

Blaum, Joaquin; Esposito, Federico; Heise, Sebastian

Working Paper

Input sourcing under supply chain risk: Evidence from U.S. manufacturing firms

Staff Reports, No. 1141

Provided in Cooperation with:

Federal Reserve Bank of New York

Suggested Citation: Blaum, Joaquin; Esposito, Federico; Heise, Sebastian (2025) : Input sourcing under supply chain risk: Evidence from U.S. manufacturing firms, Staff Reports, No. 1141, Federal Reserve Bank of New York, New York, NY, <https://doi.org/10.59576/sr.1141>

This Version is available at:

<https://hdl.handle.net/10419/315924>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

NO. 1141
FEBRUARY 2025

Input Sourcing Under Supply Chain Risk: Evidence from U.S. Manufacturing Firms

Joaquin Blaum | Federico Esposito | Sebastian Heise

Input Sourcing Under Supply Chain Risk: Evidence from U.S. Manufacturing Firms

Joaquin Blaum, Federico Esposito, and Sebastian Heise

Federal Reserve Bank of New York Staff Reports, no. 1141

February 2025

<https://doi.org/10.59576/sr.1141>

Abstract

We study the effect of climate risk on how firms organize their supply chains. We use transaction-level data on U.S. manufacturing imports to construct a novel measure of input sourcing risk based on the historical volatility of ocean shipping times. Our measure isolates the unexpected component of shipping times that is induced by weather conditions along more than 40,000 maritime routes. We first document that unexpected shipping delays induced by weather shocks have significant negative effects on importers' revenues, profits, and employment. We then show that more exposed firms actively diversify the risk of weather delays by using more routes and sourcing from more foreign suppliers, although their total imports decline. To rationalize these findings, we introduce shipping time risk into a general equilibrium model of importing with firm heterogeneity. Our quantitative analysis predicts substantial costs for the U.S. economy associated with different sources of supply chain risk.

JEL classification: F10, F15, Q54

Key words: supply chains, climate shocks, shipping time risk, input sourcing

Heise: Federal Reserve Bank of New York (email: sebastian.heise@ny.frb.org). Blaum: Boston University (email: jblaum@bu.edu). Esposito: Tufts University (email: federico.esposito@tufts.edu). We thank our discussants Davin Chor, Simona Giglioli and Shafaat Yar Khan, as well as Laura Alfaro, David Atkin, Teresa Fort, Matthew Gudgeon, Fadi Hassan, Georg Schaur, Peter Schott and participants at numerous seminars and conferences for their valuable comments. The 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 and 2407 (CBDRB-FY21-P1883-R8915, CBDRB-FY23-P1883-R10453, CBDRB-FY23-P1883-R10672, CBDRB-FY25-0124).

This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the author(s) and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the author(s).

To view the authors' disclosure statements, visit
https://www.newyorkfed.org/research/staff_reports/sr1141.html.

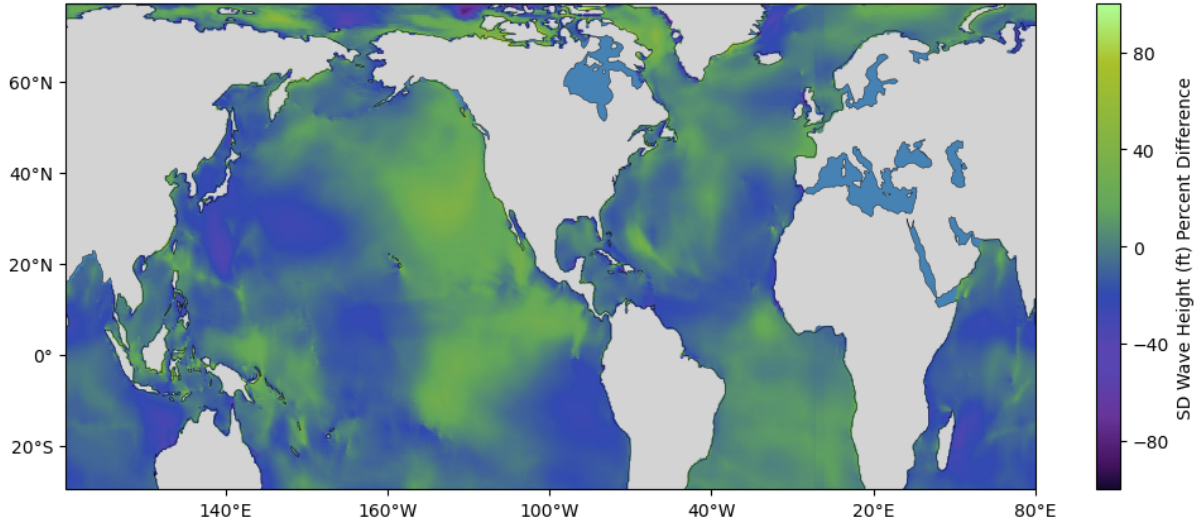
1 Introduction

The past decades have seen a dramatic transformation in the international organization of production, with intermediate inputs accounting for two-thirds of global trade and complex global value chains spanning multiple countries ([Johnson and Noguera \(2012\)](#), [Antràs and Chor \(2022\)](#)). For many firms, the timely delivery of their inputs is a crucial element of the production process ([Hummels and Schaur \(2010\)](#), [Hummels and Schaur \(2013\)](#)). However, the increased reliance on imported inputs has exposed firms to a host of supply chain risks that can adversely impact the timeliness of their inputs. Salient recent examples include the increased frequency of extreme weather events associated with climate change, the geopolitical risk stemming from attacks by Houthi militias in the Red Sea, and the strain on port infrastructure that followed the Covid pandemic (e.g., [Brancaccio et al. \(2024\)](#)). How do these and other supply chain risks impact firms' import behavior? Do firms adapt their supply chains to hedge the delay risk stemming from these shocks? Answering these questions is challenging due to the inherent difficulty in developing credible measures of firm-level risk.

We shed light on these questions by focusing on a specific but important source of risk: weather shocks. We start by establishing that weather conditions have a significant effect on the ocean shipping times of U.S. imports. To do so, we rely on transaction-level import data on ocean shipments provided by the U.S. Census Bureau as well as detailed data on oceanic wave conditions along more than 40,000 maritime routes. We exploit this relationship to measure the component of shipping times that is induced by weather, which we interpret as unexpected by U.S. importers given the unpredictability of high frequency ocean conditions.

Armed with this measure, we establish two key empirical results. We first show that unexpected shipping delays induced by weather shocks have large and disruptive effects on U.S. importers' production levels and profit margins. Second, we build a measure of risk based on the volatility of the weather-induced shipping times. As Figure 1 shows, the standard deviation of wave height has increased in many locations over the past decade. We show that firms systematically respond to this type of weather risk along different margins of adjustment. More exposed firms rely on more routes and foreign suppliers, and they lower both their imports and the concentration of expenditure across routes and suppliers. We next incorporate risky shipping times into a quantitative model of firm-level importing, and calibrate the model to match salient features of the data. We use our framework to quantify the impact of three scenarios of heightened risk: climate change, geopolitical tensions in the Red Sea, and port congestion. Overall, we find that these shocks trigger an important risk diversification response by importers, but nevertheless reduce U.S. real income by 0.4% to 1.33%.

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



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

The cornerstone of our analysis is the U.S. Census Bureau’s Longitudinal Firm Trade Transactions Database (LFTTD), which provides transaction-level data recording the identity of the U.S. importer and its foreign supplier, as well as information about the product, quantity, and value transacted for the universe of U.S. imports. Importantly, the data record the delivery time between the foreign port of exit and the U.S. port of entry and, for ocean shipments, the vessel identity. Since our customs data do not contain details on each vessel’s journey across the ocean, we propose an algorithm that uses the vessel name, foreign port stops, 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 program.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).

Our methodology relies crucially on the measurement of the components of shipping times that are unexpected to importers. We use a rich set of fixed effects and controls to remove components that are presumably known at the time the inputs are bought, including the

¹Ganapati et al. (2024) 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.

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 view this assumption as plausibly satisfied in the data. On the one hand, most maritime shipments to the U.S. involve multi-week ocean crossings, and import orders are placed typically many weeks before production finalizes and goods are shipped (see [Deloitte \(2024\)](#)). On the other hand, weather forecasts are reasonably accurate for about 7 days into the future, and only general patterns can be predicted beyond 2 weeks—with ocean wave height being particularly hard to predict given the chaotic nature of ocean dynamics ([Alley et al. \(2019\)](#), [Zhang et al. \(2022\)](#) and [Mishra et al. \(2022\)](#)).

We analyze the effects of shipping delays induced by weather shocks on firms’ outcomes. We identify for each year the shipments that were extremely delayed, which we define as having a weather-induced delivery time larger than the 95th 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 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 weather-delayed 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 supply chain disruptions on firms’ production, and suggest that firms are typically not able to fully hedge their supply chain risk with insurance or financial instruments. We next study whether U.S. importers adjust their sourcing strategy and import demand ex-ante to reduce the potential impact of weather shocks.

To explore whether importers hedge against weather shocks, we build a measure of risk based on the volatility of weather-induced shipping times. In particular, we measure the riskiness of each foreign supplier-route-product combination as the standard deviation of the weather-induced shipping times over 3-year rolling windows. We construct a shift-share exposure to risk for each importer as a weighted average of the risk of its suppliers and routes over the previous 3 years, using pre-determined import shares as weights. We then estimate, 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 weather-induced risk along the extensive and intensive margins. Going from the 25th to the 75th percentile of the shipping risk distribution increases

the number of routes used and the number of foreign suppliers by 7.7% and 4.9%, respectively. Moreover, it reduces the total value imported by 5.1%. Thus, importers with ex-ante riskier supply chains spread their input expenditures among more routes and foreign suppliers, and import less overall. Importantly, the negative effect of risk on imports is estimated conditional on the negative effect that longer shipping times have on import demand, indicating that uncertainty has an additional detrimental impact on international trade.

To rationalize these findings, we incorporate shipping risk into a standard model of importing with firm heterogeneity, along the lines of [Blaum et al. \(2018\)](#), [Gopinath and Neiman \(2014\)](#), and [Halpern et al. \(2015\)](#). Firms can source their inputs domestically or from foreign suppliers. We follow [Hummels and Schaur \(2013\)](#) in their treatment of timeliness by assuming that input qualities are reduced when inputs take longer to arrive, for example due to spoilage, absence of key inputs, etc. The key departure from the literature is that firms are uncertain about shipping times at the time of placing orders. While firms are risk-neutral, the presence of market power with elastic demand introduces curvature in revenues, making expected revenues fall with more volatile input qualities.² Firms can diversify their shipping time risk by sourcing from multiple foreign suppliers, or equivalently by using multiple routes, albeit this strategy is limited by per-supplier fixed costs. We provide conditions under which firms increase their number of foreign suppliers and reduce their import values after a mean-preserving spread to supplier qualities.

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

We assess the impact of various risk-related scenarios that have recently received significant attention, namely, climate change, geopolitical tensions in the Red Sea, and port congestion. The volatility of ocean wave height has increased on average by 0.34% per year between 2011-2023, consistent with work suggesting an increasing likelihood of extreme wave heights

²The imperfect substitutability between labor and material inputs, and between domestic and foreign materials, also introduce curvature in the revenue function.

([Young et al. \(2011\)](#)). We evaluate the effects of an increase in the volatility of ocean wave heights that continues along this trend over the next 50 years on the U.S. economy in our model. In a second exercise, we investigate how the re-routing of commercial ships following the Houthi attacks around the Suez Canal affects the U.S. economy, through an increase in both the average and volatility of navigation time. In a third exercise, we consider the greater variability of waiting times at ports associated with the rise in port congestion that took place in the post Covid period of 2021-2022. For all exercises, we find that, despite a strong risk diversification response along the extensive margin, imports fall substantially, as firms reduce their risk exposure by substituting towards domestic production. This shift increases production costs and prices, reducing U.S. real income by 0.4% to 1.33%.

Related Literature. Our paper contributes to several strands of the literature. First, it relates to work that investigates 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 \(2013\)](#)). While these seminal papers focus on the role of the level of shipping times, i.e., their first moment, we study the effect of the variance of shipping times, i.e., the second moment. Our empirical results show that uncertainty around shipping times has an additional negative effect on import demand. We propose a theory of the firm that incorporates this mechanism in a way that is both tractable and amenable to quantitative analysis.

Second, we contribute to a broader literature that analyzes the impact of uncertainty on firms. Most of the international trade literature on this topic has focused on exports and FDI (e.g., [Ramondo et al. \(2013\)](#), [Fillat and Garetto \(2015\)](#), [Esposito \(2022\)](#), [Baley et al. \(2020\)](#) and [De Sousa et al. \(2020\)](#)). In contrast, we analyze risk on the input side and how it affects firms' sourcing decisions. Only a few papers have studied the effects of sourcing uncertainty on international trade (e.g., [Gervais \(2018\)](#), [Grossman et al. \(2023\)](#), and [Handley et al. \(2024\)](#)). Our contribution to this literature is to develop a novel and plausibly exogenous measure of firm-level shipping time risk using weather shocks, which we use to study the causal impact of risk on importers in the United States. We combine weather data with comprehensive firm-level administrative data and show that importers actively adjust the intensive and extensive margins of importing in response to weather risk.³ Complementary to our work are [Balboni et al. \(2023\)](#) and [Castro-Vincenzi et al. \(2024\)](#), who study how firms diversify their sourcing locations in Pakistan and India, respectively. In contrast to our focus on maritime shipping risk and international trade, these works focus on

³The diversification mechanism we highlight is complementary to firms' use of inventories, as shown by [Alessandria and Ruhl \(2021\)](#) and [Carreras-Valle \(2021\)](#).

adaptation to flood risk and on domestic trade.

Third, we contribute to work that studies the effects of supply chain disruptions on firms (Carvalho et al. (2021), Barrot and Sauvagnat (2016), Boehm et al. (2019), Khanna et al. (2022), Alessandria et al. (2023), Lafrogne-Joussier et al. (2023)).⁴ Relative to this literature, we provide a new way to identify supply shocks using readily available weather data, rather than large, aggregate shocks—such as the Japanese earthquake or the Covid lockdowns. Our measure therefore lends itself to a wide range of applications that require exogenous shocks to firms.

Finally, we use our empirical findings to calibrate a model to shed light on the long-run implications of climate risk. Existing models of firm-level input sourcing typically abstract from supplier risk considerations—e.g., Gopinath and Neiman (2014), Halpern et al. (2015), Antras et al. (2017) and Blaum et al. (2018). Our contribution is to extend a sourcing model to allow for risk and to quantify the impact of risk in general equilibrium. Therefore, our calibrated model can serve as a laboratory to estimate the impact of any type of sourcing risk on U.S. importers.

The remainder of the paper proceeds as follows. Section 2 describes our data 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 in Section 5. Section 6 concludes.

2 Data Construction

In this section, we describe how we measure shipping times, routes, and weather conditions for every import transaction headed to the U.S. Our novel methodology allows us to infer vessels’ shipping routes from Census data and to combine these with the observed weather conditions at granular locations in the oceans for over 40,000 distinct shipping routes. These will be the building blocks of our measure of shipping time risk, which we exploit in the empirical analysis of Section 3.

⁴Also related is work that studies the effects of climate shocks on firms, e.g., Pankratz and Schiller (2024) and Dunbar et al. (2023). More broadly, we contribute to a large literature that studies firm-to-firm relationships, see e.g., Dhyne et al. (2021), Bernard et al. (2019), Esposito and Hassan (2023), Heise (2024). This literature does not typically focus on uncertainty.

2.1 Census Data

Our empirical analysis 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 HS-10 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 and an indicator of whether the transaction is between related parties.⁵ We calculate prices as the value of the shipment divided by the quantity shipped and keep both related party and arm’s-length transactions.

The LFTTD contains several additional variables that are critical to construct our measure of risk. First, each customs record reports the export and the import dates (after customs are cleared), which allows us to construct shipping times. Second, for seaborne imports, we also observe the foreign port of departure, the U.S. port of arrival, and the vessel name. We use this information to construct shipping routes, as explained below.⁶ 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. 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.⁷ Given our focus on supply chains, we restrict our analysis to firms that operate in the manufacturing sector, whose imports are most likely intermediates into production. We also obtain firms’ total sales, cost of materials, and employee compensation from the Census of Manufactures (CMF) in census years (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 payroll.

⁵The foreign seller is identified by a Manufacturer ID (MID), which is an alphanumeric code that combines information on the seller’s country, name, street address, and city. We follow [Kamal et al. \(2015\)](#) and [Kamal and Monarch \(2018\)](#) in combining sellers with the same 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.

⁶Less information is available for other modes of transportation. For non-vessel imports we only have the 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.

⁷We 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 in the first row are reported in billions of 2009 dollars.

Table 1 reports summary statistics of our 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 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

Shipping Times For all shipments, irrespective of their mode of transportation, we calculate the shipping time as the difference, in days, between the import date in the U.S. and the export date from the foreign country. We show statistics of the distribution of shipping times in Table 2. Vessel shipments take on average 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 shipping times. There is less dispersion for other modes of transportation, which have a median shipping time of zero.⁸

Routes and Journeys For seaborne shipments, we develop an algorithm to construct ocean shipping routes and vessels' journeys between ports from the information on the port of entry

⁸Of course, predictable factors such as the origin country or the time of year affect vessel shipping times to the U.S. We show additional statistics on shipping times and their determinants in Appendix A.2.

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.

and port of origin. We assign each transaction in the customs data to a *trip*, defined as a journey of a vessel that begins with the loading of cargo at a foreign port and ends (possibly after some intermediate stops) with the unloading of cargo at a U.S. port. As a starting point, we sort all transactions involving a given vessel by their foreign departure date. We then take all the vessel’s transactions and assign them to a single trip (“Trip 1”). Next, we find the earliest arrival date of the vessel in the U.S. for this trip. If there exists any transaction of the same vessel with an export departure date abroad that is later than this earliest arrival date in the U.S., we assign these transactions to a new trip (“Trip 2”). We continue splitting trips into sub-trips until no further splits are possible.⁹ 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 path taken by the vessel across the ocean on any route segment using Eurostat’s SeaRoute program. This program computes the shortest maritime paths using the network of global shipping lanes and observed vessel movements from satellite data.¹⁰ Our sample includes around 10,500 route segments and 43,000 routes, i.e., distinct ordered sets of route segments ending in any U.S. port. We show that, for a selected sample, these routes closely follow actual vessel movements reported by AIS data (Appendix A.3).¹¹ The upper panel of Figure 2 illustrates the route segments in our data.

For shipments arriving with modes of transportation other than vessel, we only know the country of departure rather than the precise departure port. We therefore approximate the

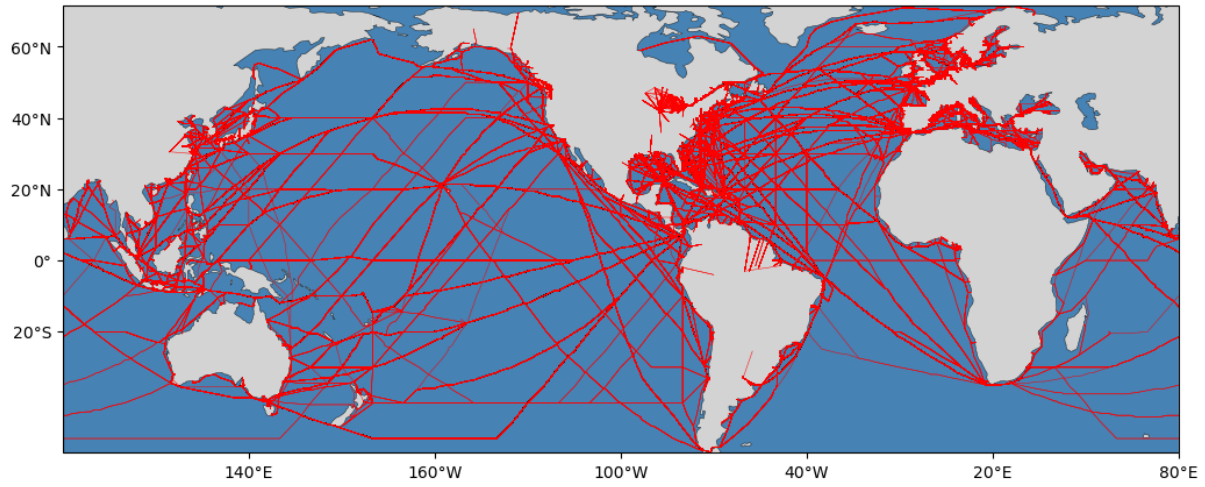
⁹In some instances, the arrival date may be misreported. In Appendix A.1 we explain how we identify such cases and how we refine our algorithm to redefine the trips.

¹⁰The shipping lanes are from the Oak Ridge National Labs CTA Transportation Network Group, Global Shipping Lane Network: <http://geocommons.com/datasets?id=25>.

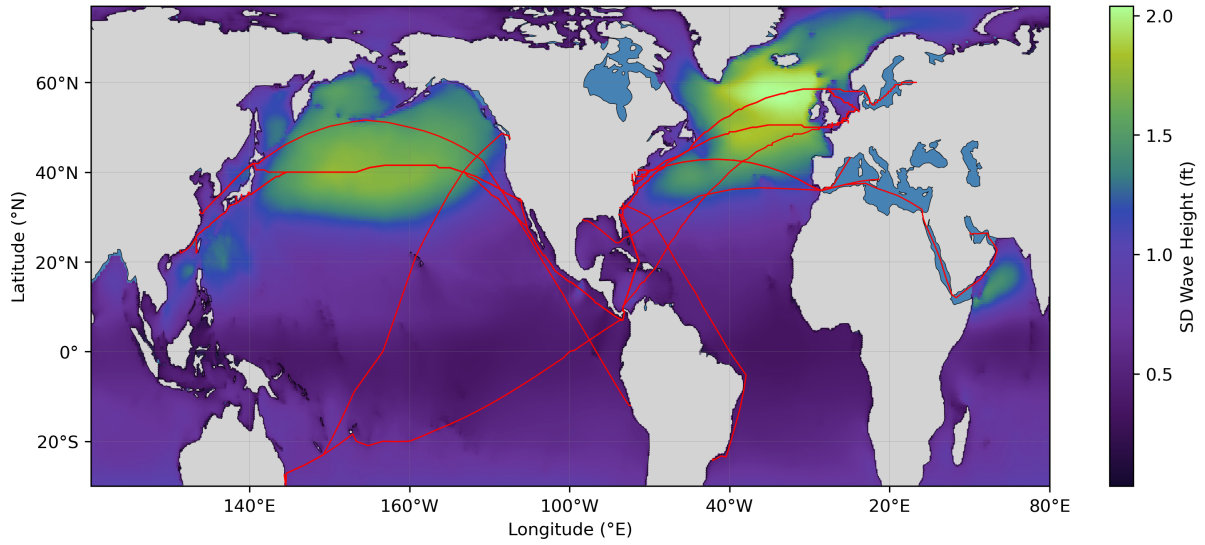
¹¹This evidence is consistent with [Ganapati et al. \(2024\)](#), who show that vessels typically follow the minimum-distance routes fairly closely. In addition, two-thirds of world trade in manufacturing travels on container ships, which typically follow fixed itineraries (i.e. the so-called “bus system”, see [Brancaccio et al. \(2020\)](#) and [Heiland et al. \(2022\)](#)), which are more likely to be picked up by the Eurostat’s SeaRoute program.

Figure 2: Weather Conditions and Routes

(a) Network of Shipping Routes



(b) Volatility of Wave Height



Notes: The upper panel shows the network of shipping routes constructed in our data. The bottom panel shows the standard deviation of wave height across all days from 2011-2016 and some selected shipping routes.

shipping route as a foreign country of origin and a U.S. entry point (e.g., airport or border crossing) pair.¹² Since we cannot compute the weather conditions for these transactions, we will assume that their shipping risk is zero for the empirical analysis we do in Section 3.¹³

Our measures of shipping times and routes are the building blocks of our empirical analysis. The key advantage of relying on the U.S. Census transaction-level data to obtain these measures is its comprehensive nature and extreme detail, which allows for a systematic analysis of the role of shipping risk in the full economy. However, there are 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. Hence we do not capture the risks associated with those parts of the trip. However, typically goods spend several weeks on the vessel to the U.S. and thus this part of the journey is likely a large fraction of the total travel time. Second, we do not know whether goods are reloaded from one vessel to another, i.e., “trans-shipped” (Ganapati et al., 2024).¹⁴ 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.

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 2016.¹⁵ 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).¹⁶ We aggregate the hourly information to the daily level and compute the daily average of

¹²In our sample, there are 11,500 distinct non-seaborne routes.

¹³We will analyze below to what extent importers use other modes of transportation, in particular shipments by airplane, to mitigate shipping risk.

¹⁴This 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. Note that any intermediate stops the vessel makes on its journey are captured in our customs data, as explained above, as long as some cargo that is bound for the U.S. is loaded.

¹⁵Significant 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>.

¹⁶Ocean 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.

significant wave height and direction for each geo-location in the oceans.¹⁷

We combine the weather information with the route segments for seaborne shipments constructed earlier. Since the effect of waves on vessel speed depends on the direction of travel, we compute for each route coordinate a relative wave direction. This relative direction is computed by taking the absolute difference between the direction of the waves and the estimated direction of vessels at that point. We estimate vessels' direction using the latitude and longitude of subsequent route coordinates. A greater relative direction means that the waves are less aligned with vessels' likely course.¹⁸ For each route segment, we compute the average weather (i.e., wave height and relative direction) for each day by averaging across all segment coordinates.

In the final step, we merge the route and weather information with the trade transaction data. For each day a vessel spends on a segment, we merge in the corresponding segment-level average weather. We then take an average across the day-level weather measures for each transaction. Our final dataset thus contains, for each transaction, an average wave height and an average relative wave direction along the entire route.¹⁹

To illustrate the source of our exogenous variation, the blue and green shading in the bottom panel of Figure 2 report the standard deviation of the average daily wave height for each grid point in our data. The red lines indicate some shipping routes used by U.S. importers. There are significant differences across locations. Routes across the Atlantic and Pacific have higher wave height volatility than routes along the coast of South America. Importantly, there is variation even across routes that are relatively close to each other. Vessels traversing the Northern Atlantic Ocean on their way to the East Coast face a significantly higher standard deviation of wave height than vessels that travel further South.²⁰

¹⁷We aggregate the data from the hourly to the daily level to reduce computational requirements. If information is unavailable at a geo-location, we take a simple average of the weather in the surrounding grid points. 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.

¹⁸For 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.

¹⁹Since we do not know a vessel's precise location on each day, we cannot use the actual weather in a vessel's vicinity and instead use the average segment-level weather. In robustness analysis below, we impute a vessel's location and use the weather only near the vessel's imputed location. The results are similar.

²⁰Appendix A.4 provides some summary statistics on wave height. We also show that there are significant differences in weather in the Northern Atlantic and Northern Pacific across seasons.

Table 3: Summary Statistics on Foreign Sourcing

	Mean	St. Dev.	P50	P95
Number of Routes	2.23	4.08	1	6.77
Number of Suppliers	1.90	3.84	1	5.39
HHI across Suppliers	0.88	0.21	1	1
HHI across Suppliers and Routes	0.81	0.27	1	1

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.

3 Empirical Analysis

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 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 Multi-Route and Multi-Supplier Sourcing

We show in Table 3 that firms rely on multiple routes to source the same HS-10 product within the same year. The average firm uses 2.2 routes per product and year, across all modes of transportation. The large standard deviation compared to the mean indicates that there is substantial heterogeneity across buyers. While the median importer uses only one route for a given product, firms at the 95th percentile use nearly 7. Table 3 shows that firms also rely on multiple foreign suppliers to source the same HS-10 products. Firms sourcing a given product from multiple suppliers account for almost 89% of imports. We argue below that one reason to use multiple routes and suppliers is to hedge against shipping delay risk (Section 3.4). In addition, although input expenditures are, on average, highly concentrated among routes and suppliers, the degree of this concentration varies substantially.

3.2 Measuring Unexpected Shipping Times

A central goal of our methodology is to measure shipping delays, that is, instances where goods arrive later than expected. However, we do not observe the shipping times expected by the importers, only the realized ones. To isolate the component of shipping times that is

unexpected, we propose a two-step approach. In the first, we regress the observed shipping times on a rich set of fixed effects and observables to capture determinants of the shipping times that are likely anticipated by importers, such as the route, season, supplier, or shipping charges. To remove any additional unobserved determinants of shipping times, in a second step, we regress these “residualized” shipping times on the weather conditions observed along the maritime route. We treat the predicted effects from this regression as the weather-induced unexpected shipping times. We include the residualization in the first step to remove predictable effects of weather on shipping times. For example, shippers could systematically use different vessels on routes on which weather conditions are more severe.²¹

Our measurement of unexpected shipping times relies on the assumption that, at the time of placing orders, importers do not fully anticipate the weather conditions along the entire maritime route beyond the usual seasonal patterns, which are picked up by the route-month fixed effects in the residualization. We believe that this is a reasonable assumption as import orders are placed typically many weeks before production finalizes and goods are shipped (see [Deloitte \(2024\)](#)). Moreover, most shipments to the U.S. require multi-week ocean crossings where weather conditions cannot be perfectly predicted, even by shipping companies relying on sophisticated weather forecasting technology.²² Importers may also be uncertain about the exact shipping date of their orders.

Step I: Residualization. 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 arm’s 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.²³ The time it takes for the shipment to arrive in the U.S., $T_{xhrtvfa}^s$, is a stochastic variable with the following law of motion:

²¹Another advantage of the residualization is that we can perform it on the full sample from 1992 while we have weather data only for 2011-2016. In a robustness exercise below, we use the risk measure obtained from the residualization step, and show that our results carry over to the longer sample starting in 1992.

²²Forecasts are reasonably accurate only until around 7 days into the future, and only general weather trends can be predicted beyond 2 weeks (see [Alley et al. \(2019\)](#) and [Ritchie \(2024\)](#)). [Zhang et al. \(2022\)](#) and [Mishra et al. \(2022\)](#) argue that despite advancements in machine learning and predictive modeling, accurately forecasting ocean wave height remains a difficult problem due to the chaotic and non-linear nature of ocean dynamics.

²³The rationale to include the shipment weight and freight rate is that they have a significant effect on shipping times, as we show in Table A.2 in Appendix A.2.

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

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, e.g. the average time it takes to unload a shipment and clear customs.²⁴ 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. 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).

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 even conditional on shipping charges. For example, shipping contracts could involve negotiations over dimensions beyond freight cost, such as promises of future business, which could be correlated with shipping times in a known way. We therefore focus on the variation in navigation times induced by weather along the shipping route.

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:

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

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, and Direction^s is the average relative wave direction in degrees, relative to the vessel's direction of travel.

Table 4 presents the results from the regression. The first column includes only wave

²⁴Similarly, $\bar{\xi}_v$ 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.

Table 4: Effect of Weather on Shipping Times

Dep. Var:	\tilde{t}^s	\tilde{t}^s	\tilde{t}^s
Wave Height ^s	-0.025*** (0.000)	-0.026*** (0.000)	-0.029*** (0.001)
Direction ^s		-0.010*** (0.000)	-0.020*** (0.000)
Wave Height ^s × Direction ^s			0.003*** (0.000)
R-Squared	0.013	0.013	0.013
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. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

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 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 (Appendix A.3). These data do not rely on an imputation of routes and report information on waves and vessel speed at exact vessel locations in the ocean, indicating that our results are not driven by our imputation methodology. The results are also in line with earlier findings that have shown a positive effect of wave height on speed (Filtz et al. (2015)), and could be consistent with vessels increasing cruising speed when passing through areas with bad weather. The predicted values from the regression in column (3), $\tilde{t}^{s,weather}$, constitute our measure of unexpected shipping times due to weather.

3.3 Shipping Delays and Importers Performance

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

$$FracDelayed_{ft}^{weather} = \frac{\sum_s \mathbb{D}_{ft}^{s,weather} \cdot \text{Imp Value}_{ft}^s}{\text{Total Input costs}_{ft}}, \quad (3)$$

Table 5: Effect of Extreme Delays on Firms' Outcomes

	(1)	(2)	(3)
Weather Shocks			
Dependent Variable (in logs):	Sales	Profits	Employees
Frac Delayed	-2.434*** (0.513)	-1.298** (0.502)	-0.390* (0.231)
Importer FE	Y	Y	Y
Year FE	Y	Y	Y
R-Squared	0.97	0.92	0.98
Observations	40,500	40,500	40,500
Residualized Only			
Dependent Variable (in logs):	Sales	Profits	Employees
Frac Delayed	-2.270*** (0.386)	-1.197*** (0.405)	-0.348** (0.152)
Importer FE	Y	Y	Y
Year FE	Y	Y	Y
R-Squared	0.94	0.92	0.98
Observations	40,500	40,500	40,500

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. Mean of $FracDelayed_{ft}^{weather}$ is 0.0038 and standard deviation is 0.0266. R^2 is the overall fit inclusive of the fixed effects. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

where $\mathbb{D}_{ft}^{s,weather}$ is an indicator that shipment s to importer f was delayed in year t , Imp Value_{ft}^s is the import value of such shipment, and $\text{Total Input costs}_{ft}$ are the importer's total production costs (materials, including domestically sourced inputs, plus labor) in the year.²⁵ Intuitively, this measures the share of an importer's input expenditures that are subject to extremely long shipping times in a given year. We then estimate the following panel regression for the years 2011-2016:

$$\ln(Y_{ft}^o) = \alpha + \beta_1 FracDelayed_{ft}^{weather} + \gamma_f + \delta_t + \epsilon_{ft}, \quad (4)$$

where Y_{ft}^o is either the sales, operating profits (sales minus materials and labor costs), or number of employees, and γ_f and δ_t are firm and year fixed effects, respectively. Our identifying assumption is that, conditional on firm fixed effects, the fraction of inputs that is subject to extreme shipping delays due to weather is orthogonal to any unobservable characteristics that may affect an importer's post-delay performance. Our construction of weather shocks in the previous section aims to satisfy this assumption. The top panel of Table 5 reports the results. Standard errors are clustered at the firm-level.

Shipping delays significantly disrupt production levels and profit margins. Increasing the

²⁵Data 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.

fraction of delayed shipments by one standard deviation (2.66 percentage points, which is almost a ten-fold increase from the average fraction delayed, 0.38%) is associated with a drop in sales by 6.5%, a fall in profits by 3.5%, and in employment by 1.0%. Therefore, extreme unexpected delays have a large and significant impact on U.S. importers. Our large effects of shipping delays are also consistent with evidence that the shipping risk is borne primarily by the buyer (see [Herghelegiu and Monastyrenko \(2020\)](#) and [Eurosender \(2023\)](#)) and that insurance for supply disruption events is limited and expensive ([Heckmann \(2016\)](#)).²⁶

In the bottom panel of Table 5 we re-run the regressions using the residualized shipping times from Step 1 instead of the weather-induced shipping times from Step 2 to construct delays. This specification picks up all types of delays, including those unrelated to weather such as port delays or strikes, but requires a stronger identification assumption. Specifically, we now have to assume that any delays we capture with the residuals \tilde{t}^s are unanticipated by the importers. While this assumption is not as clearly satisfied as for weather shocks, we find relatively similar coefficients. A one standard deviation increase in the fraction of delayed shipments is associated with a drop in sales by 6%, a fall in profits by 3.2%, and in the number of employees by 0.9%.

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 construct weather-induced shipping delays with this measure. We find similar results 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, we now investigate whether firms take actions to actively diversify this source of risk. To do so, we compute a measure of exposure to shipping time risk based on the volatility of the weather-induced shipping times experienced by importers. We then document how sourcing patterns are affected by exposure to such risk.

We start by computing the standard deviation of the weather-induced residualized shipping

²⁶Note that most importers are relatively small. Our estimates of the impacts of shipping delays on sales are in line with the effects of other supply chain disruptions found in a recent literature. [Carvalho et al. \(2021\)](#) estimate an elasticity of sales of -3.6% following a shock hitting a domestic supplier. [Barrot and Sauvagnat \(2016\)](#) find that when one of their suppliers is hit by a major natural disaster, firms experience an average drop 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.

times over three-year rolling windows, denoted by $\widehat{StdTime}_{xhrt-3,t-1}$, at the supplier-product-route-year (x, h, r, t) level.²⁷ While weather risk varies at the route-level, we compute our measure at the supplier-product-route level to allow for variation in risk across suppliers and products within the same route.²⁸ We aggregate this risk measure at the importer-product-year level by taking a weighted average over the importer’s suppliers and routes over the previous three years, i.e.,

$$\widehat{StdTime}_{fht-3,t-1} \equiv \sum_{x,r} \omega_{fxhr,t-3,t-1} \widehat{StdTime}_{xhrt-3,t-1}, \quad (5)$$

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

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

$$\ln(Y_{fht}) = \alpha + \beta_1 \ln(\widehat{StdTime}_{fht-3,t-1}) + \beta_2 X_{fht} + \gamma_f + \mu_h + \delta_t + \epsilon_{fht}, \quad (6)$$

where Y_{fht} is an import demand outcome of importer f in year t for product h . This regression analyzes whether risk faced by the importer in the previous three years ($t - 3$ to $t - 1$) affects the importer’s sourcing in the current year t . Our shift-share measure of risk helps alleviate concerns of reverse causality, that is, the endogeneity of the risk measure through the importers’ choice of routes and suppliers. Given the well-known stickiness in buyer-supplier relationships (e.g., [Martin et al. \(2023\)](#), [Heise \(2024\)](#)), our panel specification with firm fixed effects exploits variation over time in importers’ exposure to risk driven by within-route changes in the volatility of weather shocks.²⁹

We consider the following dimensions of sourcing as dependent variable: (i) the number of routes, (ii) the number of foreign suppliers, (iii) the concentration of imports across routes as measured by the Herfindhal–Hirschman index (HHI), (iv) the HHI of imports across suppliers, and (v) the total value imported. X_{fht} is a vector of controls, and γ_f , μ_h and δ_t are importer,

²⁷For non-vessel transactions, we set $\widehat{StdTime}_{xhrt-3,t-1} = 0$ for each (x, h, r, t) cell. Cells with fewer than 10 transactions are dropped.

²⁸For example, if suppliers use different shipping companies that vary in their ability to predict weather conditions, then the supplier’s risk would not be correctly picked up by a route’s average risk level. Similarly, a timely arrival could be more important for some products than others, leading firms to choose different shipping companies or to invest more in forecasting.

²⁹Note also that since the sum of the shares $\omega_{fxhr,t-3,t-1}$ for each firm-product is one, we do not need to control for the sum of shares interacted with period fixed effects, as discussed in [Borusyak et al. \(2024\)](#).

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	0.127*** (0.010)	0.080*** (0.009)	-0.075*** (0.003)	-0.053*** (0.003)	-0.084*** (0.012)
Importer FE	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y
R-squared	0.73	0.69	0.44	0.46	0.90
Observations	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. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

product, and year fixed effects, respectively.³⁰

Our controls X_{fht} include the importer’s unexpected shipping time $\tilde{t}^{s,weather}$, averaged over the previous three years. This variable accounts for the direct negative effect of shipping times on import demand, as documented in [Hummels and Schaur \(2013\)](#). The average shipping time also controls for the fact that suppliers located in countries further away may mechanically have more volatile shipping times purely because they have more scope for delays. We also include the average unit value paid by the importer for product h in year t . This variable controls for the fact that riskier suppliers may sell cheaper inputs, confounding the relationship between risk and import demand we aim to estimate. Finally, we control for importers’ size (proxied by the total imports of product h over the previous 3 years), and for suppliers’ size (proxied by the total exports of the product over the previous 3 years) since larger importers or exporters have more shipments, which could mechanically increase their risk.³¹

Table 6 presents the findings, using the volatility of weather-induced shipping times for the sample period 2011-2016. Standard errors are clustered at the firm level. Column 1 documents a positive and significant relationship between the number of routes used and shipping risk. An increase in risk from the 25th to the 75th percentile of the weather risk

³⁰In our baseline specification, we omit firm-product pairs with only one foreign supplier in year t since the purchase volume may be too small to make diversification viable or the product may be too specialized. We show below that our results are robust to including such firms.

³¹For example, exporters shipping greater volumes in a given period may need to send more shipments with more vessels, which may introduce a predictable correlation between risk and exporter size. This effect would not be picked up by the exporter fixed effects in the residualization step. The reasoning is similar for importers.

distribution (61 log points) increases the number of shipping routes used by 7.7%. Column 2 shows that there is also a positive relationship between the number of foreign suppliers and risk. An increase from the 25th to the 75th percentile of the risk distribution increases the number of suppliers used by 4.9%.³² The fact that the coefficient on risk is larger in column 1 than in column 2 is consistent with importers relying more on additional routes than on additional suppliers as a vehicle to diversify weather risk, as this type of risk is largely determined at the route level, and different suppliers may rely on the same routes.

We next look at the relationship between shipping time risk and the concentration of import value across an importer’s routes and suppliers. Column 3 shows a negative and significant relation between shipping risk and the HHI over routes, suggesting that importers with riskier routes feature a more diversified pattern of expenditure across their routes. Column 4 shows that this effect is similar when we look at the concentration across suppliers. Lastly, in column 5, we find a negative and statistically significant relationship between our risk measure and total imports in each year. Quantitatively, we find that going from the 25th to the 75th percentile of the risk distribution decreases the route HHI by 4.6%, the supplier HHI by 3.2%, and total imports by 5.1%.

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

3.4.1 Selection Bias, Robustness, and Diversification via Air Shipping

Our empirical results show that firms that are more exposed to shipping time risk feature a more diversified structure of import demand. The measure of risk exposure, however, takes the firm’s set of suppliers and routes as given, raising the concern of selection bias. We now discuss various forms in which this selection could affect our results. Consider first the case where importers differ in their risk aversion. To the extent that more risk averse

³²Note that we consider all the suppliers used in a year by each importer. Since the empirical analysis is 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.

Table 7: Shipping Time Risk and Import Demand, Robustness

	(1)	(2)	(3)	(4)	(5)
Dep. Var.:	Number of Routes	Number of Suppliers	HHI over Routes	HHI over Suppliers	Value Imported
Std Time	0.071*** (0.005)	0.040*** (0.004)	-0.042*** (0.001)	-0.028*** (0.001)	-0.091*** (0.006)
Importer FE	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y
R-squared	0.66	0.62	0.34	0.37	0.83
Observations	328,000	328,000	328,000	328,000	328,000

Notes: Table reports the coefficients from running specification (6) using the standard deviation of the residualized shipping times from Step I as our measure of risk, for the entire period 1992-2016. 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. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

importers feature safer suppliers/routes and also more suppliers/routes, this selection works against our empirical findings. That is, it produces a negative relationship between shipping time risk and the number of suppliers or routes. Consider next the role of firm size. In the presence of fixed costs to adding suppliers and routes, larger firms would feature more suppliers/routes. If in addition larger firms feature riskier suppliers and routes, this selection could produce relationships as the ones documented in the previous section. We address this issue by including firm fixed effects and controlling for past imports. Moreover, in our sample we find a negative and significant correlation (-0.12) between size of the firm (proxied with log sales) and our risk measure.

We next perform a number of robustness exercises with our weather risk measure. First, in Table 6, we report the results using the volatility of residualized shipping times obtained from Step I of our methodology. This measure has the advantage that we can use the whole sample period 1992-2016 but, as discussed earlier, it comes at the expense of stronger required exogeneity assumptions. Results are similar both qualitatively and quantitatively. An increase from the 25th to the 75th percentile of the risk distribution (92 log points) increases the number of routes used by 6.6% and the number of suppliers by 3.7%. Moreover, it decreases the route HHI, supplier HHI, and total imports by 3.9%, 2.5%, and 10.5%, respectively.

Next, in Appendix B.2, we document that using an alternative measure of weather shocks, which relies on predictions of where the vessel is on the route in each day and uses local weather conditions, produces similar results. Our results also hold when we include firm-time fixed effects, which control for firm-level shocks that may affect production choices, and also

Table 8: Shipping Time Risk and Import Demand with Air Shipments

Dep. Var.:	Air Shipments
Std Time	0.017*** (0.004)
Avg Time	-0.720*** (0.334)
Importer FE	Y
Product FE	Y
Year FE	Y
Controls	Y
R-Squared	0.60
Observations	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. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

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 rather than a weighted average across suppliers. 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.

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 estimate a variant of our main specification, equation (6),

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

where d_{fht} is a dummy that is equal to one if firm f uses air shipments for HS10 h in year t , and $\widehat{StdTime}_{fht-3,t-1}$ is the same weather-based risk measure as before. The controls X_{fht} are identical to the ones used before, except that the importer’s total log value of imports is now split up into imports by vessel and imports by airplane to account for the relative importance of both. Table 8 documents that higher shipping risk is associated with a higher likelihood of using air shipments. An increase in risk from the 25th to the 75th percentile of the weather risk distribution increases the likelihood of using air shipments by 1.0%. While the effect is small since firms likely use air shipments for many other reasons, for example to get seasonal goods quickly, our findings suggest that firms use air transportation to hedge

ocean shipping risk.³³

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

4 A Model of Input Sourcing with Shipping Risk

To rationalize the empirical evidence on shipping time risk and import demand and to quantify the aggregate implications of risk, 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. \(2018\)](#), and [Gopinath and Neiman \(2014\)](#). The key departure from this literature is that inputs' shipping times are a component of input quality, thus affecting production levels in the spirit of [Hummels and Schaur \(2013\)](#), and that such shipping times are stochastic. Imported inputs lower production costs due to production complementarities and differences in qualities and prices, but require payment of fixed costs. In addition to these standard forces, firms have incentives to increase the number of foreign inputs to mitigate the impact of shipping time risk on expected revenues.

Section 4.1 outlines the environment of the model, while Section 4.2 characterizes the firm's problem. Section 4.3 provides theoretical results that describe the impact of risk on import demand. Finally, Section 4.4 closes the model in equilibrium.

4.1 Environment

We consider a small open economy populated by a fixed mass of firms that produce differentiated manufacturing varieties which are sold locally. Firms buy inputs from foreign suppliers to whom they are exogenously matched. Each of these suppliers corresponds to a different route in the empirical analysis of Section 3. At the time of placing orders, firms are uncertain about the time it will take the inputs to arrive and, crucially, inputs arriving late are less productive. Firms can diversify the risk of late deliveries by having multiple foreign suppliers.³⁴ For tractability, we assume that the suppliers of any given firm are

³³This result is in line with the findings in [Hummels and Schaur \(2010\)](#), which show that firms use air shipping to smooth demand volatility on international markets.

³⁴Our model abstracts from firms' use of inventories, which have been shown to be an important diversification mechanism by [Alessandria and Ruhl \(2021\)](#) and [Carreras-Valle \(2021\)](#), among others. This is because we document in Section 3 that our main results on the response of import demand to shipping risk

ex-ante identical—but may differ *ex-post* in their realized shipping times—implying that the extensive margin of trade can be summarized by the number of suppliers. While suppliers are identical for a given firm, firms differ in the riskiness of their pool of foreign suppliers—a feature that will be instrumental in Section 5 below for the model to come to terms with the empirical patterns documented above.³⁵

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

$$y_f = \varphi_f l^{1-\gamma} \left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + \left(\sum_{i=1}^N \alpha_i x_i \right)^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\gamma \frac{\varepsilon}{\varepsilon-1}} \quad (7)$$

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. 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 . The N suppliers are taken from a pool of unlimited foreign suppliers to whom the firm is matched to. 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.³⁶

A central element of our theory is the assumption that shipping delays are detrimental to production. In particular, we assume that (i) longer shipping times reduce the qualities of inputs, similarly to [Hummels and Schaur \(2013\)](#), and (ii) shipping times are stochastic and unknown to firms at the time of placing orders. We parametrize the relationship between supplier quality and shipping time as follows:

$$\alpha_i = \begin{cases} \bar{\alpha}_i & \text{if } d_i \leq \mathbb{E}[d_i] \\ e^{-\tau \cdot d_i} & \text{if } d_i > \mathbb{E}[d_i], \end{cases} \quad (8)$$

where d_i are the number of days it takes to ship the input of supplier i to firm f , $\bar{\alpha}_i = e^{-\tau \cdot \mathbb{E}[d_i]}$, and $\mathbb{E}[d_i]$ is the expected shipping time. This formulation implies that if an input arrives

are large and significant even after controlling for firms' inventories.

³⁵Our theory takes the assignment between importers and their suppliers as exogenously given. In the quantitative exercise of Section 5 below, we make this assignment to replicate the negative correlation between firm size and supplier riskiness observed in the data and described above. A micro-foundation where firms search for suppliers, and finding safer suppliers is more costly, would deliver this pattern.

³⁶The production structure thus far is standard in the literature—it corresponds to the one in [Gopinath and Neiman \(2014\)](#) or [Blaum et al. \(2018\)](#) when foreign inputs are perfect substitutes. We abstract from love-of-variety effects to focus on the extensive margin of importing as a channel of risk diversification. For tractability, we assume that the fixed costs of adding suppliers are not supplier-specific. For a treatment of the case with supplier-specific fixed costs in a deterministic setting see [Antras et al. \(2017\)](#).

earlier than or just as expected, i.e. $d_i \leq \mathbb{E}[d_i]$, it has a constant level of quality $\bar{\alpha}_i$. 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 .³⁷

We assume that, for each importer, the shipping days d_i are i.i.d. and denote their CDF by $G_f(\cdot)$. This distribution is known to the firm at the time of placing input orders. We allow this distribution to be firm-specific as, in our quantitative exercise below, firms differ in the riskiness of their foreign suppliers.

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

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

$$C = \left(\int c_f^{\frac{\sigma-1}{\sigma}} df \right)^{\frac{\sigma}{\sigma-1}}, \quad (9)$$

where $\sigma > 1$ and c_f denotes final consumption of the good produced by firm f .³⁸ 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 (9).³⁹

4.2 Firm's Problem under Risk

We next describe the firm's problem of choosing the quantity of imports and the number of suppliers in the presence of risk. The total sales of firm f , which include demand from both

³⁷The specification in equation (8) imposes that inputs arriving earlier than expected do not increase production. This structure effectively limits the positive effect of very low shipping times on revenues, ruling out risk loving behavior. This is useful to come to terms with the empirical evidence of Section 3. Note also that the specification in equation (8) 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.

³⁸For simplicity, we abstract from exporting, importing of final goods, and consumption of non-tradable goods. Incorporating these elements, as in [Blaum \(2024\)](#), is feasible but would complicate the analysis without adding additional insights.

³⁹The 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. See also [Blaum et al. \(2018\)](#); [Gopinath and Neiman \(2014\)](#); [Adão et al. \(2020\)](#) for a similar assumption.

consumers and other firms, are given by

$$R_f = y_f^{\frac{\sigma-1}{\sigma}} P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma}, \quad (10)$$

where y_f is given by equation (7), P is the price index associated with (9) and S denotes total spending (including demand by firms and consumers). 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

$$\max_{x_D, x, N} \chi_f \mathbb{E} \left[\left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + \left(\sum_{i=1}^N \alpha_i(d_i) \right)^{\frac{\varepsilon-1}{\varepsilon}} x^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\psi} \right] - p_D x_D - N p_M x - w N F, \quad (11)$$

where χ_f and ψ depend on firm efficiency, general equilibrium variables and parameters, and $\alpha(d_i)$ is given by (8).⁴⁰ Note that the expectation operator is taken over the possible realizations of d_i and thus depends on the distribution of shipping times $G_f(\cdot)$.

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

4.3 The Workings of Risk

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

⁴⁰See Section C.1 of the Appendix for details.

After maximizing out labor and the foreign inputs, the firm's problem simplifies to:

$$\max_N \underbrace{\tilde{\chi}_f \left(\mathbb{E} [\bar{\alpha}^\psi] \right)^{\frac{1}{1-\psi}}}_{=\tilde{R}} - wNF, \quad (12)$$

where $\bar{\alpha} \equiv \frac{1}{N} \sum_{i=1}^N \alpha_i$ is the average supplier quality, $\psi \equiv \frac{\gamma(\sigma-1)}{1+\gamma(\sigma-1)} \in (0, 1)$, $\tilde{\chi}_f$ is a constant that depends on firm efficiency, general equilibrium objects and parameters, and \tilde{R} is expected revenues net of labor and foreign input variable costs. Expression (12) makes it clear that volatility in the average supplier quality lowers expected revenues as $\psi < 1$. Relying on a second order approximation of $\bar{\alpha}^\psi$ around $\mathbb{E}[\alpha]$, the firm problem can be written as:

$$\max_N \tilde{\chi} \left((\mathbb{E}[\alpha])^\psi - \psi \frac{(1-\psi)}{2} (\mathbb{E}[\alpha])^{\psi-2} \frac{1}{N} \mathbb{V}[\alpha] \right)^{\frac{1}{1-\psi}} - wNF, \quad (13)$$

where $\mathbb{E}[\alpha]$ and $\mathbb{V}[\alpha]$ are the mean and variance of the supplier-level quality.⁴¹ This expression highlights the role of the mean and the variance of supplier quality, as well as of the number of suppliers, in shaping expected revenues. In particular, dispersion in input qualities reduces expected revenues. By increasing the number of suppliers N , the firm lowers the variance of the *average* supplier quality $\mathbb{V}[\bar{\alpha}] = \mathbb{V}[\alpha]/N$, thus reducing the amount of risk faced and mitigating its effects on expected revenues. The following result formalizes the effect of increased supplier risk for the case where N is continuous.

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

$$\frac{1}{N^*} \frac{\mathbb{V}[\alpha]}{(\mathbb{E}[\alpha])^2} < \frac{2}{\psi}. \quad (14)$$

That is, under condition (14), a higher $\mathbb{V}[\alpha]$ leads to an increase in N^* . Furthermore, $\epsilon_{N^*, \mathbb{V}[\alpha]} < 1$ regardless of whether (14) holds. It follows that a higher $\mathbb{V}[\alpha]$ always reduces the import value.

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

The first part of Proposition 1 states that a mean preserving spread in the variance of

⁴¹Taking a second-order Taylor approximation of expected utility (in our case, the expectation of a concave function) has a long-standing tradition in the finance literature, since [Markowitz \(1952\)](#) and [Samuelson \(1970\)](#), and it has been recently used also in the literature on trade under uncertainty (see e.g. [De Sousa et al. \(2020\)](#) and [Esposito \(2022\)](#)).

supplier-level quality increases the number of suppliers, if condition (14) holds. There are two opposite forces at work. When the variance of quality is higher, holding constant the expected value, an increase in the number of suppliers leads to a larger reduction in the variance of average supplier quality, $\mathbb{V}[\alpha]/N$, which is what matters for expected profits (13). This force increases the returns to adding suppliers. At the same time, higher variance in supplier-level quality reduces expected revenue, leading to a lower return to adding suppliers. When supplier risk is a small part of expected revenue, as ensured by condition (14), the negative level effect on expected revenues is dominated by the stronger reduction in the variance of average quality. As a result, a mean preserving spread in the variance of supplier quality increases the number of suppliers.

In the second part of the proposition, we turn our attention to import values. To understand this result, note that import value $N^*x^*p_M$, which is proportional to expected revenue, is a decreasing function of the variance of average supplier quality $\mathbb{V}[\alpha]/N^*$ —see (13).⁴² An increase in the variance of supplier-level quality therefore lowers import value if N^* either decreases or if it increases less than proportionally with $\mathbb{V}[\alpha]$. The proposition establishes that $\epsilon_{N^*, \mathbb{V}[\alpha]} < 1$ and hence that import value necessarily falls with more volatility in supplier quality.

Finally, we investigate the heterogeneity across importers in their response to risk.

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

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

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

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

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

⁴²As shown in Appendix C.2, total import value is given by

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

Lastly, Proposition 2 also implies that larger firms feature a smaller decline in sales when faced with heightened risk.

In connecting these results to our findings of Section 2, it should be noted that in the empirical analysis we measure the volatility of suppliers' log shipping times, not of input qualities, which are unobservable. In our theory, a mean-preserving spread to the distribution of shipping days affects both the variance and the mean of input qualities. In particular, the expected input qualities can increase or decrease depending on parameters.⁴³ In turn, the effect of a given change in expected quality on the returns to adding suppliers also depends on parameters.⁴⁴ Ultimately, whether the theory can come to terms with our empirical findings on shipping time volatility and import demand documented above is a quantitative matter which we tackle in Section 5 below. There, we consider a calibrated version of the model with a domestic input, general equilibrium, and firm heterogeneity. Before turning to this analysis, we next show how the model is closed in equilibrium.

4.4 Equilibrium

Thus far we have studied the effects of increased shipping time risk for a firm when prices are kept fixed. In going forward, we allow risk to affect aggregate equilibrium variables. We abstract from aggregate risk by assuming that there is a continuum of firms of unit mass within each type f .⁴⁵ We consider an equilibrium where firms maximize profits, the consumer maximizes utility, and goods markets clear. An equilibrium is fully characterized by the aggregate domestic spending S and the price index P associated to consumer utility. Note that the price of the domestic input bundle is given by $p_D = P$ as the domestic input aggregator is identical to consumer preferences. We now describe how S and P are determined.

Consumer expenditure is given by:

$$PC = wL + \Pi, \tag{15}$$

⁴³Recall that input quality is a function of shipping days given by expression (8). When all the mass of the shipping days distribution lies in the exponential region of the quality function, a mean preserving spread to log shipping days leads to a mean preserving spread in log qualities. Consequently, both the variance and the mean of input quality increase. When some mass lies in the flat region of the quality function, expected quality may fall with a mean preserving spread to log days.

⁴⁴Intuitively, a higher expected input quality affects the returns to adding suppliers in two ways. There is a direct positive effect as suppliers are more productive. However, there is an additional effect as a higher expected quality reduces the effective importance of the variance in expected revenue, thus reducing the incentives to adding suppliers. For ease of exposition, we relegate a formal treatment of changes in expected quality to Section C.3 of the Appendix, which provides a result characterizing this comparative static. We will come back to this issue in the quantitative Section 5.

⁴⁵In the quantitative exercise below, we assume that a firm's type is determined by its efficiency and the riskiness of its suppliers.

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

$$S = PC + p_D \int x_{Df} df, \quad (16)$$

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

$$P = \left(\int p_f^{1-\sigma} df \right)^{\frac{1}{1-\sigma}}, \quad (17)$$

where p_f is the expected price set by firm type f . An equilibrium is attained whenever (15), (16), and (17) are satisfied and firms maximize profits.^{46,47}

5 Climate Change, Geopolitical, and Infrastructure Risk

The theory developed in the previous section establishes that, when faced with greater shipping time risk, firms may increase their number of suppliers and reduce their import values, patterns that are aligned with the empirical evidence of Section 3. We now discipline the model's parameters with key moments of the data and establish that the model can quantitatively come to terms with our empirical evidence (Section 5.1). We then employ the model to quantify the consequences of climate change, geopolitical risks in the Red Sea, and strains on port infrastructure (Section 5.2). These exercises highlight how our quantitative model can be used as a laboratory to assess the impact of any risk event affecting ocean shipping on U.S. imports and welfare.

5.1 Calibration

Our calibration strategy requires the model to replicate various moments related to suppliers' shipping time risk. To assess whether the model can be consistent with the empirical evidence

⁴⁶To compute an equilibrium, given a guess of (S, P) we solve the firm's profit maximization problem and obtain input choices and expected prices and profits. We then find the level of spending and price index (S', P') implied by the right hand sides of expressions (16) and (17), respectively. An equilibrium is attained whenever $S = S'$ and $P = P'$.

⁴⁷We do not impose labor market clearing and, as a result, trade balance may not be attained. In particular, the trade balance is given by $TB = -(L - L_d)$, where L_d is the total labor demand. In other words, the manufacturing sector can finance a trade deficit by being a net supplier of labor to the rest of the economy. We thus normalize the wage to 1.

of Section 3, we target our estimate of the effect of shipping time risk on the extensive margin of importing.⁴⁸ We also require that the model matches the sensitivity of sales to shipping times. Finally, given our focus on aggregate effects in the counterfactuals below, we target the joint distribution of firm size and supplier risk across importers. We next describe how we parameterize such a distribution.

Parametrization of Firm Heterogeneity and Shipping Days To generate cross-sectional variation in importer size and exposure to supplier risk, we allow for firm types to be heterogeneous in two dimensions: efficiency φ_f and the standard deviation of shipping days of their suppliers σ_{df} . We assume that firm efficiency φ_f is drawn from a log-normal distribution with standard deviation σ_φ (we normalize average efficiency). The distribution of suppliers' shipping days d_i is also assumed to be log normal with type-specific standard deviation σ_{df} , and the shipping days are i.i.d. across a firm's suppliers. For computational simplicity, we assume that there are two types of suppliers, low and high risk: $\sigma_{df} \in \{\sigma_{dL}, \sigma_{dH}\}$.⁴⁹ When $\sigma_{df} = \sigma_{dH}$, firm f is assigned to high risk suppliers. We allow the risk type and efficiency to be correlated by assuming that a firm gets the high risk type with probability p which satisfies:

$$p = \kappa_\sigma + \rho_{\varphi\sigma}\varphi, \quad (18)$$

where $\kappa_\sigma, \rho_{\varphi\sigma}$ are parameters that control the prevalence of the high risk type and its correlation with efficiency, respectively.⁵⁰

Parameters, Moments, and Identification We directly measure the standard deviation of log shipping days for each type, σ_{dL} and σ_{dH} , using the standard deviation of the weather-predicted residualized log shipping times over three-year rolling windows, that is, the measure of risk used in our empirical analysis of Section 3. We compute the average of this measure within the groups of firms above and below the median risk. To isolate the role of supplier risk, we set expected log shipping days μ_{df} for each type to match a common average shipping

⁴⁸In this way, the nature of the quantitative exercise is to assess whether the model is able to match our empirical estimates of Section 3. Note, however, that we only target the estimate of risk on the extensive margin of importing; we do not target the effect of risk on import values.

⁴⁹Having more than two risk types is feasible but would complicate the numerical approach we use to solve and calibrate the model. When solving the firm's problem, expected revenues are computed for each combination of input choices (the domestic input, the foreign input, and the number of suppliers) and risk type.

⁵⁰To ensure that p is bounded between 0 and 1 we impose:

$$p = \max \{ \min \{ \kappa_\sigma + \rho_{\varphi\sigma}\varphi, 1 \}, 0 \}.$$

time.

The dispersion in firm efficiency σ_φ , the risk type assignment parameters κ_σ and $\rho_{\varphi\sigma}$, the fixed cost of adding foreign suppliers F , the elasticity of input quality to shipping time τ , and the price of imported inputs p_M are chosen to match the following moments of the data: (i) the coefficient of variation of log sales, (ii) a share of 50% of high risk firms, (iii) the correlation between log sales and our risk measure, (iv) the elasticity of the extensive margin of importing with respect to shipping risk, (v) the elasticity of sales to shipping times, and (vi) the aggregate import share. Given our weather-based measure of risk, we rely on the number of routes as our measure of the extensive margin of importing in (iv).⁵¹ All moments are measured from the same sample of U.S. manufacturing importers used in the empirical analysis of Section 3.

While each moment is affected by all parameters in equilibrium, intuitively, σ_φ controls the dispersion in firm size, κ_σ controls the share of firms of high risk, and $\rho_{\varphi\sigma}$ regulates the correlation between firm size and risk. By affecting the cost of adding suppliers, F controls the elasticity of the number of suppliers with respect to risk. Intuitively, a higher fixed cost F makes it more costly to diversify risk by increasing the number of foreign suppliers to source from, lowering the elasticity of N with respect to risk (Figure 3, left panel). The parameter τ affects how qualities, and thus revenues, fall with longer shipping times (Figure 3, center panel). Importantly, this negative association between shipping times and sales predicted by the model is verified in the data. A panel regression of log sales on the average shipping time of the firm’s suppliers yields a negative and statistically significant coefficient.⁵² By affecting the relative price of imported inputs, p_M controls firms’ expenditure on foreign inputs and thus the aggregate import share (Figure 3, right panel).

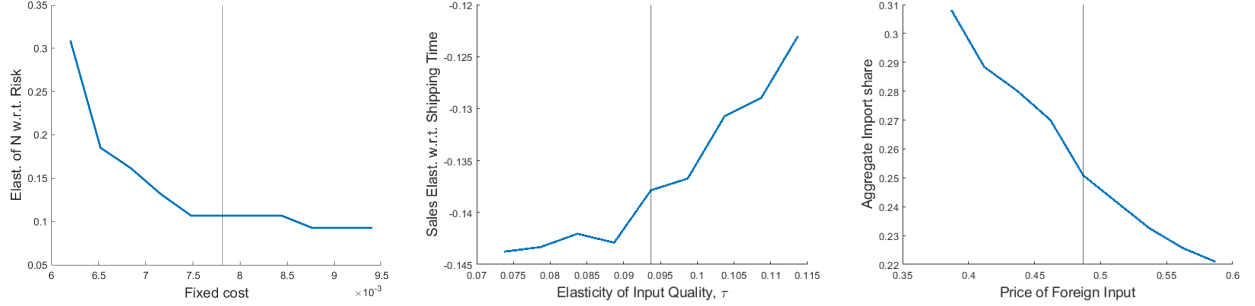
Lastly, we take three parameters from the literature. We set the output elasticity of materials to $\gamma = 0.6$ and the elasticity of substitution between domestic and foreign inputs to $\varepsilon = 2.38$ as in [Blaum et al. \(2018\)](#); we set the demand elasticity to $\sigma = 5$, in line with the average estimate for manufacturing differentiated goods in [Broda and Weinstein \(2006\)](#).

Calibration Results We report the calibrated parameter values in Table 9. The model is able to closely match the targeted moments. The average fixed costs paid are \$124,000

⁵¹Results 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, since its value is close to the elasticity of the number of routes with respect to risk.

⁵²The regression includes firm and year fixed effects and yields an estimated coefficient for average shipping times of -0.28 with a standard error of 0.14 clustered at the firm level. In the model, we run this regression for each of 1,000 simulated states of the world, where in each state we draw realizations of shipping days for each supplier. We then take an average across states of the estimated regression coefficient for shipping times.

Figure 3: Identification of Parameters



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

dollars, which are line with the literature (Fielier et al. (2018); Antras et al. (2017)).⁵³ An important feature of our calibration is that the model matches perfectly the elasticity of imports with respect to risk estimated in Section 3 (Table 6), despite not targeting this moment directly (see bottom of Table 9). This feature, together with the close match of the elasticity of the number of routes with respect to risk, implies that our quantitative model is able to come to terms to the key empirical findings of Section 3.

5.2 Counterfactual Analysis

Armed with the calibrated model, we assess the impacts of three prominent risk-related scenarios. We consider increases in risk associated with climate change, Red Sea attacks, and port congestion. We also study the consequences of a complete removal of shipping time risk.

Climate Change. We simulate an increase in weather volatility due to climate change over the next 50 years under the assumption that future weather conditions will continue to follow their historical trend. For this exercise, we use the matched dataset of shipping routes and weather conditions used in our empirical analysis, extended to 2023 to pick up a trend over a longer period. We compute the standard deviation of wave height across the days of each

⁵³We back out the average fixed costs from the ratio of average sales to average fixed costs and the assumption that average sales are the same as in the data.

Table 9: Calibrated Parameters and Targeted Moments

Parameter			Moment	Model	Data
Fixed Cost per Supplier	F	0.008	Elast. of N w.r.t. Risk	0.11	0.12
Elasticity of Input Quality	τ	0.094	Sales Elast. w.r.t. Ship. Time	-0.14	-0.28
Foreign Input Price	p^*	0.487	Aggregate Import Share	0.25	0.23
Std. Dev. Log Efficiency	σ_φ	0.02	Coef. of Variation Log Sales	0.35	0.24
Prevalence of High Risk Type	κ_σ	0.50	Share of High Risk Firms	0.56	0.50
Corr. High Risk and Efficiency	$\rho_{\varphi\sigma}$	-0.02	Corr. Log Sales and Risk	-0.19	-0.12
Std. Dev. of Shipping Times (High)	σ_{dH}	0.29	Avg. $\widehat{StdTime}$ above median	0.29	0.29
Std. Dev. of Shipping Times (Low)	σ_{dL}	0.09	Avg. $\widehat{StdTime}$ below median	0.09	0.09
Expected Log Shipping Days (High)	μ_{dH}	2.73	Average Shipping Days	16	16
Expected Log Shipping Days (Low)	μ_{dL}	2.77	Average Shipping Days	16	16
<i>Not calibrated:</i>			<i>Not targeted:</i>		
Elasticity Domestic-Foreign Inputs	ε	2.38	Elast. of Imports w.r.t. Risk	-0.08	-0.08
Demand Elasticity	σ	5			
Output Elast. w.r.t. Materials	γ	0.6			

Notes: The elasticities are coefficient estimates of cross-sectional regressions of log N , log sales, and log import value on log of the standard deviation of shipping days (controlling for firm efficiency), or log shipping times. The aggregate import share is the fraction of material expenditure accounted by foreign inputs. $\widehat{StdTime}$ is the measure of risk defined in Section 3.4. The moments in the data are measured in the 2011-2016 period. Sources: U.S. Census Bureau and authors' calculations.

year for each location on a shipping route, and then compute the annual growth rate of these standard deviations between 2011 and 2023 (see Figure 1 in the Introduction). The average annual growth rate across all locations is 0.34%. This number is in the ballpark of the annual growth rate estimated in [Young et al. \(2011\)](#) for 1985-2008, and is more broadly consistent with work suggesting an increasing likelihood of extreme wave heights (e.g., [Shi et al. \(2024\)](#)).⁵⁴ Compounding this growth over 50 years, we obtain a long-run growth rate in the standard deviation of wave height of 18.5%.

Table 10 (first column) summarizes the aggregate effects of climate change. If current climate trends persist and everything else is held fixed, the average number of suppliers will be 24% higher 50 years from now. This happens as firms diversify the increased risk of international shipping delays across multiple foreign suppliers (or routes, equivalently).

⁵⁴The oceanography literature typically focuses on the mean and the 99th percentile of wave height as a proxy for large events, rather than on predicting the standard deviation of wave height. [Young et al. \(2011\)](#) find a near neutral trend in wave height at the mean, an annual increase of 0.25% at the 90th percentile, and an increase of 0.50% per year at the 99th percentile of wave height. These patterns strongly suggest an increase in the standard deviation of wave height. Under the assumption that weather conditions are log normally distributed, these patterns suggest an increase in the standard deviation of wave height of 0.24% per year, reasonably close to our estimate.

Table 10: Counterfactuals

Variable	Climate Change	Red Sea Attacks	Port Congestion	Removing Risk
Average N	23.72%	2.79%	54.88%	-6.98%
Total Import Value	-2.89%	-42%	-8%	20.89%
Import Share	-2.40%	-33.95%	-7.37%	13.31%
Price Index	0.43%	-	1.46%	-0.13%
Total Spending	0.19%	-	0.8%	1.96%
Real Income	-0.39%	-	-1.33%	2.07%

Notes: The first column reports aggregate statistics after we shock the model with an increase in σ_{df} by 18.5%. In the second column we shock the model with an increase in μ_{df} by 27% and in σ_{df} by 14.85%. In the third column we shock the model with an increase σ_{df} by 14%. In the last column we set $\sigma_{df} = 0$.

The increase in N is larger for firms that had riskier suppliers to begin with, i.e. firms with $\sigma_{df} = \sigma_{dH}$, which on average raise N by 45%. Instead, most of the small, low-risk firms have no change in N because of the fixed costs of adding suppliers. The increase in the number of suppliers is also larger for more productive firms (the correlation between the growth in N and efficiency is 0.41 among high risk firms), consistent with our theoretical results in Proposition 2.

Climate change also lowers manufacturing imports by 2.9%, corresponding to a decline in U.S. imports of 46 billion dollars.⁵⁵ Consistent again with Proposition 2, the reduction in imports is smaller for larger firms (the correlation between import growth and efficiency is -0.14), as these firms are better diversified than small firms to begin with. The decline in imports is also stronger for high risk firms (-3.2%) than for low risk firms (-0.23%), as expected. As importers substitute riskier foreign inputs with domestic ones, the heightened risk implies higher production costs and prices (the price index increases by 0.43%) and lower output. Despite an increase in total domestic spending stemming from the production of the domestic input, U.S. real income is reduced by 0.4%.

Red Sea Attacks. In the second exercise, we evaluate the effects of the ongoing attacks by Houthi militias against commercial ships traveling along the Red Sea waterway that started in October 2023 (see [Rodriguez-Diaz et al. \(2024\)](#)). To incorporate the effects of the attacks on shipping times into our model, we first compare the time it takes to travel to the U.S. through the Suez Canal to the time it takes to travel via the Cape of Good Hope, avoiding the Suez Canal. Since the routes used in our empirical analysis are imputed, we cannot use

⁵⁵In 2016, the U.S. imported 2.2 trillion dollars of goods, of which \$590 billion were consumer goods (see <https://www.bea.gov/index.php/system/files/2017-12/trad1216.xls>, sheet 8). Classifying the remaining \$1.6 trillion as intermediate imports and using the decrease of 2.9% from our counterfactual, we obtain a change in intermediate imports of -\$46 billion.

them to study the difference in shipping times between two routes that start and end at the same port. Instead, we use information from Searates, a shipping time calculator, to compute the implied change in shipping time for each of the top-10 routes that head to the U.S. through the Suez Canal.⁵⁶ Shipping times increase by 27% on average across routes, with the largest change for the routes starting from the Arabian Peninsula (a growth rate of around 60%) and the smallest one for routes from South-East Asia to the East Coast (a growth of around 6%). Since Searates reports only average shipping times, we obtain the change in the standard deviation by relying on how the variability of shipping times changes with the average length of the trip in our Census data. We find that the correlation between the (log) mean and the (log) standard deviation of shipping times is 0.61, which reflects the fact that delay shocks accumulate with longer distance. Using this estimate, a change in the average shipping time μ_{df} of 27% should be associated to a change in the standard deviation σ_{df} of 16.5%.

We feed both the change in the mean and in the standard deviation of shipping times into our model. We find substantial impacts of the Red Sea attacks on total imports and the import share, which contract by 42% and 34%, respectively (Table 10, column 2). We find a moderate effect on the average number of suppliers, which grows by 2.79%. In assessing these results, it should be noted that in this exercise we are increasing the average shipping time, on top of increasing its variance. For import values, the increase in the mean and in the variance of shipping times both tend to reduce import value (Proposition 1)—but they have opposite effects on the number of suppliers, with longer shipping times reducing N . As in the previous counterfactual, the import reduction is larger for riskier firms (-51.47%), but since the shock also increases the average shipping time, the drop in imports is larger for more efficient firms. Lastly, since the share of goods that are imported into the U.S. and travel through the Suez Canal is only 2% of the total imports in our sample period, we assume for this exercise that the aggregate variables determined in general equilibrium are unaffected.⁵⁷

⁵⁶Searates can be accessed at <https://www.searates.com/distance-time>. We find the top-10 maritime routes by U.S. import value in 2011-2016. These are Jamnagar, India - New York; Kaohsiung, Taiwan - Chicago; Yantian, China - Chicago; Singapore - New York; Singapore - New Orleans; Yantian - Detroit; Hong Kong - Chicago; and three routes starting in the Middle East (the Census Bureau does not allow us to reveal the identity of these three routes due to disclosure concerns).

⁵⁷We note, however, the share of U.S. imports affected by a Suez Canal closure is likely higher: 3 of the top 10 routes by import value going through the Suez Canal are from ports in the Middle East, which ship mostly mineral fuels (HS27). It is likely that many of these fuel shipments originated in the Gulf are transshipped through hubs such as Antwerp or Rotterdam. Other products significantly affected by a Suez Canal closure are palm oil and coconut oil (64% of U.S. imports in 2011-2016 are via the Suez Canal) and TV reception apparatuses (21%).

Port Congestion. In the third exercise, we evaluate the economic effects of the greater variability of waiting times at ports due to the rise of port congestion that occurred globally in the aftermath of the Covid pandemic. We capture congestion using the Average Congestion Rate (ACR) measure developed by [Bai et al. \(2024\)](#), which reflects the average number of hours a container ship waits at port before docking at the berth for the top-50 container ports worldwide. Both the level and the volatility of wait times increased in the post-Covid period, which raises the volatility of shipping times. In particular, we find that the standard deviation of ACR went up by 14% in the post-Covid period of 2021-2022 relative to the 2017-2020 period. We feed this shock into our model as a change in the standard deviation σ_{df} of 14% for all firms, holding the mean fixed to isolate the effect of volatility.

The results in column 3 of Table 10 show a large effect of permanently higher port congestion, with imports and the import share contracting by 8% and 7.4%, respectively, and the number of suppliers (routes) increasing by 55%. The magnitude of these effects follows from the large shock to risk that is fed into the model, and is consistent with firms being able to use new routes by substituting towards additional ports, as in [Brancaccio et al. \(2024\)](#). The results are driven mostly by the response of larger firms: firms above the median efficiency more than double the number of foreign suppliers (routes), as they exploit their scale to diversify risk along the extensive margin.

Removing Shipping Time Risk. We conclude this section by quantifying the effects of a complete removal of shipping time risk (Table 10, column 4). When risk is removed, there is a 7% reduction in the average number of foreign suppliers used by importers. This happens because, without risk, foreign suppliers are as safe as the domestic ones, making it no longer necessary to diversify the risk of international shipping delays across multiple foreign producers. The reduction in N is larger for firms that had riskier suppliers to begin with, i.e. firms with $\sigma_{df} = \sigma_{dH}$, which on average reduce N by 18%. The decline in the number of suppliers when risk is removed is stronger for larger firms (the correlation between the growth in N and efficiency is -0.87 among high risk firms), consistent with Proposition 2.

The removal of shipping uncertainty also implies a substantial increase of 21% in aggregate imports and of 13% in the aggregate import share. If shipping risk was removed in 2016, U.S. manufacturing imports would go up by 334 billion dollars. Consistent again with Proposition 2, the increase in imports is larger for smaller firms (the correlation between import growth and efficiency is -0.78), and it is stronger for high risk compared to low risk firms (29.32% vs 16.17%). Overall, removing shipping risk lowers production costs and prices, resulting in a 2% increase in U.S. real income.

6 Conclusions

In this paper, we use U.S. Census shipment-level data combined with information on wave height and direction from NOAA to construct a novel measure of supply chain risk based on weather shocks. We document substantial negative 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 routes and suppliers, and a lower concentration of expenditure across routes and suppliers, which indicates that firms actively diversify this source of risk. To rationalize this evidence, we introduce risky delivery times into a standard quantitative model of firm-level importing. We show that an increase in shipping risk reduces the import share and can substantially lower aggregate income.

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

References

- Adão, R., Arkolakis, C., Esposito, F., 2020. General equilibrium indirect effects in space: Theory and measurement. Technical Report.
- 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* 146, 103788.
- Alessandria, G., Ruhl, K.J., 2021. Mitigating international supply-chain risk with inventories and fast transport Mimeo.
- Alley, R.B., Emanuel, K.A., Zhang, F., 2019. Advances in weather prediction. *Science* 363, 342--344.
- Antràs, P., Chor, D., 2022. Global value chains. *Handbook of International Economics* 5, 297--376.
- Antras, P., Fort, T.C., Tintelnot, F., 2017. The Margins of Global Sourcing: Theory and Evidence from US Firms. *American Economic Review* 107.
- Bai, X., Fernández-Villaverde, J., Li, Y., Zanetti, F., 2024. The causal effects of global supply chain disruptions on macroeconomic outcomes: evidence and theory. CESifo Working Paper No. 10930.
- Balboni, C., Boehm, J., Waseem, M., 2023. Firm adaptation in production networks: evidence from extreme weather events in pakistan Mimeo.
- Baley, I., Veldkamp, L., Waugh, M., 2020. Can global uncertainty promote international trade? *Journal of International Economics* 126, 103347.
- Barrot, J.N., Sauvagnat, J., 2016. Input specificity and the propagation of idiosyncratic shocks in production networks. *The Quarterly Journal of Economics* 131, 1543--1592.
- Bartik, T.J., 1991. Who benefits from state and local economic development policies? .
- Bernard, A.B., Moxnes, A., Saito, Y.U., 2019. Production networks, geography, and firm performance. *Journal of Political Economy* 127, 639--688.
- Blaum, J., 2024. Global Firms in Large Devaluations. *Quarterly Journal of Economics* Forthcoming.
- 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.
- Borusyak, K., Hull, P., Jaravel, X., 2024. A practical guide to shift-share instruments. Technical Report. National Bureau of Economic Research.
- Brancaccio, G., Kalouptsi, M., Papageorgiou, T., 2020. Geography, transportation, and endogenous trade costs. *Econometrica* 88, 657--691.
- Brancaccio, G., Kalouptsi, M., Papageorgiou, T., 2024. Investment in Infrastructure and Trade: The Case of Ports. Technical Report. National Bureau of Economic Research.
- 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. Mimeo.
- Carvalho, V.M., Nirei, M., Saito, Y.U., Tahbaz-Salehi, A., 2021. Supply Chain Disruptions: Evidence from the Great East Japan Earthquake. *The Quarterly Journal of Economics* 136, 1255--1321.
- Castro-Vincenzi, J., Khanna, G., Morales, N., Pandalai-Nayar, N., 2024. Weathering the storm: Supply

- chains and climate risk. Technical Report. National Bureau of Economic Research.
- De Sousa, J., Disdier, A.C., Gaigné, C., 2020. Export decision under risk. *European Economic Review* 121, 103342.
- Deloitte, 2024. Restructuring the supply base: Prioritizing a resilient, yet efficient supply chain. URL: <https://www2.deloitte.com/us/en/insights/industry/manufacturing/global-supply-chain-resilience-amid-disruptions.html>.
- Dhyne, E., Kikkawa, A.K., Mogstad, M., Tintelnot, F., 2021. Trade and domestic production networks. *The Review of Economic Studies* 88, 643–668.
- Dunbar, G.R., Steingress, W., Tomlin, B., 2023. Climate Variability and International Trade. Technical Report. Bank of Canada.
- 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* Mimeo.
- Eurosender, 2023. Incoterms for sea freight and ocean transport. URL: <https://www.eurosender.com/en/incoterms/sea-freight>.
- Evans, C.L., Harrigan, J., 2005. Distance, time, and specialization: Lean retailing in general equilibrium. *American Economic Review* 95, 292–313.
- Fieler, A.C., Eslava, M., Xu, D.Y., 2018. Trade, Quality Upgrading, and Input Linkages: Theory and Evidence from Colombia. *American Economic Review* 108, 109–146.
- Fillat, J., Garetto, S., 2015. Risk, return and multinational production. *The Quarterly Journal Economics* 130, 2027–2073.
- Filtz, E., De La Cerda, E.S., Weber, M., Zirkovits, D., 2015. Factors affecting ocean-going cargo ship speed and arrival time, in: *International Conference on Advanced Information Systems Engineering*, Springer. pp. 305–316.
- Fort, T.C., Klimek, S.D., 2018. The Effects of Industry Classification Changes on US Employment Composition. U.S. Census Bureau Center for Economic Studies Working Paper 18-28.
- Ganapati, S., Wong, W.F., Ziv, O., 2024. Entrepot: Hubs, scale, and trade costs. *American Economic Journal: Macroeconomics* 16, 239–278.
- Gervais, A., 2018. Uncertainty, risk aversion and international trade. *Journal of International Economics* 115, 145–158.
- Gopinath, G., Neiman, B., 2014. Trade Adjustment and Productivity in Large Crises. *The American Economic Review* 104, 793–831.
- Grossman, G.M., Helpman, E., Lhuillier, H., 2023. Supply chain resilience: Should policy promote international diversification or reshoring? *Journal of Political Economy* 131, 3462–3496.
- Halpern, L., Koren, M., Szeidl, A., 2015. Imported Inputs and Productivity. *American Economic Review* 105, 3660–3703.
- Handley, K., Limão, N., Ludema, R.D., Yu, Z., 2024. Firm input choice under trade policy uncertainty. *Journal of International Economics* 150, 103909.
- Heckmann, I., 2016. Supply chain risk analytics. *Towards Supply Chain Risk Analytics: Fundamentals, Simulation, Optimization*, 131–162.
- Heij, C., Knapp, S., 2015. Effects of wind strength and wave height on ship incident risk: regional trends and seasonality. *Transportation Research Part D: Transport and Environment* 37, 29–39.

- Heiland, I., Moxnes, A., Ulltveit-Moe, K.H., Zi, Y., 2022. Trade from space: Shipping networks and the global implications of local shocks. Mimeo.
- Heise, S., 2024. Firm-to-firm relationships and the pass-through of shocks: Theory and evidence. *Review of Economics and Statistics* , forthcoming.
- Herghelegiu, C., Monastyrenko, E., 2020. Risk- and cost-sharing in firm-to-firm trade. DEM Discussion Paper Series 20-24, Department of Economics at the University of Luxembourg.
- Hummels, D.L., Schaur, G., 2010. Hedging price volatility using fast transport. *Journal of International Economics* 82, 15--25.
- Hummels, D.L., Schaur, G., 2013. Time as a Trade Barrier. *American Economic Review* 103, 2935--2959.
- Johnson, R.C., Noguera, G., 2012. Accounting for intermediates: Production sharing and trade in value added. *Journal of International Economics* 86, 224--236.
- Kamal, F., Krizan, C.J., Monarch, R., 2015. Identifying Foreign Suppliers in US Merchandise Import Transactions. Mimeo.
- Kamal, F., Monarch, R., 2018. Identifying foreign suppliers in u.s. import data. *Review of International Economics* 26, 117--139.
- Khanna, G., Morales, N., Pandalai-Nayar, N., 2022. Supply chain resilience: Evidence from indian firms. NBER Working Paper No. 30689.
- Lafragne-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.
- Markowitz, H., 1952. Portfolio selection. *The journal of finance* 7, 77--91.
- Martin, J., Mejean, I., Parenti, M., 2023. Relationship stickiness, international trade, and economic uncertainty. *Review of Economics and Statistics* , forthcoming.
- Mishra, J.P., Singh, K., Chaudhary, H., 2022. Intelligent ocean wave height prediction system using light gbm model. *International Journal of Systematic Innovation* 7, 64--77.
- Pankratz, N.M., Schiller, C.M., 2024. Climate change and adaptation in global supply-chain networks. *The Review of Financial Studies* 37, 1729--1777.
- Pierce, J.R., Schott, P.K., 2012. Concording u.s. harmonized system categories over time. *Journal of Official Statistics* 28, 53--68.
- Ramondo, N., Rappoport, V., Ruhl, K.J., 2013. The proximity-concentration tradeoff under uncertainty. *Review of Economic Studies* 80, 1582--1621.
- Ritchie, H., 2024. Weather forecasts have become much more accurate; we now need to make them available to everyone. *Our World in Data* <https://ourworldindata.org/weather-forecasts>.
- Rodriguez-Diaz, E., Alcaide, J.I., Garcia-Llave, R., 2024. Challenges and security risks in the red sea: Impact of houthi attacks on maritime traffic. *Journal of Marine Science and Engineering* 12, 1900.
- Samuelson, P.A., 1970. The fundamental approximation theorem of portfolio analysis in terms of means, variances and higher moments. *The Review of Economic Studies* 37, 537--542.
- Shi, J., Feng, X., Toumi, R., Zhang, C., Hodges, K.I., Tao, A., Zhang, W., Zheng, J., 2024. Global increase in tropical cyclone ocean surface waves. *nature communications* 15, 174.
- Viellechner, A., Spinler, S., 2020. Novel data analytics meets conventional container shipping: predicting delays by comparing various machine learning algorithms. *Proceedings of the 53rd Hawaii International Conference on System Sciences* .
- Young, I., Zieger, S., Babanin, A.V., 2011. Global trends in wind speed and wave height. *Science* 332, 451--455.

Zhang, S., Jiang, H., Yin, Y., 2022. Characteristics of ocean waves in the southern ocean and seasonal prediction of difficulty for the vessels crossing the westerlies. *Atmosphere* 13, 916.

Online Appendix (Not for Publication)

A Data Construction and Summary Statistics

A.1 Data Construction

In this section, we describe the steps taken to clean the 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 in 2009 prices using the U.S. GDP deflator.

Since the manufacturer ID (MID) differs across establishments of the same firm in different locations and since logistics are likely arranged at the firm-level, we consider MIDs with the same name and country component but with a different street address or city component to belong to the same exporter. Specifically, we replace the MID with a shortened identifier that contains only the country ISO code and the name portion of the ID.⁵⁸ This approach follows earlier work by [Kamal et al. \(2015\)](#) and [Kamal and Monarch \(2018\)](#). [Kamal et al. \(2015\)](#) compare the number of MIDs in the Census data to the number of foreign exporters for 43 countries from the World Bank’s Exporter Dynamics Database (EDD), which is based on foreign national government statistics and private company data. They show that the number of MIDs in the Census data matches well with the number of sellers in the EDD when the street address or the city component are omitted. [Kamal and Monarch \(2018\)](#) provide further support that the MID is a good identifier of foreign exporters.

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

⁵⁸While different establishments may have different efficiencies or distances to the port, we only observe port-to-port shipping times.

in the given year. Our final dataset includes all related-party trade, and we include the type of relationship (arms’ length or related party) as one of the dimensions in the residualization of shipping times.

Trips construction We provide some further details on how we construct vessels’ trips. As described in the main text, we sort all transactions involving a given vessel by their foreign departure date. We then take all the vessel’s transactions and split them into trips using the arrival date in the U.S. and the export departure date abroad. Specifically, for each trip we find the earliest arrival date of the vessel in the U.S., and assign transactions with a later export departure date abroad to a new trip. In some cases, however, the U.S. arrival date is possibly misreported. For example, if transactions 1-5 depart abroad on June 22 and arrive in the U.S. on July 5, transaction 6 departs on June 23 and purportedly arrives on June 24, and transactions 7-10 depart on June 25 and arrive on July 5, then the procedure described above would assign transactions 1-6 to one trip and transactions 7-10 to another, even though almost all shipments arrive on the same day in the U.S.. It seems likely that the arrival date for transaction 6 is misreported. We therefore re-combine some of the previously separated trips. For each of the trips assigned in the first step, we compare the latest importation date in the U.S. to the earliest departure date abroad of the *next* trip. If the earliest departure date abroad of the next trip is before the latest importation date of the earlier trip, then the two trips must have been part of the same journey and we recombine these trips into one. We again iterate through this procedure until no more trips can be combined. Our resulting final dataset contains trips with completely non-overlapping foreign departure and importation dates for each vessel.

A.2 The Determinants of Shipping Times

In this section we examine the factors affecting vessel-borne shipping times to motivate our residualization procedure to construct shipping risk for ocean shipments.

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

Table A.1: 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 next investigate the role of the shipping route more formally by regressing 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. 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.

We finally analyze the role played by the season of the year, related party status, shipping weight, and charges, conditional on the shipping route, by running regressions of log shipping time on these characteristics. We present the results in Table A.2. We include fixed effects for the route in all regressions. 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 nearly 3% shorter in the summer quarters of the northern hemisphere, highlighting the potential role of weather in affecting shipping routes. In column 2, we find that related party transactions have slightly longer shipping times relative to arms-length transactions. The next columns find a positive relationship between shipping time and shipment weight, and a negative one between shipping charges and delivery times, conditional on weight.

A.3 Analyzing Vessel Movements with AIS Data

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

Table A.2: 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*** (0.000)		0.010*** (0.000)
Log Shipping Charges				0.004*** (0.000)	−0.003*** (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.

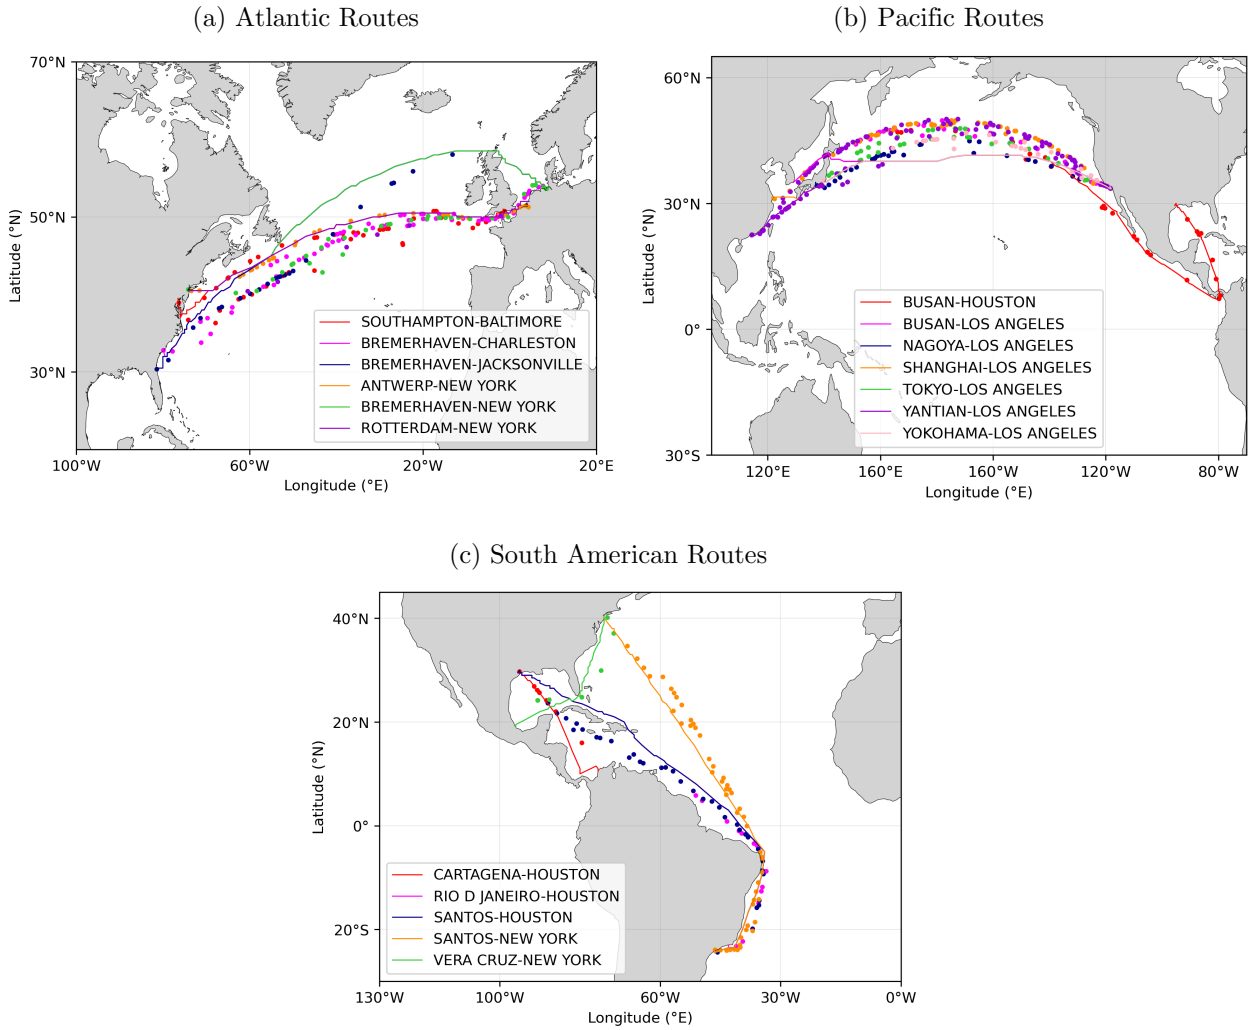
Vessel movements vs. routes As discussed in the main text, we construct the route taken by vessels between ports across the ocean using Eurostat’s SeaRoute program. For a subset of them, we compare these constructed routes to the actual vessel movements using AIS data from MarineTraffic, a provider of ship tracking and maritime analytics services.⁵⁹ These data report the precise location of vessels on the oceans based on transceiver signals of ships. We downloaded detailed geolocations with time stamps for 74 vessels traveling on 19 routes between July 10 and July 16, 2023 and between August 21 and August 28, 2023. For each vessel, we obtained origin and destination port as well as speed, direction, and weather at different locations (latitudes and longitudes) with detailed time stamps along the route. We obtained on average 121 different observations for each vessel along its route, with a range of between 13 to 373 data points. While we have historical data for each vessel, allowing us to observe each vessel from its departure port, the limited time of our data access did not permit us to observe each vessel until its arrival at the destination. On average, we observe 83% of a vessel’s full voyage from origin to destination (median: 90%).

Figure A.1a plots some routes across the Atlantic from the SeaRoute program against the observed locations of vessels traveling on these routes. The vessel locations are reasonably close to the routes, though not perfect. In particular, the SeaRoute program suggests that vessels traveling between Bremerhaven and New York mostly follow a route to the North of

⁵⁹See <https://www.marinetraffic.com>.

the United Kingdom, while the vessels we observe making the journey between these ports followed a route to the South of Britain. In contrast, some vessels traveling to Jacksonville followed the Northern route. Figure A.1b plots the routes across the Pacific. Here vessels are closer to the routes near the end points, but follow a more Northern trajectory in the middle. Finally, Figure A.1c plots vessel movements on South American routes. Overall, the analysis shows that vessels tend to broadly follow the routes from the SeaRoute program, but that there is substantial variation. This variation will introduce measurement error into our weather variable that will bias our results towards zero.

Figure A.1: Vessel Movements vs Routes



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

Vessel Speed and Weather Conditions We examine the effect of weather conditions on vessel speed in the AIS data. MarineTraffic provides for each vessel at each recorded location

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

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

where i indexes the vessel, j the location, and t is the time stamp. Here, Speed_{ijt} is the speed of the vessel at location j and time t , Height_{ijt} is the height of the waves, Direction_{ijt} is the wave direction relative to the direction of travel (where zero indicates that the waves are in the direction of travel), and γ_i are vessel fixed effects.⁶⁰ The first column of Table A.3 presents the results. As in the Census data, higher waves increase the vessel speed: a one standard deviation increase in significant wave height from the mean increases vessel speed (hence reduces shipping time) by about 4 log points. Also consistent with the Census data, a greater wave angle relative to the direction of travel has a positive effect on speed. When the waves are against the direction of travel (180 degrees), vessel speed is about 13 log points higher, reducing the shipping time.

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

⁶⁰We observe each vessel only on one route and so these are effectively vessel-route fixed effects.

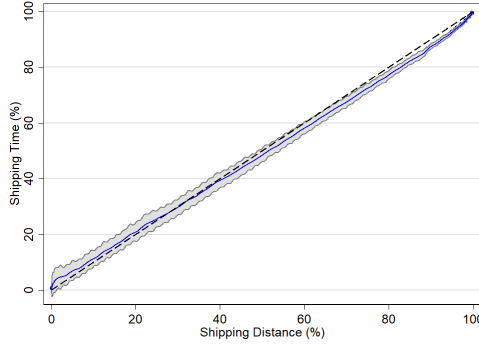
Table A.3: Effect of Weather on Shipping Times

Dep. Var:	Vessel speed	Vessel speed
Wave Height	0.0584*** (0.0122)	
Direction	0.0007*** (0.0002)	
Wave Height \times Direction ^s	-0.0002** (0.0001)	
Wind Speed		0.0024** (0.0011)
Wind Direction		0.0003** (0.0001)
Wind Speed \times 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 In the last step, we verify our assumption that vessels travel at approximately constant speed across the ocean using the AIS data, which we use to infer vessels' approximate location for our alternative weather shock measure in Appendix B.1. We use the 18 vessels for which we observe the entire journey from origin to destination port. For these vessels, we compute at each location the share of the journey completed, in terms of distance, as well as the share of the journey passed in terms of total voyage time. We then plot in Figure A.2 a bin scatter of the distance share against the share of voyage time. Overall, we find that the fit line is approximately on the 45 degree line throughout the journey, indicating that our assumption of constant speed is reasonable. Vessels are slightly slower at departure, and then make up for this delay along the journey before slowing down again near the arrival port.

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



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

A.4 Weather Summary Statistics

In this section we provide some further details on the wave height variable. Table A.4 provides some summary statistics on the mean and standard deviation of the significant wave height and its (absolute) direction across all days 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.

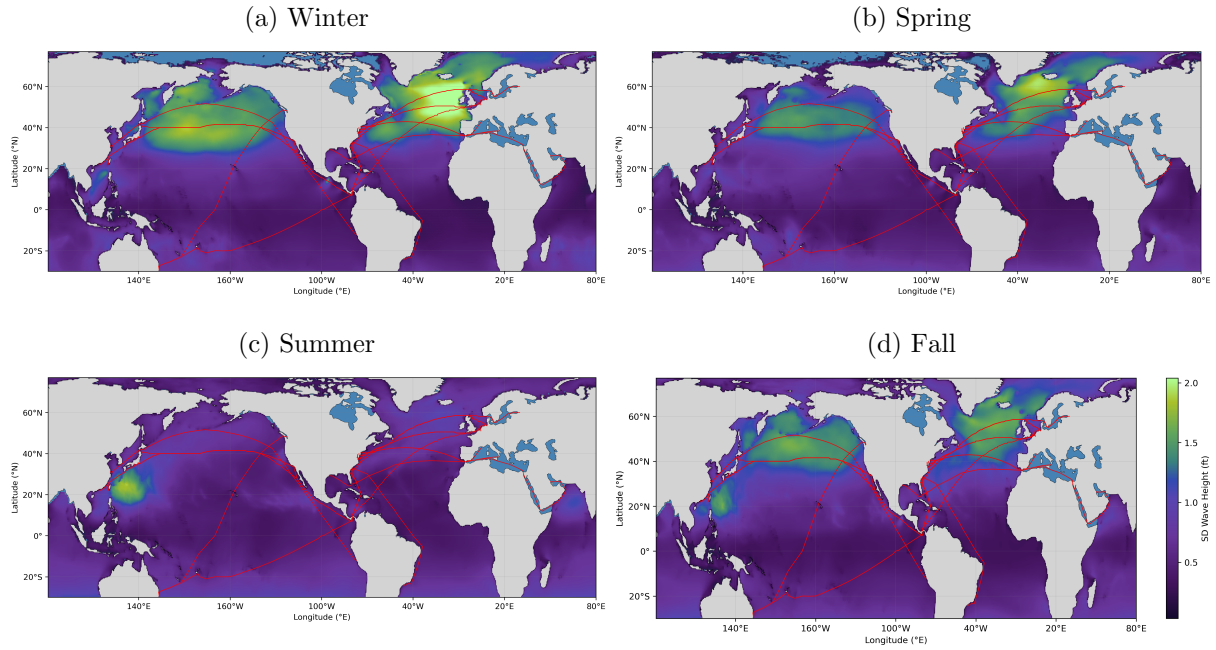
Figure A.3 further shows that there is also significant variation in the standard deviation of wave height across seasons. For example, both the northern Atlantic and the northern Pacific experience significant volatility in wave height in the fall, but very little in the summer.

Table A.4: Weather conditions: summary statistics

	Mean	Sd	p1	p5	p50	p95	p99
Significant wave height (m)							
All	2.6	1.5	0.1	0.6	2.3	5.5	7.3
North Atlantic	2.2	1.3	0.2	0.7	1.9	4.8	7.0
South Atlantic	2.7	1.3	0.3	1.1	2.4	5.2	6.8
North Pacific	2.3	1.2	0.2	0.7	2.1	4.6	6.5
South Pacific	2.9	1.4	0.3	1.1	2.6	5.7	7.4
Indian Ocean	2.9	1.6	0.1	0.6	2.7	5.9	7.7
Significant wave direction (degrees)							
All	203	80	20	52	216	318	343
North Atlantic	182	102	11	31	191	333	350
South Atlantic	216	65	33	92	220	314	342
North Pacific	186	98	17	40	195	325	344
South Atlantic	215	67	34	83	223	314	339
Indian Ocean	209	62	28	81	218	295	326

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

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



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

B Additional Results for Empirical Analysis

B.1 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.⁶¹ We then find for each day of the journey the local weather in the vessel’s current sub-segment and average these local weather conditions across the vessel’s journey. We use this variable to run regression (2) and obtain from this regression an alternative measure of weather-induced shipping times $\tilde{t}^{s,weather,alt}_{ft}$. Table B.1 shows the coefficients from this regression, analogous to Table 4 in the main text. Overall, the results are very similar.

We then construct weather-induced shipping delays analogously to Section 3.3 and re-run the regression (4):

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

where $\text{FracDelayed}_{ft}^{weather,alt}$ is the fraction of imports subject to shipping delays constructed with the alternative measure. Table B.2 shows the results. The coefficients are qualitatively very similar, but slightly smaller than with the measure in the main text.

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

Dep. Var:	$\hat{t}_{xhrtvfa}^s$	$\hat{t}_{xhrtvfa}^s$	$\hat{t}_{xhrtvfa}^s$
Wave Height ^s	−0.014*** (0.000)	−0.014*** (0.000)	−0.009*** (0.000)
Direction ^s		−0.004*** (0.000)	0.008*** (0.000)
Wave Height ^s × Direction ^s			−0.006*** (0.000)
R-Squared	0.010	0.010	0.010
Observations	5,728,000	5,728,000	5,728,000

⁶¹We use AIS data in Appendix A.3 to show that this assumption approximately holds in vessel tracking data.

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

	(1)	(2)	(3)
Weather Shocks (Alternative Measure)			
Dependent Variable (in logs):	Sales	Profits	Employees
Frac Delayed	-1.864*** (0.518)	-1.090** (0.529)	-0.369 (0.242)
Importer FE	Y	Y	Y
Year FE	Y	Y	Y
R-Squared	0.97	0.91	0.98
Observations	40,500	40,500	40,500

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. R^2 is the overall fit inclusive of the fixed effects.

B.2 Robustness of the Risk Regressions

In this section, we show that our results on the impact of shipping risk on importers' behavior from Section 3.4 are robust to a variety of alternative specifications.

First, we replace the separate firm and time fixed effects in our baseline specification (6) with firm-time fixed effects. This specification picks up changes in a firms' sourcing strategy that are common to all imported goods, and identifies our effect from within-firm variation in sourcing behavior across goods. In this specification, firms that only import a single good are therefore dropped. The results in Table B.3 are similar to the baseline and increase in magnitude for all effects except for total imports. Going from the 25th to the 75th percentile of the weather risk distribution is associated with a 10.3% increase in routes and a 7.2% increase in the number of suppliers. The route HHI, supplier HHI, and import value drop by 5.5%, 4.0%, and 2.8%, respectively.

Next, we use the alternative measure of weather-induced shipping risk introduced in Appendix B.1. As described above, we decompose each trip segment into smaller sub-segments, assume that vessels travel through these areas at constant speed, and determine for each day of a vessel's journey the local weather conditions in a vessel's current location. We then use the weather-induced shipping times based on this measure, $\tilde{t}^{s,weather,alt}$, to construct our risk measure using the same steps as described in Section 3.4. Table B.4 shows the results from running regression (6) with this alternative risk measure. The effects are similar to before.

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

$$\widehat{StdTime}_{fht-3,t-1}^{main} \equiv \sum_{r \in R(x^{main})} \omega_{fxhr,t-3,t-1}^{x^{main}} \widehat{StdTime}_{xhrt-3,t-1}, \quad (19)$$

where the main supplier of good h is defined as the one with the largest shipment value to importer f in the years $t - 3$ to $t - 1$. The weighted average is now taken only across the routes r used by the main supplier, $R(x^{main})$ and the weights $\omega_{fxhr,t-3,t-1}^{x^{main}}$ are the import shares of each route for the main supplier. The results of running specification (6) with this measure are in Table B.5. Going from the 25th to the 75th percentile of shipping risk increases the number of routes and suppliers by 5.7% and 4.8%, respectively, while the route HHI, supplier HHI, and total import value fall by 3.9%, 3.7%, and 4.5%. Since we are not able to compute shipping risk for all main suppliers (for example because they have fewer than 10 transactions), the number of observations drops relative to the other regressions.

Table B.6 includes firms with only one supplier, which are dropped in the main specification. These firms are therefore by definition not diversified in t . Here, we find that the relationship between shipping risk and the number of suppliers and routes remains significantly positive, but decreases slightly in magnitude. However, all results remain strongly significant.

Finally, in Table B.7 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.

Table B.3: Firm-Time Fixed Effects

	(1)	(2)	(3)	(4)	(5)
Dep. Var.:	Number of Routes	Number of Suppliers	HHI over Routes	HHI over Suppliers	Value Imported
Std Time	0.169*** (0.011)	0.117*** (0.010)	-0.090*** (0.004)	-0.066*** (0.004)	-0.046*** (0.013)
Importer-Year FE	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y
R-squared	0.75	0.71	0.48	0.49	0.92
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.

Table B.4: Alternative Measure of Weather Risk

	(1)	(2)	(3)	(4)	(5)
Dep. Var.:	Number of Routes	Number of Suppliers	HHI over Routes	HHI over Suppliers	Value Imported
Std Time	0.181*** (0.010)	0.122*** (0.010)	-0.091*** (0.003)	-0.069*** (0.003)	-0.038*** (0.012)
Importer-Year FE	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y
R-squared	0.76	0.71	0.49	0.49	0.92
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.

Table B.5: Risk Measure Based on Main Supplier

	(1)	(2)	(3)	(4)	(5)
Dep. Var.:	Number of Routes	Number of Suppliers	HHI over Routes	HHI over Suppliers	Value Imported
Std Time	0.094*** (0.014)	0.078*** (0.011)	-0.064*** (0.007)	-0.061*** (0.005)	-0.074*** (0.012)
Importer-Year FE	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y
R-squared	0.76	0.70	0.56	0.53	0.93
Observations	46,000	46,000	46,000	46,000	46,000

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 B.6: Including Firms with One Supplier

	(1)	(2)	(3)	(4)	(5)
Dep. Var.:	Number of Routes	Number of Suppliers	HHI over Routes	HHI over Suppliers	Value Imported
Std Time	0.159*** (0.010)	0.098*** (0.010)	-0.088*** (0.004)	-0.056*** (0.003)	-0.060*** (0.011)
Importer-Year FE	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y
R-squared	0.75	0.72	0.57	0.56	0.90
Observations	96,000	96,000	96,000	96,000	96,000

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 B.7: Shipping Time Risk and Import Demand with Inventory-Sales Ratio

	(1)	(2)	(3)	(4)	(5)
Dep. Var.:	Number of Routes	Number of Suppliers	HHI over Routes	HHI over Suppliers	Value Imported
Std Time	0.129*** (0.011)	0.083*** (0.009)	-0.077*** (0.004)	-0.056*** (0.004)	-0.080*** (0.0013)
Inventory- Sales Ratio	-0.036** (0.015)	-0.018 (0.013)	0.005* (0.003)	0.003 (0.003)	-0.062 (0.048)
Importer FE	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y
Observations	55,000	55,000	55,000	55,000	55,000

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.

C Additional Results for Section 4

C.1 Derivation of equation (11)

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 (7) and (10), firms' profits in the second stage are given by

$$\max_{l_p} \pi_f = \varphi_f^{\frac{\sigma-1}{\sigma}} l^{(1-\gamma)\frac{\sigma-1}{\sigma}} \left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + \left(\sum_{i=1}^N \alpha_i \right)^{\frac{\varepsilon-1}{\varepsilon}} x^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\gamma \frac{\varepsilon}{\varepsilon-1} \frac{\sigma-1}{\sigma}} P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma} - p_D x_D - N p_M x - w l - w N F, \quad (20)$$

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:

$$l_f = \left[\tilde{\gamma} \varphi_f^{\frac{\sigma-1}{\sigma}} \left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + \left(\sum_{i=1}^N \alpha_i \right)^{\frac{\varepsilon-1}{\varepsilon}} x^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\gamma \frac{\varepsilon}{\varepsilon-1} \frac{\sigma-1}{\sigma}} P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma} w^{-1} \right]^{\frac{1}{1-\tilde{\gamma}}} \quad (21)$$

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 (20), expected profits are

$$\max_{x_D, x, N} \chi_f \mathbb{E} \left[\left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + \left(\sum_{i=1}^N \alpha_i \right)^{\frac{\varepsilon-1}{\varepsilon}} x^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\psi} \right] - p_D x_D - N p_M x - w N F, \quad (22)$$

where $\chi_f \equiv \left(\varphi_f^{\frac{\sigma-1}{\sigma}} P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma} \right)^{\frac{1}{1-\tilde{\gamma}}} w^{-\frac{\tilde{\gamma}}{1-\tilde{\gamma}}} \left[(\tilde{\gamma})^{\frac{\tilde{\gamma}}{1-\tilde{\gamma}}} - (\tilde{\gamma})^{\frac{1}{1-\tilde{\gamma}}} \right]$ and $\psi \equiv \gamma \frac{\varepsilon}{\varepsilon-1} \frac{\sigma-1}{\sigma} \frac{1}{1-\tilde{\gamma}}$.

C.2 Derivations for Section 4.3

Derivation of expression (12). In the case without the domestic input, given the realization of input qualities, firm profits are

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

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

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

Solving for labor and plugging back into profits yields:

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

Letting

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

and ψ be defined as in the main text, the ex-ante firm profit maximization problem is:

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

The first order condition with respect to x gives:

$$\psi x^{\psi-1} \chi_f \mathbb{E} \left(\sum_{i=1}^N \alpha_i \right)^\psi - Np_M = 0 \quad (23)$$

Using this condition to eliminate x yields expression (12) in the main text, where $\tilde{\chi} \equiv \left(\psi^{\frac{\psi}{1-\psi}} - \psi^{\frac{1}{1-\psi}} \right) p_M^{-\frac{\psi}{1-\psi}} \chi^{\frac{1}{1-\psi}}$.

Derivation of expression (13). A second-order approximation of the function $\bar{\alpha}^\psi$ around $\mathbb{E}[\alpha]$ yields

$$\bar{\alpha}^\psi \approx (\mathbb{E}[\alpha])^\psi + \psi (\mathbb{E}[\alpha])^{\psi-1} (\bar{\alpha} - \mathbb{E}[\alpha]) + \psi \frac{(\psi-1)}{2} (\mathbb{E}[\alpha])^{\psi-2} (\bar{\alpha} - \mathbb{E}[\alpha])^2 \quad (24)$$

Taking expectations on both sides gives

$$\mathbb{E} [\bar{\alpha}^\psi] \approx (\mathbb{E} [\alpha])^\psi + \psi \frac{(\psi - 1)}{2} (\mathbb{E} [\alpha])^{\psi-2} \mathbb{E} [(\bar{\alpha} - \mathbb{E} [\alpha])^2] \quad (25)$$

since $\mathbb{E} [\bar{\alpha}] = \mathbb{E} [\alpha]$. Noting that $\mathbb{V} [\bar{\alpha}] = \mathbb{E} [(\bar{\alpha} - \mathbb{E} [\alpha])^2]$ and that $\mathbb{V} [\bar{\alpha}] = \mathbb{V} [\alpha] / N$, we plug this into equation (12) to obtain expression (13).

Proof of Propositions 1 and 2. For continuous N , the first order condition with respect to N is given by:

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

Using the implicit function theorem:

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

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

$$\epsilon_{N^*, \mathbb{V}[\alpha]} = \frac{\partial N^*}{\partial \mathbb{V}[\alpha]} \frac{\mathbb{V}[\alpha]}{N^*} = - \frac{\left(1 - \frac{\frac{\psi^2}{2} (\mathbb{E}[\alpha])^{\psi-2} \mathbb{V}[\alpha]}{N (\mathbb{E}[\alpha])^\psi - \psi \frac{(1-\psi)}{2} (\mathbb{E}[\alpha])^{\psi-2} \mathbb{V}[\alpha]} \right)}{\left(\frac{\frac{\psi^2}{2} (\mathbb{E}[\alpha])^{\psi-2} \mathbb{V}[\alpha]}{N (\mathbb{E}[\alpha])^\psi - \psi \frac{(1-\psi)}{2} (\mathbb{E}[\alpha])^{\psi-2} \mathbb{V}[\alpha]} - 2 \right)} < 1 \quad (26)$$

Note that the denominator is negative as $G_N < 0$ follows from the second order condition at the optimal N . The numerator is positive whenever condition (14) in the main text holds; this establishes the first part of the proposition.

Relying on the first order condition with respect to x in (23), import values are given by:

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

where the second line uses the approximation in (24). It follows that the effect of a mean preserving spread on import values depends on its effect on the variance of average supplier

quality $\mathbb{V}[\alpha]/N$. In turn, this effect is given by:

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

as $\epsilon_{N^*, \mathbb{V}[\alpha]} < 1$ follows from 26. It follows that, regardless of whether the optimal N falls or increases after a mean preserving spread, $\mathbb{V}[\alpha]/N$ increases and import values fall.

To Proposition 2, note first that

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

as $G_\varphi > 0$. That is, more efficient firms select a greater number of suppliers. Note next that $\epsilon_{N^*, \mathbb{V}[\alpha]}$ depends on efficiency only through N . Using (26), it is straightforward to show that:

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

which establishes the first part of the proposition. Note next that import value $M \equiv Nxp^*$ is proportional to expected revenues, so that:

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

Note next that, relying on the approximation in (25), we have that:

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

It follows that:

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

Note that the right hand side of this expression increases with efficiency (as both N and $\epsilon_{N^*, \mathbb{V}[\alpha]}$ increase). This completes the proof of Proposition 2.

C.3 The Effects of Changes in Expected Quality

In the main text, we study the effects of a mean preserving spread to input qualities. Because our empirical exercises involve days and not qualities, it is necessary to also explore the effect of changes in expected quality. A higher expected quality affects expected revenue in

two ways. There is a direct positive effect that stems from output being increasing in input qualities and is captured by the first term in expression (13). A second effect is that a higher expected quality makes a given variance less important for expected revenue, as captured by the second term in expression (13). In other words, the importance of the variance of quality for expected revenues is mediated by the mean quality. By making a given variance less detrimental, this effect reduces the incentive to adding suppliers. This negative effect turns out to be dominated by the positive direct effect whenever the variance of quality is large relative to its mean, as captured by the coefficient of variation. The following result formalizes this argument.

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

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

whenever:

$$\frac{1}{N^*} \frac{\mathbb{V}[\alpha]}{(\mathbb{E}[\alpha])^2} > \frac{2-3\psi}{(1-\psi)(2-\psi)} \frac{2}{\psi}. \quad (27)$$

Note that, when $\psi < 2/3$, the right hand side of condition (27) is lower than the right hand side of condition (14) in the main text.⁶² It follows that in the region of parameters where a mean preserving spread to quality increases the returns to adding suppliers (i.e., condition (14) is satisfied), there is a sub-region where an increase in expected quality increases the returns to adding suppliers and a sub-region where the opposite happens.

Proof. Note that (13) can be written as:

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

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

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

⁶²This follows from the fact that

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

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

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

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

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

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