speed on wind speed and relative wind direction. Direction of zero means that the waves or wind are in the direction of travel.

Vessel Speed In the robustness analysis of our main results, we consider an alternative weather shock measure that is based on vessels' predicted location on the route segments, assuming vessels travel at approximately constant speed. We verify this assumption using the AIS data. We use the 18 vessels for which we observe the entire journey from origin to destination port. For these vessels, we compute at each location the share of the journey completed, in terms of distance, as well as the share of the journey passed in terms of total voyage time. We then plot in Figure A.2 a bin scatter of the distance share against the share of voyage time. Overall, we find that the fit line is approximately on the 45 degree line throughout the journey, indicating that our assumption of constant speed is reasonable. Vessels are slightly slower at departure, and then make up for this delay along the journey before slowing down again near the arrival port.

Figure A.2: Vessel Distance Covered vs Voyage Time Elapsed



Source: MarineTraffic and authors' calculations. Notes: The figure plots the share of the distance completed against the share of voyage time elapsed for 18 vessels for which we have complete voyage information.

A.4 Weather Summary Statistics

In this section we provide some further details on the wave height variable. Table A.4 provides some summary statistics on the mean and standard deviation of the significant wave height and its (absolute) direction across all days in the data. We find that there is substantial variation across both height and direction variable. For example, the mean wave height is 2.6 meters, but at the 95th percentile the wave height is 5.4 meters.

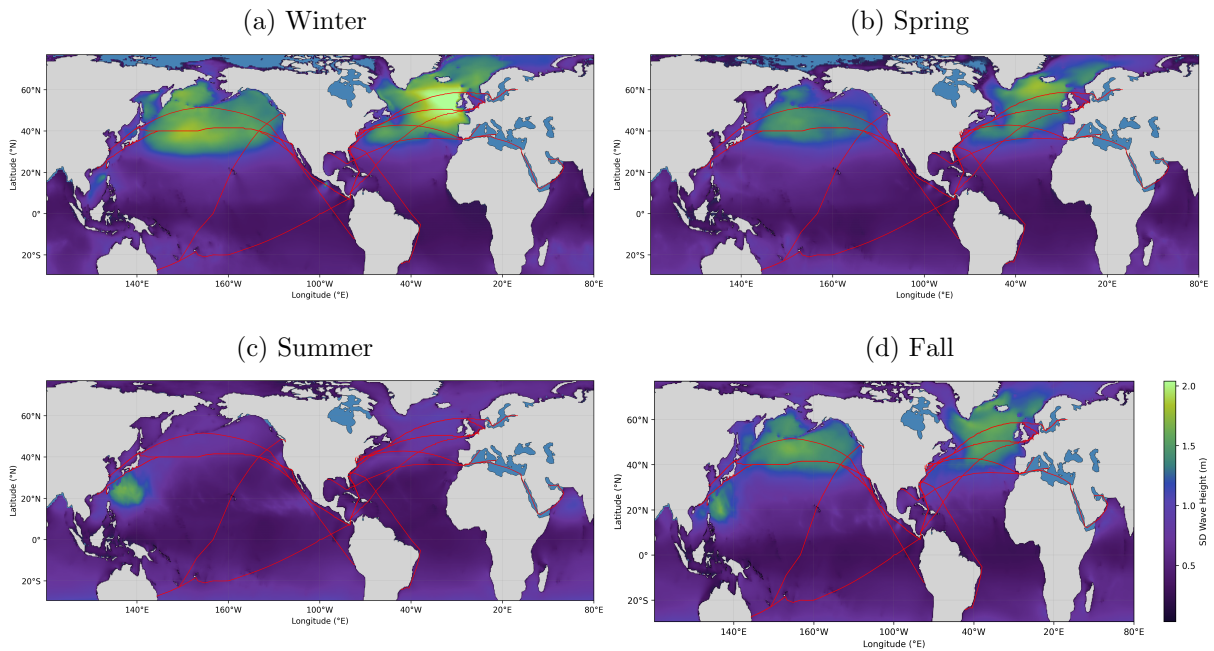
Figure A.3 further shows that there is also significant variation in the standard deviation of wave height across seasons. For example, both the northern Atlantic and the northern Pacific experience significant volatility in wave height in the fall and winter, but less in the summer.

Table A.4: Weather conditions: summary statistics

	Mean	Sd	p1	p5	p50	p95	p99
Significant wave height (m)							
All	2.6	1.5	0.2	0.7	2.3	5.4	7.2
North Atlantic	2.2	1.3	0.3	0.8	1.9	4.8	6.9
South Atlantic	2.7	1.3	0.3	1.1	2.4	5.2	6.8
North Pacific	2.3	1.2	0.2	0.8	2.1	4.5	6.5
South Pacific	2.9	1.4	0.3	1.2	2.6	5.7	7.3
Indian Ocean	2.9	1.6	0.2	0.7	2.7	5.9	7.7
Significant wave direction (degrees)							
All	203	81	20	52	217	319	343
North Atlantic	183	102	11	31	193	333	350
South Atlantic	217	65	31	91	220	316	343
North Pacific	185	98	17	40	194	325	344
South Pacific	216	67	33	83	223	315	339
Indian Ocean	210	63	27	80	219	297	328

Notes: The table shows summary statistics for the weather variables across all days in the data.

Figure A.3: Standard Deviation of Wave Height in Different Seasons



Notes: The figure shows the standard deviation of wave height across all days from 2011-2019 for different seasons. Seasons are based on the Northern Hemisphere, i.e. December, January and February are winter months, March, April, and May are spring months, and so on.

B Additional Results for Empirical Analysis

B.1 Additional Results for Importer Performance

One concern with our baseline methodology of averaging the weather across all locations of each trip segment on each day is that some of these locations may be very far away from the vessel’s current location. We use an alternative approach that estimates vessels’ location on each segment and uses weather only from the surrounding area. Specifically, we decompose each trip segment into smaller sub-segments of 1,000 km of length and assume that vessels travel through these areas at constant speed.³⁵ We then find for each day of the journey the local weather in the vessel’s current sub-segment and average these local weather conditions across the vessel’s journey. We use this variable to run regression (3) of shipping times on weather conditions, and show the results from this regression in Table B.1. We find that the relationships between weather conditions and shipping times are less strong than under our baseline specification, although still highly significant.

We construct an alternative measure of weather-induced shipping times $\tilde{t}^{s,weather,alt}$ using equation (4) as in the main text, and construct weather-induced shipping delays as in Section 3.1. We then re-run regression (6):

$$\ln(Y_{ft}^o) = \alpha + \beta_1 \text{FracDelayed}_{ft}^{alt} + \gamma_f + \delta_t + \epsilon_{ft},$$

where $\text{FracDelayed}_{ft}^{alt}$ is the fraction of imports subject to shipping delays constructed with the alternative measure. Table B.2 shows the results. The coefficients are similar as with the measure in the main text. A one standard deviation increase in delayed shipments (2.04 percentage points) is associated with a drop in sales of 5.4%, a decline in profits of 3.8%, and a fall in employment of 0.9%.

As an additional robustness check, we re-run regression (3) without the weather terms, and then use the residual from this regression as our measure of unanticipated shipping times to construct delays. This measure picks up delays beyond those due to weather, such as delays due to port delays or strikes, but requires that the delays we capture with these residuals are unanticipated by importers. We find relatively similar coefficients to the main text, but larger overall effects due to the greater standard deviation of this delay measure. A one standard deviation increase in the fraction of delayed shipments, 0.031, is associated with a drop in sales by 8.3%, a fall in profits by 4.7%, and in the number of employees by 1.6%.

³⁵We use AIS data in Appendix A.3 to show that this assumption approximately holds in vessel tracking data.

Table B.1: Effect of Weather on Shipping Times – Alternative Weather Conditions

Dep. Var:	$\ln(T_s)$		
Wave Height ^s	-0.029*** (0.000)	-0.029*** (0.000)	-0.026*** (0.000)
Direction ^s		0.006*** (0.000)	0.013*** (0.000)
Wave Height ^s × Direction ^s			-0.004*** (0.000)
R-Squared	0.788	0.788	0.788
Observations	9,892,000	9,892,000	9,892,000

Notes: The table reports the coefficients on the weather terms from estimating specification (3), using the weather at a vessel’s imputed location. Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. The variable wave height is expressed in meters, while the variable direction is expressed in hundreds of degrees. Importer, HS10 product, exporter, route-time, vessel, and related party fixed effects, as well as log charges and log weight, are included in the regression but not reported in the table for brevity. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

Table B.2: Effect of Extreme Delays on Firms’ Outcomes (Alternative Measure)

	(1)	(2)	(3)
Weather Shocks (Alternative Measure)			
Dependent Variable (in logs):	Sales	Profits	Employees
Frac Delayed	-2.627*** (0.309)	-1.862*** (0.354)	-0.444*** (0.153)
Importer FE	Y	Y	Y
Year FE	Y	Y	Y
R-Squared	0.97	0.90	0.98
Observations	78,000	78,000	78,000

Notes: The table reports the coefficients on the fraction of imports delayed, $FracDelayed_{ft}^{alt}$, from specification (6), using the alternative measure of weather-induced shipping times, $\tilde{i}^{s,weather,alt}$, to compute the fraction delayed. Fixed effects for the firm and the year are included. Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the firm level. Mean of $FracDelayed_{ft}^{weather}$ is 0.0031 and standard deviation is 0.0204. R^2 is the overall fit inclusive of the fixed effects. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

Table B.3: Effect of Extreme Delays on Firms' Outcomes (Without Weather)

	(1)	(2)	(3)
Dependent Variable (in logs):	Sales	Profits	Employees
Frac Delayed	-2.665*** (0.219)	-1.511*** (0.248)	-0.515*** (0.117)
Importer FE	Y	Y	Y
Year FE	Y	Y	Y
R-Squared	0.97	0.90	0.98
Observations	78,000	78,000	78,000

Notes: The table reports the coefficients on the fraction of imports delayed, $FracDelayed_{ft}$, from specification (6), using the alternative measure of shipping times based on the residual from specification (3) without the weather terms to compute the fraction delayed. Fixed effects for the firm and the year are included. Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the firm level. Mean of $FracDelayed_{ft}$ is 0.0060 and standard deviation is 0.0310. R^2 is the overall fit inclusive of the fixed effects. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

B.2 Robustness of the Risk Regressions

In this section, we show that our results on the impact of shipping risk on importers' behavior from Section 3.2 are robust to a variety of alternative specifications.

First, we use the alternative measure of weather-induced shipping times introduced in Appendix B.1. As described above, we decompose each trip segment into smaller sub-segments, assume that vessels travel through these areas at constant speed, and determine for each day of a vessel's journey the local weather conditions in a vessel's current location. We then compute the average weather a vessel experienced on its journey and use this variable to run regression (3) of shipping times on weather conditions. We use the weather-induced shipping times based on this regression, $\tilde{t}^{s,weather,alt}$, to construct our risk measure following the same steps as described in Section 3.2. Table B.4 shows the results from running regression (8) with this alternative risk measure. The effects are similar to before: going from the 25th to the 75th percentile of the risk distribution (61 log points) increases the number of shipping routes used by 9.9%, the number of suppliers by 6.4%, and reduces total imports by 5.7%.

Next, we analyze a broader risk measure that is constructed by taking the standard deviation of the residual from running equation (3) without the weather terms. Table B.5 shows the results of the risk regression (8) with this alternative risk measure. An increase from the 25th to the 75th percentile of the risk distribution (91 log points) increases the number of routes used by 8.4% and the number of suppliers by 5.2%. Moreover, it decreases the route HHI, supplier HHI, and total imports by 4.3%, 3.2%, and 5.6%, respectively.

Third, we replace the separate firm and time fixed effects in our baseline specification (8) with firm-time fixed effects. This specification picks up changes in a firms' sourcing strategy

that are common to all imported goods, and identifies our effect from within-firm variation in sourcing behavior across goods. In this specification, firms that only import a single good are therefore dropped. The results in Table B.6 are similar to the baseline and increase in magnitude for all effects except for total imports. Going from the 25th to the 75th percentile of the weather risk distribution is associated with a 11.5% increase in routes and a 8.0% increase in the number of suppliers. The route HHI, supplier HHI, and import value drop by 5.2%, 4.4%, and 1.5%, respectively.

Fourth, we show that our results also hold when we use only the riskiness of the importers' main supplier to construct our measure of risk, rather than a weighted average across all suppliers. Specifically, we compute the risk measure

$$\widehat{StdTime}_{fht-3,t-1}^{main} \equiv \sum_{r \in R(x^{main})} \omega_{fxhr,t-3,t-1}^{x^{main}} \widehat{StdTime}_{xhrt-3,t-1}, \quad (22)$$

where the main supplier of good h is defined as the one with the largest shipment value to importer f in the years $t - 3$ to $t - 1$. The weighted average is now taken only across the routes r used by the main supplier, $R(x^{main})$ and the weights $\omega_{fxhr,t-3,t-1}^{x^{main}}$ are the import shares of each route for the main supplier. The results of running specification (8) with this measure are in Table B.7. Going from the 25th to the 75th percentile of shipping risk (42 log points) increases the number of routes and suppliers by 3.5% and 2.5%, respectively, while the route HHI, supplier HHI, and total import value fall by 2.1%, 2.1%, and 5.0%. Since we are not able to compute shipping risk for all main suppliers (for example because they have fewer than 10 transactions), the number of observations drops relative to the other regressions.

Table B.8 additionally includes firms with only one supplier in our main specifications, which are dropped in the baseline. These firms are therefore by definition not diversified in t . Here, we find that the relationship between shipping risk and the number of suppliers and routes remains significantly positive, and increases slightly in magnitude. All results remain strongly significant.

Finally, in Table B.9 we include an additional control for the inventory-sales ratio. We obtain the beginning of year value of the total inventory of materials for all firms in each census year from the CMF, and for a subset of firms from the ASM in all other years. These inventories contain domestically sourced supplies, and are therefore only a proxy of the inventory of imported inputs. We find that the relationship between shipping risk and the number of suppliers and routes strengthens once we include the inventory control. However, conditional on shipping risk, a higher inventory-sales ratio decreases the number of suppliers

used.

Table B.4: Alternative Measure of Weather Risk

	(1)	(2)	(3)	(4)	(5)
Dep. Var.:	Number of Routes	Number of Suppliers	HHI over Routes	HHI over Suppliers	Value Imported
Std Time	0.162*** (0.009)	0.105*** (0.008)	-0.083*** (0.003)	-0.067*** (0.003)	-0.093*** (0.011)
Importer FE	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y
R-squared	0.72	0.67	0.46	0.41	0.88
Observations	129,000	129,000	129,000	129,000	129,000

Notes: The table shows the coefficients on the standard deviation of the weather-induced shipping times, $\widehat{StdTime}_{fht-3,t-1}^{weather,alt}$ from specification (8). The risk measure is constructed using the local weather conditions around a vessel's predicted location on each day, as described in the text. We consider several dependent variables: the log of the number of routes (col. 1), the log of the number of foreign suppliers (col. 2), log HHI across routes (col. 3), log HHI across suppliers (col. 4), and the log total value imported (col. 5). Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the firm level. Regression includes controls for three-year average of weather-induced shipping times, average unit value paid, the importer's total imports over the previous three years, and the suppliers' total exports to the U.S. over the previous three years. For brevity, the table does not report the coefficients of these regressors. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

Table B.5: Broader Risk Measure

	(1)	(2)	(3)	(4)	(5)
Dep. Var.:	Number of Routes	Number of Suppliers	HHI over Routes	HHI over Suppliers	Value Imported
Std Time	0.093*** (0.006)	0.057*** (0.006)	-0.048*** (0.002)	-0.035*** (0.002)	-0.062*** (0.007)
Importer FE	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y
R-squared	0.71	0.67	0.45	0.41	0.88
Observations	130,000	130,000	130,000	130,000	130,000

Notes: The table shows the coefficients on the standard deviation of the residualized shipping times, $\widehat{StdTime}_{fht-3,t-1}^{alt}$ from specification (8). The risk measure is computed using the standard deviation of the residuals obtained from estimating equation (3) without the weather terms. We consider several dependent variables: the log of the number of routes (col. 1), the log of the number of foreign suppliers (col. 2), log HHI across routes (col. 3), log HHI across suppliers (col. 4), and the log total value imported (col. 5). Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the firm level. Regression includes controls for three-year average of weather-induced shipping times, average unit value paid, the importer's total imports over the previous three years, and the suppliers' total exports to the U.S. over the previous three years. For brevity, the table does not report the coefficients of these regressors. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

Table B.6: Firm-Time Fixed Effects

	(1)	(2)	(3)	(4)	(5)
Dep. Var.:	Number of Routes	Number of Suppliers	HHI over Routes	HHI over Suppliers	Value Imported
Std Time	0.227*** (0.010)	0.158*** (0.009)	-0.103*** (0.004)	-0.087*** (0.004)	-0.030*** (0.010)
Importer-Year FE	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y
R-squared	0.75	0.69	0.51	0.45	0.90
Observations	114,000	114,000	114,000	114,000	114,000

Notes: The table shows the coefficients on the standard deviation of the weather-induced shipping times, $\widehat{StdTime}_{fht-3,t-1}$, from specification (8) for different dependent variables: the log of the number of routes (col. 1), the log of the number of foreign suppliers (col. 2), log HHI across routes (col. 3), log HHI across suppliers (col. 4), and the log total value imported (col. 5). The specification is run using firm-time fixed effects, rather than firm and time fixed effects separately, in addition to product fixed effects. Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the firm level. Regression includes controls for three-year average of weather-induced shipping times, average unit value paid, the importer's total imports over the previous three years, and the suppliers' total exports to the U.S. over the previous three years. For brevity, the table does not report the coefficients of these regressors. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

Table B.7: Risk Measure Based on Main Supplier

	(1)	(2)	(3)	(4)	(5)
Dep. Var.:	Number of Routes	Number of Suppliers	HHI over Routes	HHI over Suppliers	Value Imported
Std Time	0.085*** (0.013)	0.061*** (0.009)	-0.051*** (0.005)	-0.051*** (0.003)	-0.120*** (0.011)
Importer FE	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y
R-squared	0.72	0.65	0.51	0.45	0.89
Observations	94,000	94,000	94,000	94,000	94,000

Notes: The table shows the coefficients on the standard deviation of the weather-induced shipping times, $\widehat{StdTime}_{fht-3,t-1}^{main}$, from specification (8) for different dependent variables: the log of the number of routes (col. 1), the log of the number of foreign suppliers (col. 2), log HHI across routes (col. 3), log HHI across suppliers (col. 4), and the log total value imported (col. 5). The specification uses as risk measure the standard deviation of weather-induced shipping times of the main supplier, as described in equation (22). Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the firm level. Regression includes controls for three-year average of weather-induced shipping times, average unit value paid, the importer's total imports over the previous three years, and the suppliers' total exports to the U.S. over the previous three years. For brevity, the table does not report the coefficients of these regressors. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

Table B.8: Including Firms with One Supplier

	(1)	(2)	(3)	(4)	(5)
Dep. Var.:	Number of Routes	Number of Suppliers	HHI over Routes	HHI over Suppliers	Value Imported
Std Time	0.183*** (0.009)	0.116*** (0.008)	-0.090*** (0.003)	-0.067*** (0.003)	-0.073*** (0.009)
Importer FE	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y
R-squared	0.71	0.69	0.49	0.51	0.86
Observations	192,000	192,000	192,000	192,000	192,000

Notes: The table shows the coefficients on the standard deviation of the weather-induced shipping times, $\widehat{StdTime}_{fht-3,t-1}$, from specification (8) for different dependent variables: the log of the number of routes (col. 1), the log of the number of foreign suppliers (col. 2), log HHI across routes (col. 3), log HHI across suppliers (col. 4), and the log total value imported (col. 5). The specification includes importers with only one supplier for a given product per year, in addition to importers with multiple suppliers. Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the firm level. Regression includes controls for three-year average of weather-induced shipping times, average unit value paid, the importer's total imports over the previous three years, and the suppliers' total exports to the U.S. over the previous three years. For brevity, the table does not report the coefficients of these regressors. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

Table B.9: Shipping Time Risk and Import Demand with Inventory-Sales Ratio

	(1)	(2)	(3)	(4)	(5)
Dep. Var.:	Number of Routes	Number of Suppliers	HHI over Routes	HHI over Suppliers	Value Imported
Std Time	0.201*** (0.011)	0.128*** (0.010)	-0.097*** (0.004)	-0.077*** (0.004)	-0.088*** (0.013)
Inventory-Sales Ratio	-0.068 (0.065)	-0.104** (0.042)	0.002 (0.016)	0.029*** (0.010)	-0.306** (0.128)
Importer FE	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y
R-squared	0.71	0.67	0.44	0.41	0.89
Observations	91,500	91,500	91,500	91,500	91,500

Notes: The table shows the coefficients on the standard deviation of the weather-induced shipping times, $\widehat{StdTime}_{fht-3,t-1}$, from specification (8) for different dependent variables: the log of the number of routes (col. 1), the log of the number of foreign suppliers (col. 2), log HHI across routes (col. 3), log HHI across suppliers (col. 4), and the log total value imported (col. 5). The specification includes the inventory-to-sales ratio as additional control, computed as beginning of year inventories divided by sales for the firm in year t from the CMF or ASM. Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the firm level. Regression includes controls for three-year average of weather-induced shipping times, average unit value paid, the importer's total imports over the previous three years, and the suppliers' total exports to the U.S. over the previous three years. For brevity, the table does not report the coefficients of these regressors. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

C Additional Results for Section 4

C.1 Derivation of Expression (14)

Firms choose their inputs in two stages. In the first, firms choose N , x and x_D under uncertainty about the shipping times of their foreign inputs. After the uncertainty is realized, firms choose the level of labor conditional on the choices for N , x and x_D . Combining the production function (10) and the expression for revenues (13), realized firm profits in the second stage are given by

$$\hat{\pi}_f = \varphi_f^{\frac{\sigma-1}{\sigma}} l^{(1-\gamma)\frac{\sigma-1}{\sigma}} \left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + \left(\sum_{i=1}^N \hat{\alpha}_i \right)^{\frac{\varepsilon-1}{\varepsilon}} x^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\gamma \frac{\varepsilon-1}{\varepsilon} \frac{\sigma-1}{\sigma}} P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma} - p_D x_D - N p_M x - w l - w N F, \quad (23)$$

where we have used that in stage 1 firms choose $x_i = x$ since suppliers are ex-ante symmetric, and $\{\hat{\alpha}_i\}$ are the realized supplier qualities which depend on the realized shipping times. Given input choices from the first stage (x_D, x, N) , and realizations of qualities $\{\alpha_i\}$, the optimal choice of labor is:

$$\hat{l}_f = \left[\tilde{\gamma} \varphi_f^{\frac{\sigma-1}{\sigma}} \left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + \left(\sum_{i=1}^N \hat{\alpha}_i \right)^{\frac{\varepsilon-1}{\varepsilon}} x^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\gamma \frac{\varepsilon-1}{\varepsilon} \frac{\sigma-1}{\sigma}} P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma} w^{-1} \right]^{\frac{1}{1-\tilde{\gamma}}} \quad (24)$$

where $\tilde{\gamma} \equiv (1-\gamma)\frac{\sigma-1}{\sigma}$. The associated optimal profits are given by:

$$\chi_f \left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + \left(\sum_{i=1}^N \hat{\alpha}_i \right)^{\frac{\varepsilon-1}{\varepsilon}} x^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\psi} - p_D x_D - N p_M x - w N F, \quad (25)$$

where $\chi_f \equiv \left(\varphi_f^{\frac{\sigma-1}{\sigma}} P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma} \right)^{\frac{1}{1-\tilde{\gamma}}} w^{-\frac{\tilde{\gamma}}{1-\tilde{\gamma}}} \left[(\tilde{\gamma})^{\frac{\tilde{\gamma}}{1-\tilde{\gamma}}} - (\tilde{\gamma})^{\frac{1}{1-\tilde{\gamma}}} \right]$ and $\psi \equiv \gamma \frac{\varepsilon-1}{\varepsilon} \frac{\sigma-1}{\sigma} \frac{1}{1-\tilde{\gamma}}$. Taking expectations of this expression yields equation (14) in the main text.

C.2 Derivations for Section 4.3

Derivation of expression (15). In the case without the domestic input, given the realization of input qualities, firm profits are

$$\pi = P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma} \left(\varphi_f l^{1-\gamma} \left(\sum_{i=1}^N \alpha_i x \right)^\gamma \right)^{\frac{\sigma-1}{\sigma}} - wFN - Np_M x - wl.$$

Letting $\tilde{\gamma} \equiv (1 - \gamma) \frac{\sigma-1}{\sigma}$, the first order condition with respect to labor is

$$\tilde{\gamma} P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma} \varphi_f^{\frac{\sigma-1}{\sigma}} l^{\tilde{\gamma}-1} \left(\sum_{i=1}^N \alpha_i x \right)^{\gamma \frac{\sigma-1}{\sigma}} = w$$

Solving for labor and plugging back into profits yields:

$$\pi = \left(\tilde{\gamma}^{\frac{\tilde{\gamma}}{1-\tilde{\gamma}}} - \tilde{\gamma}^{\frac{1}{1-\tilde{\gamma}}} \right) \left(P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma} \varphi_f^{\frac{\sigma-1}{\sigma}} \right)^{\frac{1}{1-\tilde{\gamma}}} \left(\sum_{i=1}^N \alpha_i x \right)^{\gamma \frac{\sigma-1}{\sigma} \frac{1}{1-\tilde{\gamma}}} w^{-\frac{\tilde{\gamma}}{1-\tilde{\gamma}}} - wFN - Np_M x.$$

Letting again

$$\chi_f \equiv \left(\tilde{\gamma}^{\frac{\tilde{\gamma}}{1-\tilde{\gamma}}} - \tilde{\gamma}^{\frac{1}{1-\tilde{\gamma}}} \right) \left(\varphi_f^{\frac{\sigma-1}{\sigma}} P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma} \right)^{\frac{1}{1-\tilde{\gamma}}} w^{-\frac{\tilde{\gamma}}{1-\tilde{\gamma}}}$$

and $\tilde{\psi}$ be defined as in the main text, the ex-ante firm profit maximization problem is:

$$\max_{x, N} \chi_f x^{\tilde{\psi}} \mathbb{E} \left(\sum_{i=1}^N \alpha_i \right)^{\tilde{\psi}} - Np_M x - wNF.$$

The first order condition with respect to x gives:

$$\tilde{\psi} x^{\tilde{\psi}-1} \chi_f \mathbb{E} \left(\sum_{i=1}^N \alpha_i \right)^{\tilde{\psi}} - Np_M = 0 \tag{26}$$

Using this condition to eliminate x yields expression (15) in the main text, where $\tilde{\chi}_f \equiv \left(\tilde{\psi}^{\frac{\tilde{\psi}}{1-\tilde{\psi}}} - \tilde{\psi}^{\frac{1}{1-\tilde{\psi}}} \right) p_M^{-\frac{\tilde{\psi}}{1-\tilde{\psi}}} \chi_f^{\frac{1}{1-\tilde{\psi}}}$.

Derivation of expression (16). A second-order approximation of the function $\bar{\alpha}^{\tilde{\psi}}$ around $\mathbb{E}[\alpha]$ yields

$$\bar{\alpha}^{\tilde{\psi}} \approx (\mathbb{E}[\alpha])^{\tilde{\psi}} + \tilde{\psi} (\mathbb{E}[\alpha])^{\tilde{\psi}-1} (\bar{\alpha} - \mathbb{E}[\alpha]) + \tilde{\psi} \frac{(\tilde{\psi} - 1)}{2} (\mathbb{E}[\alpha])^{\tilde{\psi}-2} (\bar{\alpha} - \mathbb{E}[\alpha])^2 \quad (27)$$

Taking expectations on both sides gives

$$\mathbb{E}[\bar{\alpha}^{\tilde{\psi}}] \approx (\mathbb{E}[\alpha])^{\tilde{\psi}} + \tilde{\psi} \frac{(\tilde{\psi} - 1)}{2} (\mathbb{E}[\alpha])^{\tilde{\psi}-2} \mathbb{E}[(\bar{\alpha} - \mathbb{E}[\alpha])^2] \quad (28)$$

since $\mathbb{E}[\bar{\alpha}] = \mathbb{E}[\alpha]$. Noting that $\mathbb{V}[\bar{\alpha}] = \mathbb{E}[(\bar{\alpha} - \mathbb{E}[\alpha])^2]$ and that $\mathbb{V}[\bar{\alpha}] = \mathbb{V}[\alpha]/N$, we plug this into equation (15) to obtain expression (16).

Proof of Proposition 1. For continuous N , the first order condition with respect to N is given by:

$$G(N, \mathbb{V}[\alpha]) \equiv \tilde{\chi} \left((\mathbb{E}[\alpha])^{\tilde{\psi}} - \tilde{\psi} \frac{(1 - \tilde{\psi})}{2} (\mathbb{E}[\alpha])^{\tilde{\psi}-2} \frac{1}{N} \mathbb{V}[\alpha] \right)^{\frac{\tilde{\psi}}{1-\tilde{\psi}}} \frac{\tilde{\psi}}{2} (\mathbb{E}[\alpha])^{\tilde{\psi}-2} N^{-2} \mathbb{V}[\alpha] - wF = 0.$$

Using the implicit function theorem:

$$\frac{\partial N^*}{\partial \mathbb{V}[\alpha]} = - \frac{G_{\mathbb{V}[\alpha]}}{G_N},$$

where $G_y \equiv \partial G / \partial y$. Taking these partial derivatives and plugging them back into the previous expression yields:

$$\epsilon_{N^*, \mathbb{V}[\alpha]} = \frac{\partial N^*}{\partial \mathbb{V}[\alpha]} \frac{\mathbb{V}[\alpha]}{N^*} = - \frac{\left(1 - \frac{\frac{\tilde{\psi}^2}{2} (\mathbb{E}[\alpha])^{\tilde{\psi}-2} \mathbb{V}[\alpha]}{N (\mathbb{E}[\alpha])^{\tilde{\psi}-\tilde{\psi} \frac{(1-\tilde{\psi})}{2}} (\mathbb{E}[\alpha])^{\tilde{\psi}-2} \mathbb{V}[\alpha]} \right)}{\left(\frac{\frac{\tilde{\psi}^2}{2} (\mathbb{E}[\alpha])^{\tilde{\psi}-2} \mathbb{V}[\alpha]}{N (\mathbb{E}[\alpha])^{\tilde{\psi}-\tilde{\psi} \frac{(1-\tilde{\psi})}{2}} (\mathbb{E}[\alpha])^{\tilde{\psi}-2} \mathbb{V}[\alpha]} - 2 \right)} < 1 \quad (29)$$

Note that the denominator is negative as $G_N < 0$ follows from the second order condition at the optimal N . The numerator is positive whenever condition (17) in the main text holds; this establishes the first part of the proposition.

Relying on the first order condition with respect to x in (26), import values are given by:

$$\begin{aligned} Nxp_M &= (p_M)^{-\frac{\tilde{\psi}}{1-\tilde{\psi}}} \left(\tilde{\psi} \chi_f \mathbb{E} \left[(\bar{\alpha})^{\tilde{\psi}} \right] \right)^{\frac{1}{1-\tilde{\psi}}} \\ &\approx (p_M)^{-\frac{\tilde{\psi}}{1-\tilde{\psi}}} \left(\tilde{\psi} \chi_f \left((\mathbb{E}[\alpha])^{\tilde{\psi}} - \tilde{\psi} \frac{(1-\tilde{\psi})}{2} (\mathbb{E}[\alpha])^{\tilde{\psi}-2} \frac{1}{N} \mathbb{V}[\alpha] \right) \right)^{\frac{1}{1-\tilde{\psi}}}, \end{aligned}$$

where the second line uses the approximation in (27). It follows that the effect of a mean preserving spread on import values depends on its effect on the variance of average supplier quality $\mathbb{V}[\alpha]/N$. In turn, this effect is given by:

$$\frac{\partial \mathbb{V}[\alpha]/N}{\partial \mathbb{V}[\alpha]} = \frac{1 - \epsilon_{N^*, \mathbb{V}[\alpha]}}{N} > 0$$

as $\epsilon_{N^*, \mathbb{V}[\alpha]} < 1$ follows from (29). It follows that, regardless of whether the optimal N falls or increases after a mean preserving spread, $\mathbb{V}[\alpha]/N$ increases and import values fall.

C.3 Heterogeneity in the Response to Risk

In this section, we investigate how the importers' response to risk depends on productivity. We establish the following proposition.

Proposition 2. (*Heterogeneity of the Effect of Risk*) *Consider the case without a domestic input in production. Let M denote import value. Then:*

$$\frac{\partial}{\partial \varphi} \epsilon_{N^*, \mathbb{V}[\alpha]} > 0 \text{ and } \frac{\partial}{\partial \varphi} \epsilon_{M, \mathbb{V}[\alpha]} \geq 0.$$

Thus, more efficient firms feature a greater change in the number of suppliers, and a smaller reduction in total imports, after an increase in $\mathbb{V}[\alpha]$.

Proof. Note first that

$$\frac{\partial N^*}{\partial \varphi} = -\frac{G_\varphi}{G_N} \geq 0$$

as $G_\varphi > 0$. That is, more efficient firms select a greater number of suppliers. Note next that $\epsilon_{N^*, \mathbb{V}[\alpha]}$ depends on efficiency only through N . Using (29), it is straightforward to show that:

$$\frac{\partial \epsilon_{N^*, \mathbb{V}[\alpha]}}{\partial N^*} > 0,$$

which establishes the first part of the proposition. Since import value $M \equiv Nxp_M$ is

proportional to expected revenues, we have that:

$$\epsilon_{M, \mathbb{V}[\alpha]} \equiv \frac{\partial \log M}{\partial \log \mathbb{V}[\alpha]} = \frac{1}{1 - \tilde{\psi}} \frac{\partial \log \left(\mathbb{E} \bar{\alpha}^{\tilde{\psi}} \right)}{\partial \log \mathbb{V}[\alpha]}$$

Moreover, relying on the approximation in (28), we have that:

$$\frac{\partial \left(\mathbb{E} \bar{\alpha}^{\tilde{\psi}} \right)}{\partial \mathbb{V}[\alpha]} = -\tilde{\psi} \frac{(1 - \tilde{\psi})}{2} (\mathbb{E}[\alpha])^{\tilde{\psi}-2} \left\{ \frac{1 - \epsilon_{N^*, \mathbb{V}[\alpha]}}{N} \right\}$$

It follows that:

$$\epsilon_{M, \mathbb{V}[\alpha]} = -\frac{1}{1 - \tilde{\psi}} \frac{\tilde{\psi} \frac{(1 - \tilde{\psi})}{2} (\mathbb{E}[\alpha])^{\tilde{\psi}-2} \mathbb{V}[\alpha]}{N (\mathbb{E}[\alpha])^{\tilde{\psi}} - \tilde{\psi} \frac{(1 - \tilde{\psi})}{2} (\mathbb{E}[\alpha])^{\tilde{\psi}-2} \mathbb{V}[\alpha]} \{1 - \epsilon_{N^*, \mathbb{V}[\alpha]}\} \leq 0$$

Note that the right-hand side of this expression increases with efficiency (as both N and $\epsilon_{N^*, \mathbb{V}[\alpha]}$ increase), proving the result. \square

Intuitively, larger firms can more easily afford the fixed costs of adding suppliers when faced with greater risk, and feature a larger increase (or a smaller decrease) in the number of suppliers when faced with greater volatility in supplier quality. As a result, they manage to attenuate the resulting increase in exposure to risk and see their import values fall by less. Lastly, Proposition 2 also implies that larger firms feature a smaller decline in sales when faced with heightened risk.

C.4 The Effects of Changes in Expected Quality

In the main text, we study the effects of a mean preserving spread to input qualities. Because our empirical exercises involve days and not qualities, it is necessary to also explore the effect of changes in expected quality. A higher expected quality affects expected revenue in two ways. There is a direct positive effect that stems from output being increasing in input qualities and is captured by the first term in expression (16). A second effect is that a higher expected quality makes a given variance less important for expected revenue, as captured by the second term in expression (16). In other words, the importance of the variance of quality for expected revenues is mediated by the mean quality. By making a given variance less detrimental, this effect reduces the incentive to adding suppliers. This negative effect turns out to be dominated by the positive direct effect whenever the variance of quality is large relative to its mean, as captured by the coefficient of variation. The following result

formalizes this argument.

Proposition 3. *Consider the case without a domestic input in production. Let N^* be the optimal number of suppliers. An increase in the mean of the supplier-level quality $\mathbb{E}[\alpha]$, holding constant its variance $\mathbb{V}[\alpha]$, increases the returns to adding suppliers*

$$\left. \frac{\partial^2 \tilde{R}}{\partial N \partial \mathbb{E}[\alpha]} \right|_{N=N^*} > 0$$

whenever:

$$\frac{1}{N^*} \frac{\mathbb{V}[\alpha]}{(\mathbb{E}[\alpha])^2} > \frac{2 - 3\tilde{\psi}}{(1 - \tilde{\psi})(2 - \tilde{\psi})} \frac{2}{\tilde{\psi}}. \quad (30)$$

Note that, when $\tilde{\psi} < 2/3$, the right hand side of condition (30) is lower than the right hand side of condition (17) in the main text.³⁶ It follows that in the region of parameters where a mean preserving spread to quality increases the returns to adding suppliers (i.e., condition (17) is satisfied), there is a sub-region where an increase in expected quality increases the returns to adding suppliers and a sub-region where the opposite happens.

Proof. Note that (16) can be written as:

$$\frac{\partial \tilde{R}}{\partial N} = \tilde{\chi} \frac{\tilde{\psi}}{2} N^{-2} \left((\mathbb{E}[\alpha])^{\frac{3\tilde{\psi}-2}{\tilde{\psi}}} - (\mathbb{E}[\alpha])^{\frac{\tilde{\psi}-2}{\tilde{\psi}}} \tilde{\psi} \frac{(1-\tilde{\psi})}{2} \frac{1}{N} \mathbb{V}[\alpha] \right)^{\frac{\tilde{\psi}}{1-\tilde{\psi}}} \mathbb{V}[\alpha]$$

The marginal effect of a change in the mean of input quality, holding its variance constant, on the returns to adding suppliers is:

$$\begin{aligned} \frac{\partial^2 \tilde{R}}{\partial N \partial \mathbb{E}[\alpha]} &\propto \frac{\tilde{\psi}}{1-\tilde{\psi}} \left((\mathbb{E}[\alpha])^{\frac{3\tilde{\psi}-2}{\tilde{\psi}}} - (\mathbb{E}[\alpha])^{(\tilde{\psi}-2)\frac{1}{\tilde{\psi}}} \tilde{\psi} \frac{(1-\tilde{\psi})}{2} \frac{1}{N} \mathbb{V}[\alpha] \right)^{\frac{\tilde{\psi}}{1-\tilde{\psi}}-1} \times \\ &\left(\frac{3\tilde{\psi}-2}{\tilde{\psi}} (\mathbb{E}[\alpha])^{\frac{3\tilde{\psi}-2}{\tilde{\psi}}-1} - (\tilde{\psi}-2) (\mathbb{E}[\alpha])^{(\tilde{\psi}-2)\frac{1}{\tilde{\psi}}-1} \frac{(1-\tilde{\psi})}{2} \frac{1}{N} \mathbb{V}[\alpha] \right). \end{aligned}$$

³⁶This follows from the fact that

$$\frac{2 - 3\tilde{\psi}}{(1 - \tilde{\psi})(2 - \tilde{\psi})} < 1$$

for $\tilde{\psi} \in (0, 2/3)$.

Because the first term on the right hand side is positive, the sign of this crossed derivative depends on the sign of the second term, which is positive whenever:

$$\frac{1}{N} \frac{\mathbb{V}[\alpha]}{(\mathbb{E}[\alpha])^2} > \frac{2 - 3\tilde{\psi}}{(1 - \tilde{\psi})(2 - \tilde{\psi})} \frac{2}{\tilde{\psi}}$$

If $\tilde{\psi} > 2/3$, then the right hand side of this expression is negative and the inequality holds. When $\tilde{\psi} < 2/3$, then the above condition is required for the crossed derivative to be positive. \square

C.5 Additional Details on Equilibrium

After shipping times have been realized, the realized output of a firm of type f is given by

$$\hat{y}_f = \varphi_f^{\frac{1}{1-\tilde{\gamma}}} \left[\tilde{\gamma} P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma} w^{-1} \right]^{\frac{1-\tilde{\gamma}}{1-\tilde{\gamma}}} \left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + \left(\sum_{i=1}^N \hat{\alpha}_i \right)^{\frac{\varepsilon-1}{\varepsilon}} x^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\gamma \frac{\varepsilon}{\varepsilon-1} \frac{1}{1-\tilde{\gamma}}}$$

where $\hat{\alpha}_i$ is the realized input quality. Because there is a continuum of identical firms of each type, aggregate output is deterministic and given by $y_f \equiv \mathbb{E}[\hat{y}_f]$. Prices are then determined by the demand from consumers and the domestic input producers, $p_f = y_f^{-\frac{1}{\sigma}} P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma}$, who face no risk. Aggregate labor demand by firms of type f is $l_f \equiv \mathbb{E}[\hat{l}_f]$ where \hat{l}_f is given by (24). Aggregate profits are given by:

$$\Pi = \int_f (p_f y_f - p_D x_{Df} - N_f p_M x_f - w l_f - N_f F) df,$$

where x_{Df} , x_f and N_f are the optimal input choices in stage 1 of type f firms. An equilibrium (S, P) is attained whenever the associated firm input choices, prices, and aggregate profits are consistent with equations (18), (19) and (20).

C.6 Additional Results for Counterfactuals

Table C.1 contains the effects of the climate change and port congestion counterfactuals on average sales, profits, and labor. The calibrated model predicts increases of about 0.3-0.5% in these variables for the climate change counterfactual, and 0.6-0.8% for the port congestion (columns 1 and 4). These effects are driven primarily by general equilibrium forces: greater risk increases the cost of foreign inputs and hence the price level P , and increases domestic

spending S , as firms substitute towards domestically produced inputs. Instead, when S and P are kept fixed, the model predicts decreases in sales, profits and labor of about 1-1.2% for the climate counterfactual and 2.2-2.6% for port congestion (columns 2 and 5), reflecting the effects of increased volatility in supplier quality—see Section 4.3.

We next compare these counterfactuals to a back-of-the-envelope calculation of the changes in sales, profits, and labor using the reduced-form estimates of the effects of delays due to weather shocks from Table 4. We start by mapping the increase in the standard deviation of shipping times in each of the two counterfactuals into an increase in the share of costs which are delayed—i.e., the variable in the right-hand-side of the reduced-form regressions in Table 4. Holding input choices constant, the change in the expected share of delayed inputs after a change in the distribution of shipping times is given by:

$$\Delta \mathbb{E} [FracDelayed] = s_{IM} \Delta \mathbb{E} [s_N],$$

where s_N is the share of suppliers that are delayed and s_{IM} is the share of expenditure on imported inputs in total costs. We perform our calculation for a representative firm with $s_{IM} = 0.16$ —resulting from the observed import share in materials of 0.28 and the share of materials in total costs of 0.58 in our data. Following the empirical analysis of Section 3.1, we consider a supplier’s shipment to be delayed whenever its shipping time exceeds the 95th percentile of shipping times in the pre-counterfactual equilibrium. Assuming that shipping times are log-normally distributed, we simulate 1,000 draws of shipping times for each supplier after the counterfactual increase in risk. We compute the counterfactual $\mathbb{E} [s_N]$ as the average fraction of suppliers delayed ($s_N = 0.05$ in the pre-counterfactual equilibrium). We perform this calculation for each risk type and take the average.

Armed with the change in the expected fraction of delayed costs for each counterfactual, we use the coefficients of Table 4 to compute the predicted changes in sales, profits, and labor. We find reductions in these variables of about 0.2-1.2% for the climate counterfactual, and of 0.5-2.6% for the port congestion (columns 3 and 6 in Table C.1). Because the reduced-form estimates of the cost of delays do not capture changes in general equilibrium variables, which are absorbed by the time fixed effects, we compare the back-of-the-envelope calculations to the predictions of the model in partial equilibrium. The model’s predictions of sales and profits line up closely to those implied by the reduced-form estimates. The predictions for labor are somewhat less aligned, with the model over-predicting the magnitude of the decline.

Table C.1: Additional Outcomes After Counterfactuals

	Climate Change			Port Congestion		
	(1)	(2)	(3)	(4)	(5)	(6)
Growth in (%)	Model	Model P.E.	Back-of-the-envelope	Model	Model P.E.	Back-of-the-envelope
Sales	0.32	-1.09	-1.16	0.64	-2.38	-2.56
Profits	0.36	-1.18	-0.76	0.60	-2.58	-1.69
Labor	0.49	-1.00	-0.2	0.84	-2.18	-0.44

Notes: The table reports the percentage change in average sales, profits and labor associated with (i) an increase in the standard deviation of log shipping times of 18.5% (columns 1-3) and (ii) an increase in the standard deviation of shipping times of 52% (columns 4-6). The Model and Model P.E. columns report the predictions of the calibrated model, with Model P.E. keeping S, P constant. The back-of-the-envelope columns report predictions using the reduced-form coefficients of Table 4.