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

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

May 2023

Abstract

We use transaction data on U.S. manufacturing imports to construct a novel measure of supply chain risk at the firm-level: the historical volatility of ocean shipping times. We isolate the unexpected component of this risk that is induced by weather conditions along the maritime shipping routes. We first establish that shipping delays are associated with lower employment and revenues. We then show that importers actively diversify this source of risk: firms with more volatile supply chains have lower imports, a larger number of foreign suppliers, and a lower concentration of expenditure across suppliers. A quantitative model of importing with shipping time risk and importer risk aversion can rationalize our empirical findings. The model implies that U.S. importers would prefer to have certain but up to 25% longer shipping times rather than face supply chain risk.

*Blaum: 270 Bay State Road, Boston, MA 02215, email: jblaum@bu.edu. Esposito: Joyce Cummings Center, Room 663, 177 College Avenue, Medford, MA 02155, email: federico.esposito@tufts.edu. Heise: 33 Liberty Street, New York, NY 10045, email: sebastian.heise@ny.frb.org.

We thank our discussant Davin Chor, as well as Laura Alfaro, Georg Schaur and participants at numerous seminars and conferences for their valuable comments. The views and opinions expressed in this work do not necessarily represent the views of the Federal Reserve Bank of New York or the Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 1883 (CBDRB-FY21-P1883-R8915, CBDRB-FY23-P1883-R10453).

1 Introduction

The past decades have seen a dramatic transformation in the international organization of production, with supply chains spanning an increasing number of countries (Antras and Chor (2021)). These changes had important consequences for the pattern of international trade, with intermediate inputs now accounting for roughly two-thirds of global trade (Johnson and Noguera (2012)). At the same time, the increasing reliance on imported inputs has exposed firms to novel sources of risk, such as foreign natural disasters, supplier bankruptcy, or port backlogs (Baldwin and Freeman (2021); Carvalho et al. (2021)). According to a recent survey of U.S. manufacturing firms by the Federal Reserve Bank of New York, more than 50% of the respondents were affected by supply chain disruptions in 2020 (Federal Reserve Bank of New York (2021)). While a growing literature in international trade focuses on foreign input sourcing (Antras et al. (2017a), Blaum et al. (2018b), Halpern et al. (2015)), little is known about how importers cope with supply chain risks and how this affects international trade patterns. An important challenge for answering these questions is the lack of comprehensive measures of supply chain risk at the firm level.

We shed light on this topic by focusing on a specific source of risk: the volatility of shipping times. Using transaction-level import data on ocean shipments provided by the U.S. Census Bureau, we construct measures of shipping time risk at the supplier-product-route level. We then estimate the effect of shipping time volatility on U.S. manufacturing firms' performance and sourcing behavior. Our results indicate that U.S. importers that are more exposed to shipping time risk feature lower imports, a larger number of foreign suppliers, and a lower concentration of expenditure across suppliers, suggesting that firms actively diversify this source of risk. We incorporate risky shipping times into an otherwise standard quantitative model of firm-level importing and assess its ability to replicate our empirical findings. We show that a model where firms maximize expected profits cannot account for the negative association between imports and shipping time risk we document in the data. In contrast, when firms are risk averse, the model can rationalize our empirical findings.

Our methodology relies crucially on the measurement of the components of shipping times that are unexpected to the importers. We make progress on this front in two ways. First, we exploit the granularity of the U.S. Census Bureau's Longitudinal Firm Trade Transactions Database (LFTTD). The LFTTD provides, for each import transaction since 1992, the identity of the U.S. importer and its overseas supplier, the product at the HS-10 level, value and quantity shipped, the mode of transportation and, importantly, the shipping time between the foreign port of exit and the U.S. port of entry, after customs are cleared.[1]

[1]Since most shipments by airplane, rail, or truck arrive in the U.S. within one day, we focus on vessel

With this detailed data at hand, we use a rich set of fixed effects to remove the components of the shipping times that are presumably known to the importers at the time the inputs are bought, such as shipping charges, the shipment weight, or the seasonality. The volatility of the resulting "residualized" shipping times is our baseline measure of supply chain risk. This measure captures several sources of uncertainty, such as weather-related delays, suppliers' ability to comply with regulation, or delays at the port of entry, due to backlogs or unexpected inspections.

Second, we refine our risk measure by focusing on the unanticipated variation in shipping times that is induced by weather conditions along the shipping routes. Weather can have a sizable effect on shipping times: Filtz et al. (2015) shows that greater wind speeds in the direction of travel or higher waves significantly reduce shipping times. Our identifying assumption is that *realized* weather conditions along the entire maritime route are not anticipated by the importer at the time the orders are made. To this end, we first use customs data on the vessel name, foreign port of exit, and U.S. port of entry to determine all the intermediate stops a vessel made on its journey to the U.S. We then construct the shipment route by finding the shortest ocean distance for each trip segment of the vessel's journey, using data from Eurostat's SeaRoute. We finally compute the weather conditions along each shipment route using detailed data on the average height and direction of waves from the National Oceanic and Atmospheric Administration (NOAA).

We start our analysis by assessing the effects of longer unexpected shipping times on firms' economic performance. We identify for each importer in a given year the shipments that were "delayed", i.e. which had a travel time larger than the 95th percentile of the shipping times distribution for a given route. We document that U.S. importers with a higher share of delayed inputs experienced significant declines in sales and employment in the same year. These effects are stronger if we focus only on the delays triggered by weather shocks, and for importers that do not rely on backup suppliers. These results highlight the negative effects of supply chain disruptions on firms production, similarly to recent empirical evidence, such as Boehm et al. (2019) and Carvalho et al. (2021), and provide the basis for our methodology to measure supply chain risk.

We next study whether U.S. importers adjust their sourcing strategy and import demand to cope with this source of risk, and explore the different margins of adjustment. We compute the standard deviation of the unexpected shipping times over 5-year rolling windows for each foreign supplier and product. We measure the risk exposure of an importer for a given product with the volatility of its suppliers averaged over all the routes used in the

shipments, but also explore whether firms use air transportation to manage risk in a robustness analysis. Vessel-based imports accounted for nearly half of all U.S. imports by value in the past decade. See Table 1.

previous 5 years. We analogously construct a measure of risk that is based only on the travel times predicted by weather conditions along the route. We estimate panel regressions at the importer-product-year level and include a rich set of fixed effects and controls. Our results indicate that U.S. importers that are more exposed to shipping time risk feature lower imports, a larger number of foreign suppliers, and a lower concentration of expenditure across suppliers. This suggests that U.S. importers actively diversify this source of risk along different margins of adjustment. The economic magnitudes are large. We find that going from the 25th to the 75th percentile of the supply chain risk distribution reduces the total value imported by 5.3-8.2 percent depending on the specification.

Motivated by these findings, we incorporate supply chain risk into a standard model of firm-level importing along the lines of Blaum et al. (2018a); Halpern et al. (2015); Gopinath and Neiman (2014). Firms can source their inputs domestically or from a number of foreign suppliers and there is imperfect substitutability between suppliers. Firms can choose the number of foreign suppliers, which is limited by per-supplier fixed costs. The key departure from the literature is the presence of stochastic input qualities. To map the theory to our empirical setting, we follow Hummels and Schaur (2013) and assume that slow delivery times reduce the effective quality of inputs. We assume that importers know the probability distribution function of their supplier delivery times and choose the number of suppliers and the values imported before the realizations of the shipping times.

We consider a quantitative version of the model where firms are heterogeneous both in their productivity and in the shipping risk they face. We calibrate the model to match salient features of the data. In particular, we set the fixed cost of adding foreign suppliers to match the aggregate import share, and choose the elasticity of inputs quality to delivery times to match the observed elasticity of sales to shipping times. Intuitively, this captures the negative relationship between longer shipping times and final sales we document. We then examine whether the calibrated model can reproduce the observed relationships between inputs sourcing and supply chain risk. We find that a model where firms maximize expected profits fails to do so, as it predicts a positive relationship between shipping time risk and imports. To rationalize the evidence, we introduce risk aversion in the firm's problem following Gervais (2018) and Esposito (2022). We calibrate the risk aversion to exactly match the elasticity of total imports with respect to delivery risk that we estimate. The implied calibration suggests that U.S. importers are moderately risk averse: they would prefer to have certain but up to 25% longer shipping times rather than face supply chain risk.[2]

[2]To be precise, we find that U.S. importers would choose, in order to avoid supply chain risk, to increase the average shipping times from 16 days – the observed average shipping time – to 16-20 days, i.e. an increase on average of 6%, but up to 25%.

Most of the literature on risk and international trade has focused on the impact of uncertainty on production and exports, i.e., on the output side — see e.g. Ramondo et al. (2013), Fillat and Garetto (2015), Esposito (2022), Baley et al. (2020) and De Sousa et al. (2020). In contrast, we focus on the risk faced by importers. On the empirical front, we provide a novel comprehensive measure of supply chain risk at the firm-level and assess whether this source of risk has significant implications for the design of firms' production processes. Our empirical findings suggest that firms actively use the extensive and intensive margins of importing to reduce the impact of supply chain disruptions on their production processes. On the theory side, existing models of firm-level input sourcing typically abstract from supplier risk considerations — see e.g. Gopinath and Neiman (2014), Halpern et al. (2015), Antras et al. (2017b) and Blaum et al. (2018b). We contribute to this literature by exploring how supplier risk affects our understanding of input trade with heterogeneous firms.

There is a small but growing literature that empirically investigates how firms manage supply chain risk, such as Clark et al. (2014), Huang (2017), Gervais (2018) and Charoenwong et al. (2020).[3] Our paper departs from the approach used by this literature in that we use transaction-level data of the universe of U.S. importers and their trade relationships with foreign suppliers, and exploit variation in shipping times across different routes, allowing to perform an accurate firm-to-firm analysis. On this front, we also connect to a recent literature that studies firm-to-firm relationships, see e.g. Miyauchi (2018), Tintelnot et al. (2018), Heise (2019), Bernard et al. (2019) and Esposito and Hassan (2023). In contrast to this strand of literature, we focus on the risk associated with buyer-supplier relationships. Using highly granular data on weather conditions along the shipping routes allows us to identify plausibly unexpected shipping times which we use to study the impact of risk on firms' production process.

We also connect to an empirical literature that studies the effect of supply chain disruptions on firms (Carvalho et al. (2016), Barrot and Sauvagnat (2016) and Boehm et al. (2019)), and a related strand of works that investigate the importance of shipping times for international trade (Evans and Harrigan (2005), Hummels and Schaur (2010) and Hummels and Schaur (2013)). We contribute to this body of work by analyzing the effect of risk in the supply chain on firms' import demand.

Finally, we contribute to a strand of literature in management that advocates for the use of multiple suppliers as a way to reduce supply chain risk, see e.g. Zsidisin et al. (2004), Kleindorfer and Saad (2005) and Chopra and Sodhi (2014). Our contribution is to conduct a systematic analysis of the risk diversification behavior of importers by exploiting the richness of U.S. Census data.

[3] Handley et al. (2020) study the effect of trade policy uncertainty on firms' sourcing strategies.

The remainder of the paper proceeds as follows. Section 2 lays out our conceptual framework for measuring supply chain risk. Section 3 constructs our measure of risk, while Section 4 discusses our empirical results. Section 5 presents the model, which we calibrate in Section 6 to perform our quantitative analysis. Section 7 concludes.

2 Measurement of Supply Chain Risk

The first goal of our analysis is to measure supply chain risk. To this end, we propose a simple model of stochastic shipping times that will guide our methodology. Based on the characteristics of our dataset, we consider a buyer f that orders a shipment s of product h from seller x in time period t. The seller can either be a related party or at arms' length, and this is captured by the index a. The shipment arrives to the U.S. on vessel v via route r, which consists of a combination of the ports of origin and destination and possibly any intermediate stops made by the vessel. The shipment has a weight of W^s and the buyer spends C^s dollars on freight costs and insurance for shipment s—we refer to these expenses as shipping charges. The time it takes for the shipment to arrive in the U.S., $T^s_{xhrtvfa}$, is random and follows:

$$\begin{aligned}\ln(T^s_{xhrtvfa}) = {} & \bar{\pi}_x + \bar{\alpha}_h + \bar{\gamma}_r + \bar{\theta}_t + \bar{\xi}_v + \bar{\delta}_f + \bar{\omega}_a \\ & + \pi_x + \alpha_h + \gamma_r + \theta_t + \xi_v + \delta_f + \omega_a + \eta \ln (C^s) + \rho \ln (W^s)\end{aligned} \tag{1}$$

where $\bar{\pi}_x, \bar{\alpha}_h, \bar{\gamma}_r, \bar{\theta}_t, \bar{\xi}_v, \bar{\delta}_f, \bar{\omega}_a$ capture deterministic components associated with, respectively, a supplier, product, route, time period, vessel, buyer, and related party status, which are known to the buyer. They capture the average effect of each characteristic on the shipping times. For instance, the supplier-level average $\bar{\pi}_x$ may reflect the ability of a supplier to comply with regulation, while the product-level deterministic component $\bar{\alpha}_h$ may capture the fact that some products are harder to ship or take longer to get cleared at customs. The route-specific term $\bar{\gamma}_r$ may capture route characteristics, such as ocean distance to the US; or characteristics of the ports of departure or entry, such as the average time it takes to unload a shipment. The time component $\bar{\theta}_t$ may capture the effect of seasonality, as induced by weather, while $\bar{\xi}_v$ 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. Finally, 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. We further assume that $\pi_x, \alpha_h, \gamma_r, \theta_t, \xi_v, \delta_f, \omega_a$ are mean-zero random variables, each with a known distribution. Finally, shipping charges and shipment weight log-linearly affect shipping times.

Equation (1) implies that the population variance of log shipping times for a tuple (x, h, r, t, v, f, a), given charges C^s and weight W^s is given by:

$$\sigma^2_{xhrtvfa} = \mathbb{V}\left(T^s_{xhrtvfa}\right) = \mathbb{V}\left(\pi_x\right) + \mathbb{V}\left(\alpha_h\right) + \mathbb{V}\left(\gamma_r\right) + \mathbb{V}\left(\theta_t\right) + \mathbb{V}\left(\xi_v\right) + \mathbb{V}\left(\delta_f\right) + \mathbb{V}\left(\omega_a\right). \quad (2)$$

We assume that equation (1) is the data generating process of the shipping times we observe in the data. As our goal is to measure the uncertainty around delivery times, a key step is to isolate the unexpected components from the deterministic ones. We argue that this depends crucially on the level of disaggregation at which we measure risk. Consider the most disaggregated measure of risk, at the (x, h, r, t, v, f, a) level. The observed shipping times within such a tuple contain shipments made with different charges C^s and weights W^s. This variation is expected by the buyer and should not be taken into account when measuring risk. To take out this variation, we regress the observed log shipping times on charges and weight and identify η and ρ, and then define the shipping times net of these controls as $\tilde{t}^s_{xhrtvfa} \equiv \ln(T^s_{xhrtvfa}) - \hat{\eta}C^s - \hat{\rho}W^s$. In what follows, we will work with these net shipping times. For a tuple (x, h, r, t, v, f, a), an unbiased estimator of the population variance $\sigma^2_{xhrtvfa}$ is the sample variance of the observed shipping times within that tuple, denoted by $\hat{\sigma}^2_{xhrtvfa}$, as all the deterministic components drop out.

Consider now measuring risk at a more aggregated level. For example, at the supplier-product (x, h) level. The sample variance of the net-of-charges shipping times includes also the deterministic terms $\bar{\gamma}_r, \bar{\theta}_t, \bar{\xi}_v, \bar{\delta}_f$ and $\bar{\omega}_a$, as they are not constant within each (x, h) pair. In other words, holding the supplier and product constant, the realized net shipping times include variation induced by deterministic components that vary by route-time, vessel, and buyer. As these are reasonably known to the importers, they affect the estimated sample variance, leading to a mis-measurement of supply chain risk.

This issue arises more generally when one measures risk at a more aggregated level than the tuple (x, h, r, t, v, f, a). To control for the deterministic components associated with the omitted dimensions, we perform a residualization step. Specifically, we regress the observed log shipping times on fixed effects for those dimensions that we aggregate over. This identifies the set of constants for these dimensions. In our analysis below, our baseline measure of risk is at the supplier-product-route level, (x, h, r). Therefore, we first estimate the coefficients for C^s and W^s and the fixed effects $\bar{\theta}_t$, $\bar{\xi}_v$, $\bar{\delta}_f$, and $\bar{\omega}_a$, and then subtract them from the shipping times to obtain the residualized shipping times:

$$\hat{t}^s_{xhrtvfa} \equiv \ln(T^s_{xhrtvfa}) - \hat{\eta}\ln(C^s) - \hat{\rho}\ln(W^s) - \bar{\theta}_t - \bar{\xi}_v - \bar{\delta}_f - \bar{\omega}_a. \tag{3}$$

By construction, the stochastic components of the time, vessel, buyer, and relationship characteristics are included in $\hat{t}^s_{xhrtvfa}$, while the deterministic components are not.[4] Our measure of supply chain risk is the sample variance of $\hat{t}^s_{xhrtvfa}$ within (x, h, r), denoted by $\hat{\sigma}^2_{xhr}$. In the following section, we implement our methodology using U.S. transaction-level import data.

3 Constructing Supply Chain Risk

In this section we describe how we construct our measure of supply chain risk. We first discuss the U.S. Census datasets we employ in our methodology. We then describe how we construct the maritime route taken by each shipment combining Census data with data on ocean distance from Eurostat's SeaRoute, and how we merge data on weather conditions from NOAA. We finally provide some stylized facts on the factors that affect shipping times.

3.1 Data

Our empirical analysis relies mostly 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 occurred during the period 1992-2016. Each transaction is associated with an identifier of the U.S. importer, the 10-digit Harmonized System (HS10) product code traded, the mode of transportation (vessel, air, etc.), the vessel name for ocean shipments, as well as the value, weight, and quantity shipped. We use the concordance by Pierce and Schott (2012) to transform the HS10 codes into time-consistent product codes, and calculate prices as the value of shipment divided by the quantity shipped. Importantly, we also observe the date of departure abroad (the export date) and the date of importation in the U.S. (after customs are cleared). We compute (raw) shipping times as the difference, in days, between the importation date in the U.S. and the export date from the foreign port. For vessel-borne imports, we also observe the foreign port of departure and the U.S. port of arrival. We use this information to construct shipping routes, as explained below.

[4] To be precise, the shipping times $\hat{t}^s_{xhrtvfa}$ are the realization of the random variable $\hat{T}^s_{xhrtvfa} = (\bar{\pi}_x + \bar{\alpha}_h + \bar{\gamma}_r) + \pi_x + \alpha_h + \gamma_r + \theta_t + \xi_v + \delta_f + \omega_a$. Note that the deterministic components $\bar{\pi}_x, \bar{\alpha}_h, \bar{\gamma}_r$ do not need to be removed because they are constant within each (x, h, r) tuple.

The LFTTD contains two additional variables that are important for our purposes. First, the data report an identifier of the seller in the form of a Manufacturer ID (MID), an alphanumeric code that combines information on the seller's country, name, street address, and city.[5] Because shipping negotiations most likely take place at the overall 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, since plants of the same firm located in different locations have a different MID.[6] This approach follows earlier work by Kamal et al. (2015) and Kamal and Monarch (2018), who show that the number of MIDs in the Census data matches well with official firm counts once the street address or the city component are omitted. Second, the LFTTD data also contain information on whether the transaction is between related parties. We keep track of related-party status since it could affect shipping times. We classify an importer-exporter pair as related if it had at least one related-party transaction in a given year. Since in some cases an import transaction spans multiple customs records, we collapse the data to the supplier (x) - HS10 (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 more detail the data cleaning process in Appendix A.

Table 1 reports the summary statistics of our baseline import dataset. The first column presents all manufacturing imports for 1992-2016 and the second column shows vessel-borne trade only. Our dataset covers about 10.5 trillion dollars of imports (in 2009 dollars), of which about 40 percent are by vessel. We also 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 our measure, as explained below.

The second dataset we use is the Longitudinal Business Database (LBD), which reports in each year U.S. establishments' employment and a time-consistent six-digit North American Industrial Classification (NAICS) 2007 industry code constructed by Fort and Klimek (2018). We collapse 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 plant with the highest employment. We focus our analysis on all importers that are in the manufacturing sector, since these firms are most likely to source intermediates into production that require an on-time arrival. We also obtain firms' total sales, cost of materials, and employee compensation from the Census

[5]Specifically, the MID consists of the two-digit ISO country code of origin of the good, the first three characters of the first word of the exporter's name, the first three characters of the second word of the exporter's name, the first four numbers of the street address of the foreign exporter, and the first three letters of the foreign exporter's city.

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

Table 1: U.S. Import Transaction Summary Statistics

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

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

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 materials and employee costs. We merge these data into the transaction-level import dataset.[7]

3.2 Construction of Shipping Routes

As the Census data do not provide direct information on shipping routes, we develop an algorithm to construct them. We assign each transaction to a *trip*, defined as a journey of a vessel that begins with the loading of cargo at a foreign port and ends (possibly after some intermediate stops) with the unloading of cargo at a U.S. port. As a starting point, we sort all transactions involving a given vessel by their foreign departure date. We then take all the vessel's transactions and assign them to a single trip ("Trip 1"). Next, we find the earliest arrival date of the vessel in the U.S. for this trip. If there exist 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"). Since the departures are after the vessel initially arrived in the U.S., they cannot have been from the same trip, and must have been from a new trip that began after the vessel left the U.S. again. We continue splitting trips into sub-trips until no further splits are possible.

In the second step, we re-combine some of the previously separated trips. This step deals with cases in which the U.S. arrival date is possibly misreported.[8] For each of the

[7]Since the LBD , ASM, and CMF data are only available annually, each transaction in the same year will have the same employment, sales, and costs.

[8]For example, if transactions 1-5 depart abroad on 6/22 and arrive in the U.S. on 7/5, transaction 6 departs on 6/23 and purportedly arrives on 6/24, and transactions 7-10 depart on 6/25 and arrive on 7/5, then the procedure would assign transactions 1-6 to one trip and transactions 7-10 to another, even though almost all shipments arrive on the same day in the U.S.. It seems likely that the arrival date for transaction

trips assigned in the previous step, we compare the latest importation date in the U.S. to the earliest departure date abroad of the *next* trip. If the earliest departure date abroad of the next trip is before the latest importation date of the earlier trip, then the two trips must have been part of the same journey and we recombine these trips into one. We again iterate through this procedure until no more trips can be combined. The resulting dataset contains trips with completely non-overlapping foreign departure and importation dates for each vessel.

We construct for each trip an ordered sequence of foreign and U.S. ports. We order the foreign ports using their mean export dates, i.e., foreign ports with an earlier mean export date were reached earlier. Similarly, we order the U.S. ports by their mean dates of importation. Putting together this information, we construct the sequence of ports visited by each shipment, for example, Le Havre - Birmingham - New York - Newport News. We refer to a leg of the trip between two ports as *trip segment.* We next determine the exact route taken by the vessel by taking the shortest ocean distance for each trip segment, using data from Eurostat's SeaRoute program.[9]

3.3 Construction of Weather Conditions

Our final source of data is 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 contain the height and direction (in degrees) of significant waves, swell, and wind waves, respectively, at hourly or three-hourly frequency for geo locations at 0.5 degree distances in the oceans during the period 2011 to 2017. We focus on the direction and height of significant waves. We compute the daily average of significant wave height and direction for each geo location in the oceans on each day. If information is unavailable at a geo location, we interpolate the weather using the information of neighboring geo locations.

For each route segment that we have previously constructed, we obtain the geo coordinates (longitudes and latitudes) of each point in 0.5 degree increments. We then merge these points with the information on significant wave height and direction on each day to obtain the average weather on each route segment on each day. We translate the wave direction into a relative direction using the likely bearing of the vessel on the route.[10] A relative direction of zero means that the significant waves are in the same direction as the vessel's travel. We finally use the trip information to obtain for each transaction the days the good spent

6 is misreported.

[9] This shortest sea route is not necessarily the shortest air distance, since vessels may have to pass, e.g., through the Panama Canal to reach their destination.

[10] For example, a wave direction of 75 degrees for a vessel traveling at direction 90 degrees would be translated into a wave direction of 15 degrees relative to the vessel's path of travel.

Figure 1: Volatility of Wave Height on Selected Routes

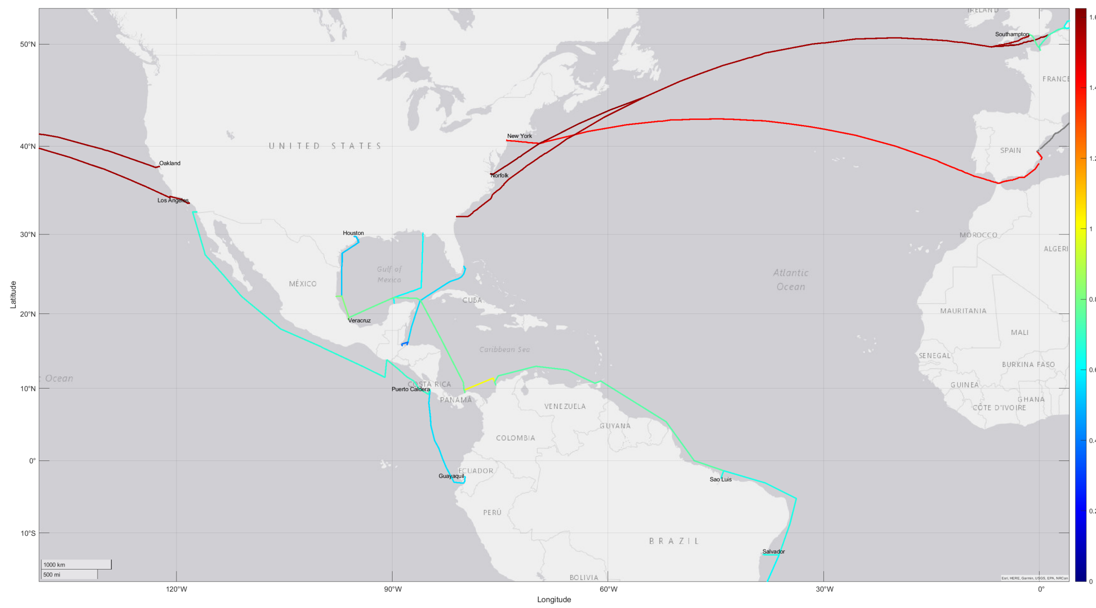

Notes: The figure shows the standard deviation of wave height across all days from 2011-2017 for selected shipping routes into the U.S.

on each route segment.[11] We average the wave height and relative direction across these days to obtain for each transaction in our dataset the average weather conditions on the shipment's journey. Table 10 in Appendix A 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.

To illustrate the source of our exogenous variation, Figure 1 reports the standard deviation of wave height, one of our measures of weather conditions, for some of the main shipping routes into the U.S. There are significant differences across routes. Routes across the Atlantic and Pacific have higher volatility in wave height than routes along the coast of South America. Importantly, there is variation even across routes that are relatively close to each other. Compared to the Valencia - New York route, the route from Southampton - Savannah, which crosses the Northern Atlantic Ocean, has a standard deviation of wave height that is more than 10 percent higher.

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

Table 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^{***}		0.010^{***}
			(0.000)		(0.000)
Log Shipping Charges				0.004^{***}	-0.003^{***}
				(0.000)	(0.000)
R^2	0.616	0.616	0.616	0.616	0.617
Route FE	Y	Y	Y	Y	Y
Observations (thousands)	35,480	35,480	35,480	35,480	35,480

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

3.4 Factors Affecting Shipping Times

In our analysis, we focus on vessel-borne shipments because the other modes of transportation (train, truck and airplane) feature virtually no dispersion in shipping times (see Table 11 in the Appendix). Vessel shipments are also substantially slower than all other modes of transportation, as they take 16 days on average to arrive in the U.S., while air and truck shipments arrive in the U.S. on average within the same day, while train shipments arrive on average in four days. Importantly, Table 11 also shows that vessel is the most important mode of transportation measured by total import value.[12]

We now examine the factors affecting vessel-borne shipping times to motivate the variables we will include in our residualization procedure. We present the results in Table 2. All regressions include fixed effects for the combination of ports used. In column (1), we test whether seasonality affects shipping times by adding dummies for each quarter of the year. Shipping times for a given port pair are about $2-3\%$ shorter in the summer quarters of the northern hemisphere, highlighting the role of weather in affecting shipping routes. In column (2), we find that related party transactions have slightly longer shipping times relative to arms-length transactions. The next columns find a positive relationship between shipping time and shipment weight, and a negative one between shipping charges and delivery times, conditional on weight.

[12] We present some additional statistics on shipping times by geographic region in Appendix B.

Table 3: Attributes of Importer-Product-Year Tuples

	All		Vessel Only	
	Mean	Standard Deviation	Mean	Standard Deviation
Suppliers per Importer-Product-Year	1.90	3.84	1.83	2.81
Suppliers per Importer-Product-Country-Year	1.39	1.56	1.39	1.39
Dep. Port-Entry Port Pairs per Importer-Product-Year			2.18	3.18
Dep. Port-Entry Port Pairs per Importer-Product-Country-Year			1.75	1.91
Vessels per Port Combination-Year			16.78	44.94

Source: LFTTD and authors' calculations. Table reports the mean and standard deviation across importer (f) by product (h) by year or importer (f) by product (h) by country (c) by year tuples during our 1992 to 2016 sample period.

4 Shipping Risk and Import Demand

In this Section, we first document some characteristics of firms' sourcing behavior (Section 4.1) and outline our procedure to residualize the shipping times from deterministic components that are presumably known to the importer (Section 4.2). We then present two key pieces of empirical evidence. First, we establish that long shipping times, which we informally refer to as ''delays'', are associated with lower employment and revenues of U.S. importers (Section 4.3). This finding is the basis for using the volatility of shipping times as a measure of risk. Second, we document how firms adjust their sourcing strategy and import demand in response to shipping time risk (Section 4.4).

4.1 Characteristics of Firms' Sourcing Behavior

We first document some key characteristics of firms' sourcing behavior and show that firms use multiple suppliers and routes even for the same HS10 product. The first row of Table 3 shows that the average importer has 1.9 suppliers for the same HS10 product per year across all modes of transportation. The large standard deviation compared to the mean indicates that there is significant heterogeneity across buyers. Figure 2a shows the tail of the distribution of the number of suppliers used by importers within HS10-year. While the 60th percentile of importers uses only one supplier for a given product, some firms in the tail of the distribution have more than 50 suppliers. Overall, 29 percent of importer-product pairs source from two or more suppliers. These importers tend to be large: firms sourcing a given product from at least two suppliers account for almost 90% of imports. In Section 4, we will argue that one reason to use multiple suppliers is to hedge against shipping delay risk.

The second row of Table 3 shows that many firms have multiple suppliers of the same

Figure 2: Distribution of the Average Number of Suppliers per Importer

(a) Within HS10 Product

(b) Within HS10 Product-Country

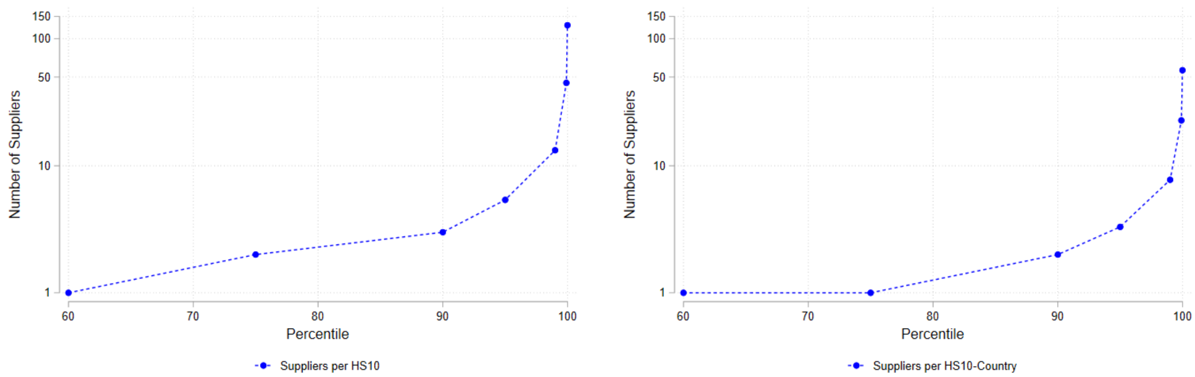

HS10 even within the same country. The average number of suppliers within importer-HS10-country-year is 1.4. Figure 2b plots the tail of the distribution of the number of suppliers used by importers within HS10-country-year. About 21 percent of within-country importer-product pairs source from multiple suppliers, but these firms account for roughly three quarters of imports.

Firms may not only use multiple suppliers to hedge against delays but can also, for vessel shipments, use different ports. The third and fourth rows of Table 3 show that importers on average use two departure port - entry port combinations to source a given HS10 per year, even within the same country. The final row shows that the average number of vessels supplying the U.S. on any port combination is 17, with a large standard deviation.

4.2 Residualization of shipping times

Motivated by the evidence shown in Section 3.4, we implement our residualization procedure described in Section 2. We construct shipping times net of the deterministic season, vessel, importer, and arm's-length components, by running the following regression:

$$\ln(T^s_{xhrtvfa}) = \bar{\theta}_t + \bar{\xi}_v + \bar{\delta}_f + \bar{\omega}_a + \eta \ln(C^s) + \rho \ln(W^s) + \epsilon^s_{xhrtvfa}, \tag{4}$$

where $\bar{\theta}_t, \bar{\xi}_v, \bar{\delta}_f, \bar{\omega}_a$ are season-year, vessel, importer, and relationship status fixed effects, and C^s and W^s are the observed charges and shipment weight, respectively. We obtain from this regression the residuals $\hat{t}^s_{xhrtvfa} \equiv \ln(T^s_{xhrtvfa}) - \hat{\eta}\ln(C^s) - \hat{\rho}\ln(W^s) - \bar{\theta}_t - \bar{\xi}_v - \bar{\delta}_f - \bar{\omega}_a$ as defined above. These residuals contain the random components of shipping times as well as the deterministic exporter, product, and route components.

4.3 Effect of Shipping Delays

Using the residualized shipping times, we identify shipping delays in our dataset in four steps. First, we compute the average shipping time residual, $\bar{\hat{t}}^s_{xhrtvfa}$, within each product by route.[13] This variable measures how long it takes on average to ship a specific product from a specific origin to destination, conditional on season, vessel, importer, relationship status, charges, and weight. We do not residualize with respect to the exporter component, and hence delays due to on average slower suppliers will be included in our measure. Second, we compute for each shipment the log deviation of the observed $\hat{t}^s_{xhrtvfa}$ from the associated average. Third, we set a delay indicator variable $\mathbb{D}^s_{xhrtvfa}$ equal to one if the log deviation of a tuple is larger than the 95th percentile of deviations in the data. Thus, we flag as "extreme delays" shipments which took far longer than normal within a given product and port of departure - port of entry combination. Finally, we multiply $\mathbb{D}^s_{xhrtvfa}$ by the value of affected shipments, $ShipDelayed_{xhrtvfa} = \mathbb{D}^s_{xhrtvfa} \cdot Value^s_{xhrtvfa}$. We then aggregate this variable to the importer-year level by summing across all transactions of the importer in the year, and scale by he importer's total production costs (materials plus labor) from the manufacturing census or the ASM. Our measure thus provides the fraction of the importer's total inputs that are subject to the delay. We use our delay variable to estimate the following regression of firms' outcomes:

$$\ln(Y_{ft}) = \beta_1 FracDelayed_{ft} + \gamma_f + \rho_t + \epsilon_{ft}, \tag{5}$$

where Y_{ft} is either number of employees, sales, or operating profits (sales minus materials and labor costs), $FracDelayed_{ft}$ is the share of firm f's costs that is subject to import delays, and γ_f and ρ_t are firm and time fixed effects, respectively.

Columns (1)-(3) of Table 4 present the results from an OLS regression. Shipping delays are significantly negatively associated with all three variables. Increasing the fraction of delayed shipments in costs by one percent is associated with a drop in sales by 2% and a fall in profits by 0.9%. The number of employees falls by 0.4%. Therefore, extreme delays in shipments, which are reasonably unexpected by U.S. importers, have a large and significant negative impact on importers' economic outcomes.

A concern with the OLS design is reverse causality. For example, importers that expect a decline in sales in the current year could request shipments that are relatively slower. We use weather shocks to isolate a plausibly exogenous source of variation in shipping delays. Filtz et al. (2015) shows that higher significant waves or greater wind speeds in the direction of travel can reduce shipping times. Our identifying assumption is that realized weather conditions along the entire maritime route are not anticipated by the importer at the time

[13] To limit the number of cells, we approximate the route by the port of departure - port of entry combination.

Table 4: Effect of Extreme Delays on Firms' Outcomes

	(1)	(2)	(3)	(4)	(5)	(6)
		OLS			Weather IV	
Dependent Variable (in logs):	Sales	Profits	Employees	Sales	Profits	Employees
Frac Delayed	-1.982^{***}	-0.869^{**}	-0.371^{**}	-6.131^{***}	-3.307^{**}	-0.816^{*}
	(0.387)	(0.351)	(0.173)	(2.056)	(1.651)	(0.420)
Importer FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
F-Stat				19.53	19.53	19.53
Observations	142,000	142,000	142,000	142,000	142,000	142,000

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 country level.

the orders are made. To this end, we modify the residualization procedure as follows:

$$\begin{aligned} \ln(T^s_{xhrtvfa}) = & \left(\bar{\pi}_x + \bar{\alpha}_h + \bar{\gamma}_r + \bar{\theta}_t + \bar{\xi}_v + \bar{\delta}_f + \bar{\omega}_a\right) \\ & + \pi_x + \alpha_h + \left(\tilde{\gamma}_r + \tilde{\theta}_t + Weather_{rt}\right) + \xi_v + \delta_f + \omega_a + \eta \ln\left(C^s\right) + \rho \ln\left(W^s\right) \end{aligned} \quad (6)$$

where we decompose the route component γ_r and the season-year effect θ_t into a component that is due to the weather, $Weather_{rt}$, and the residual terms $\tilde{\gamma}_r$ and $\tilde{\theta}_t$. To construct our instrument, we first find the residualized shipping time $\hat{t}^s_{xhrtvfa}$ using equation (4) as before. We then regress $\hat{t}^s_{xhrtvfa}$ on significant wave height, relative direction, and the interaction between the two, to obtain the predicted shipping time due to weather, $\hat{t}^{s,weather}_{xhrtvfa}$. This variable reflects the component of the shipping time that is due to unanticipated weather conditions, and is unrelated to the other variables. Table 13 in Appendix B shows that higher waves increase vessel speed (consistent with Filtz et al. (2015)), and tail wind marginally increases vessel speed. We then use the same procedure as above to compute a delay indicator $\mathbb{D}^{s,weather}_{xhrtvfa}$ based on the weather-induced shipping times, and generate $FracDelayed^{weather}_{ft}$, the share of the firm's costs that are subject to import delays due to weather. We then use this variable as an instrument for our delay variable.

Columns (4)-(6) of Table 4 show that our results strengthen with the IV specification. Increasing the fraction of delayed shipments in costs by one percent is associated with a drop in sales by 6.1%, a fall in profits by 3.3%, and a decline in the number of employees by 0.8%.

We next investigate whether firms with alternative back-up suppliers perform better when they face extreme delivery delays. For this purpose, we split our delay indicator into two variables $\mathbb{D}^{s,alt}_{xhrtvfa}$ and $\mathbb{D}^{s,noalt}_{xhrtvfa}$, where $\mathbb{D}^{s,alt}_{xhrtvfa}$ is equal to one if the shipment is delayed, as defined before, but the importer has more than one supplier for the affected product, and $\mathbb{D}^{s,noalt}_{xhrtvfa}$ is equal to one if the shipment is delayed and there is no alternative supplier. Thus,

$\mathbb{D}^{s}_{xhrtvfa} = \mathbb{D}^{s,alt}_{xhrtvfa} + \mathbb{D}^{s,noalt}_{xhrtvfa}$. We then compute the fraction of affected inputs as before and run

$$\ln(Y_{ft}) = \beta_1 FracDelayed^{alt}_{ft} + \beta_2 FracDelayed^{noalt}_{ft} + \gamma_f + \rho_t + \epsilon_{ft},$$

where $FracDelayed^{alt}_{ft}$ and $FracDelayed^{noalt}_{ft}$ are the fraction of imports affected by delays with and without an alternative supplier, respectively. We instrument for these variables using the analogous variables computed using the weather-induced shipping times.

Table 14 in Appendix B shows that the effect on firms' outcomes is significantly more negative when firms do not have back-up suppliers. A one percent increase in the fraction of delayed shipments in cost is associated with a decline in sales of 8.3 percent for firms without alternate suppliers, more than double the decline in sales for those with alternative suppliers. Similarly, firms without alternative suppliers exhibit a larger response in terms of profits and employees - a one percent increase in the fraction delayed leads to a 5.8 percent decrease in sales and a 1.3 percent decrease in the number of employees, relative to 1.6 and 0.4 percent, respectively, for firms with alternative suppliers. In the next section, we analyze how firms set up their sourcing strategy to reduce the risk imposed by shipping delays.

4.4 Shipping Time Risk and Import Demand

We next examine how U.S. importers adjust their sourcing strategy in response to the riskiness of their supply chains. To obtain a measure of risk, we use our baseline residualized shipping times $\hat{t}^{s}_{xhrtvfa}$ constructed above and compute their mean and standard deviation within each supplier-HS10-route-year cell using five-year rolling windows, starting from 1992-1996. Cells with fewer than 10 transactions are dropped. The resulting standard deviations $\widehat{StdTime}_{xhrt-5,t-1}$ provide a measure of shipping risk of the supplier for the given route and product. We then estimate the following empirical specification for each importer-product-year:

$$\ln(Y_{fht}) = \beta_1 \ln(\widehat{StdTime}_{fht-5,t-1}) + \beta_2 X_{fht} + \gamma_h + \gamma_t + \epsilon_{fht}, \tag{7}$$

where Y_{fht} is an importer choice variable. We consider multiple dimensions of sourcing: (i) the number of suppliers, (ii) the number of routes, (iii) the concentration of suppliers, (iv) concentration of supplier-routes, (v) the value per supplier per year, and (vi) the total value imported per year. The variable $\widehat{StdTime}_{fht-5,t-1}$ is a weighted average over the risk of the importer's suppliers and routes over the past five years, where the weights are the importer's purchases across all modes of transportation. For shipments by modes other than vessel we assign a risk of zero, since they mostly arrive within one day. Thus, an importer using the same supplier and route for vessel imports could face a lower shipping time risk dependent

Table 5: Shipping Time Risk and Import Demand

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.:	Number of Suppliers	Number of Routes	HHI over Suppliers	HHI over Supplier-Routes	Value per Supplier-Year	Value Imported per Year
Std Time	0.053***	0.096***	−0.036***	−0.054***	−0.151***	−0.098***
	(0.004)	(0.006)	(0.002)	(0.002)	(0.006)	(0.007)
Importer FE	Y	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Observations	328,000	328,000	328,000	328,000	328,000	328,000

Notes: Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the firm level.

on the share of imports by vessel. Shipments by vessel for which we cannot compute the standard deviation since they are not associated with enough transactions are assigned the average risk of the remaining imports. The variable X_{fht} are controls, and γ_h and γ_t are product and year fixed effects, respectively.

Our controls include the suppliers' mean shipping time over the previous five years. This accounts for the fact that suppliers located in countries further away may have a higher standard deviation of shipping times purely because they have more scope for delays. We also control for the average unit value charged to the importer in year t. Importers could face a trade-off between risk and cost, with low input prices compensating for higher riskiness. We explicitly take into account this potential mechanism by controlling for the average unit value. Third, we control for the weighted average of the suppliers' shipments used to compute the risk measure and for the total imports of the importer-product over the past five years. Larger shipments tend to be associated with a greater weight, which is correlated with more volatile shipments.

Table 5 presents our results. Column (1) documents a positive and significant relationship between the number of suppliers used and shipping risk. An increase from the 25th to 75th percentile of the risk distribution (0.84 log points) increases the number of suppliers by 4.5%. Column (2) shows the relationship between the number of routes and shipping risk. Again, we find a positive and significant relationship between the two. An increase from the 25th to the 75th percentile of the risk distribution increases the number of routes used by 8.1%.

We next look at the relationship between shipping time risk and the concentration of suppliers for each importer. Column (3) shows a negative and significant relation between shipping risk and the HHI over suppliers, suggesting that importers with riskier suppliers diversify their sources more. Column (4) shows that this effect is larger for supplier-route combinations. In column (5), we look at the relationship between our risk measure and log value per supplier in each year. We find that going from the 25th to the 75th percentile of the risk distribution decreases value per supplier-year by 12.7%. Finally, column (6) investigates the relationship between shipping time risk and firms' total imports in a given year. A one log point increase in riskiness over the past five years is associated with a decline in the total imports by around 8.2%.

As highlighted in the previous section, some of the remaining variation in shipping times could be still anticipated by the importer. To alleviate this concern, we refine our risk measure by focusing on the variation in shipping times induced by weather along the shipping route of each transaction. The starting point to construct this risk measure is our weather-induced shipping time $\hat{t}^s_{xhrtvfa}$, constructed similarly to the previous section. We then compute the mean and the standard deviation of these weather-induced shipping times. Since we only have weather information for the period 2011 to 2017, we compute our measures over three-year rolling windows only. We use the standard deviation $\widehat{StdTime}^{weather}_{xhrt-3,t-1}$ as a measure of weather-induced shipping risk for the supplier and route. As before, we then compute a weighted average over an importer's suppliers and routes to arrive at the importer-specific risk measure $\widehat{StdTime}^{weather}_{fht-3,t-1}$. We then estimate the empirical specification:

$$\ln(Y_{fht}) = \beta_1 \ln(\widehat{StdTime}^{weather}_{fht-3,t-1}) + \beta_2 X_{fht} + \gamma_h + \gamma_t + \epsilon_{fht} \tag{8}$$

where the controls are the same as before.

Our results in Table (6) are consistent with the previous results using the overall risk measure. Going from the 25th to the 75th percentile of the risk distribution increases the number of suppliers by 4.7% percent and the number of routes by 5.86%. Again we find a negative and significant relationship between weather-related shipping risk and both HHI over suppliers and supplier-routes. Columns (5) and (6) show that weather-related shipping risk is associated with lower imports per supplier and total value imported, with inter-quartile range effects of -9.8% and -5.3%, respectively.

Table 6: Shipping Time Risk and Import Demand (Weather Risk)

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.:	Number of Suppliers	Number of Routes	HHI over Suppliers	HHI over Supplier-Routes	Value per Supplier-Year	Value Imported per Year
Std Time	0.077***	0.123***	−0.052***	−0.074***	−0.163***	−0.086***
	(0.008)	(0.009)	(0.003)	(0.003)	(0.011)	(0.012)
Importer FE	Y	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Observations	72,500	72,500	72,500	72,500	72,500	72,500

Notes: Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the firm level.

4.4.1 Robustness

In Appendix B.1, we document that our results are robust to omitting firm fixed effects, including firms with only one supplier, and alternative sets of controls. We also show that our results hold when we use only the risk of the main supplier, rather than a weighted average across all suppliers. 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 over the past year, and regress it on our measure of risk, constructed for vessel shipments alone. We include the baseline controls, and additionally include controls for the value imported by air and vessel separately over the past five years. Table 20 in Appendix B.1 indicates that higher shipping risk is associated with a higher likelihood of using air shipments, suggesting that firms use air transportation to hedge the ocean shipping risk.

Overall, our findings suggest that U.S. importers systematically react to supplier risk. In equilibrium, importers matched to riskier suppliers feature more suppliers and lower imports. An important caveat is that our analysis is silent on the matching between importers and foreign suppliers. Note, however, that the positive relation between the risk measure and the number of suppliers holds even with the addition of firm fixed effects, suggesting that selection may not fully explain this relation.

5 A Model of Input Sourcing with Shipping Risk

In this Section, we lay out a theoretical framework to study the effects of shipping time risk on import demand. The model builds on the standard models of importing with firm heterogeneity in Halpern et al. (2015); Blaum et al. (2018b); Gopinath and Neiman (2014). Similarly to Hummels and Schaur (2013), we assume that inputs' shipping times are a measurable component of input quality. The key departure from the literature is the assumption that such shipping times and thus input qualities are stochastic. In Section 6, we evaluate the ability of this model of importing with shipping time risk to rationalize the facts documented in Section 4.

We start by describing the model environment in Section 5.1 and then turn to characterizing the solution to the firm's problem in Section 5.2. We describe the equilibrium of the model in Section 5.3.

5.1 Environment

We consider a small open economy populated by a mass of firms that produce differentiated varieties. Firms combine local and foreign inputs according to the following nested structure:

$$\begin{aligned} y_f &= \varphi_f l^{1-\gamma} x^{\gamma} \\ x &= \left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + x_I^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}} \\ x_I &= \left(\sum_{i=1}^{N} (\alpha_i x_i)^{\frac{\kappa-1}{\kappa}} \right)^{\frac{\kappa}{\kappa-1}} , \end{aligned} \tag{9}$$

where f denotes a firm, $\gamma \in (0,1)$ and $\varepsilon, \kappa > 1$. The firm combines intermediate material inputs x with labor l with a Cobb-Douglas aggregator, with an heterogeneous efficiency φ_f drawn from a known distribution with CDF $H(\cdot)$. In turn, materials are a CES aggregator of a domestic bundle x_D and a foreign bundle x_I. The firm sources its foreign inputs from N suppliers with quantities denoted by x_i and qualities by α_i. We assume that each firm is matched with a pool of potential foreign suppliers. The foreign inputs are combined with a CES aggregator with elasticity of substitution κ. This production structure is standard in the literature, see e.g. Gopinath and Neiman (2014).

We assume that qualities are a function of the time it takes the inputs to be shipped from the foreign country to the domestic market. Inputs that arrive in a timely manner are more productive – see Hummels and Schaur (2013) for a similar treatment of time. In particular, we assume that

$$\alpha_i = d_i^{-\tau}, \tag{10}$$

where d_i are the number of days it takes to ship foreign input i to firm φ. An additional key assumption is that shipping times d_i are stochastic and unknown to the firm when it makes production choices. For simplicity, we assume that the input shipping times are ex-ante symmetric and independent. That is, the $d_i's$ are i.i.d. across input suppliers. We denote their CDF by $G_f(\cdot)$ and assume it is known to the firm. We allow this distribution to be firm-specific as, in our quantitative exercise below, firms differ in the riskiness of their foreign suppliers.

In terms of market structure, we assume that firms are price takers in input markets. Thus, firms can use any quantity of the foreign inputs, the domestic bundle and labor at prices p^*, p_D and w, respectively. We assume that foreign inputs prices p^* are the same across all suppliers. Any variable trade costs are embedded in these prices p^*. As standard in the literature, we assume that the extensive margin of trade is limited by fixed costs. In particular, each additional foreign supplier entails payment of a fixed cost F units of domestic labor.[14]

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

$$U = \left(\int c_f^{\frac{\sigma-1}{\sigma}} df \right)^{\frac{\sigma}{\sigma-1}}, \tag{11}$$

where $\sigma > 1$ and c_f denotes final consumption of the good produced by firm f. In addition, we assume a structure of roundabout production by which firms use the output of all other domestic firms as inputs. In particular, we assume that the domestic bundle x_D is produced according to a CES aggregator given by (11).[15] Finally, we assume that firms compete under monopolistic competition on the output market.

5.2 Firm's problem

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

$$R_{Df} = y_f^{1-1/\sigma} P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma}, \tag{12}$$

where P is the CES price index associated with (11) and S denotes total spending. Both P and S are endogenous variables determined in general equilibrium.

Firms are risk-neutral and maximize profits in two stages. In the first stage, 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. In the second stage, profits are given by

$$\max_{l_p} \ \pi_f = \varphi_f^{\frac{\sigma-1}{\sigma}} l^{(1-\gamma)\frac{\sigma-1}{\sigma}} \left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + \left(\sum_{i=1}^{N} \alpha_i^{\frac{\kappa-1}{\kappa}} \right)^{\frac{\kappa}{\kappa-1}\frac{\varepsilon-1}{\varepsilon}} x^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\gamma \frac{\varepsilon}{\varepsilon-1}\frac{\sigma-1}{\sigma}} P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma} - p_D x_D - N p^* x - wl - NF, \tag{13}$$

where α_i are the quality shocks which depend on the realized shipping times. The optimal

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

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

choice of production labor is given by

$$l_f = \left[\tilde{\gamma} \varphi_f^{\frac{\sigma-1}{\sigma}} \left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + \left(\sum_{i=1}^{N} \alpha_i^{\frac{\kappa-1}{\kappa}} \right)^{\frac{\kappa}{\kappa-1} \frac{\varepsilon-1}{\varepsilon}} x^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\gamma \frac{\varepsilon}{\varepsilon-1} \frac{\sigma-1}{\sigma}} P^{\frac{\sigma-1}{\sigma}} S^{1/\sigma} w^{-1} \right]^{\frac{1}{1-\tilde{\gamma}}} \tag{14}$$

where $\tilde{\gamma} \equiv (1-\gamma) \frac{\sigma-1}{\sigma}$. In the first stage, taking l_{pf} as given, the firm maximizes *expected* profits. By plugging the expression for l_f into equation (13), expected profits are

$$\max_{x_D, x, N} \ \mathbb{E}[\pi_f] = \chi_f \mathbb{E} \left[\left(x_D^{\frac{\varepsilon-1}{\varepsilon}} + \left(\sum_{i=1}^{N} (d_i)^{-\frac{\tau(\kappa-1)}{\kappa}} \right)^{\frac{\kappa}{\kappa-1} \frac{\varepsilon-1}{\varepsilon}} x^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\psi} \right] - p_D x_D - N p^* x - NF, \tag{15}$$

where χ_f and ψ are functions of firm efficiency, equilibrium objects and parameters.[16] Note that the expectation operator is taken over the possible realizations of d_i and thus depends on the distribution of shipping times $G_f(\cdot)$. We now turn to defining the equilibrium of the economy.

5.3 Equilibrium

We consider an equilibrium where firms maximize profits, the consumer maximizes utility, and goods markets clear. An equilibrium can be fully characterized by the pair (S, P), where S is the level of total domestic spending and P is the price index of the locally-produced manufacturing bundle. Note that the price of the domestic input bundle is given by $p_D = P$ due to the symmetry between the domestic input aggregator and consumer utility. The price of foreign inputs is exogenously given by p^*. We now turn to the characterization of the equilibrium.

The consumer's budget constraint is:

$$PC = wL + \Pi,$$

where $\Pi \equiv \int \pi_f df$ are total firms' profits. Standard calculations imply that $P = \left(\int p_f^{1-\sigma} df \right)^{\frac{1}{1-\sigma}}$. Note that both prices and profits depend on the actual realization of the shipping days:

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

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

and

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

where $\{d_i\}$ are the realized shipping times for each supplier i of firm f.[17] Thus, the price index and total profits are given by

$$P = \left(\int\int p\left(\{d_i\}; S, P, w; \varphi\right)^{1-\sigma} dG\left(d\right) dH\left(\varphi\right)\right)^{\frac{1}{1-\sigma}} \tag{16}$$

$$\Pi = \int\int \pi\left(\{d_i\}; S, P, w; \varphi\right) dG\left(d\right) dH\left(\varphi\right). \tag{17}$$

Aggregate spending must satisfy

$$S = wL + \Pi + p_D \int x_D\left(S, P, w; \varphi\right) dH\left(\varphi\right), \tag{18}$$

where $x_D\left(S, P, w; \varphi\right)$ is the demand for the domestic input.

The consumer budget constraint and the goods market clearing conditions imply that the trade balance is given by

$$TB = -\left(L - L_d\right), \tag{19}$$

where L_d is the total labor demand and L is the labor endowment of the consumer. Because the equilibrium does not impose labor market clearing, the manufacturing sector can be a net supplier of labor to the rest of the economy and thus attain a trade deficit.

An equilibrium consists of the vector (S, P) such that equations (16) and (18) are satisfied and firms maximize profits. It is relatively straightforward to augment the model with the inclusion of a non-tradeable sector and international trade in final manufacturing goods, as in Blaum (2019).

6 Quantitative Analysis

In this section, we assess whether the theory outlined above can rationalize the empirical findings of Section 4. To this end, we first calibrate the parameters of the model to match salient features of the data. We then simulate the model and reproduce the risk regressions we estimate in the data. We show how the model fails to match the empirical elasticities. We further show that introducing risk aversion in the firms' maximization problem allows the model to match the elasticities.

[17] Note that we have written both p_f and π_f explicitly as a function of the equilibrium variables S, P and w (as the choices N, x and x_D ultimately depend on those and parameters).

Table 7: Calibration of Baseline Model

Parameters			Moments	Model	Data
Fixed Cost per Supplier	F	0.05	Aggregate Import Share	0.23	0.23
Inputs Quality Elasticity	τ	0.47	Revenue Elasticity	-0.06	ND

Notes: The import share corresponds to the fraction of materials expenditures accounted by imported inputs in the manufacturing sector. Revenue elasticity is the elasticity of firm revenues to the average shipping times. ND stands for "not-disclosed": we have submitted the disclosure request to the U.S. Census, but it has not been approved yet.

6.1 Calibration

We take the following parameters from the literature. As in Blaum et al. (2018b), we set the output elasticity of materials to $\gamma = 0.6$ and the elasticity of substitution between domestic and foreign inputs to $\varepsilon = 2.38$. The elasticity of substitution across foreign suppliers is set to $\kappa = 8$ from Eaton and Kortum (2002), and the demand elasticity to $\sigma = 5$, in line with the average estimate for manufacturing differentiated goods in Broda and Weinstein (2006).

We assume that firm efficiency φ is drawn from a log-normal distribution. We set the mean and variance of efficiency to match the mean and variance of log-sales observed in the data. Shipping days are also drawn from a log-normal distribution, with mean and variance taken from our dataset. To introduce cross-sectional variation in the supply chain risk faced by importers, we assume that there are two types of risk, low and high. In other words, half of the firms in our economy draw the shipping days from a distribution with low variance, and the other half from a distribution with high variance. We set these variances as the averages of the baseline importer-level risk measure we use in our empirical analysis (the standard deviation of the shipping times $\widehat{StdTime}_{fht-5,t-1}$) within the groups of firms above and below the median risk measure.

We then calibrate the fixed cost of adding foreign suppliers, F, and the elasticity of input quality to shipping time, τ, to match the observed aggregate import share and the elasticity of revenue with respect to average shipping times.[18] Intuitively, this captures the negative relationship between longer shipping times and final sales we have documented in the empirical analysis. Table 7 summarizes the results of the calibration.

6.2 The Effect of Supply Chain Risk

With the calibrated model at hand, we now examine whether the model can match the observed relationships between firm choices and supply chain risk. In particular, we simulate a large number of importers and run the same regressions we have estimated in the data.

[18] We normalize the foreign input price p^* and the wage w to 1.

Table 8: Risk Elasticities

Variable	Data	Baseline Model	Extended Model
Number of suppliers	0.08	0.05	0.09
Value imported per supplier	-0.16	0.05	-0.16
Total value of imports	-0.08	0.1	-0.07

Notes: The table reports the elasticity of each variable with respect to the standard deviation of shipping times predicted by weather, $\widehat{StdTime}^{weather}_{fht-3,t-1}$. The empirical elasticities are reported in Table 6, columns 1, 5 and 6.

We regress the number of suppliers, the value imported per supplier, and the total value of imports on the importer-level measure of supply chain risk, which is the standard deviation of shipping times. As shown in Table 8, the model delivers elasticities that are different from the ones estimated in Table 6: the elasticity of the number of suppliers is positive but a bit lower than in the data, and the intensive margin elasticities have the wrong sign. Therefore, the standard model of importing augmented with shipping time risk is unable to reproduce the negative relationship between supplier risk and import values documented above.

A Model with Importer Risk Aversion To rationalize the data, we modify the baseline model and introduce risk aversion in the firm's problem. In particular, we follow Gervais (2018) and Esposito (2022) and assume that each firm maximizes the following function:

$$\max \quad E\left(\pi_f\right) - \frac{\rho}{2} Var\left(\pi_f\right) \tag{20}$$

where $\rho \geq 0$ is the coefficient of absolute risk aversion. This expression can be micro-founded with a second-order approximation of a CARA utility function. Importantly, this is a simple and intuitive way to introduce (constant) risk aversion in the firms' decision making process.

Next, we discipline the risk aversion parameter ρ with the estimated elasticities of the number of suppliers and the imports per supplier with respect to shipping risk. In particular, we jointly calibrate the model parameters F, τ, ρ to match the same moments as in the baseline calibration as well as the estimated risk elasticities.

Our results suggest a risk aversion of 9 and values for F and τ similar to the baseline calibration — see Table 9. To give economic significance to this estimate, we compute the Certainty Equivalent (CE) implied by our calibration. We define CE_f as the minimum number of certain shipping days such that firm f would choose to avoid supply chain risk. We can compute CE_f such that the the expected utility the firm derives when facing supply chain risk equals the utility the firm obtains when there is no uncertainty and shipping times are always equal to CE_f:

Table 9: Calibration, extended table

Panel A: Calibrated Parameters			Baseline Model	Extended Model
Fixed Cost per Supplier	F		0.05	0.12
Quality Elasticity	τ		0.47	0.43
Risk aversion	ρ		-	9
Panel B: Targeted Moments		Data	Baseline Model	Extended Model
Import Share		0.23	0.23	0.21
Revenue Elasticity		ND	-0.06	-0.05
Elasticity of number of suppliers		0.08	-	0.09
Elasticity of imports per supplier		-0.16	-	-0.16
Panel C: Fixed Parameters			Baseline Model	Extended Model
Demand elasticity	σ		5	5
Elasticity between domestic and foreign	ε		2.38	2.38
Materials share	γ		0.6	0.6
Elasticity across foreign suppliers	κ		8	8
Average shipping time (days)			16	16
Variance of shipping times (high type)			ND	ND
Variance of shipping times (low type)			ND	ND
Average productivity			2	2
Variance of productivity			0.23	0.23

Notes: The import share corresponds to the fraction of materials expenditures accounted by imported inputs in the manufacturing sector. Revenue elasticity is the elasticity of firm revenues to the average shipping times. ND stands for "not-disclosed": we have submitted the disclosure request to the U.S. Census, but it has not been approved yet.

$$E\left(\pi_f\right) - \frac{\rho}{2} Var\left(\pi_f\right) = \chi_f \left(\left(x_D\right)^{\frac{\varepsilon-1}{\varepsilon}} + \left(x\right)^{\frac{\varepsilon-1}{\varepsilon}} \left(\left(CE_f\right)^{-\tau} \right)^{\frac{\varepsilon-1}{\varepsilon}} \left(N\right)^{\frac{\kappa}{\kappa-1}\frac{\varepsilon-1}{\varepsilon}} \right)^{\psi} - p_D x_D - Npx - NF$$

By construction, the certainty equivalent is larger than the average shipping time if the firm is risk averse, and it is heterogeneous across importers, as they vary in their productivity and risk type. Our calibrated model implies that CE_f ranges between 16-20 days, with an average of 17. This means that U.S. importers would choose, in order to avoid supply chain risk, to increase their shipping times from 16 (the average number of shipping days in the data) to 16-20 days, i.e. an increase on average of 6%, but up to 25%.

7 Conclusions

In this paper, we use U.S. Census shipment-level data to construct a novel measure of supply chain risk. We then explore how exposure to shipping time risk correlates with the pattern of import demand of U.S. manufacturing firms at the intensive and extensive margins. Our results suggest that U.S. importers that are more exposed to shipping time volatility feature lower imports, a larger number of foreign suppliers, and a lower concentration of expenditure across suppliers, which indicates that firms actively diversify this source of risk. To rationalize this evidence, we introduce risky delivery times into a standard quantitative model of firm-level importing. We show that a model where firms maximize expected profits cannot account for the negative relationship between imports and shipping time risk we document in the data. In contrast, when firms are risk averse, the model can rationalize our empirical findings. To match the data, the model implies that U.S. importers have a moderate degree of risk aversion. An important caveat of our current approach is that we take the matching between importers and suppliers as exogenously given. An avenue for future research is to improve our understanding of this matching process, and evaluate its implications for the aggregate economy.

References

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

Antras, P., Fort, T.C., Tintelnot, F., 2017a. The margins of global sourcing: Theory and evidence from us firms. American Economic Review 107, 2514--64.

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

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

Baley, I., Veldkamp, L., Waugh, M., 2020. Can global uncertainty promote international trade? Journal of International Economics 126, 103347.

Barrot, J.N., Sauvagnat, J., 2016. Input specificity and the propagation of idiosyncratic shocks in production networks. The Quarterly Journal of Economics 131, 1543--1592.

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

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

Blaum, J., Lelarge, C., Peters, M., 2018a. Firm Size and the Intensive Margin of Import Demand .

Blaum, J., Lelarge, C., Peters, M., 2018b. 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.

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

Carvalho, V.M., Nirei, M., Saito, Y., Tahbaz-Salehi, A., 2016. Supply chain disruptions: Evidence from the great east japan earthquake. Columbia Business School Research Paper .

Carvalho, V.M., Nirei, M., Saito, Y.U., Tahbaz-Salehi, A., 2021. Supply Chain Disruptions: Evidence from the Great East Japan Earthquake*. The Quarterly Journal of Economics 136, 1255--1321. URL: https://doi.org/10.1093/qje/qjaa044, doi:10.1093/qje/qjaa044.

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

Chopra, S., Sodhi, M., 2014. Reducing the risk of supply chain disruptions. MIT Sloan management review 55, 72--80.

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

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

Eaton, J., Kortum, S., 2002. Technology, Geography, and Trade. Econometrica 70, 1741--1779. URL: http://doi.wiley.com/10.1111/1468-0262.00352, doi:10.1111/1468-0262.00352.

Esposito, F., 2022. Demand risk and diversification through international trade. Journal of International Economics 135, 103562.

Esposito, F., Hassan, F., 2023. Import competition, trade credit, and financial frictions in general equilibrium. Trade Credit, and Financial Frictions in General Equilibrium (February 18, 2023) .

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

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

Fillat, J., Garetto, S., 2015. Risk, return and multinational production. The Quarterly Journal Economics 130, 2027--2073.

Filtz, E., De La Cerda, E.S., Weber, M., Zirkovits, D., 2015. Factors affecting ocean-going cargo ship speed and arrival time, in: International Conference on Advanced Information Systems Engineering, Springer. pp. 305--316.

Fort, T.C., Klimek, S.D., 2018. The Effects of Industry Classification Changes on US Employment Composition U.S. Census Bureau Center for Economic Studies Working Paper 18-28.

Gervais, A., 2018. Uncertainty, risk aversion and international trade. Journal of International Economics 115, 145--158.

Gopinath, G., Neiman, B., 2014. Trade Adjustment and Productivity in Large Crises. The American Economic Review 104, 793--831. URL: http://www.ingentaconnect.com/content/aea/aer/2014/00000104/00000003/art00003.

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

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

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

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

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

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

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

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

Kamal, F., Monarch, R., 2018. Identifying foreign suppliers in u.s. import data. Review of International Economics 26, 117--139.

Kleindorfer, P.R., Saad, G.H., 2005. Managing disruption risks in supply chains. Production and operations management 14, 53--68.

Miyauchi, Y., 2018. Matching and agglomeration: Theory and evidence from japanese firm-to-firm trade. Technical Report. Working Paper.

Pierce, J.R., Schott, P.K., 2012. Concording u.s. harmonized system categories over time. Journal of Official Statistics 28, 53--68.

Ramondo, N., Rappoport, V., Ruhl, K.J., 2013. The proximity-concentration tradeoff under uncertainty. The Review of Economic Studies , rdt018.

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

Zsidisin, G.A., Ellram, L.M., Carter, J.R., Cavinato, J.L., 2004. An analysis of supply risk assessment techniques. International Journal of Physical Distribution & Logistics Management .

A Appendix: Data

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 (HS10) codes, and calculate prices as unit values by dividing the value of shipment by the quantity shipped. Fourth, we translate the nominal shipment values into real values using the U.S. GDP deflator.

Since the manufacturer ID (MID) differs across establishments of the same firm in different locations and since logistics are likely arranged at the firm-level, we replace the MID with a shortened identifier that contains only the country ISO code and the name portion of the ID, as described in the main text. External validation of the ability of the MIDs to identify foreign firms is provided in Kamal et al. (2015). They 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. Kamal et al. (2015) show that the number of MIDs in the Census data matches well with the number of sellers in the EDD when the street address or the city component are omitted. Kamal and Monarch (2018) provide further support that the MID is a good identifier of foreign exporters as follows. First, errors due to manual data entry are likely low because most firms use customs brokers for their official customs invoice and nearly all entries are filed electronically using specialized customs software. Second, the MID is used for regulatory purposes, such as enforcing anti-dumping measures or tracking compliance with U.S. restrictions for textile shipments, which provide an incentive for the U.S. government to ensure that the MIDs are correct. Third, as an external validation, Kamal and Monarch (2018) assess whether the MID can distinguish between distinct exporters using Chinese data: they construct artificial MIDs from exporter names and addresses in the Chinese Annual Survey of Industrial Firms, and show that they tend to be unique firm identifiers within sectors.

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

Table 10: Weather conditions: summary statistics

	Mean	Sd	p5	p50	p95
Significant wave height (m)	1.9	1.3	0.4	1.7	4.5
Significant wave direction (degrees)	187	90	44	197	318

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

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

Table 10 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 1.6 meters, but at the 95th percentile the wave height is 4.2 meters.

Table 11: Shipping Times by Mode of Transportation

	(1) Avg. Time	(2) Std. Time	(3) P5	(4) P25	(5) P50	(6) P75	(7) P95	(8) Total Value ($Bill.)
Vessel	16.38	23.54	3.49	10.00	13.46	20.48	33.32	4,250
Train	4.42	6.256	0.00	0.00	0.00	8.48	16.86	1,450
Truck	0.05	0.36	0.00	0.00	0.00	0.00	0.00	2,210
Airplane	0.49	0.89	0.00	0.00	0.00	1.00	2.298	1,610
Other	11.94	74.70	0.00	0.00	0.00	2.47	24.18	1,020

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.

Table 12: Shipping Times by Region

	(1) Avg. Time	(2) Std. Time	(3) 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.

B Appendix: Additional Results

Table 11 reports the percentiles of the distribution of shipping times for different modes of transportation. The table clearly shows how vessel-borne shipments are substantially slower and volatile than other types of transportation.

Table 12 shows average shipping times and their standard deviation for vessel-based shipments by region of origin. Shipments from Latin America and Canada tend to arrive fastest in the U.S., while shipments from Oceania and Africa take the longest. There is a large standard deviation of shipping times for all source countries.

Table 13 reports the results of a regression of residualized shipping times on weather conditions, identified with wave height and wind direction. The table shows a negative and significant effect of these variables on shipping times. This constitutes the first stage of our

Table 13: Effect of Weather on Shipping Times

Dep. Var:	$\hat{t}^{s}_{xhrtvfa}$
Wave Height^s	-0.0330^{***}
	(0.0006)
Direction^s	-0.0002^{***}
	(0.0000)
Wave Height$^s\times$ Direction^s	0.00003^{***}
	(0.0000)
Route FE	Y
Observations	5,774,000

Table 14: Effect of Extreme Delays on Firms' Outcomes, II

	(1)	(2)	(3)
		IV	
Dependent Variable:	Sales	Profits	Employees
Frac Delayed (Alt)	-4.024^{*}	-1.579	-0.396
	(2.062)	(1.359)	(0.431)
Frac Delayed (NoAlt)	-8.317^{**}	-5.755^{*}	-1.271^{**}
	(3.403)	(3.344)	(0.623)
Importer FE	Y	Y	Y
Year FE	Y	Y	Y
F-Stat (Alt)	8.359	8.359	8.359
F-Stat (No-Alt)	3.53	3.53	3.53
Observations	142,000	142,000	142,000

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 country level.

main regressions.

Table 14 regresses firm-level economic outcomes on the fraction of inputs that were delayed, but, differently from the baseline table, it differentiates between U.S. importers that were sourcing a given product from only one supplier and importers which instead had alternative back-up suppliers. We can see that the latter type of firms perform significantly better than the former, highlighting the benefits of diversifying inputs sourcing across multiple suppliers.

B.1 Robustness

In this section, we show that our results on the impact of shipping risk on importers' behavior are robust to a variety of alternative specifications. First, in Table 15, we run our baseline specification without firm fixed effects and find similar results as before. An increase in the standard deviation of shipping times has a positive and significant effect on both the number of suppliers and routes. Similarly, we find that the relationship between shipping risk and importers' concentration of suppliers remains significant, and the effect on value imported even strengthens.

Table 15: Shipping Time Risk and Import Demand with Different Fixed Effects

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.:	Number of Suppliers	Number of Routes	HHI over Suppliers	HHI over Supplier-Routes	Value per Supplier-Year	Value Imported per Year
Std Time	0.063^{***}	0.101^{***}	-0.038^{***}	-0.055^{***}	-0.188^{***}	-0.125^{***}
	(0.011)	(0.011)	(0.003)	(0.002)	(0.012)	(0.008)
Importer FE	–	–	–	–	–	–
Product FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Observations	328,000	328,000	328,000	328,000	328,000	328,000

Notes: Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the firm level.

Next, in Table 16 we do not include the unit value in the regression and find results consistent with our baseline.

Table 16: Shipping Time Risk and Import Demand Excluding Price Controls

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.:	Number of Suppliers	Number of Routes	HHI over Suppliers	HHI over Supplier-Routes	Value per Supplier-Year	Value Imported per Year
Std Time	0.055^{***}	0.097^{***}	-0.037^{***}	-0.054^{***}	-0.153^{***}	-0.098^{***}
	(0.004)	(0.006)	(0.002)	(0.001)	(0.006)	(0.007)
Importer FE	Y	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Controls	Ex. Price	Ex. Price	Ex. Price	Ex. Price	Ex. Price	Ex. Price
Observations	328,000	328,000	328,000	328,000	328,000	328,000

Notes: Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the firm level.

Table 17 includes firms with only one supplier. Here, we find that the relationship between shipping risk and the number of suppliers and routes remains significantly positive, but decreases slightly in magnitude. Compared to the baseline, the inclusion of firms with one supplier weakens the association between shipping risk and HHI over suppliers slightly. Value per supplier and total value imported remains significantly negatively correlated with shipping risk.

In Table 18 we include an additional control for the inventory-sales ratio and find that the relationship between shipping risk and the number of suppliers and routes strengthens.

Table 17: Shipping Time Risk and Import Demand Including Firms with One Supplier

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.:	Number of Suppliers	Number of Routes	HHI over Suppliers	HHI over Supplier-Routes	Value per Supplier-Year	Value Imported per Year
Std Time	0.047***	0.090***	−0.031***	−0.054***	−0.147***	−0.100***
	(0.004)	(0.006)	(0.002)	(0.002)	(0.006)	(0.007)
Importer FE	Y	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Observations	502,000	502,000	502,000	502,000	502,000	502,000

Notes: Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the firm level.

However, conditional on shipping risk, a higher inventory-sales ratio decreases the number of routes used. Supplier concentration and value imported remain significantly negatively correlated with shipping risk.

Table 18: Shipping Time Risk and Import Demand with Inventory-Sales Ratio

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.:	Number of Suppliers	Number of Routes	HHI over Suppliers	HHI over Supplier-Routes	Value per Supplier-Year	Value Imported per Year
Std Time	0.061***	0.114***	−0.039***	−0.059***	−0.157***	−0.095***
	(0.005)	(0.007)	(0.002)	(0.002)	(0.008)	(0.008)
Inventory-Sales Ratio	−0.018	−0.036**	0.000	0.005	−0.063	−0.082
	(0.013)	(0.015)	(0.003)	(0.003)	(0.048)	(0.060)
Importer FE	Y	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Observations	237,000	237,000	237,000	237,000	237,000	237,000

Notes: Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the firm level.

In Table 19, we also find that our results hold when we use the risk of the main supplier as our risk measure, rather than a weighted average across all suppliers. A one log point increase in our measure of shipping risk increases the number of suppliers used by 4 percent and the number of routes by 6 percent. In columns (3) and (4), we see that having one percent higher shipping risk is associated with a 4 percent lower HHI over suppliers and supplier-route combinations. Finally, columns (5) and (6) show that increasing shipping risk by one percent leads to lower imports.

Lastly, Table 20 documents that higher shipping risk is associated with a higher likelihood of using air shipments, suggesting that firms use air transportation to hedge the ocean shipping risk.

Table 19: Shipping Time Risk and Import Demand with Main Supplier

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.:	Number of Suppliers	Number of Routes	HHI over Suppliers	HHI over Supplier-Routes	Value per Supplier-Year	Value Imported per Year
Std Time	0.043***	0.059***	−0.039***	−0.039***	−0.153***	−0.111***
	(0.004)	(0.006)	(0.002)	(0.003)	(0.007)	(0.006)
Importer FE	–	–	–	–	–	–
Product FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Controls	Y	Y	Y	Y	Y	Y
Observations	237,000	237,000	237,000	237,000	237,000	237,000

Notes: Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the firm level.

Table 20: Shipping Time Risk and Import Demand with Air Shipments

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

Notes: Number of observations has been rounded to the nearest 1000 as per U.S. Census Bureau Disclosure Guidelines. Standard errors are clustered at the firm level.