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

Joaquin Blaum — Boston University
Federico Esposito — Tufts University
Sebastian Heise — NY Fed

March 2024

Abstract

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 40,000 maritime routes. We first document that unexpected shipping delays induced by weather shocks have negative effects on importers' revenues, profits and employment. We then show that firms actively diversify this source of risk: importers facing greater shipping volatility use more routes and source from more foreign suppliers, while they reduce imports. To rationalize these findings, we introduce shipping time risk into a general equilibrium model of importing with firm heterogeneity. Our quantitative analysis predicts that an increase in weather volatility in line with historical trends would reduce the U.S. manufacturing import share by 9% and real income by 2% over the next 50 years.

*Blaum: 270 Bay State Road, Boston, MA 02215, email: jblaum@bu.edu. Esposito: Joyce Cummings Center, Room 663, 177 College Avenue, Medford, MA 02155, email: federico.esposito@tufts.edu. Heise: 33 Liberty Street, New York, NY 10045, email: sebastian.heise@ny.frb.org. We thank our discussants Davin Chor, Simona Giglioli and Shafaat Yar Khan, as well as Laura Alfaro, 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-FY21-P1883-R8915, CBDRB-FY23-P1883-R10453, CBDRB-FY23-P1883-R10672).

1 Introduction

The past decades have seen a dramatic transformation in the international organization of production, with intermediate inputs accounting for two-thirds of global trade (Johnson and Noguera (2012), Antras and Chor (2021)). The increasing reliance on imported inputs has exposed firms to more sources of disruption risk, such as natural disasters, port backlogs, or supply chain bottlenecks (Baldwin and Freeman (2021); Carvalho et al. (2021)). In the post-Covid years, shortages and shipping delays have put an additional significant strain on supply chains. According to a recent survey of U.S. manufacturing firms, more than 50% of respondents were affected by supply chain disruptions, and the average delivery time for inputs increased by 35 days in 2021 (Federal Reserve Bank of New York (2021)). Moreover, climate change will likely increase the frequency and severity of extreme weather events, with negative consequences for shipping times (Young et al., 2011, Allan et al. (2023)).

While there has been a growing focus on the effects of supply chain disruptions in academic and policy circles (Barrot and Sauvagnat (2016), Carvalho et al. (2021), White House, 2021 and U.S. Department of Commerce, 2023), little is known about how importers cope with supply chain risks, and how these risks affect international trade patterns. Important challenges for answering these questions are the need for detailed firm-to-firm data and the difficulty to identify exogenous sources of risk faced by importers.

We shed light on these questions by focusing on a specific but important source of risk: weather shocks. We provide a novel measure of shipping time risk using transaction-level import data on ocean shipments provided by the U.S. Census Bureau, which we combine with detailed data on oceanic wave conditions along more than 40,000 maritime routes to isolate the unexpected component of shipping time volatility. We estimate the effect of weather-induced shipping time volatility on U.S. manufacturing firms' performance and sourcing behavior. Our results indicate that, first, unexpected shipping delays induced by weather shocks have large and disruptive effects on U.S. importers' production levels and profit margins. We then show that firms systematically respond to shipping risk induced by weather conditions along different margins of adjustment: more exposed firms feature lower imports, a larger number of foreign suppliers, a lower concentration of expenditure across suppliers, and a greater reliance on more routes and modes of transportation. We next incorporate risky shipping times into an otherwise standard quantitative model of firm-level importing, and calibrate the model to match salient features of the data. Our model suggests that if the volatility of weather conditions continues to increase in line with historical trends, the U.S. manufacturing import share in 50 years will be 9% lower and real income will fall by 2% compared to today.

Our methodology relies crucially on the measurement of the components of shipping times that are unexpected to importers. We exploit the granularity of the U.S. Census Bureau's Longitudinal Firm Trade Transactions Database (LFTTD), which provides for each transaction the identity of the U.S. importer and its foreign supplier, the product at the ten-digit Harmonized System (HS-10) level, value and quantity shipped, and the mode of transportation 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 our customs data do not contain details on each vessel's journey, we propose an algorithm that uses the vessel name, foreign port of departure, and U.S. port of entry to determine the intermediate stops a 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.[1] We compute the weather conditions along each shipment's route using detailed hourly data at the 0.5 degree level on oceanic wave conditions, measured with the average wave height and direction from the National Oceanic and Atmospheric Administration (NOAA).

We use a rich set of fixed effects and controls to remove the components of the shipping times that are presumably known to the importers at the time the inputs are bought, including the identity of the supplier, the route, the vessel, the month, and the shipping charges. We then isolate the variation in the residualized shipping times that is induced by weather conditions, i.e., which is explained by the realized wave height and direction observed along the route of each transaction. To interpret the variation in the weather-induced shipping times as unexpected, our identifying assumption is that the *realized* weather conditions along the entire maritime route are not anticipated by the importers when they make the orders, beyond seasonal patterns that are picked up by route-month fixed effects.

We start our analysis by assessing the effects of extreme shipping delays induced by weather shocks on firms' outcomes. To this end, we identify for each year the shipments that were extremely delayed due to weather shocks, which we define as having a weather-induced delivery time larger than the 95^{th} percentile of its distribution for a given route. We estimate panel regressions for the years 2011-2016 and document that U.S. importers with a higher share of extremely delayed inputs due to weather experienced significant declines in sales, profits and employment. A one standard deviation increase in the share of input costs that are subject to weather delays reduces firms' sales by 6.5%, profits by 3.5% and employment by 1% in the same year. These large negative effects highlight the substantial impact of

[1] Ganapati et al. (2023) show that vessels on average follow the optimal maritime routes very closely. Moreover, we confirm, using AIS tracking data, that the major routes we construct are close to the actual routes that vessels follow.

supply chain disruptions on firms' production. Given the large economic costs associated with extreme shipping delays, we next study whether U.S. importers adjust their sourcing strategy and import demand to reduce the potential impact of this source of risk.

We measure the riskiness of each foreign supplier-route-product as the standard deviation of the weather-induced shipping times over 3-year rolling windows, and then construct the risk exposure of each importer as a weighted average of the risk of its suppliers and routes over the previous 3 years, using the pre-determined import shares as weights. We then estimate, for the years 2011-2016, 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 this weather-induced risk along the extensive and intensive margins. Going from the 25^{th} to the 75^{th} percentile of the shipping risk distribution increases the number of routes used and the number of foreign suppliers by 10.4% and 7.3%, respectively. Moreover, it reduces the total value imported by 3%. Thus, importers matched with riskier suppliers and routes increase the number of both, have less concentrated input expenditures, and import less overall. Importantly, the negative effect of risk on import demand we find is estimated *conditional* on the level effect that longer shipping times have on import demand, as already documented by Hummels and Schaur (2013), indicating that uncertainty amplifies the detrimental impact of long shipping times on international trade.

Motivated by these findings, we incorporate shipping risk into a standard model of firm-level importing 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. The key departure from the literature is that, at the time of placing orders, firms are uncertain about the time it will take the foreign inputs to arrive, while the delivery time for domestic inputs is deterministic. We assume that inputs that arrive later than expected are less productive. Firms can diversify this shipping time risk by sourcing from multiple foreign suppliers, albeit this strategy is limited by per-supplier fixed costs.

We consider a quantitative version of the model where firms are heterogeneous both in their productivity and in the shipping risk they face. Larger importers are assigned to safer foreign suppliers, as we observe in the data. We calibrate the model to match salient features of the data. In particular, the elasticity of the number of suppliers with respect to shipping risk disciplines the per-supplier fixed cost, while the relationship between sales and shipping times pins down the elasticity of input quality with respect to shipping time.

Armed with the calibrated model, we assess the quantitative importance of shipping time risk. First, we show that removing supply chain risk entirely implies, on average, a reduction in the number of foreign suppliers and an increase is U.S. manufacturing imports by 86 billion dollars. In addition, removing risk would increase the import share by 7.1%, as the demand for the "safe" domestic input is reduced, and would decrease the price index by 2.6%. In a second counterfactual, we assume that the standard deviation of wave height continues to increase over the next 50 years in line with the trend we observe during our sample period. This trend implies an increase in the volatility of wave height of 81% over the next 50 years, consistent with an increasing frequency of severe weather events due to climate change (e.g., Young et al., 2011). We find that this counterfactual generates a decline in the import share by 9.2% and a reduction in U.S. real income by 1.8%. Firms respond to the higher risk by spending resources to match with more foreign suppliers and by sourcing more domestically.

Related Literature Most of the literature on risk and international trade has focused on the impact of uncertainty on exports and FDI, i.e., on the output side — see e.g. Ramondo et al. (2013), Fillat and Garetto (2015), Esposito (2022), Baley et al. (2020) and De Sousa et al. (2020). In contrast, our work focuses on the risk faced by importers.[2] On the empirical front, we provide a novel comprehensive measure of shipping risk at the firm-level and show that firms actively use the extensive and intensive margins of importing to reduce the potential impact of shipping delays on their production processes. On the theory side, existing models of firm-level input sourcing typically abstract from supplier risk considerations — see e.g. Gopinath and Neiman (2014), Halpern et al. (2015), Antras et al. (2017) and Blaum et al. (2018). We contribute to this literature by exploring how delivery risk affects our understanding of input trade with heterogeneous firms, and quantify its impact in general equilibrium.

There is a small but growing literature that empirically investigates how firms manage supply chain risk, such as Clark et al. (2014), Huang (2017), Charoenwong et al. (2020), and Ersahin et al. (2023).[3] Our paper departs from the approach used by this literature in that we use transaction-level data of the universe of U.S. importers and their trade relationships with foreign suppliers, and exploit variation in shipping times across different routes, allowing us to perform firm-to-firm analysis. Alessandria and Ruhl (2021) and Carreras-Valle (2021) have recently focused on the role of inventories for the firms' response to supply chain risk, while our paper studies a different margin of adjustment to such exposure.[4]

We also connect to a growing literature that studies firm-to-firm relationships, see e.g. Tintelnot et al. (2018), Heise (2019), Bernard et al. (2019) and Esposito and Hassan (2023),

[2]Some notable exceptions are Gervais (2018) and Handley et al. (2020).

[3]There is also a literature that studies this matter theoretically from a policy perspective, such as Handley et al. (2020) and Grossman et al. (2023).

[4]It is worth noting, however, that our results are robust to controlling for firms' inventories.

and in particular 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)). Our paper provides a new way to identify supply shocks, using weather-induced shipping delays, and estimate their effects on importers. Interestingly, our results are in the same ballpark of the findings of this literature, but the effects are identified in panel regressions using firm-level shocks that are relatively frequent (i.e. the shipping delays), rather than using large, aggregate episodes – 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 that goes beyond specific episodes. Closely related to our work, Pankratz and Schiller (2022), Dunbar et al. (2023) and Balboni et al. (2023) study the impact of weather shocks in production decisions. We contribute to these papers by analyzing the effect of shipping time risk, highlighting the importance of second-order moments for firms' import demand. We demonstrate that firms adjust their sourcing behavior in response to shipping time risk by relying on more suppliers and decreasing the exposure to each of them.[5]

Finally, there is a strand of works that investigate the importance of shipping times for international trade, both in theory and in the data (Evans and Harrigan (2005), Hummels and Schaur (2010) and Hummels and Schaur (2013)). Our empirical results show that uncertainty around shipping times amplifies the negative direct impact of timeliness on import demand, and propose a theory of the firm that incorporates this mechanism in a tractable way.

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

2 Empirical Measurement

In this section we describe how we combine U.S. Census transaction-level import data with detailed data on weather conditions to measure shipping delays and shipping time risk.

[5]In addition, we are also the first in economics to use data on oceanic wave conditions to identify weather shocks, while the literature on the effects of weather and climate change on the economy has focused on other variables , such as temperature, wind and precipitation (see e.g. Eichenbaum et al. (2020)).

2.1 Data

Our empirical analysis mostly relies on the Longitudinal Firm Trade Transactions Database (LFTTD) of 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 1992-2016. Each transaction is associated with an identifier of the U.S. importer, the HS10 product code traded, the mode of transportation (vessel, air, etc.), as well as the value, weight, and quantity shipped. The data also report an identifier of the foreign seller in the form of a Manufacturer ID (MID), an alphanumeric code that combines information on the seller's country, name, street address, and city, and an indicator of whether the transaction is between related parties.[6] We calculate prices as the value of 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 allow 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.[7] 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.[8] 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 industry of each U.S. establishment.[9] We restrict our analysis to all importers that are in the manufacturing sector, since these firms are most likely to source intermediates into production that require an on-time arrival. We also

[6] We follow Kamal et al. (2015) and Kamal and Monarch (2018) in combining sellers with the same street address and city 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.

[7] Less information is available for other modes of transportation. For non-vessel imports we only have the generic shipping company name rather than the name of the individual truck, train, or plane, and we only know the country of departure rather than the precise departure location. For these modes of transportation, we therefore approximate the shipment route simply as a foreign country of origin and U.S. entry point (e.g., airport or border crossing) pair.

[8] For transactions that are not by vessel, the vessel name is either the name of the shipping company (e.g., Delta Airlines) or blank.

[9] We prepare the LBD by collapsing these data to the firm-level, and construct the 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).

Table 1: U.S. Import Transaction Summary Statistics

	All	**Seaborne**
Total Imports	10,540	4,250
Unique Importers (f)	171,400	92,300
Unique Exporters (x)	815,000	407,400
Number of Transactions (millions)	109	35.8
Number of U.S. Ports of Entry (p_i)		302
Number of Foreign Ports (p_e)		1,795
Number of Origin-Destination Port Pairs		43,080
Unique Vessels (v)		401,700

Source: LFTTD and authors' calculations. Table summarizes U.S. imports from 1992 to 2016. Values are reported in billions of 2009 dollars.

obtain firms' total sales, cost of materials, and employee compensation from the Census of Manufactures (CMF) in census years (1992, 1997, etc.) and from the Annual Survey of Manufacturers (ASM) for non-census years. We construct profits as sales minus cost of materials and employee costs.

Table 1 reports the summary statistics of our baseline dataset. The first column considers all manufacturing imports over the period 1992-2016. The second column shows the sample of seaborne trade only, which we will use to construct our measure of shipping risk below. Our dataset covers about 10.5 trillion dollars of imports (in 2009 dollars), of which about 40 percent are by vessel. For vessel-based transactions, we observe 302 U.S. ports and nearly 1,800 foreign ports, as well as more than 400,000 unique vessels, which are crucial pieces of information to construct shipping routes, to which we turn to next.

2.2 Construction of Shipping Times and Routes

We now describe how we compute shipping times and routes. First, for all shipments, irrespective of their mode of transportation, we calculate the shipping time as the difference, in days, between the importation date in the U.S. and the export date from the foreign country. We show statistics on the distribution of shipping times in Table 2. Vessel shipments on average take 16 days to arrive to the U.S., which is substantially more than 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 4 days. Importantly, there is a high degree of dispersion in vessel shipments. Shipping times are less dispersed for other modes of transportation, which have a median shipping time of zero. Given such small dispersion in

Table 2: Shipping Times by Mode of Transportation

	Avg. Time	Std. Time	P5	P25	P50	P75	P95	Total Value
Vessel	16.4	23.5	3.5	10	13.5	20.5	33.3	4,250
Train	4.4	6.2	0	0	0	8.5	16.9	1,450
Truck	0.1	0.4	0	0	0	0	0	2,210
Airplane	0.5	0.9	0	0	0	1	2.3	1,610

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

shipping times, we will focus on risk from seaborne shipments going forward and assume that all other modes of transportation have zero risk.[10]

Second, for seaborne shipments, we develop an algorithm to construct ocean shipping routes and vessels' journey between ports. We assign each transaction 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.[11] We then use the dates of import and export to construct the sequence of ports visited by each shipment, e.g. Le Havre - Birmingham - New York - Newport News. We refer to a leg of the trip between two ports as *route segment.*

We determine the route taken by the vessel across the ocean between any two ports using Eurostat's SeaRoute program. This program computes the shortest maritime routes from a network of lines following the most frequent maritime routes. Our initial sample includes around 32,000 route segments – see summary statistics in Appendix A.3. Once we merge these routes with the transaction-level import data, we are left with 10,500 route segments and 43,000 distinct routes.

Our measures of shipping times and routes are the building blocks of our empirical analysis. The key advantage of deriving such measures from U.S. Census transaction-level data is that these data are comprehensive and extremely detailed, allowing us to conduct a systematic

[10]We show some additional statistics on shipping times and their determinants in Appendix A.2.

[11]In some instances, the arrival date may be misreported. In Appendix A.1 we explain how we identify such cases and how we refine our algorithm to redefine the trips.

analysis of the role that shipping time volatility has for firms' sourcing strategy. However, the construction of these variables features a number of limitations, due to the nature of the available data. First, we do not have information 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. However, typically goods spend several weeks on the vessel to the U.S. crossing the Atlantic or Pacific Ocean, and thus this part of the journey is likely a large fraction of the total travel time. Second, we do not know whether goods are reloaded from one vessel to another, i.e., "trans-shipped" (Ganapati et al., 2023). This is because the port of exportation that we observe in the data is the foreign port where goods are loaded onto a vessel headed to the U.S. Therefore, for the products that are trans-shipped we only observe the journey on the last vessel to the U.S., implying that we underestimate the overall shipping delay risk faced by importers. Third, we do not observe the actual maritime route taken by the vessels, but assume that the ships follow the shortest ocean-distance routes. We show in Appendix A.5 that this assumption is not too far from reality: our routes typically follow very closely the actual vessel movements reported by AIS data for a selected sample of routes. In addition, two-thirds of world trade in manufacturing travels on containerships, which typically follow fixed itineraries (i.e. the so-called "bus system", see Brancaccio et al. (2020) and Heiland et al. (2022)), which mitigates this concern.[12]

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 during the period 2011 to 2016.[13] 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).[14] We aggregate the hourly information to the daily level and compute the daily average of significant wave height and direction for each geo-location in the oceans. If information is

[12] In addition, Ganapati et al. (2023) show that vessels typically follow the minimum-distance routes pretty closely.

[13] Significant 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. If both swell and wind-waves are present, it equals the square root of the sum of the squares of the swell and wind-wave heights. See https://www.ndbc.noaa.gov/waveobs.shtml.

[14] Ocean 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.

unavailable at a geo-location, we interpolate the weather using the information of neighboring geo locations.[15]

We obtain the geo coordinates (longitudes and latitudes) of each route segment in 0.5 degree increments and merge these coordinates with daily weather information. Since the effect of wave direction on vessel speed depends on the direction of travel, we compute for each coordinate the waves' relative direction by taking the absolute difference between the wave direction and the estimated bearing of vessels at that point. We compute vessels' bearing using the latitude and longitude of subsequent route coordinates. A greater relative direction means that the waves are less aligned with vessels' likely course.[16] To merge this information into the trade transaction data, we obtain for each import transaction the days the vessel spent on each trip segment. We then compute the segment-level average weather for each day the vessel spent on each segment, and average across days and segments.[17]

To illustrate the source of our exogenous variation, the blue and green shading in Figure 1 reports the standard deviation of the average daily wave height for each grid point in our data. The red lines indicate some of the major shipping routes used by U.S. importers. There are significant differences across locations. Routes across the Atlantic and Pacific have higher volatility of wave height 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.[18]

3 Shipping Risk and Import Demand

In this section we investigate how U.S. importers cope with the risk stemming from shipping delays. We start by documenting that importers rely on multiple foreign suppliers and shipping routes to source their products. We then establish that extremely long shipping

[15]Note that while we have weather conditions in the oceans, we do not have weather information for some of the more interior bodies of water such as the Great Lakes, the Mediterranean sea, and the Baltic Sea. Consequently, for trip segments in these regions the weather is missing.

[16]For example, a wave direction of 75 degrees for a vessel traveling at direction 90 degrees would be translated into a wave direction of $\text{abs}(90-75) = 15$ degrees. When this absolute difference exceeds 180, we subtract 180 to get the minimum distance. For instance, if a vessel travels North and the waves go West, the relative direction would be $\text{abs}(0-270) - 180 = 90$ degrees.

[17]For example, for a good loaded in Le Havre on June 23 onto a vessel that made an intermediate stop in Southampton on June 25 and that arrived in New York on July 4, we would average over the weather on the Le Havre - Southampton segment from June 23 to June 25 and the Southampton to New York segment on June 25 to July 4.

[18]Appendix A.4 provides some summary statistics on this variable. We also show that there significant differences in weather in the Northern Atlantic and Northern Pacific across seasons.

Figure 1: Volatility of Wave Height on Selected Routes

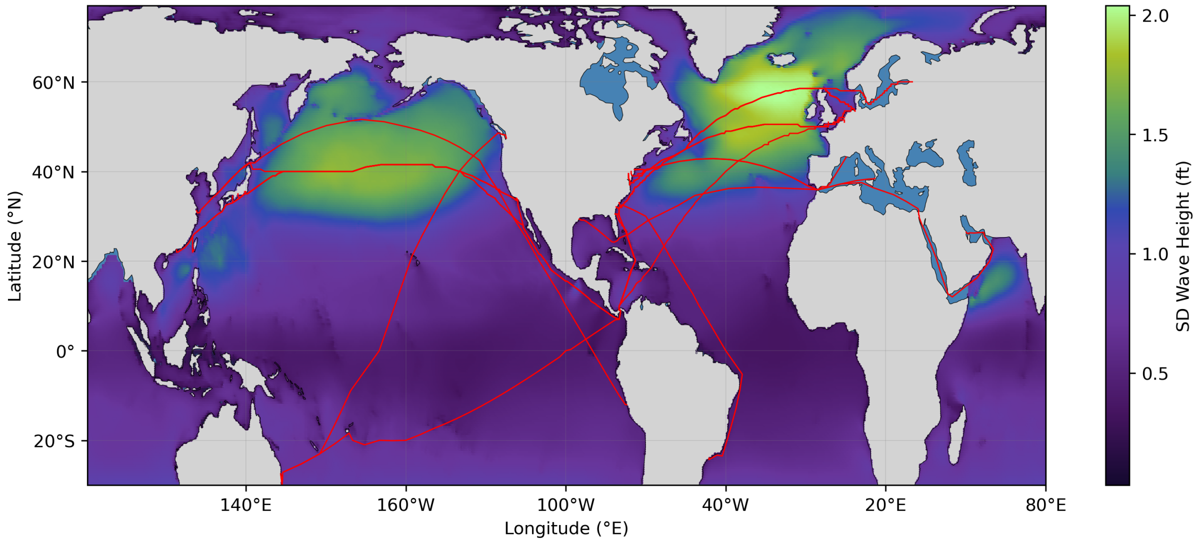

Notes: The figure shows the standard deviation of wave height across all days from 2011-2016 for selected shipping routes into the U.S.

times induced by weather conditions (“shipping delays”) have negative and significant consequences on importers’ performance. Lastly, we document that U.S. manufacturing importers adjust their sourcing strategy and import demand in response to shipping risk.

3.1 Suppliers, Routes and Sourcing Behavior

We show in Table 3 that firms use multiple foreign suppliers and routes even for the same HS-10 product within the same year. The first row of Table 3 shows that the average importer has 1.9 suppliers for the same product per year, across all modes of transportation. The large standard deviation compared to the mean indicates that there is significant heterogeneity across buyers. While the median importer uses only one foreign supplier for a given product, firms at the 95^{th} percentile use more than five. Firms sourcing a given product from multiple suppliers account for almost 90% of imports. In Section 3.4 below, we argue that one reason to use multiple suppliers is to hedge against shipping delay risk.

The second row of Table 3 shows that firms may not only use multiple foreign suppliers to hedge against delays but also use different routes. The average firm uses 2.2 routes per product and year across all modes of transportation, and the 95^{th} percentile is nearly seven routes. The third and fourth rows of Table 3 highlight that there is substantial heterogeneity across importers in the concentration of input expenditures across different suppliers. While the average HHI is 0.88, indicating that firms’ sourcing of a given product is concentrated on

Table 3: Summary Statistics on Foreign Sourcing

	Mean	St. Dev.	P50	P95
Number of Suppliers	1.90	3.84	1	5.39
Number of Routes	2.23	4.08	1	6.77
HHI across Suppliers	0.88	0.21	1	1
HHI across Suppliers and Routes	0.81	0.27	1	1
Imports per Supplier	425	1,370	15	833
Imports	1,659	58,900	19	2,092

Source: LFTTD and authors' calculations. Table reports the mean and standard deviation across importer-product-year tuples in our 1992-2016 sample period. Values are expressed in thousands of 2009 dollars.

a main supplier, the standard deviation is 0.21. The last two rows show that the average firm imports $425,000 per supplier for a given product and year and has total imports of that product of on average \$1.7 million. There is a substantial tail of very large importers, as evidenced by the standard deviation of \$59 million.

3.2 Measuring Unexpected Shipping Times

One of the goals of our methodology if to measure "unexpected" shipping delays: production processes are likely disrupted if one or more inputs arrive later than expected. However, we cannot observe the shipping times expected by the importers, only the realized ones. To isolate unexpected shipping delays, we propose a two-step approach. In the first step, we regress the observed shipping times on a rich set of fixed effects and observables to capture any determinants of the shipping times that may be fully anticipated by the importers, and treat the residuals from such regression as the stochastic (i.e. unexpected) component of the shipping times. In the second step, we regress these "residualized" shipping times on the weather conditions observed along the maritime route, and treat the predicted effects from this regression as weather-induced unexpected shipping times.

We first residualize the shipping times with respect to a rich set of fixed effects and observables for two reasons. First, weather conditions may vary systematically across routes and seasons, or impact shipping times in a systematically different way depending on the vessel type. If we regressed reported shipping times directly on weather, some of the effect of weather on shipping times that we estimate could in fact be anticipated. The residualization step seeks to remove such sources of predictable variation. Second, we only have data on weather conditions for the limited sample period of 2011-2016. Since we can perform the residualization procedure on the full sample from 1992, we also construct a risk measure using only the residualized shipping times as robustness, and show that our results carry over to the full sample.

Step I: Residualization. We propose a residualization procedure to remove the components of the shipping times that are presumably known to the importers at the time of making orders. As discussed, we focus on vessel shipments only and treat shipments using all other modes as riskless. Consider a buyer f that orders a vessel shipment s of product h from seller x in time period t. The seller can either be a related party or at arms' length, and this is captured by the index a. The shipment arrives to the U.S. on vessel v via route r, which consists of a combination of the port of origin and destination. A given shipment s has a weight of W^s and total shipping charges (freight costs plus insurance) are C^s dollars.[19] The time it takes for the shipment to arrive in the U.S., $T^s_{xhrtvfa}$, is a stochastic variable with the following law of motion:

$$\begin{aligned}\ln(T^s_{xhrtvfa}) = {} & \bar{\pi}_x + \bar{\alpha}_h + \bar{\gamma}_{rt} + \bar{\xi}_v + \bar{\delta}_f + \bar{\omega}_a \\ & + \pi_x + \alpha_h + \gamma_{rt} + \xi_v + \delta_f + \omega_a + \eta \ln\left(C^s\right) + \rho \ln\left(W^s\right).\end{aligned} \tag{1}$$

The terms with upper bars capture a long list of deterministic components that might be known to the buyer at the time of ordering. For instance, $\bar{\pi}_x$ may reflect the ability of a supplier to arrange logistics with shipping companies, while $\bar{\alpha}_h$ may capture the fact that some products are harder to ship or take longer to get cleared at customs. $\bar{\gamma}_{rt}$ captures route characteristics in a given month t (e.g., April 2015), such as weather on the route or characteristics of the ports of departure or entry such as the average time it takes to unload a shipment and clear customs.[20] Shipping times are also determined by random shocks $(\pi_x, \alpha_h, \gamma_{rt}, \xi_v, \delta_f, \omega_a)$ which are realized after import orders are placed. As a first step to measure plausibly unexpected shipping times, we compute the residuals $\tilde{t}^s$ from regressing the observed shipping times on the set of fixed effects and observables that are specified in equation (1).

[19]The rationale to include the shipment weight and freight rate is that they have a significant effect on shipping times, as we show in Table 10 in Appendix A.2.

[20]Similarly, $\bar{\xi}_v$ may capture the speed or weight of a vessel. The buyer component $\bar{\delta}_f$ captures buyer characteristics that may affect shipping times, such as its ability to arrange logistics with the supplier. The relationship status component $\bar{\omega}_a$ may capture that it is easier to arrange transport when the partners are related rather than at arms' length.

Step II: Weather Conditions. While our residualization removes various plausibly known components of shipping times, some of the remaining variation could still be anticipated by the importers. For example, firms with low import demand could request shipments that are relatively slower. We therefore focus on the variation in shipping times induced by weather along the shipping route. Our identification relies on the assumption that, at the time the inputs orders are made, 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 residualization. We believe this is a reasonable assumption, as most shipments to the U.S. require multi-week ocean crossings where weather conditions cannot be perfectly predicted, even by shipping companies relying on sophisticated weather forecasting technology.

Following the literature on ocean shipping, we measure weather conditions with significant wave height and relative direction (Filtz et al. (2015)). We regress the transaction-level residualized shipping times (obtained from the previous step) on these variables and their interaction, for the years 2011-2016 for which we have weather data:

$$\tilde{t}^s = \beta_1 \text{Height}^s + \beta_2 \text{Direction}^s + \beta_3 \text{Height}^s \cdot \text{Direction}^s + \epsilon^s, \tag{2}$$

where, for each shipment s in the LFTTD, $\tilde{t}^s$ is the (residualized) log shipping time, Height^s is the average wave height along the shipment's route in meters, as constructed above, and Direction^s is the average relative wave direction (in degrees, relative to the vessel's direction of travel). As explained above, Height^s and Direction^s are both measured along the entire maritime route taken by shipment s.

Table 4: Effect of Weather on Shipping Times

Dep. Var:	$\tilde{t}^s$	$\tilde{t}^s$	$\tilde{t}^s$
Wave Height^s	-0.025^{***}	-0.026^{***}	-0.029^{***}
	(0.000)	(0.000)	(0.001)
Direction^s		-0.0001^{***}	-0.0002^{***}
		(0.0000)	(0.0000)
Wave $\text{Height}^s \times$ Direction^s			0.00003^{***}
			(0.00001)
Route FE	Y	Y	Y
Observations	5,728,000	5,728,000	5,728,000

Source: LFTTD and authors' calculations. 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.

Table 4 presents the results from the regression. The first column includes only wave height as regressor, the second column adds the relative wave direction, and the third column adds the interaction term. Higher waves reduce shipping times: as shown in the final column, a one standard deviation increase in wave height (1.5m) reduces shipping time by nearly 4 log-points . Shipping times are also slightly 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 4 log points. While the 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 in Appendix A.5. 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. The predicted values from the regression in column (3), $\tilde{t}^{s,weather}$, constitute our measure of unexpected shipping times.

3.3 Shipping Delays and Importers Performance

We identify weather-induced delays as cases where a shipment's unexpected shipping time, $\tilde{t}^{s,weather}$, is above the 95^{th} percentile of its distribution within a given product-route.[21] For each importer, we then compute the weighted share of these weather-delayed inputs as:

$$FracDelayed_{ft} = \frac{\sum_s \mathbb{D}^s_{ft} \cdot \text{Imp Value}^s_{ft}}{\text{Total Input costs}_{ft}}, \tag{3}$$

where $\mathbb{D}^s_{ft}$ is an indicator that shipment s to importer f was delayed in year t, Imp Value$^s_{ft}$ is the import value of such shipment, and Total Input costs$_{ft}$ are the importer's total production costs (materials, including domestically sourced inputs, plus labor) in the year.[22] Intuitively, this measures the share of an importer's input expenditures that are subject to extremely long shipping times, i.e. delays, in a given year. We then estimate the following panel regression for the years 2011-2016:

$$\ln(Y^o_{ft}) = \alpha + \beta_1 FracDelayed_{ft} + \gamma_f + \delta_t + \epsilon_{ft}, \tag{4}$$

where Y^o_{ft} 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. Table 5 reports the results. Standard errors are clustered at the firm-level.

[21] Since we analyze delays within a given product, we omit the product fixed effect from the first-step residualization and compute $\tilde{t}^s \equiv \ln(T^s) - \hat{\eta}\ln(C^s) - \hat{\rho}\ln(W^s) - \hat{\gamma}_{rt} - \hat{\xi}_v - \hat{\delta}_f - \hat{\omega}_a - \bar{\pi}_x$ in the first stage.

[22] Data on production costs is taken from the manufacturing census or the ASM. Note that for non-census years we only have this information for the subset of firms that are in the ASM.

Table 5: Effect of Extreme Delays on Firms' Outcomes

	(1)	(2)	(3)
Dependent Variable (in logs):	Sales	Profits	Employees
Frac Delayed	-2.434^{***}	-1.298^{**}	-0.390^{*}
	(0.513)	(0.502)	(0.231)
Importer FE	Y	Y	Y
Year FE	Y	Y	Y
Observations	$40,500$	$40,500$	$40,500$
R-Squared	0.973	0.914	0.982

Notes: 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 0.0038 and standard deviation is 0.0266. R^2 is the overall fit inclusive of the fixed effects.

Shipping delays significantly disrupt production levels and profits margins. Increasing the fraction of delayed shipments by one standard deviation (2.66%, which is almost a ten-fold increase from the average fraction delayed, which is 0.38%) is associated with a drop in sales by 6.5%, a fall in profits by 3.5%, and in the number of employees by 1.0%. Our identifying assumption is that, conditional on firm fixed effects, the fraction of inputs that is subject to weather-induced shipping delays is orthogonal to any unobservable characteristics that may affect an importer's post-delay performance. Since realized weather conditions are not under control of the importer or the foreign supplier, this condition is plausibly satisfied. We find that unexpected delays have a large and significant negative impact on U.S. importers' economic outcomes. Such large disruptive effects also suggest that our weather-induced measure of shipping times captures a meaningful component of the variation in delivery times.

It is worth noting that the elasticity of firms' sales with respect to the shipping delays is similar to the elasticity of sales with respect to supply chain disruptions estimated by a recent literature. In fact, 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 by 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. There are two relevant differences, however, relative to this literature. First, our identification arises from idiosyncratic shocks that are relatively frequent and widespread (the shipping delays), while the literature has identified the effects from large, aggregate episodes – such as the Japanese earthquake or the Covid lockdowns. Our measure therefore lends itself to a wide range of applications that require exogenous firm-level shocks. Second, our elasticity should be interpreted as an intensive margin response, rather than an extensive margin one: in other words, our shock occurs when a foreign input arrives later than expected, rather than an input never arriving at destination.

In Appendix B.1, we construct an alternative measure of weather shocks. Instead of averaging over the weather conditions of the entire route, we predict where on the route the vessel is on each day and only use local weather conditions around this location. We then follow the same steps as before to construct weather-induced shipping delays with this measure. We find that our results are similar using this alternative risk measure.

3.4 Shipping Time Risk and Import Demand

Having shown that shipping delays have large negative consequences on U.S. importers' production, we now investigate whether and how firms respond to this source of risk. We measure the exposure of U.S. importers to shipping time risk for narrowly defined products. To do so, we first define a risk measure for vessel shipments at the supplier-product-route-year (x, h, r, t) level, denoted by $\widehat{StdTime}_{xhrt-3,t-1}$. This measure is computed as the standard deviation of the weather-predicted residualized shipping times computed over three-year rolling windows.[23] For non-vessel transactions, we set $\widehat{StdTime}_{xhrt-3,t-1} = 0$ for each (x, h, r, t) cell. We then compute the 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. Formally, we compute

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

where the weights are firm f's past import shares of product h from each supplier-route. Our measure is akin a shift-share exposure measure, where the supplier-route-product level standard deviations are the "shift", and the import shares are the pre-determined "shares".

Armed with this measure of risk, we estimate the following panel specification for the years 2011-2016:

$$\ln(Y_{fht}) = \alpha + \beta_1 \ln(\widehat{StdTime}_{fht-3,t-1}) + \beta_2 X_{fht} + \gamma_{ft} + \mu_h + \epsilon_{fht}, \tag{5}$$

where Y_{fht} is an importer f's choice variable in year t for product h. We consider the following dimensions of sourcing as dependent variable: (i) the number of routes, (ii) the number of

[23]The residualized shipping times are constructed as $\tilde{t}^s_{xhrtvfa} \equiv \ln(T^s_{xhrtvfa}) - \hat{\eta}\ln(C^s) - \hat{\rho}\ln(W^s) - \hat{\gamma}_{rt} - \hat{\xi}_v - \hat{\delta}_f - \hat{\omega}_a$. Cells with fewer than 10 transactions are dropped.

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	0.17^{***}	0.12^{***}	-0.09^{***}	-0.07^{***}	-0.05^{***}
	(0.01)	(0.01)	(0.00)	(0.00)	(0.01)
Importer-Year FE	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y
Observations	64,000	64,000	64,000	64,000	64,000

Notes: 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.

suppliers, (iii) the concentration of import value across supplier-routes as measured by the Herfindahl–Hirschman index (HHI), (iv) the HHI of import value across suppliers, and (v) the total value imported per year. X_{fht} is a vector of controls, and γ_{ft}, μ_h are importer-year and product fixed effects, respectively.[24]

Our controls include the suppliers' (residualized) average shipping time over the previous 3 years. This variable accounts for the direct effect of shipping time on import demand (as in Hummels and Schaur (2013)). It 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 control for the average unit value charged to the importer in year t. This variable addresses the concern that riskier suppliers may sell cheaper intermediate inputs. In addition, we follow Borusyak et al. (2022) and control for the sum of the import shares $\omega_{fxhr,t-3}$ for each importer-product-year, i.e. the total imports made in the previous 3 years for a given input. This addresses the concern that the weights we use to construct the risk measure may be correlated with the firms' choices, even though they are pre-determined at t, akin to the non-random exposure to shocks highlighted in Borusyak et al. (2022) in the context of shift-share regressions. Finally, we control for suppliers' size (proxied by the total exports for each product over the previous 3 years) to account for the role of capacity constraints in explaining sourcing behavior.

Table 6 presents our results. Standard errors are clustered at the firm level.[25] Column (1) documents a positive and significant relationship between the number of routes used and shipping risk.[26] An increase from the 25^{th} to the 75^{th} percentile of the risk distribution (0.61

[24] In our baseline specification, we omit firm-product pairs with only one 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 even if we include such firms.

[25] In unreported results, we find that standard errors are similar if we use two-way clustering to account for correlated shocks across importers and time.

[26] Note that we consider all the suppliers used in a year by each importer. Since the empirical analysis is

log points) increases the number of routes by 10.4%. Column (2) shows the relationship between the number of foreign suppliers and shipping risk. Again, we find a positive and significant relationship between the two. An increase from the 25^{th} to the 75^{th} percentile of the risk distribution increases the number of suppliers used by 7.3%.

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) shows a negative and significant relation between shipping risk and the HHI over routes, suggesting that importers with riskier routes diversify more their expenditures across different sources. Column (4) shows that this effect is similar for supplier-route combinations. Lastly, in column (5), we look at the relationship between our risk measure and total imports in each year. We find that going from the 25^{th} to the 75^{th} percentile of the risk distribution decreases imports by 3%.[27]

3.4.1 Discussion and Robustness

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/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/routes, this selection could produce relationships as the ones documented in the previous section. In our sample, however, we find a negative and significant correlation (-0.12) between size of the firm (proxied with log sales) and our risk measure, which is a feature we will use to calibrate the model in the next section.

We perform a number of robustness exercises. In Appendix B.2, we document that the findings are robust to residualizing shipping times with supplier-by-month fixed effects instead of route-by-month fixed effects, capturing that suppliers might be slower fulfilling orders at

done at the annual level, we do not look at how different suppliers are sequentially added throughout a given year. We will adopt this static approach also in our quantitative model in Section 4.

[27]In Appendix B.2, we report the results using only the residualized shipping times obtained from Step I. Results are qualitatively similar to the ones in Table 6, although they imply different interquartile range effects. Specifically, an increase from the 25^{th} to the 75^{th} percentile of the risk distribution (0.92 log points) increases the number of foreign suppliers by 3.7% and the number of routes used by 6.6%, while it reduces the imports per supplier by 12.1% and total imports by 8.4%.

different times of the year, or they might choose different routes. Our results also hold for different combinations of fixed effects, and also when we include in the sample firms sourcing from a single supplier, or when we focus only on the risk of the importer's main supplier. Importantly, we show that controlling 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 (2021)), does not change our main findings. We also show that results are similar when we calculate the weather conditions along the route in a different way. 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 (see Table 1). 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 regress it on our measure of risk, constructed using vessel shipments alone.[28] Table 22 indicates that higher shipping risk is associated with a higher likelihood of using air shipments, suggesting that firms use air transportation to hedge the ocean shipping risk.

Overall, our findings suggest that U.S. importers systematically react to shipping risk along different margins of adjustment. In equilibrium, importers with riskier suppliers or routes feature i) more suppliers and routes, ii) less concentrated inputs expenditures, iii) lower imports, and iv) use multiple modes of transportation.

4 A Model of Input Sourcing with Shipping Risk

To rationalize the empirical evidence on shipping time risk and import demand and quantify its aggregate implications, we lay out a theoretical framework that builds on the standard models of importing with firm heterogeneity in Halpern et al. (2015); Blaum et al. (2018); Gopinath and Neiman (2014). The key departure from the literature is that inputs' shipping times are a component of input quality, as in Hummels and Schaur (2013), and that such shipping times are stochastic.

We start by describing the model environment in Section 4.1 and then turn to characterizing the solution to the firm's problem in Section 4.2. We describe the equilibrium of the model in Section 4.3. In Section 4.4, we calibrate the model and evaluate its ability to rationalize the facts documented in Section 3. Lastly, we simulate a number of counterfactuals in Section 4.5.

[28]We include the baseline controls, but replace the importer's total purchases over the past five years by the value imported by air and vessel separately over the past five years.

4.1 Environment

We consider a small open economy populated by a fixed mass of firms that produce differentiated varieties which are sold locally.[29] Firms buy inputs from foreign suppliers to whom they are exogenously matched. At the time of placing orders, firms are uncertain about the time it will take the inputs to arrive and, crucially, inputs which arrive late are less productive. Firms can diversify the risk of late deliveries by having multiple foreign suppliers. For tractability, we assume that the suppliers of any given firm are ex-ante identical, implying that extensive margin of trade can be summarized by the number of suppliers.[30] Different firms, however, can be matched to suppliers of different riskiness — this will be instrumental in Section 4.4 below for the model to come to terms with the empirical patterns documented above.

Firms combine labor, domestic and foreign inputs according to the following nested structure:

$$y_f = \varphi_f l^{1-\gamma} \left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + \left(\sum_{i=1}^{N} \alpha_i x_i \right)^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\gamma \frac{\varepsilon}{\varepsilon-1}} \tag{6}$$

where f denotes a firm, $\gamma \in (0,1)$ and $\varepsilon > 1$. The firm combines intermediate inputs with labor l using a Cobb-Douglas aggregator, where efficiency φ_f is firm-specific and drawn from a known distribution with CDF $H(\cdot)$. The intermediate inputs, in turn, are a CES aggregator of a domestic input 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, we assume that the extensive margin of trade is limited by fixed costs. In particular, each additional foreign supplier entails the payment of a fixed cost F in units of domestic labor.[31] This production structure corresponds to the one in Gopinath and Neiman (2014) or Blaum et al. (2018) when foreign inputs are perfect substitutes.

We assume that qualities are a function of the time it takes the inputs to be shipped from the foreign country to the domestic market. In particular, we assume that:

$$\alpha_i = \begin{cases} \bar{\alpha}_i & if \quad d_i \leq \mathbb{E}\left[d_i\right] \\ e^{-\tau \cdot d_i} & if \quad d_i > \mathbb{E}\left[d_i\right] \end{cases} \tag{7}$$

where d_i are the number of days it takes to ship foreign input i to firm f. This formulation

[29]While incorporating exporting in the model is feasible, it would make the firms' problem more involved without adding additional insights.

[30]We abstract from the more general problem where firms choose from a set of suppliers of varying riskiness.

[31]For tractability, we assume that these fixed costs are not supplier-specific. For a treatment of the case with supplier-specific fixed costs in a deterministic setting see Antras et al. (2017).

implies that if an input arrives earlier than or just as expected, i.e. $d_i \leq \mathbb{E}[d_i]$, it has a constant level of quality $\bar{\alpha}_i$.[32] Instead, if an input arrives *later* than expected, i.e. if $d_i > \mathbb{E}[d_i]$, quality falls with shipping time at an elasticity given by τ (as in Hummels and Schaur (2013)).[33]

A key assumption is that shipping times d_i are stochastic and unknown to firms when choosing input quantities. In addition to assuming that suppliers are ex-ante symmetric, we assume that their shipping times are independent, so that the d_i's are i.i.d.. We denote their CDF by $G_f(\cdot)$ and assume it is known to the firm. We allow this distribution to be firm-specific as, in our quantitative exercise below, firms differ in the riskiness of their foreign suppliers.

In terms of market structure, we assume that firms are price takers in input markets. Thus, firms can source any quantity of the domestic and foreign inputs and labor at prices p_D, p^*, and w, respectively. We assume that foreign input prices p^* are the same across all suppliers and exogenously given. Any variable trade costs are embedded in these prices p^*. In output markets, firms are assumed to compete under monopolistic competition.

There is a representative consumer that is endowed with L units of labor, owns the firms, and consumes the locally produced goods with preferences given by

$$C = \left(\int c_f^{\frac{\sigma-1}{\sigma}} df \right)^{\frac{\sigma}{\sigma-1}}, \tag{8}$$

where $\sigma > 1$ and c_f denotes final consumption of the good produced by firm f. In addition, we assume a structure of roundabout production by which firms use the output of all other domestic firms as inputs. In particular, we assume that the domestic bundle x_D is produced using the same CES aggregator as in (8).[34]

4.2 Firm's problem

The total sales of firm f, which include demand from both consumers and other firms, are given by

[32] An input arriving early could increase production, but it can also increase inventory costs. The specification in equation (7) implicitly assumes that these two forces offset each other.

[33] We set $\bar{\alpha}_i \equiv e^{-\tau \cdot \mathbb{E}[d_i]}$, such that the overall α_i function in continuous. Note that specification in equation (7) implies that the input qualities are bounded between 0 and 1. We can therefore interpret α_i as the fraction of the input quantity x_i that is effectively used in production.

[34] The assumption that the CES aggregators for the domestic bundle and consumer utility coincide is made for tractability. Under this assumption, we do not need to treat sales to the consumer and to other firms separately in the firm's problem.

$$R_f = y_f^{\frac{\sigma-1}{\sigma}} P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma}, \tag{9}$$

where y_f is given by equation (6), P is the CES price index associated with (8) and S denotes total spending. Both P and S are endogenous variables determined in general equilibrium.

Firms are risk-neutral and 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 exercise below. Due to the ex-ante symmetry of foreign suppliers, 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 is given by

$$\max_{x_D,x,N} \quad \chi_f \mathbb{E}\left[\left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + \left(\sum_{i=1}^{N} \alpha_i(d_i)\right)^{\frac{\varepsilon-1}{\varepsilon}} x^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\psi}\right] - p_D x_D - Np^* x - wNF, \tag{10}$$

where χ_f and ψ are functions of firm efficiency, general equilibrium objects and parameters, and $\alpha(d_i)$ is given by (7).[35] 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 risk of late deliveries vs payment of the fixed costs. When exposed to higher shipping risk, firms increase the number of foreign suppliers and tilt their expenditure towards the domestic input, which is assumed to be safe. We now turn to the definition of the equilibrium of the economy.

4.3 Equilibrium

We consider an equilibrium where firms maximize profits (10), the consumer maximizes utility (8), and goods markets clear. Note that the price of the domestic input bundle is given by $p_D = P$ due to the symmetry between the domestic input aggregator and consumer utility. The price of foreign inputs is exogenously given by p^*. We now turn to the characterization of the equilibrium.

The consumer's budget constraint is:

$$PC = wL + \Pi,$$

where $\Pi \equiv \int \pi_f df$ are total firms' profits. Standard calculations imply that $P = \left(\int {p_f}^{1-\sigma} df\right)^{\frac{1}{1-\sigma}}$,

[35] See Section B.3 of the Appendix for details.

where $p_f = (P)^{\frac{\sigma-1}{\sigma}} S^{\frac{1}{\sigma}} (y_f)^{-\frac{1}{\sigma}}$.[36] Note that both prices and profits depend on the *realization* of the shock, i.e. on the actual shipping days:

$$\pi_f = \pi\left(\{d_i\}; S, P, w; \varphi_f\right)$$

and

$$p_f = p\left(\{d_i\}; S, P, w; \varphi_f\right)$$

where $\{d_i\}$ are the realized shipping times for each supplier i of firm f.[37] Aggregating across all firms, the price index and total profits are then given by

$$P = \left(\int\int p\left(\{d_i\}; S, P, w; \varphi\right)^{1-\sigma} dG\left(d\right) dH\left(\varphi\right)\right)^{\frac{1}{1-\sigma}} \tag{11}$$

$$\Pi = \int\int \pi\left(\{d_i\}; S, P, w; \varphi\right) dG\left(d\right) dH\left(\varphi\right). \tag{12}$$

Aggregate spending must satisfy

$$S = wL + \Pi + p_D \int x_D\left(S, P, w; \varphi\right) dH\left(\varphi\right), \tag{13}$$

where $x_D\left(S, P, w; \varphi\right)$ is the demand for the domestic input. The consumer budget constraint and the goods market clearing conditions imply that the trade balance is given by

$$TB = -\left(L - L_d\right), \tag{14}$$

where L_d is the total labor demand and L is the labor endowment of the consumer. We normalize the wage to 1, and thus the equilibrium does not impose labor market clearing. This implies that the manufacturing sector can be a net supplier of labor to the rest of the economy and thus attain a trade deficit.[38]

An equilibrium consists of the vector (S, P) such that equations (11) and (13) are satisfied

[36]Replacing equations (7) and (16) into (6), output equals to:

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

which depends on the realization of the shipping days.

[37]Note that we have written both p_f and π_f explicitly as a function of the equilibrium variables S, P and w, as the choices N, x and x_D ultimately depend on those and parameters, but not on the realized shipping days.

[38]Note that there are no exports in our small open economy.

and firms maximize profits. It is relatively straightforward to augment the model with the inclusion of a non-tradeable sector and international trade in final manufacturing goods, as in Blaum (2019).

4.4 Calibration

In this section, we assess whether the theory outlined above can rationalize the empirical findings of Section 3. To this end, we first calibrate the parameters of the model to match salient features of the data. We then simulate the model and reproduce the risk regressions we estimate in the data. As in Blaum et al. (2018), we set the output elasticity of materials to $\gamma = 0.6$ and the demand elasticity to $\sigma = 5$, in line with the average estimate for manufacturing differentiated goods in Broda and Weinstein (2006).

We assume that firm efficiency φ is drawn from a log-normal distribution. We set the mean and variance of efficiency to match the coefficient of variation of log-sales observed in the data.[39] Shipping days are also drawn from a log-normal distribution, with mean and variance taken from our dataset. To introduce cross-sectional variation in the shipping risk faced by importers, we assume that there are two types of suppliers: low and high risk, both with the same mean. We assign these two types to the importers such that the correlation between firm sales and risk is -0.12, as the within route correlation that we observe in the data.[40]

We set the two types of risk as the averages of the importer-level risk measure we use in our empirical analysis (i.e., the standard deviation of shipping times $\widehat{StdTime}_{fht-3,t-1}$) within the groups of firms above and below the median risk measure.[41] We report all the parameter values in Table 7.

We then calibrate the fixed cost of adding foreign suppliers, F, the elasticity of input quality to shipping time, τ, and the elasticity of substitution between domestic and foreign inputs, ε, to match salient moments from the data. Specifically, we calibrate the fixed cost F to match the elasticity of N with respect to risk that we estimate in Table 6.[42] Intuitively, a

[39]The average log-sales across importers and years in our sample is 7.308, while the standard deviation is 1.727. The implied coefficient of variation is 0.24. We normalize the average productivity to 2, so the variance that matches the coefficient of variation in the data is 0.23.

[40]While we take the assignment between suppliers and importers as exogenous, it can be micro-founded with a model where firms search for suppliers, and finding safer suppliers is more costly. This would imply that larger importers have safer suppliers, as in the data.

[41]The average shipping time across importers and years in our sample is 16 days (see Table 2). The standard deviation of (log) residualized times in the two groups is 0.09 and 0.29, respectively, and their average is 0. This implies that, once we set the mean to 16, the variance of shipping days (in levels) is 2.3 and 26.3, respectively.

[42]Note that, as explained at the beginning of the section, in our model routes and suppliers are the same

Table 7: Calibration

Parameter			**Target Moment**	**Model**	**Data**
Fixed Cost per Supplier	F	0.01	Elasticity of N	0.17	0.17
Elasticity of Inputs Quality	τ	0.07	Revenue Elasticity	-0.23	-0.28
Elasticity b/w domestic and foreign inputs	ε	4.06	Trade Elasticity	-2.34	-2
Average shipping time (days)				16	16
Variance of shipping times (high type)				26.3	26.3
Variance of shipping times (low type)				2.3	2.3
CV of productivity				0.24	0.24
Correlation between sales and risk				-0.12	-0.12

higher fixed cost F makes it more costly to diversify risk by increasing the number of foreign suppliers chosen, and thus lowers the elasticity of N with respect to risk. We estimate this elasticity in the model with a cross-sectional regression of the log of firm-level optimal N on the log of the standard deviation of shipping days, just as we do for Table 6. The negative relationship between F and the elasticity of N is shown in Figure 5 in Appendix B.3.

Second, as a greater value of τ makes revenues more elastic to changes in shipping times, we calibrate this parameter to capture the negative relationship between longer shipping times and final sales that we estimate in the data.[43] We estimate this elasticity in the model as follows: we simulate 1,000 states of the world, where in each state we draw realizations of shipping days for each importer-supplier pair from the same calibrated distribution. For each state we estimate a cross-sectional regression of the log of firm-level revenues on the log of average shipping days across all suppliers. We then take an average of this elasticity across all states. The positive relationship between τ and the revenue elasticity is shown in Figure 5 in Appendix B.3.

Third, we calibrate the elasticity of substitution between domestic and foreign inputs, ε, to match the aggregate elasticity of imports with respect to trade costs. We find this elasticity in our model by increasing the price of foreign inputs by 1% and recomputing the implied aggregate imports in general equilibrium. The log-change in imports from the calibrated level divided by the log-change in the price of foreign inputs delivers the trade elasticity (see also Edmond et al. (2015) for a similar exercise). Our target is a trade elasticity

thing. For this reason, we calibrate F to match the elasticity of the number of routes with respect to risk, reported in column (1) of Table 6. Results are similar if, instead, we target the elasticity of the number of foreign suppliers with respect to risk, reported in column (2) of Table 6.

[43] In particular, we estimate a panel regression of firm-level log sales on the average shipping time of each firm's suppliers, with firm and year fixed effects. The estimated coefficient is -0.28, with standard error 0.14 clustered at the firm level.

Table 8: Counterfactuals

Variable	No risk	Climate Change
Average N	-10.8%	42.9%
Average x_D	-11.1%	12.5%
Total imports	5.4%	-3.5%
Import share	7.1%	-9.3%
Price Index	-2.6%	2.7%
Real income	2.2%	-1.8%

of -2, as recently estimated by Boehm et al. (2023). The positive relationship between ε and the trade elasticity is shown in Figure 5 in Appendix B.3. Table 7 summarizes the results of our calibration.

4.5 Counterfactuals

Armed with the calibrated model, we perform two counterfactual exercises. In the first, we start from the calibrated trade equilibrium and remove shipping time risk, by setting the standard deviation of shipping days to 0 for all firms. Table 8 (first column) summarizes the aggregate effects of this exercise. Removing supply chain risk implies a 10.8% reduction in the average number of foreign suppliers used by domestic firms, together with a substantial increase of 5.4% in aggregate imports.[44] This happens because now foreign inputs are as safe as the domestic ones, making it no longer necessary to diversify the risk of international shipping delays across multiple foreign suppliers. Indeed, this mechanism implies a large reduction in the average quantity of the domestic input of 11.1% and an increase in the aggregate import share of 7.1%. Overall, such increase in imported inputs implies that domestic firms produce more and lower their output price, thus lowering the price index by 2.6%. Real income increases by 2.2%, with aggregate spending decreasing in general equilibrium acting as a mitigating force due to the reduction in the demand for the domestic input.

In a second counterfactual exercise, we simulate an increase in weather volatility due to climate change, by assuming that future weather conditions will follow the recent historical trend. In particular, we compute the growth rate of the standard deviation of wave height between 2011 and 2016 (the initial and final years of our weather dataset) for each coordinate.

[44] As the U.S. in 2016 was importing 1,598,965 million dollars worth of intermediates and capital goods, total imports would go up by 86 billion dollars.

We then assume that the average 5-year growth rate across all coordinates, 6.11%, will remain constant throughout the following 50 years. This would imply an increase in the standard deviation by $1.0611^{10} = 1.8095$, i.e., by 81%. We then simulate, in our model, an increase in the volatility of shipping times, for both low and high risk suppliers, of 81%. Results are shown in Table 8 (second column). We find that this increase in the volatility of weather would increase the average number of foreign suppliers by 42.9% and reduce aggregate imports and the aggregate import share by 3.5% and 9.3%, respectively. This climate change shock has negative welfare effects as it reduces the U.S. real income by 1.8%.

5 Conclusions

In this paper, we use U.S. Census shipment-level data to construct a novel measure of supply chain risk, using weather to generate exogenous variation. We then explore the effects of shipping delays on firms' sales 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 suggest that U.S. importers that are more exposed to shipping time volatility feature lower imports, a larger number of foreign suppliers, and a lower concentration of expenditure across suppliers, which indicates that firms actively diversify this source of risk. To rationalize this evidence, we introduce risky delivery times into a standard quantitative model of firm-level importing. We show that weather-induced shipping risk has relevant implications for the aggregate economy.

References

Alessandria, G., Khan, S.Y., Khederlarian, A., Mix, C., Ruhl, K.J., 2023. The aggregate effects of global and local supply chain disruptions: 2020--2022. Journal of International Economics , 103788.

Alessandria, G., Ruhl, K.J., 2021. Mitigating international supply-chain risk with inventories and fast transport .

Allan, R.P., Arias, P.A., Berger, S., Canadell, J.G., Cassou, C., Chen, D., Cherchi, A., Connors, S.L., Coppola, E., Cruz, F.A., et al., 2023. Intergovernmental panel on climate change (ipcc). summary for policymakers , 3--32.

Antras, P., Chor, D., 2021. Global Value Chains. Working Paper 28549. National Bureau of Economic Research. URL: https://www.nber.org/papers/w28549, doi:10.3386/w28549. series: Working Paper Series.

Antras, P., Fort, T.C., Tintelnot, F., 2017. The Margins of Global Sourcing: Theory and Evidence from US Firms. American Economic Review 107.

Balboni, C., Boehm, J., Waseem, M., 2023. Firm adaptation in production networks: evidence from extreme weather events in pakistan .

Baldwin, R., Freeman, R., 2021. Risks and global supply chains: What we know and what we need to know. Working Paper 29444. National Bureau of Economic Research. URL: https://www.nber.org/papers/w29444, doi:10.3386/w29444. series: Working Paper Series.

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.

Bernard, A.B., Moxnes, A., Saito, Y.U., 2019. Production networks, geography, and firm performance. Journal of Political Economy 127, 639--688.

Blaum, J., 2019. Global firms in large devaluations .

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.

Boehm, C.E., Levchenko, A.A., Pandalai-Nayar, N., 2023. The long and short (run) of trade elasticities. American Economic Review 113, 861--905.

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

Brancaccio, G., Kalouptsidi, M., Papageorgiou, T., 2020. Geography, transportation, and endogenous trade costs. Econometrica 88, 657--691.

Broda, C., Weinstein, D., 2006. Globalization and the gains from variety. Quarterly Journal of Economics 121, 541--585.

Carreras-Valle, M.J., 2021. Increasing inventories: The role of delivery times .

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. URL: https://doi.org/10.1093/qje/qjaa044, doi:10.1093/qje/qjaa044.

Charoenwong, B., Han, M., Wu, J., 2020. Not coming home: Trade and economic policy uncertainty in american supply chain networks. Available at SSRN 3533827 .

Clark, D.P., Kozlova, V., Schaur, G., 2014. Supply chain uncertainty in ocean transit as a trade barrier. Proceedings, working paper, Department of Economics, University of Tennessee .

U.S. Department of Commerce, T., 2023. Indo-pacific economic framework for prosperity (ipef) supply chain agreement .

De Sousa, J., Disdier, A.C., Gaigné, C., 2020. Export decision under risk. European Economic Review 121, 103342.

Dunbar, G.R., Steingress, W., Tomlin, B., 2023. Climate Variability and International Trade. Technical Report. Bank of Canada.

Edmond, C., Midrigan, V., Xu, D.Y., 2015. Competition, markups, and the gains from international trade. American Economic Review 105, 3183--3221.

Eichenbaum, M.S., Hurst, E., Parker, J.A., et al., 2020. Nber macroeconomics annual 2019, volume 34. NBER Books .

Ersahin, N., Giannetti, M., Huang, R., 2023. Supply Chain Risk: Changes in Supplier Composition and Vertical Integration. Technical Report. National Bureau of Economic Research.

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. Trade Credit, and Financial Frictions in General Equilibrium (February 18, 2023) .

Evans, C.L., Harrigan, J., 2005. Distance, time, and specialization: Lean retailing in general equilibrium. American Economic Review 95, 292--313.

Federal Reserve Bank of New York, 2021. Supplemental Survey Report. URL: https://www.newyorkfed.org/medialibrary/media/survey/business_leaders/2021/2021_02supplemental.pdf?la=en.

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., 2023. Entrepot: Hubs, scale, and trade costs. Forthcoming, American Economic Journal: Macroeconomics.

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. URL: http://www.ingentaconnect.com/content/aea/aer/2014/00000104/00000003/art00003.

Grossman, G.M., Helpman, E., Lhuillier, H., 2023. Supply chain resilience: Should policy promote international diversification or reshoring? .

Halpern, L., Koren, M., Szeidl, A., 2015. Imported Inputs and Productivity. American Economic Review 105, 3660--3703. URL: https://www.aeaweb.org/articles?id=10.1257/aer.20150443&within%5Btitle%5D=on&within%5Babstract%5D=on&within%5Bauthor%5D=on&journal=1&q=halpern&from=j, doi:10.1257/aer.20150443.

Handley, K., Limão, N., Ludema, R.D., Yu, Z., 2020. Firm Input Choice Under Trade Policy Uncertainty. Technical Report. National Bureau of Economic Research.

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., 2022. Trade from space: Shipping networks and the global implications of local shocks .

Heise, S., 2019. Firm-to-firm relationships and the pass-through of shocks: Theory and evidence. FRB of New York Staff Report .

Huang, H., 2017. Germs, roads and trade: Theory and evidence on the value of diversification in global sourcing. Available at SSRN 3095273 .

Hummels, D.L., Schaur, G., 2010. Hedging price volatility using fast transport. Journal of International Economics 82, 15--25. URL: https://linkinghub.elsevier.com/retrieve/pii/S0022199610000565, doi:10.1016/j.jinteco.2010.05.002.

Hummels, D.L., Schaur, G., 2013. Time as a Trade Barrier. American Economic Review 103, 2935--2959. URL: http://pubs.aeaweb.org/doi/10.1257/aer.103.7.2935, doi:10.1257/aer.103.7.2935.

Johnson, R.C., Noguera, G., 2012. Accounting for intermediates: Production sharing and trade in value added. Journal of International Economics 86, 224--236. URL: http://www.sciencedirect.com/science/article/pii/S002219961100122X, doi:10.1016/j.jinteco.2011.10.003.

Kamal, F., Krizan, C.J., Monarch, R., 2015. Identifying Foreign Suppliers in US Merchandise Import Transactions URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2596589.

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. Technical Report. National Bureau of Economic Research.

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.

Pankratz, N.M., Schiller, C.M., 2022. Climate change and adaptation in global supply-chain networks .

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. The Review of Economic Studies , rdt018.

Tintelnot, F., Kikkawa, A.K., Mogstad, M., Dhyne, E., 2018. Trade and domestic production networks. Technical Report. National Bureau of Economic Research.

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 .

White House, T., 2021. Building resilient supply chains, revitalizing american manufacturing, and fostering broad-based growth: 100-day reviews under executive order 14017. A Report by The White House .

Young, I., Zieger, S., Babanin, A.V., 2011. Global trends in wind speed and wave height. Science 332, 451--455.

A Appendix: Data Construction and Summary Statistics

A.1 Data Construction

In this section, we describe the steps taken to clean the LFTTD data. First, we drop 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 shipment by the quantity shipped. Fourth, we translate the nominal shipment values into real values 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.[45] 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 as follows. First, errors due to manual data entry are likely low because most firms use customs brokers for their official customs invoice and nearly all entries are filed electronically using specialized customs software. Second, the MID is used for regulatory purposes, such as enforcing anti-dumping measures or tracking compliance with U.S. restrictions for textile shipments, which provide an incentive for the U.S. government to ensure that the MIDs are correct. Third, as an external validation, Kamal and Monarch (2018) assess whether the MID can distinguish between distinct exporters using Chinese data: they construct artificial MIDs from exporter names and addresses in the Chinese Annual Survey of Industrial Firms, and show that they tend to be unique firm identifiers within sectors.

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

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 classify an importer-exporter pair as related if it had a related-party flag for any transaction in the given year.

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, and split them into trips using the arrival date in the U.S. and the export departure date abroad. In some cases, however, the U.S. arrival date is possibly misreported. For example, if transactions 1-5 depart abroad on 6/22 and arrive in the U.S. on 7/5, transaction 6 departs on 6/23 and purportedly arrives on 6/24, and transactions 7-10 depart on 6/25 and arrive on 7/5, then the procedure 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 previous 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. The resulting dataset contains trips with completely non-overlapping foreign departure and importation dates for each vessel.

A.2 The Determinants of Shipping Times

First, Table 9 shows average shipping times and their standard deviation for vessel-based shipments by region of 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 9: Shipping Times by Region

	(1)	(2)	(3)
	Avg. Time	Std. Time	Total Value ($Bill.)
Canada	8.015	25.95	67
Latin America	5.014	25.91	257
South America	19.08	29.6	254
Europe	15.29	20.72	1,160
Asia	17.32	23.7	2,330
Oceania	26.75	25.2	53
Africa	27.37	26.89	113
Other	16.66	21.38	20

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

We then examine the factors affecting vessel-borne shipping times to motivate our residualization procedure to construct shipping risk for ocean shipments. We regress each transaction's log shipping time separately on fixed effects for the foreign port of departure (p_e), U.S. port of entry (p_i), and the port combination, which we refer to as route.[46] These regressions yield an R^2 of 0.24, 0.43, and 0.63, respectively, indicating that the route explains nearly two thirds of the variation. Replacing the route with route-by-month fixed effects raises the R^2 to 0.71, and further adding buyer fixed effects increases the R^2 to 0.73. Thus, the identity of the buyer matters for shipping times even on the same route.

We next analyze more thoroughly the role played by the season of the year, related party status (a), shipping weight, and charges, and present the results in Table 10. All regressions include fixed effects for the route. In column (1), we test whether seasonality affects shipping times by adding dummies for each quarter of the year. Shipping times for a given port pair are about $2-3\%$ shorter in the summer quarters of the northern hemisphere, highlighting the 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 a negative one between shipping charges and delivery times, conditional on weight.

[46]We use a coarser definition of route for the following sections, and define it as a combination of a foreign port of exit and a U.S. port of entry. In our analysis of weather risk below, we instead differentiate routes also by the intermediate stops made between the exit and entry ports to more accurately capture the weather conditions affecting the vessel on its journey.

Table 10: Factors Affecting Shipping Times

Dep. Var.: Log Shipping Times	(1)	(2)	(3)	(4)	(5)
Q2	-0.028^{***}				
	(0.000)				
Q3	-0.029^{***}				
	(0.000)				
Q4	-0.020^{***}				
	(0.000)				
Related-Party		0.014^{**}			
		(0.000)			
Log Shipment Weight			0.008^{***}		0.010^{***}
			(0.000)		(0.000)
Log Shipping Charges				0.004^{***}	-0.003^{***}
				(0.000)	(0.000)
R^2	0.616	0.616	0.616	0.616	0.617
Route FE	Y	Y	Y	Y	Y
Observations (thousands)	35,480	35,480	35,480	35,480	35,480

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.

A.3 Route Summary Statistics

Table 11 presents some summary statistics on our route segments. The first row shows that our sample covers 32,236 route segments (origin - destination port pairs), of which 28,782 have a foreign origin and a U.S. destination port and the remainder are between foreign ports. The ports are selected using the lists of U.S. and foreign ports provided by the U.S. Census.[47] The following rows show that of the segments ending in the U.S. about a third start in Europe, around a quarter in Asia, and the remaining segments mostly start in Canada or in South America. Relatively few segments between two foreign ports start in Canada, as most shipments from stop in the U.S. The last five columns present some statistics on the distance of the route segments. The average segment ending in the U.S. is significantly longer (10,752km) than the average segment between foreign ports (4,642km), since most shipments to the U.S. have to cross either the Atlantic or the Pacific Ocean.

[47]See https://www.census.gov/foreign-trade/schedules/d/distcode.html and https://www.cbp.gov/sites/default/files/assets/documents/2017-Feb/appendix_f_0.pdf.

Table 11: Route Segments Summary Statistics

	(1) Foreign to U.S.	(2) Between Foreign Ports
Number of Segments	28,782	3,454
Starting in...		
- Europe	8,984 (31.2%)	781 (22.6%)
- Asia	8,020 (27.9%)	1,265 (36.6%)
- South and Latin America	4,636 (16.1%)	1,206 (34.9%)
- Canada	5,550 (19.3%)	59 (1.7%)
- Other	1,592 (5.5%)	143 (4.1%)
Distance (km)		
Mean	10,752	4,642
Standard Deviation	5,856	5,260
P5	2,030	176
P50	9,556	2,192
P95	20,613	17,088

Source: Eurostat SeaRoute program. The first row presents the total number of route segments. The following rows split out the number of segments by origin continent. Percentages are as a total of all segments. The final five rows show some statistics on the route segment distance.

A.4 Weather Summary Statistics

Table 12 provides some summary statistics on the mean and standard deviation of the significant wave height and its (absolute) direction across all days and route segments 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 5.5 meters.

Figure 2 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, but very little in the summer.

Table 12: Weather conditions: summary statistics

	Mean	Sd	p5	p50	p95
Significant wave height (m)					
All	2.6	1.5	0.6	2.3	5.5
North Atlantic	2.2	1.3	0.7	1.9	4.5
South Atlantic	2.7	1.3	1.1	2.4	5.2
North Pacific	2.3	1.2	0.7	2.1	4.6
South Pacific	2.9	1.4	1.1	2.6	5.7
Indian Ocean	2.9	1.6	0.6	2.7	5.9
Significant wave direction (degrees)					
All	202	80	52	216	318
North Atlantic	181	102	31	191	333
South Atlantic	216	65	92	220	314
North Pacific	186	98	40	195	325
South Atlantic	215	67	83	223	314
Indian Ocean	209	62	81	218	295

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

Figure 2: Standard Deviation of Wave Height in Different Seasons

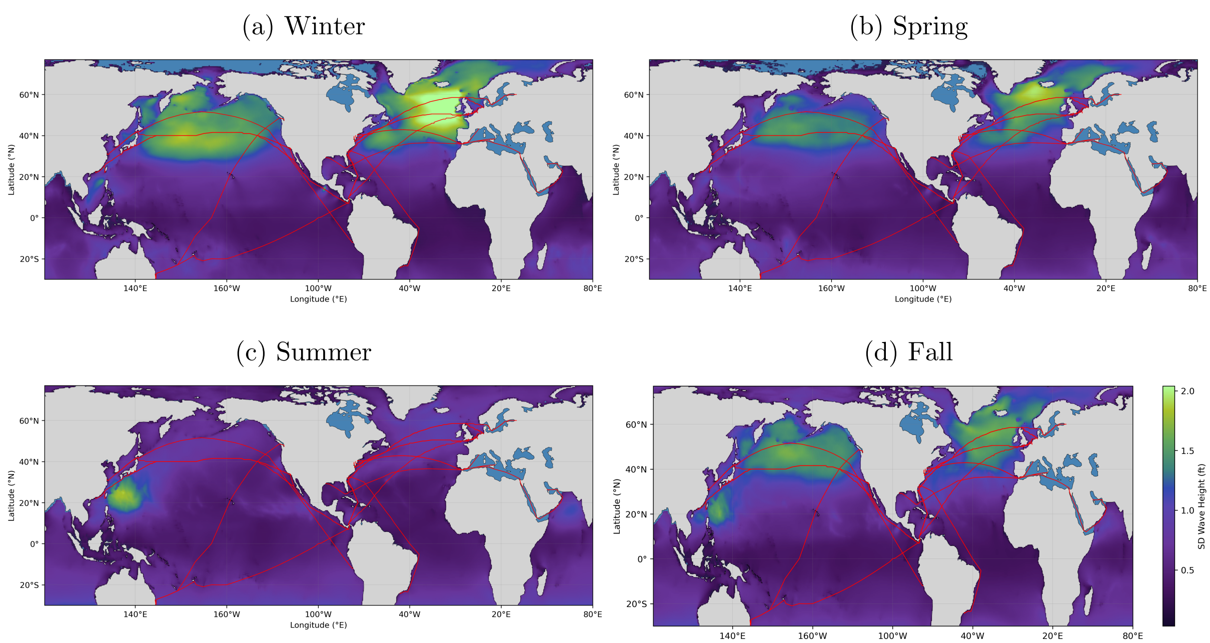

Notes: The figure shows the standard deviation of wave height across all days from 2011-2016 for different seasons. Seasons are based on the Northern Hemisphere, i.e. December, January and February are winter months, etc.

A.5 Analyzing Vessel Movements with AIS Data

Vessel movements vs. routes We construct the route taken by vessels between ports across the ocean using Eurostat's SeaRoute program. This program computes the shortest maritime routes from a network of lines based on the most frequent maritime routes.[48] For a subset of them, we compare our constructed routes to the actual vessel movements using AIS data from MarineTraffic, a provider of ship tracking services.[49] 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 3a 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 3b 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 3c 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.

[48] See https://github.com/eurostat/searoute.

[49] https://www.marinetraffic.com.

Figure 3: Vessel Movements vs Routes

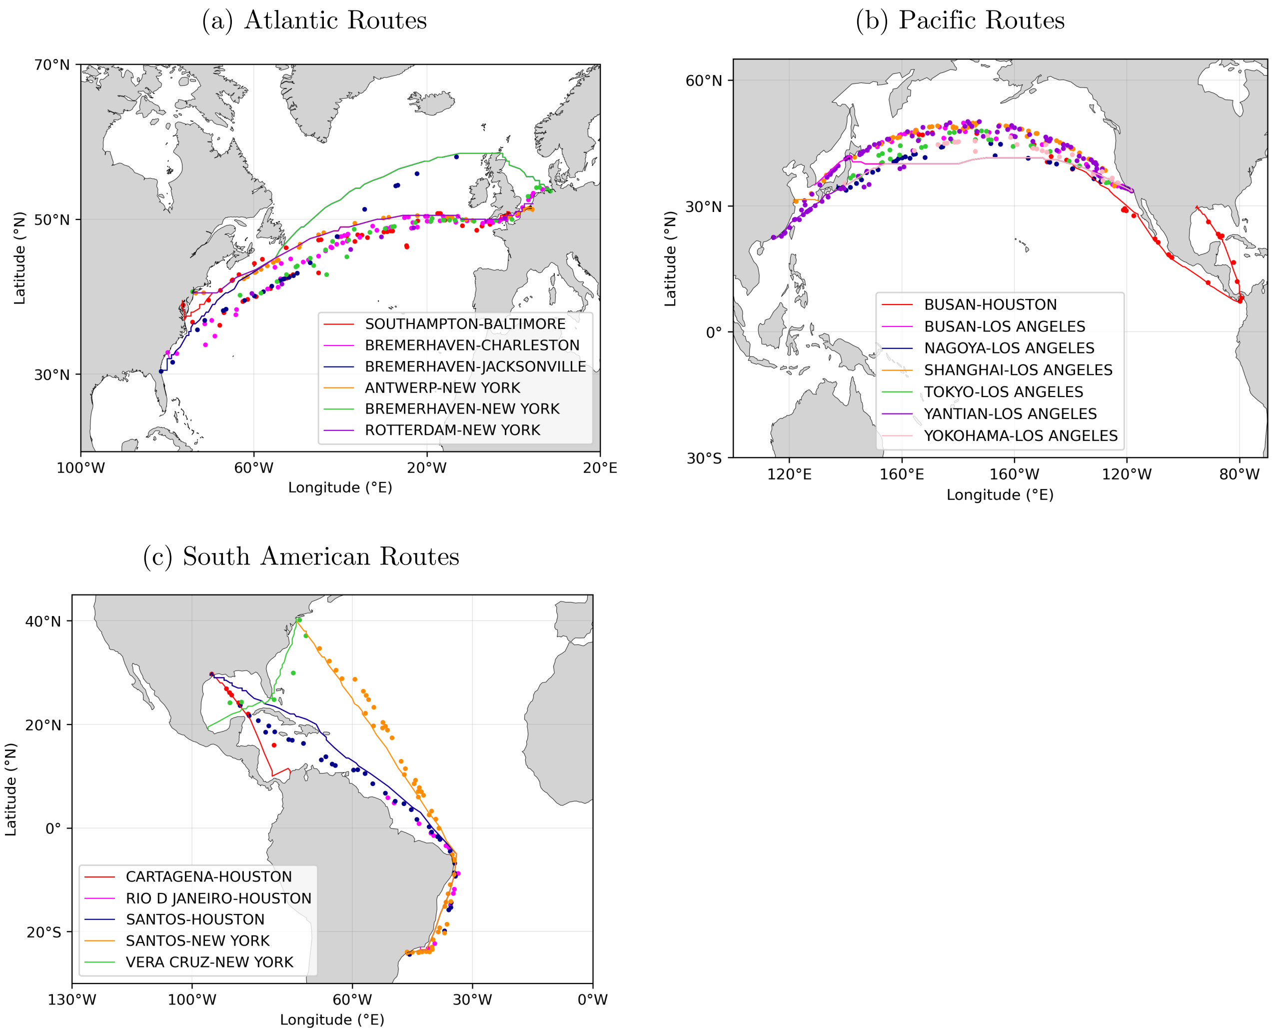

Source: MarineTraffic and authors' calculations. Notes: The figure shows the locations of vessels traveling between selected ports against the routes from the SeaRoute program used for the analysis.

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.[50] The first column of Table 13 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 13 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.

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

Table 13: Effect of Weather on Shipping Times

Dep. Var:	Vessel speed	Vessel speed
Wave Height	0.0584***	
	(0.0122)	
Direction	0.0007***	
	(0.0002)	
Wave Height × Direction^s	−0.0002**	
	(0.0001)	
Wind Speed		0.0024**
		(0.0011)
Wind Direction		0.0003**
		(0.0001)
Wind Speed × Wind Direction^s		−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 We verify our assumption that vessels travel at approximately constant speed across the ocean 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 4 a bin scatter of the distance share against the share of voyage time. We find that the fit line is approximately on the 45 degree line throughout the journey. 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 4: Vessel Distance Covered vs Voyage Time Elapsed

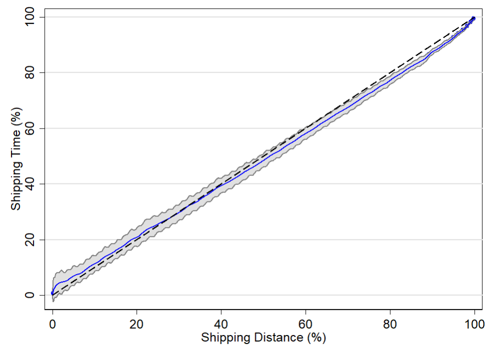

Source: MarineTraffic and authors' calculations. Notes: The figure plots the share of the distance completed against the share of voyage time elapsed for 18 vessels for which we have complete voyage information.

B Appendix: Additional Results

B.1 Effect of Weather on Shipping Times

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.[51] 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. Our results from the weather regressions are similar to before. Table 14 shows the results using such alternative approach. The results are similar but coefficients are smaller in absolute value and less significant.

Table 14: Effect of Weather on Shipping Times – Alternative Weather Conditions

Dep. Var:	$\hat{t}^s_{xhrtvfa}$	$\hat{t}^s_{xhrtvfa}$	$\hat{t}^s_{xhrtvfa}$
Wave Height^s	-0.014^{***}	-0.014^{***}	-0.009^{***}
	(0.000)	(0.000)	(0.000)
Direction^s		-0.004^{***}	0.008^{***}
		(0.000)	(0.000)
Wave Height$^s\times$ Direction^s			-0.006^{***}
			(0.000)
Route FE	Y	Y	Y
Observations	5,728,000	5,728,000	5,728,000

[51]We use AIS data in Appendix A.5 to show that this assumption approximately holds in vessel tracking data.

Table 15: Shipping Time Risk and Import Demand – Residualized Times

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.:	Number of Suppliers	Number of Routes	HHI over Suppliers	HHI over Supplier-Routes	Value per Supplier-Year	Value Imported per Year
Std Time	0.040^{***}	0.071^{***}	-0.028^{***}	-0.042^{***}	-0.131^{***}	-0.091^{***}
	(0.004)	(0.005)	(0.001)	(0.001)	(0.006)	(0.006)
Avg Time	0.015	0.077^{***}	-0.021^{***}	-0.038^{***}	0.072^{***}	-0.056^{**}
	(0.015)	(0.019)	(0.006)	(0.007)	(0.024)	(0.024)
Importer FE	Y	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Observations	328,000	328,000	328,000	328,000	328,000	328,000

Notes: 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.

B.2 Robustness of the Risk Regressions

In this section, we show that our results on the impact of shipping risk on importers' behavior are robust to a variety of alternative specifications. First, in Table 15 we repeat the analysis shown in Table 6 using only the residualized shipping times obtained from Step I. While the point estimates are smaller than the ones in Table 6, the implied differential effects are slightly larger, which likely reflects the fact that the volatility of the residualized shipping times captures additional sources of risk other than the weather conditions.

Second, in Table 16, we show that our baseline results hold when we use only the riskiness of the importers' main supplier, rather than a weighted average across all suppliers. We define the main supplier as the one with the highest value of shipments of each product over the previous five years. Going from the $25th$ to the $75th$ percentile of shipping risk increases the number of suppliers used by 2.7% and the number of routes by 3.7%. In columns (3) and (4), we see that higher shipping risk reduces the HHI over suppliers and supplier-route combinations. Finally, columns (5) and (6) show that increasing shipping risk from the $25th$ to the $75th$ percentile lowers imports per supplier by 12.2% and total imports by 9.5%.

We next examine the effect of using a different residualization for shipping times. We assume that instead of route-by-month fixed effects the residualization includes supplier-by-month fixed effects. Suppliers might be slower fulfilling orders at different times of the year, or they might choose different routes. Thus, our residualized shipping times are now

$$\tilde{t}^s_{xhrtvfa} \equiv \ln(T^s_{xhrtvfa}) - \hat{\eta}\ln(C^s) - \hat{\rho}\ln(W^s) - \hat{\gamma}_{xt} - \hat{\psi}_r - \hat{\xi}_v - \hat{\delta}_f - \hat{\omega}_a,$$

where γ_{xt} are supplier-by-month fixed effects and we additionally include route fixed effects ψ_r to pick up fixed route conditions. We then re-run our main risk regression using this alternative measure of shipping times. The results in Table 17 are similar to before.

Table 18, we run our baseline specification without firm fixed effects and find similar results as before. An increase in the standard deviation of shipping times has a positive and significant effect on both the number of suppliers and routes. Similarly, we find that the relationship between shipping risk and importers' concentration of suppliers remains significant.. The effect of risk on most variables strengthens.

Table 19 includes firms with only one supplier. Here, we find that the relationship between shipping risk and the number of suppliers and routes remains significantly positive, but decreases slightly in magnitude. Compared to the baseline, the inclusion of firms with one supplier weakens the association between shipping risk and HHI over suppliers slightly. Value per supplier and total value imported remains significantly negatively correlated with shipping risk.

In Table 20 we include an additional control for the inventory-sales ratio. We obtain the end 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 routes used. Supplier concentration and value imported remain significantly negatively correlated with shipping risk. Note that the sample size declines since we do not have inventory information for some firms.

In Table 21, we run our risk regression with weather-induced shipping risk using an alternative methodology to assign weather to each transaction. We decompose each trip segment into smaller sub-segments of 1,000km length and assume that vessels travel through these areas at constant speed. 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. The results are very similar to before. The coefficients slightly strengthen for the number of suppliers and routes, but weaken slightly for imports per supplier and overall imports.

Lastly, we estimate a variant of our main specification, equation (5),

$$d_{fht} = \beta_1 \ln(\widehat{StdTime}^{vessel}_{fht-5,t-1}) + \beta_2 X_{fht} + \gamma_f + \gamma_h + \gamma_t + \epsilon_{fht}$$

where d_{fht} is a dummy that is equal to one if firm f uses air shipments for HS10 h in year t, and $\widehat{StdTime}^{vessel}_{fht-5,t-1}$ is the weighted average standard deviation of vessel-based shipments over the last five years across all routes and suppliers used by the importer-product, computed similarly to before. In contrast to our main specification, we do not include other modes of transportation when computing this variable. The controls X_{fht} are the same as before, except that split up the importer's total log value of imports into imports by vessel and imports by airplane. Table 22 documents that higher shipping risk is associated with a higher likelihood of using air shipments, suggesting that firms use air transportation to hedge ocean shipping risk.

Table 16: Shipping Time Risk and Import Demand with Main Supplier

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.:	Number of Suppliers	Number of Routes	HHI over Suppliers	HHI over Supplier-Routes	Value per Supplier-Year	Value Imported per Year
Std Time	0.030***	0.040***	−0.028***	−0.029***	−0.132***	−0.103***
	(0.004)	(0.005)	(0.002)	(0.002)	(0.006)	(0.006)
Avg Time	0.002	0.044**	−0.017***	−0.030***	−0.062***	−0.061***
	(0.014)	(0.017)	(0.006)	(0.006)	(0.021)	(0.022)
Importer FE	–	–	–	–	–	–
Product FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Observations	237,000	237,000	237,000	237,000	237,000	237,000

Notes: 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.

Table 17: Shipping Time Risk and Import Demand (Supplier-Month Residualization)

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.:	Number of Suppliers	Number of Routes	HHI over Suppliers	HHI over Supplier-Routes	Value per Supplier-Year	Value Imported per Year
Std Time	0.041***	0.086***	−0.026***	−0.046***	−0.123***	−0.082***
	(0.004)	(0.005)	(0.002)	(0.001)	(0.006)	(0.006)
Avg Time	0.037*	0.108***	−0.023***	−0.040***	−0.015	0.022
	(0.021)	(0.022)	(0.007)	(0.008)	(0.030)	(0.028)
Importer FE	Y	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Observations	328,000	328,000	328,000	328,000	328,000	328,000

Notes: 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.

Table 18: Shipping Time Risk and Import Demand without Importer Fixed Effects

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.:	Number of Suppliers	Number of Routes	HHI over Suppliers	HHI over Supplier-Routes	Value per Supplier-Year	Value Imported per Year
Std Time	0.043^{***}	0.069^{***}	-0.029^{***}	-0.042^{***}	-0.157^{***}	-0.114^{***}
	(0.009)	(0.009)	(0.002)	(0.002)	(0.010)	(0.007)
Avg Time	-0.009	0.038	-0.014^{**}	-0.027^{***}	-0.078^{***}	-0.087^{***}
	(0.019)	(0.024)	(0.006)	(0.007)	(0.025)	(0.024)
Importer FE	–	–	–	–	–	–
Product FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Observations	328,000	328,000	328,000	328,000	328,000	328,000

Notes: 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.

Table 19: Shipping Time Risk and Import Demand Including Firms with One Supplier

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.:	Number of Suppliers	Number of Routes	HHI over Suppliers	HHI over Supplier-Routes	Value per Supplier-Year	Value Imported per Year
Std Time	0.034^{***}	0.068^{***}	-0.024^{***}	-0.042^{***}	-0.128^{***}	-0.094^{***}
	(0.004)	(0.005)	(0.001)	(0.002)	(0.005)	(0.006)
Avg Time	0.021	0.079^{***}	-0.017^{***}	-0.042^{***}	-0.066^{***}	-0.045^{*}
	(0.015)	(0.018)	(0.005)	(0.006)	(0.024)	(0.025)
Importer FE	Y	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Observations	502,000	502,000	502,000	502,000	502,000	502,000

Notes: 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.

Table 20: Shipping Time Risk and Import Demand with Inventory-Sales Ratio

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.:	Number of Suppliers	Number of Routes	HHI over Suppliers	HHI over Supplier-Routes	Value per Supplier-Year	Value Imported per Year
Std Time	0.046^{***}	0.083^{***}	-0.029^{***}	-0.045^{***}	-0.135^{***}	-0.090^{***}
	(0.004)	(0.006)	(0.002)	(0.002)	(0.007)	(0.007)
Avg Time	0.017	0.082^{***}	-0.021^{***}	-0.039^{***}	-0.076^{***}	-0.059^{**}
	(0.017)	(0.023)	(0.008)	(0.008)	(0.027)	(0.028)
Inventory-Sales Ratio	-0.018	-0.036^{**}	0.001	0.005^{*}	-0.062	-0.081
	(0.013)	(0.015)	(0.003)	(0.003)	(0.048)	(0.060)
Importer FE	Y	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Observations	237,000	237,000	237,000	237,000	237,000	237,000

Notes: 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.

Table 21: Shipping Time Risk and Import Demand – Alternative Weather Conditions

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.:	Number of Suppliers	Number of Routes	HHI over Suppliers	HHI over Supplier-Routes	Value per Supplier-Year	Value Imported per Year
Std Time	0.084^{***}	0.138^{***}	-0.055^{***}	-0.075^{***}	-0.158^{***}	-0.074^{***}
	(0.009)	(0.010)	(0.003)	(0.003)	(0.010)	(0.012)
Avg Time	9.467^{***}	12.850^{***}	-4.459^{***}	-4.701^{***}	-6.764^{***}	2.704^{**}
	(0.834)	(0.989)	(0.369)	(0.334)	(1.123)	(1.109)
Importer FE	Y	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Observations	72,500	72,500	72,500	72,500	72,500	72,500

Notes: 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.

Table 22: Shipping Time Risk and Import Demand with Air Shipments

Dep. Var.:	Air Shipments
Std Time	0.009***
	(0.001)
Avg Time	0.006
	(0.007)
Importer FE	Y
Product FE	Y
Year FE	Y
Controls	Y
Observations	328,000

Notes: 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.

B.3 Additional Results for Section 4

B.3.1 Derivation of equation (10)

The CES assumption implies that the optimal demand for variety of firm f is:

$$y_f = p_f^{-\sigma} A$$

where $A = \frac{S}{P^{1-\sigma}}$ is a demand shifter which depends on income and price index. A firm f producing a certain variety will have therefore total revenues equal to:

$$R_f = \left(\frac{y_f}{A}\right)^{-\frac{1}{\sigma}} y_f = A^{\frac{1}{\sigma}} y_f^{\frac{\sigma-1}{\sigma}} = P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma} y_f^{\frac{\sigma-1}{\sigma}}.$$

Firms maximize profits 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 optimal level of labor conditional on the choices for N, x and x_D. Combining equations (6) and (9), firms' profits in the second stage are given by

$$\max_{l_p} \quad \pi_f = \varphi_f^{\frac{\sigma-1}{\sigma}} l^{(1-\gamma)\frac{\sigma-1}{\sigma}} \left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + \left(\sum_{i=1}^{N} \alpha_i\right)^{\frac{\varepsilon-1}{\varepsilon}} x^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\gamma\frac{\varepsilon}{\varepsilon-1}\frac{\sigma-1}{\sigma}} P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma} - p_D x_D - Np^* x - wl - wNF, \tag{15}$$

where we have used the fact that $x = x_i$ since the inputs are ex-ante symmetric, and α_i are the quality shocks which depend on the realized shipping times. The optimal choice of production labor is:

$$l_f = \left[\tilde{\gamma} \varphi_f^{\frac{\sigma-1}{\sigma}} \left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + \left(\sum_{i=1}^{N} \alpha_i\right)^{\frac{\varepsilon-1}{\varepsilon}} x^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\gamma\frac{\varepsilon}{\varepsilon-1}\frac{\sigma-1}{\sigma}} P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma} w^{-1} \right]^{\frac{1}{1-\tilde{\gamma}}} \tag{16}$$

where $\tilde{\gamma} \equiv (1-\gamma)\frac{\sigma-1}{\sigma}$. In the first stage, taking l_f as given, the firm maximizes *expected* profits. Plugging the expression for l_f into equation (15), expected profits are

$$\max_{x_D,x,N} \quad \chi_f \mathbb{E}\left[\left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + \left(\sum_{i=1}^{N} \alpha_i(d_i)\right)^{\frac{\varepsilon-1}{\varepsilon}} x^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\psi} \right] - p_D x_D - Np^* x - wNF, \tag{17}$$

where χ_f and ψ are functions of firm efficiency, general equilibrium objects and parameters, and $\alpha(d_i)$ is given by (7).

Figure 5: Identification of parameters

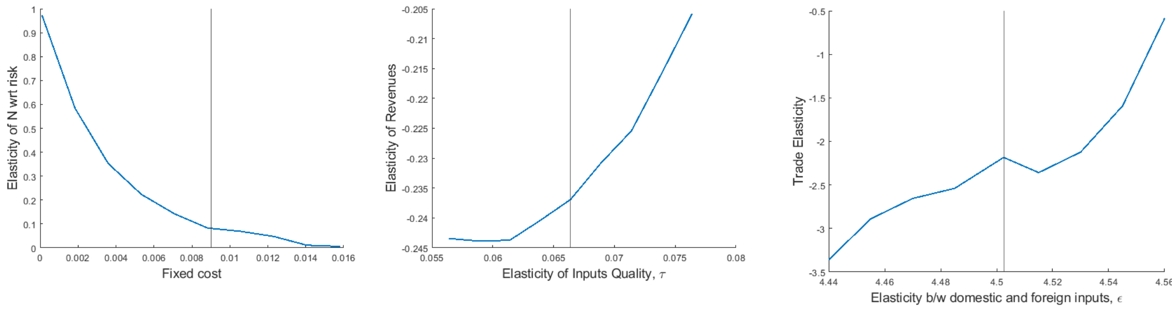

Notes: The graph on the left plots the elasticity of N with respect to risk implied by the model as a function of the fixed cost F, holding constant the other parameters at their calibrated values. The graph in the middle plots the elasticity of revenues with respect to average shipping time implied by the model as a function of τ, holding constant the other parameters at their calibrated values. The graph on the right plots the trade elasticity implied by the model as a function of ϵ, holding constant the other parameters at their calibrated values.

B.3.2 Additional Figures