

Dunbar, Geoffrey; Steingress, Walter; Tomlin, Ben

Working Paper

Climate variability and international trade

Bank of Canada Staff Working Paper, No. 2023-8

Provided in Cooperation with:

Bank of Canada, Ottawa

Suggested Citation: Dunbar, Geoffrey; Steingress, Walter; Tomlin, Ben (2023) : Climate variability and international trade, Bank of Canada Staff Working Paper, No. 2023-8, Bank of Canada, Ottawa, <https://doi.org/10.34989/swp-2023-8>

This Version is available at:

<https://hdl.handle.net/10419/297393>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

Climate Variability and International Trade

by Geoffrey Dunbar, Walter Steingress and Ben Tomlin

International Economic Analysis Department
Bank of Canada
gdunbar@bankofcanada.ca, wsteingress@bankofcanada.ca,
btomlin@bankofcanada.ca



Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

DOI: <https://doi.org/10.34989/swp-2023-8> | ISSN 1701-9397

©2023 Bank of Canada

Acknowledgements

The views presented in this paper are those of the authors and may not represent those of the Bank of Canada or its employees. We thank George Alessandria, Edouard Djeutem, Sharat Ganapati, Pierre Magontier, Nicholas Sander, Joe Steinberg and Dan Treffer as well as seminar participants at the University of Rochester, the 5th Joint Research Conference on Firm Organisation, Firm Financing and Firm Dynamics (FirmOrgDyn) at Sciences Po, the Urban Economic Association meetings, the Banque de France and the Bank of Canada for helpful comments. We thank Jamie Porthiyas for excellent research assistance, especially with the exact Earth AIS ship location data. All errors and omissions are our own.

Abstract

This paper quantifies the impact of hurricanes on seaborne international trade to the United States. Using geocoded hurricane data mapped to satellite tracking data for commercial ships, we identify hurricane intersections on sea-trade routes between U.S. and foreign ports. Matching the timing of hurricane-trade route intersections with monthly U.S. port-level trade data, we isolate the unanticipated effects of a hurricane hitting a trade route using two separate identification schemes: an event study and a local projection. Our estimates imply that a hurricane reduces route-specific monthly U.S. import flows by 5.4% to 16.0%, leading to an aggregate loss of 1.15% to 3.42% of annual U.S. west coast imports for an average storm season. We find no evidence of trade catching up in the months following a hurricane nor any evidence of rerouting to other ports or other transportation modes (e.g., air). Using our estimates in combination with climate scenarios from the Intergovernmental Panel on Climate Change, we quantify a range of costs of future hurricane disruptions that could occur if trade routes remain fixed.

Topics: Climate change; International topics

JEL codes: C22, C5, F14, F18, Q54

Résumé

Cette étude quantifie l'incidence des ouragans sur le commerce maritime international vers les États-Unis. En mettant en correspondance des données géocodées sur les ouragans et des données de suivi par satellite de navires commerciaux, nous repérons les intersections entre la trajectoire des ouragans et les routes maritimes commerciales qui relient les ports américains et étrangers. En comparant le moment où les ouragans croisent ces routes avec des données mensuelles sur le commerce des ports américains, nous isolons les effets imprévus d'un ouragan frappant une route commerciale à l'aide de deux méthodes d'identification distinctes : une étude de l'événement et une projection locale. Nos estimations impliquent qu'un ouragan réduit les flux mensuels d'importations américaines propres à une route de 5,4 à 16,0 %, ce qui représente une diminution globale de 1,15 à 3,42 % des importations annuelles de la côte ouest des États-Unis pour une saison des tempêtes moyenne. Nous ne trouvons aucune preuve de rattrapage du commerce dans les mois suivant un ouragan, ni aucune indication que les produits sont réacheminés vers d'autres ports ou par d'autres modes de transport (p. ex., l'avion). À l'aide de nos estimations combinées aux scénarios climatiques du Groupe d'experts intergouvernemental sur l'évolution du climat (GIEC), nous quantifions un éventail de coûts qui pourraient découler des ouragans futurs si les routes commerciales restent inchangées.

Sujets : Changements climatiques ; Questions internationales

Codes JEL : C22, C5, F14, F18, Q54

1 Introduction

The vast majority of international trade is conducted via the world’s oceans and seas. Although storms have long been a hazard for mariners, the growth in seaborne trade in recent decades has increased the exposure of global trade networks to extreme weather events. Extreme weather events are, themselves, a subject of increasing study because of the risk and uncertainty of climate change for extreme weather patterns. Indeed, it is widely accepted that extreme weather events can have significant impacts on economic activity. For major storms, this impact is most visible when they make landfall resulting in extensive physical damage and, in tragic circumstances, loss of life. However, many tropical cyclones never make landfall. While oceans are largely devoid of human settlement, they are not devoid of human activities—chief among them, international trade.

In this paper, we identify a new channel—seaborne trade—through which extreme weather and climate change affect economic activity. To conduct our analysis, we construct a novel dataset of trade routes, trade flows and hurricanes by combining geocoded port-to-port routes, geocoded hurricane data, and port-level trade data for US west coast ports.¹ We use this dataset and a detailed identification strategy to, first, derive an estimate for the cost of unanticipated hurricanes for route-specific trade, and second, to explore the future costs associated with projected changes in hurricane activity due to climate change. We find that hurricanes that intersect port-to-port trade routes reduce trade on those routes by 5.4% to 16.0% in the month following a hurricane, and that these losses are not offset in subsequent months, which suggests that these losses are permanent at the route level, and when aggregated across routes and months of the year, represent persistent supply shocks. Next, we use climate change scenarios from the Intergovernmental Panel on Climate Change (IPCC) (Pörtner, Roberts, Masson-Delmotte, Zhai, Tignor, Poloczanska, and Weyer, 2019) and our estimates of the trade losses from hurricanes to simulate possible future costs of extreme weather for seaborne trade. Our simulations suggest that increases in the frequency, intensity and poleward migration of hurricanes could increase annual trade losses due to hurricanes by 60%.

The range in our route-level estimates reflects two different identification strategies: an event study and a local projection. Because hurricanes are staggered and repeat occurrences on the trade routes we observe, identifying the effect of a hurricane-induced trade disruption is challenging. While the trade data we use covers the period 2003-2019, hurricanes have been naturally occurring events since before this period. As a result, determining ‘clean controls’ for an event study design is complicated and, for example, requires making explicit assumptions about the duration of hurricane events.² An additional challenge

¹Throughout this paper we use the terms hurricane, tropical cyclone and typhoon interchangeably.

²We also do not observe any trade on routes from a period prior to the existence of hurricanes, which would have permitted identification of the effects of the first hurricanes on trade routes (de Chaisemartin and D’Haultfoeuille, 2022). However, identifying such effects is not particularly salient because, as our

is that, for many countries and ports in our sample, we observe only a limited number of routes, which complicates identification of the counterfactual trade. Local projections can potentially avoid some of the difficulties that can arise from ‘unclean controls’ and estimate counterfactual trade flows using the time series dynamics of trade, but assume that the underlying data generating process follows a vector autoregression (VAR).³

The 5.4% loss in trade is estimated by the event-study design using the Poisson Pseudo-Maximum Likelihood (PPML) estimation, which includes a number of fixed effects and controls. Although we assume actual hurricanes are exogenous events, the possibility that hurricanes can occur may be anticipated because the climatic conditions that generate hurricanes are typically periodic events. To account for expected hurricane activity on trade routes, we include route-specific month dummies in our regressions (these dummies also control for other predictable monthly variation). In addition, we control for the exporting country’s supply conditions and account for heterogeneity across the type of ship (dry bulk, container, gas and liquid chemical, oil and roll-on roll-off (RoRo)) and the potential substitution across other transportation modes (e.g., air). The 5.4% estimate is robust across several event-study specifications that we consider. However, the event-study estimate is also likely to be a lower bound, because some of the variation due to unanticipated hurricanes is likely captured by the fixed effects included in our regression specification.

The upper range of our trade-loss estimates, 16%, is based on the local projection econometric design, which assumes an auto-regressive process for trade. The trade values in our data do exhibit auto-correlation, and the local projection approach leverages this feature of the data by controlling for pre-treatment values of trade and other covariates when estimating the dynamic effects of the hurricane shock. While we include route-specific month fixed effects to control for hurricane expectations, the local projection specification does not control for other contemporaneous factors that may influence trade—for example, an unobserved contemporaneous demand or supply shock that affects trade in the month the hurricane hit trade routes. However, the effects of such coincidental events for the local projection estimates may not be too severe, given that our relatively large sample size should average out such random correlations if hurricanes are exogenous events.

We note that both econometric designs yield qualitatively identical results: trade on a route falls in the month following a hurricane crossing, and the loss is not recouped in subsequent months. Our estimates imply that hurricanes lead to a permanent loss of trade on affected routes. While this is a statement about our empirical findings, it may suggest that shipping networks do not have much spare capacity.⁴ Hurricanes at sea are

identification discussion below clarifies, the object of interest is the effect of hurricanes conditional on expectations of their occurrence.

³Local projections leverage lagged values of the dependent variable to project counter-factual trade; see Montiel Olea and Plagborg-Møller (2021) and Dube, Girardi, Jordà, and Taylor (2022).

⁴In calendar day terms, our estimates suggest that trade is delayed between roughly 2 and 5 days because

also not localized events and generally affect multiple routes. Which routes are affected by a hurricane is, however, idiosyncratic to each hurricane, which implies that the aggregate trade losses from hurricanes are time-varying. Because hurricane seasons are typically concentrated in some months of the year, these aggregate trade losses exhibit persistence.⁵

To quantify the expected magnitude of the hurricane costs for aggregate US trade with the countries in our sample, we use the probability of a hurricane intersecting a route in a given calendar month weighted by the value of trade for that route. Taking the average over our 2003-2019 sample period, we find that the average hurricane season results in an annual loss of 1.15% to 3.42% of US west coast imports (using the route-specific loss estimates of 5.4% and 16%, respectively). Simulations also allow us to differentiate “good” and “bad” hurricane seasons, with a “bad” season—identified as the 90th percentile of our simulations—resulting in a 4.43% loss of imports. In terms of economic magnitude, our estimated losses of US imports lie between \$5.1 billion and \$22.5 billion per year and are comparable to the estimated physical damage of hurricanes hitting land of \$8.1 billion per year (Deryugina, 2017). We note that these costs may underestimate the total economic cost of a hurricane at sea, since the majority of goods traded by sea are intermediate goods and disruptions may affect supply chains and amplify the aggregate costs (see, for example, Barattieri and Cacciatore (2020)).

We next consider how the magnitude of these trade costs may evolve because of climate change. The IPCC has examined the scientific evidence of how changes in climatic conditions are expected to affect storm activity and emphasized three dimensions of storms: the frequency, intensity, and location (“poleward migration”) (see Pörtner, Roberts, Masson-Delmotte, Zhai, Tignor, Poloczanska, and Weyer (2019) and Pörtner, Roberts, Poloczanska, Mintenbeck, Tignor, Alegría, Craig, Langsdorf, Löschke, Möller, and Okem (2022)). However, there remains much uncertainty regarding the magnitude and even the direction of the adjustments along these dimensions, and for this reason we remain agnostic about the effects of climate change. That is, we draw on ranges of estimates from the literature on how climate change will affect storm activity and, in turn, generate cost estimates associated with these ranges. The route-specific probability of getting hit by a hurricane is adjusted to reflect changes along the three dimensions individually and their combined effects, and we present a range of possible outcomes. For example, we find that for a 1° latitude shift in poleward migration, estimates from the literature associated with different adjustments in the frequency and intensity of major storms suggest a range of hurricane-induced trade costs of 2.22% to 5.40% (or even higher), compared to the baseline

of hurricanes. This amount of delay may be too short for shippers to profitably reroute shipments (and indeed our robustness exercises suggest no such rerouting occurs). One simple example is to consider a daily passenger ferry service between two ports. If sailings are generally full, then the delay in one sailing implies a shortfall in passenger trips, assuming a fixed schedule.

⁵Our estimates do not indicate who bears this loss (e.g., insurers, wholesale firms, retailers or consumers). While an interesting question, it is beyond the scope of this paper.

estimate of 3.42%.

We contribute to the literature on the economic assessment of climate variability for the US. Mendelsohn, Nordhaus, and Shaw (1994), Schlenker, Hanemann, and Fisher (2005), Deschênes and Greenstone (2007) and Roberts and Schlenker (2010) study the effect of temperature and precipitation variability for agricultural output in the US. Dingel, Meng, and Hsiang (2019) note that both climate shocks and trade are geographically concentrated—the latter because of the gravity effect. This spatial correlation implies that climate change may increase global inequality. A similar perspective is echoed by Cruz and Rossi-Hansberg (2021), who find that spatial correlation in climate implies that international trade is unlikely to offset the costs of climate change. Our paper contributes to this literature by highlighting that trade itself is directly affected by climate variability and providing cost estimates for the effects of hurricanes on US international trade. Belasen and Polachek (2008) estimate hurricane-induced employment effects in Florida; Strobl (2011) quantifies the impact on economic growth and shows that a county’s annual economic growth rate falls significantly after a hurricane. Deryugina (2017) studies the fiscal costs of hurricanes and shows that public transfers provide considerable insurance, and that household debt even decreased in the aftermath of Hurricane Katrina (Gallagher and Hartley, 2017). A second contribution of our paper to this literature is that we show that climate variability is costly because of the trade disruption it causes, and our estimates of these costs are potentially useful in framing the anticipated costs of future climate scenarios. We make the caveat, however, that our estimates of the potential costs of future climate scenarios hold the existing structure of trade fixed and do not model potential adaptations that may occur. Thus, we view our climate scenario simulations as indicative of the opportunity costs of non-adaptation.

Another strand of literature we contribute to is the international trade literature that focuses on seaborne trade. Alessandria, Yar Khan, Khederlarian, Mix, and Ruhl (2022) show that transitory increases in shipping times (measured by the ISM Manufacturing Supplier Deliveries Index) reduce imports at impact but lead to a significant increase in the following months before reverting back to steady state. In terms of hurricane disruptions, our estimates show a negative effect on trade in the month following the hurricane shock with no subsequent significant effects up to 5 months after the shock. This effect suggests that inventory management surrounding hurricane seasons is a normal part of seaborne trade and is accounted for by our route-month fixed effects. Brancaccio, Kalouptsi, and Papageorgiou (2020) endogenize trade costs by modelling the market for seaborne transportation services and study its implications for world trade. Related to our study, this endogenous response from the transportation sector suggests that our estimates are likely to be a lower bound, as hurricanes not only reduce the value of US imports but also cause shipping delays.

Other papers, such as Heiland, Moxnes, Ulltveit-Moe, and Zi (2019) and Ganapati, Wong, and Ziv (2021), model the global shipping trade network and quantify its trade

and welfare impact. Heiland, Moxnes, Ulltveit-Moe, and Zi (2019) create a measure of optimal shipping routes by calculating the minimum travel time between different ports and quantify the welfare effect of the Panama expansion in 2016 through the reduction in travel time. In contrast, our study exploits the full geographic information of each ship’s trip and calculates optimal routes based on actual paths taken by ships. Somewhat interestingly, we find that different types of ships appear to take different routes between identical port pairs. This variation allows us to identify and estimate the effect of route-specific hurricane disruptions and provide route-specific cost estimates of climate variability. Ganapati, Wong, and Ziv (2021) document that the trade network is a hub-and-spoke system, where most trade is shipped through a few major entrepôts-hubs. Consistent with this network structure and the limited scope for substitution across ship routes, we find little evidence of ships switching destination ports to evade a hurricane. This network structure, and the fact that ports operate on fixed schedules (Brancaccio, Kalouptsidi, and Papageorgiou, 2020) with limited capacity to accommodate delayed ships, may explain our estimated permanent loss of trade after a hurricane disruption.

The rest of the paper is organized as follows. Section 2 describes the data sources and identification of ship routes and hurricane intersections. Section 3 presents the empirical framework and outlines the identification assumptions. Section 4 presents the empirical results, and Section 5 combines the route-level effects and the sample hurricane history to compute the aggregate trade costs. Section 6 describes our counterfactual climate change simulations, and the last section concludes.

2 Data

Our empirical approach estimates trade disruptions caused by hurricanes intersecting shipping routes, forcing exporters to delay or reroute their shipments. For this we require data on the trade flows associated with particular shipping routes. This data requirement can be hard to satisfy if, for example, a landlocked country that trades with the US has more than one exporting port which it could use. In this situation, we cannot confidently attribute a country’s exports to a particular port. This concern is particularly acute for transatlantic trade involving landlocked European countries. To alleviate this concern, we focus on imports to US west coast ports, which typically feature direct trade routes in the Pacific from Asian, Oceanian, Middle Eastern and South American countries.⁶

The hurricane data in our study come from the National Oceanic and Atmospheric Administration (NOAA) National Centers for Environmental Information, which maintains the International Best Track Archive for Climate Stewardship (IBTrACS)—see Knapp,

⁶Of these countries, Paraguay and Bolivia are the largest landlocked countries that export to the US; however, the US is not a main export destination for either country.

Kruk, Levinson, Diamond, and Neumann (2010).⁷ IBTrACS sources data from meteorological services worldwide and has data on tropical cyclones for the Pacific Ocean. IBTrACS provides at least daily data on the geocoded storm position (measured at the eye of the storm), the wind speed over 1-minute intervals, and the direction and speed over ground of the eye of the storm. We use the maximum recorded wind speed at every location to identify hurricane-force storms in the data.⁸

To analyze the effects of hurricanes on sea trade, we require data on the sea routes between US ports and their trading partners. For this, we use data from exactEarth for the year 2016. exactEarth provides location-based maritime vessel tracking information for ships (roughly a data point every 3 hours), using satellite automatic identification system (AIS) technology. In these data we are able to track AIS-enabled commercial ships across the globe—a total of 6,550 ships. Pertinently for our study, the data include ship positions, direction, date and 6 ship types: container, gas tanker, oil and chemical tanker, bulk carrier, RoRo and general cargo.

To identify ship routes from exporting ports to US west coast ports, we require the origin and destination ports. The AIS data do not always contain port information or port stops (ships do not always activate their AIS while in port, and ships do not always report being in port). Therefore, we define the origin port as when a ship leaves a port zone, defined as a 20km radius around a port geolocation (using the Haversine distance measure), after staying at least 24 hours in this zone.⁹ Similarly, we define a destination port when a ship enters a port zone and stays there for at least 24 hours. With this, we are able to construct a full set of geolocated trade routes for 5 major US west coast ports: San Diego, Los Angeles (LA), San Francisco/Oakland, Portland and Seattle.

Figure 1 plots the observed ship routes into the port of LA, which is the largest US port in terms of trade volume. There are two key observations. First, different types of commercial ships take different routes between the same port pairs. As an example, RoRo vehicle carriers appear to take a more direct route from Asia to LA than bulk carriers or container ships. This observation implies that ship routes are ship-type specific. Second, there is variation in routes taken for ships of the same type, which likely reflects navigation decisions taken as a result of sea-state conditions and forecasts. This second observation implies that variation in observed routes by ship type may be confounded by hurricane

⁷The data was accessed as Knapp, Diamond, Kossin, Kruk, and Schreck (2018) on 21-07-2020.

⁸We use the Saffir-Simpson index to identify hurricanes, which are storms with wind speeds greater than 64 knots per hour (knph). The Saffir-Simpson index can also be used to categorize 5 types of hurricane severity (with 5 being the strongest hurricane and 1 being the weakest hurricane). For a storm to be categorized as a category 1 hurricane, it must have a wind speed of 64 knots per hour (knph) or greater; the wind speed of category 2 hurricanes is greater than or equal to 83 knph, that of category 3 greater than or equal to 96 knph, that of category 4 greater than or equal to 113 knph, and that of category 5 greater than or equal to 137 knph.

⁹The latitude and longitude of ports comes from the World Port Index (Pub 150) of the National Geospatial-Intelligence Agency: <https://msi.nga.mil/Publications/WPI>. It contains the location and physical characteristics of major ports and terminals worldwide (approximately 3700 entries).

activity. To map hurricane paths to the route that a ship would have taken in the absence of a hurricane, we need to select from the observed routes a route that was unaffected by hurricane activity—hereafter, we refer to this as the optimal route.

We choose to measure optimal trade routes for each port-pair using one observed route in the distribution of routes observed. Because ships sometimes avoid storms by increasing speed and altering their course, using the fastest observed route risks selecting a route chosen specifically to avoid a hurricane as the optimal route. Similarly, choosing a slower than average route risks selecting a route that represents a seasonal shift or one chosen to avoid inclement weather (hurricane or not). It is probable that a route between these points of the distribution is the route a ship would choose in ideal, calm weather, conditions. We calculate shipping times for each port-to-port route for each type of commercial ship and define the optimal shipping route as the 95th percentile fastest route (or closest to but not greater than 95th for less popular routes). Figure 2 shows these optimal ship routes from exporting ports to the port of LA.

Using the coordinates of the optimal port-port trade routes, we determine if a hurricane eye intersects ship routes using the Bentley and Ottmann (1979) algorithm.¹⁰ Holding the optimal routes invariant over time, this algorithm allows us to identify the date, location and wind speed every time the eye of a storm intersects one of the optimal routes. These intersections provide a time series of hurricane activity for each optimal route. Figure 3 illustrates one hurricane, Typhoon Lionrock, that intersected a number of routes near Japan in 2016. As the figure shows, hurricanes at sea are not localized events and can potentially affect many, often spatially correlated, shipping routes.¹¹

To estimate the impact on trade flows of a hurricane intersection with a trade route, we require data on the value of imports on these routes. To calculate the value of imports by trade route, we use monthly import data, by 6-digit Harmonized System (HS) product classification and country of origin, for US ports provided by the US Census Bureau. To assign product-level trade to ship type, we use the Maritime Transport Costs (MTC) database from Korinek (2011) and code some product categories omitted from this database ourselves (for example, the MTC database does not separately identify RoRo vehicle trade).¹² We drop general-cargo ship routes from our data because these ships

¹⁰While there are other plausible definitions of a hurricane affecting a route, choosing the intersection captures the advice typically given to mariners: never cross the T(rack) of a storm.

¹¹Hurricanes that make landfall are generally more localized weather events because storms lose energy over land. We do not consider the impact of hurricanes that make landfall because our focus is to estimate the cost of trade disruption and not the impact of physical damage as does, for example, Strobl (2011) or Deryugina (2017).

¹²The US Census also records container shipments as a separate category of imports, which we link directly to container ship routes. Similarly, we map oil to oil tankers and gas and chemical liquids to gas and chemical liquid tankers. We also assign assembled passenger vehicles and other road vehicles to RoRo ships as these are the least costly method of transport for these products. Finally, we assign non-containerized grains and cereals, coal, processed wood products, potash and ores to bulk cargo carriers. We are unable to assign some trade products, such as non-containerized furniture, because

often fulfill multiple roles, such as bulk and oil or bulk and container, which we cannot differentiate. On average, our data accounts for roughly 97% of trade by value per period for the US ports. Finally, because the US trade data is not disaggregated by exporting port, we aggregate import values by country to obtain monthly bilateral trade values (exporting country to US port by ship type) for the period 2003 to 2019.¹³ We construct a hurricane dummy, $h_{i,j,s,t} = 1$, if any route from country j to US port i by type of ship s in period t is intersected by a hurricane (equal to zero otherwise).

Table 1 shows some sample statistics for our data. There are 180 trade routes across 39 exporting countries to the 5 US west coast ports. On average there are 9.39 routes per country, and the average number of ship types used per country to export is 3.32. The country with the highest number of trade routes is Japan with 18 (the highest possible number of trade routes is 25, calculated by multiplying 5 ship types by 5 US ports). The average frequency of a hurricane intersection is 0.098, though there are a few routes for which the probability of a hurricane in a given month is 100%. The average value of imports per US port and route varies by port; LA has the highest average route import value of \$524 million. The smallest US port by average route value is Portland, with an average import value of \$42.5 million. There is also substantial variation in the route import values, as shown by the minimum, maximum and standard deviation (SD).

Figure 4 shows the seasonal patterns of hurricane activity; September is the most commonly affected month. There is also significant variation across countries in the proportion of trade affected by a hurricane, largely due to the latitude of a country and its ports (see Figure 5). Finally, there appears to be correlation between import demand and hurricane activity. Figure 6 plots the monthly differences from the annual average level of imports. Imports are highest during the 4 months prior to Christmas, which also corresponds to higher than average hurricane activity.

3 Estimation and identification challenges

In this section, we start by describing what we want to estimate and then detail our identification strategy. As mentioned in the introduction, we wish to estimate the cost associated with the actual occurrence of a hurricane, separate from the expected costs of typical hurricane activity anticipated by exporters and importers. Conceptually, we assume that observed trade flows, $T_{i,j,s,t}$, from exporter j to US port i by ship type s at time t are a function of desired flows, $\tilde{T}_{i,j,s,t}$, and random shocks, $v_{i,j,s,t}$, such that $T_{i,j,s,t} = \tilde{T}_{i,j,s,t} e^{v_{i,j,s,t}}$. We assume that desired trade flows are a function of demand, $D_{i,s,t}$,

it appears these could be assigned to either general or bulk cargo ships, and we cannot differentiate between them.

¹³Public port-specific monthly trade data does not exist before 2003 and, because of the COVID pandemic, we take December 2019 as the end period of our sample.

supply, $S_{j,s,t}$, and predictable elements of route-specific trade costs, $\Theta_{i,j,s}$, which we assume includes route-level fixed effects (these account for climate-induced hurricane expectations, which are specific to individual routes).¹⁴ We also assume that the random shocks include a specific component due to unanticipated hurricanes, $h_{i,j,s,t}$: $v_{i,j,s,t} = \beta h_{i,j,s,t} + \epsilon_{i,j,s,t}$. We then model observed trade flows as:

$$T_{i,j,s,t} = \mathbf{f}(D_{i,s,t}, S_{j,s,t}, \Theta_{i,j,s}) e^{\beta h_{i,j,s,t} + \epsilon_{i,j,s,t}}$$

where $\mathbf{f}(D_{i,s,t}, S_{j,s,t}, \Theta_{i,j,s})$ is a function representing desired trade flows. Apart from assuming that this function exists, we are agnostic about the specific model of trade that it represents. We do, however, require that this function is uncorrelated with actual hurricane events, conditional on $D_{i,s,t}$, $S_{j,s,t}$ and $\Theta_{i,j,s}$. Our focus in this paper is to identify and estimate β , which is the cost of unanticipated hurricanes for trade.


3.1 Identifying hurricane disruptions

Figure 7 shows our identification scheme diagrammatically using a hypothetical example of optimal routes (the black lines) between Shanghai and three US west coast ports: LA, Oakland and Seattle. The dashed red line represents a hurricane that forms and whose eye crosses the optimal route between Shanghai and LA but does not cross either of the remaining routes. If the hurricane crossing disrupts shipping on the optimal Shanghai-LA route, exporters can either stay on the optimal route and delay shipment until the hurricane has passed, or they can re-route the shipment, for example by using the route represented by the blue dotted line. Since the blue route is longer in distance terms than the optimal route, it also entails a delay in imports from Shanghai to LA (holding ship speed constant). In both cases we expect that the hurricane-induced delay lowers imports relative to periods without hurricanes. Hurricanes also have seasonal patterns, and some routes have experienced hurricane intersections in a particular month in every year of our sample. In such instances, the blue dotted line may be the optimal route for that month. To account for such seasonal patterns, we include route-month fixed effects in $\Theta_{i,j,s}$, which adjust for seasonally-varying optimal routes in our estimation approach. Finally, we note that trade diversion may complicate identification if, for example, exporters reroute shipments from LA to Oakland or Seattle in response to a hurricane on the route to LA. We examine this issue empirically below.

As the discussion above highlights, our empirical approach requires an estimate of the counterfactual trade, $\mathbf{f}(D_{i,s,t}, S_{j,s,t}, \Theta_{i,j,s})$, which would have occurred in the absence of a hurricane in order to estimate the effect of hurricanes on trade. Because of the spatial and time series nature of our data, there are two plausible estimation strategies that we

¹⁴Note that $\Theta_{i,j,s}$ has no t subscript, which implies that the data generating process for climate is held fixed over the 17-year sample period.

Figure 7: Identification Schematic



Note: This figure presents three stylized ship routes (black lines) from the port of Shanghai to the ports of Seattle, Oakland and LA. The red dashed line represents a hypothetical hurricane, and the blue dashed line represents an alternative route taken by a ship going from Shanghai to LA to avoid the hurricane.

follow. In the first approach, an event study, we treat the occurrence of a hurricane as a random event that potentially affects trade and use information from the cross-section of routes and ports to estimate $D_{i,s,t}$, and also partly $S_{j,s,t}$. This approach assumes that trade on unaffected routes and routes intersected by hurricanes have parallel trends and that the actual occurrence of a hurricane on a route is an unanticipated event. The second assumption is consistent with meteorological forecasting, which has a typical horizon of 5 days for hurricanes.¹⁵ However, trade data can be sparse and also seasonal, which may imply that the parallel trends assumption is problematic or, at the very least, that these trends are noisy in our data. In our second approach, a local projection, we use time-series information on route-specific trade to estimate $D_{i,s,t}$ and $S_{j,s,t}$. This approach assumes that counterfactual trade values can be estimated using lagged values of trade on routes, i.e., that the trade data follow a vector autoregression (VAR) structure. This VAR assumption may over- or under-estimate trade if there are other shocks, coincidental to hurricanes, that affect trade along routes (such as a tariff shock).

The monthly port, country and ship-type trade data have several features that complicate our empirical analysis irrespective of the estimation method. First, a large fraction—almost 45%—of the monthly trade values in the data are zero, and these zero values may be either orthogonal to hurricane activity (what we refer to as structural zeros) or may arise because hurricanes disrupt trade on that route. For instance, a simple linear probability regression model suggests that hurricanes increase the probability of zero trade on a route by roughly 1.2%.¹⁶ Logarithmic transformations are infeasible for these observations, and regression specifications using logarithmic transformations would likely be biased by selection effects. Fortunately, this issue is well-studied in the empirical trade literature and,

¹⁵See, for instance, <https://www.nhc.noaa.gov/gtwo.php?basin=epac&fdays=5>.

¹⁶The estimates for the linear probability model are available upon request.

for the event study, we follow Silva and Tenreyro (2006) and use the PPML estimation.¹⁷ For the local projection, PPML is not feasible, so we follow the empirical labour literature and retain zero observations using an inverse hyperbolic sine; see Bellemare and Wichman (2020) and references therein. A second complicating feature of the data is that monthly trade is autocorrelated, which suggests that the residuals in a regression specification may be autocorrelated. Regression residuals also may be correlated by ship type because the availability of ships can affect trade flows. For example, demand for bulk carriers by one country can reduce the supply of such carriers for other exporters, Brancaccio, Kalouptsi, and Papageorgiou (2020) and Jacks and Stuermer (2021). Third, in many cases hurricanes repeatedly cross the trade routes we consider and also cross other routes to the same US ports in the neighbouring months, which may have spillover effects if, for example, port congestion affects the timing of trade. We describe our empirical strategies to mitigate these concerns below.

3.2 Event study design

The event study models each hurricane as an event that occurs on a particular trade route. We estimate the effects of the hurricane event on trade preceding, during, and following the hurricane. Identification of the hurricane shock comes from a comparison between hurricane-affected (e.g., Shanghai-LA in Figure 7) and non-affected routes (e.g., Shanghai-Seattle and Shanghai-Oakland). One challenge for the event study approach is that there are many hurricane events in our data that occur at different times for different routes. Depending on the timing of how (or if) hurricanes affect trade, we must be careful that the event study compares treated routes (affected by a hurricane) with untreated routes (routes not affected by the hurricane). In addition, if the effects of hurricanes are not purely contemporaneous, and hurricanes occur repeatedly within a few months, the dynamic effect of the previous hurricane may contaminate the control group of the current hurricane shock.

We define a hurricane event so the event-study design yields a consistent estimate of the effect of the hurricane event. Of the 31,323 route and date observations in our estimating sample, 3,062 have a hurricane intersection.¹⁸ Many of these hurricane events occur in adjacent time periods and across different routes, which implies that isolating a hurricane event poses a challenge. For example, the effect of a hurricane on one route

¹⁷An alternative approach would be to follow the two-stage selection model of Helpman, Melitz, and Rubinstein (2008). However, hurricanes affect both the selection equation (e.g. the extensive margin—whether to trade or not) as well as the return equation (e.g., the intensive margin—the value of trade conditional on trading) and thus this approach is not feasible. In the robustness section we address the presence of zeros by considering sub-samples of more frequent routes that feature fewer zeros.

¹⁸Our estimating sample includes only routes for which there has ever been a positive value of trade. We drop all routes that never recorded a value for trade, even if we observed a ship in transit on that route in our 2016 exactEarth data. This could happen if a ship travelled empty or if trade on that route was indirect.

in a given period with no other hurricanes on any other route in the adjacent period may be different from a similar hurricane that occurs during a period in which hurricanes also affect other routes in adjacent periods. For our event study design, we will define a hurricane event as a hurricane occurring on a route in a period with no other hurricanes on any route in adjacent periods. We will be specific about how we define adjacent periods below.

We label variables in our data by the quadruple $\{i, j, s, t\}$, where i denotes the US port destination, j denotes the exporter country, s denotes the type of ship (bulk carrier, container, gas, oil and RoRo) and t denotes the date. We sometimes refer to the tuple $\{i, j, s\}$ as a route by which we mean a US port and exporter country pair for a specific type of ship (recall that we find that different ships take different routes between the same port pair). Recall $h_{i,j,s,t} = 1$ is an indicator variable for a hurricane that crosses a route at time t and which potentially affects the value of imported sea trade $Y_{i,j,s,t}$ for that route. The baseline event-study specification is:

$$Y_{i,j,s,t} = \alpha_{i,j,s,m} + \gamma_{i,s,t} + \delta X_{i,j,s,t} + \sum_{p=-5}^5 \beta_p h_{i,j,s,t+p} + \sum_{p=-5}^5 \zeta_{i,j,s,t+p} + \epsilon_{i,j,s,t}, \quad (1)$$

in which $\alpha_{i,j,s,m}$ are route-month fixed effects to control for, among other seasonal factors, the expectation of a hurricane along that route in month m ; $\gamma_{i,s,t}$ are US port fixed effects interacted with type of ship and date to control for local, time-varying factors such as bottlenecks, labour disruptions, and local demand; $X_{i,j,s,t}$ are control variables such as total exports, exchange rates and exports to port i by air with effects parameterized by δ ; and $\epsilon_{i,j,s,t}$ are residuals that reflect variation in trade which are assumed orthogonal to $h_{i,j,s,t}$. The coefficients of interest are β_p which measure the percentage change in import flows due to a hurricane crossing the trade route p periods from t (β_0 is the contemporaneous effect of a hurricane crossing).¹⁹

To isolate a hurricane event we include an additional set of dummy variables, $\zeta_{i,j,s,t \pm \tau}$, in Equation (1) if that route has a hurricane $t \pm \tau$ from the date of the observed hurricane event $h_{i,j,s,t} = 1$. We consider $\tau = 5$ to follow the same lead and lag structure we use for β_p .²⁰ The inclusion of the ζ dummy variables controls for routes that have hurricanes in temporally adjacent periods, which could confound our event study design. By including these dummy variables, the estimated event study effects isolate hurricane events for which the counterfactual is not directly or indirectly affected by hurricanes in adjacent periods.

¹⁹In Equation 1 we base our estimates relative to trade in the month prior to the hurricane crossing by omitting the dummy $h_{i,j,s,t+1}$. In an alternative specification in the Appendix, we include the dummy $h_{i,j,s,t+1}$ and base our estimate relative to trade as predicted by the control variables. The results are shown in Figures 22.

²⁰We experiment with different horizons other than 5 but do not find our results are particularly sensitive in this neighbourhood.

While this approach is appropriate for our event-study design, it also suggests that our estimates may be biased away from the true hurricane effect by focusing only on a subset of hurricane events in the data. In particular, this event-study design is more likely to focus on hurricanes early or late in the season which for our data is the winter period.

This specification resembles a gravity framework for trade and naturally provokes the question of whether to include exporting country fixed effects interacted with type of ship and time (e.g. $\vartheta_{j,s,t}$ fixed effects) rather than including control variables $X_{i,j,s,t}$. Because our data spans only 5 US west coast ports, there are at most 5 data points per exporting country and type of ship in each period t and in many cases fewer than 5 because not all US ports process all types of ships. Thus, these fixed effects would be imprecisely estimated at the best of times. However, for routes affected by a hurricane the problem is even more acute since both β_p and $\vartheta_{j,s,t}$ would be estimated from at most 5 data points, and in some cases fewer. Given the relative geographic concentration of US ports, the routes for specific ship types are likely to be correlated, thus reducing the set of observations available to estimate the fixed effect.²¹ In some cases, the hurricane effect may not even be separately identifiable from the $\vartheta_{j,s,t}$. One option might be to collapse an exporting country fixed effect across type of ship. However, different ship types specialize in different traded goods, so it is not plausible to collapse the exporter-country time fixed effect across ship types because there is no reason that the demand for, or supply of, cars is, for example, related to that of rice. Thus we include $X_{i,j,s,t}$ to control for time-varying exporting country factors that may be correlated with hurricanes beyond $\alpha_{i,j,s,m}$.

Given the discussion in the previous paragraph, one may be concerned about identification of $\alpha_{i,j,s,m}$ and $\gamma_{i,s,t}$. There are 38 countries that trade with the US ports in our data and that potentially identify the $\gamma_{i,s,t}$, but the exact subset of exporting countries varies by type of ship. Across types of ships, the range of observations per period t identifying the $\gamma_{i,s,t}$ are:

Type	Minimum	Maximum
Bulk carrier	16	31
Container	28	32
Gas tanker	0	1
Oil tanker	10	28
RoRo	8	19

which also highlights that, for gas tankers, the event study will not be able to separately identify the effects of hurricanes from local demand. However, in general, there appears to be sufficient observations to identify $\gamma_{i,s,t}$ and β_p . Turning to $\alpha_{i,j,s,m}$, our data span

²¹For instance, if a hurricane crosses a container ship route from South Korea to Los Angeles, it is also likely to cross a route to San Diego and San Francisco. This situation would leave at best only 2 observations to estimate a South Korea-container-time fixed effect.

the years 2003 to 2019, which implies that we have seventeen observations per route and month for identification. Not accounting for the seasonal pattern may underestimate the effect of hurricanes because of risk-mitigating strategies that shippers may employ to avoid hurricanes. We control for expectations by including port, country, ship-type and month specific fixed effects. These fixed effects are sufficient to control for hurricane expectations if the data generating process underlying hurricane arrivals on routes is time-invariant. Since we focus on only seventeen years of data, this assumption seems reasonable even if climate change modelling suggests that the data generating process may change in future years.

3.3 Local projection design

The second estimation strategy we consider is a Jordà (2005) local projection. The local projection approach requires that the data-generating process for trade is autoregressive, such that past observations of trade are informative of future trade. As emphasized in Montiel Olea and Plagborg-Møller (2021), an advantage of local projections is that they can be augmented to include lags of the dependent variable to alleviate the bias and inconsistency in the standard errors caused by autocorrelation in the residuals. In contrast to the event study, this characteristic allows local projections to account for route-specific dynamic adjustments of trade in the control group, implying that identification relies to a lesser extent on the contemporaneous comparison of hurricane-affected and unaffected routes. The local projection assumes that the hurricanes are unanticipated shocks and then, using different horizons of the dependent variable, estimates the effect of the hurricane on trade values, which traces out the impulse response from a hurricane on trade.

The baseline estimating equation for the local projection we consider is:

$$y_{i,j,s,t+q} = \alpha_{i,j,s,m} + \sum_{p=1}^P \rho_p y_{i,j,s,t-p} + \sum_{p=1}^P G_{i,j,s,t-p} \mu_p + \beta_q h_{i,j,s,t} + v_{i,j,s,t+q} \quad (2)$$

where $y_{i,j,s,t+q}$ is the inverse hyperbolic sine of $Y_{i,j,s,t}$, $q = \{0, 1, 2, \dots, Q\}$ periods in the future; $\alpha_{i,j,s,m}$ are the route-month fixed effects which are included for the same reasons as in the event study; $y_{i,j,s,t-p}$ are the $p = \{1, 2, \dots, P\}$ lags of the dependent variable with auto-regressive parameters ρ_p ; $G_{i,j,s,t-p}$ are the remaining variables of the dynamic system with coefficient vectors μ_p ; β_q is the estimated effect of the hurricane $h_{i,j,s,t}$ in period t on $y_{i,j,s,t+q}$ and $v_{i,j,s,t}$ are the regression errors. We follow Montiel Olea and Plagborg-Møller (2021) and include enough lags to ensure that our point estimates are robust to any serial correlation. Our baseline specification sets $P = 12$, but we have experimented with smaller and larger values of P without much change to the estimates. In terms of the horizon we examine, in general we set $Q = 5$, as we find no evidence of a longer term effect of hurricanes on trade.

For some routes, there appears to be a seasonal pattern in which trade is positive for

a sequence of months and then zero for a period of time thereafter. Conceptually, with sufficient lags in the deterministic component, this pattern would be a non-issue if we have sufficient time series observations. However, given our sample size, it is possible that our local projection may not precisely capture the seasonal variation in trade, in which case we may overestimate the counterfactual trade that would have occurred, and thus overestimate the effect of the hurricane. In an attempt to address this issue, we define a dummy variable $Z_{i,j,s,t} = 1$ for all periods that have zero trade. We include Z in G for the local projection specification, which amounts to assuming an auto-regressive process for zero-trade routes. Depending on the specification, we also include the logarithm of the exchange rate to the USD, the inverse hyperbolic sine of trade by air, and/or the logarithm of total exports excluding the US as the exporting country to control for the potential correlation of hurricanes with country-specific supply shocks.

One feature of the local projection strategy is that it is agnostic about whether hurricanes occur in adjacent time periods on different routes. Including lags of the dependent variable conditions the effect of hurricanes in period t on the occurrences of hurricanes in previous periods if the local projection specification is correct. One advantage relative to the event study is that the β_q estimate the total effect of hurricanes, because there are no contemporaneous route-specific fixed effects which may absorb some of the variation due to hurricanes. However, we remind the reader that a relative disadvantage of the local projection strategy is that any contemporaneous factor that occurs at the time of the hurricane may confound the estimated effect. In a long time series, this may be less of a concern if hurricanes are conditionally random. However, given our sample size it is possible that some factors not included in the local projection specification are coincidental to hurricane events and may bias the estimated effect of hurricanes.

4 Empirical results

We begin with our baseline event study approach outlined in Equation (1). In the following subsection, we report the results from the local projection strategy presented in Equation (2). Our baseline sample contains all observations once trade has been observed on a route. We use this sample restriction as our baseline specification because it removes zero observations that plausibly reflect the absence of trade ties between that exporter and the US port. For both approaches, we evaluate our baseline specifications and then separately consider the effects of sample restrictions on zero-trade observations and the effects of route distance for our results.

4.1 Event study

Column (5) of Table 2 presents the estimation results and Figure 8 plots the event study coefficient estimates for our main specification as defined in Equation (1).²² This specification includes the logarithms of the exchange rate, total exports and imports by air to the US port for the exporting country in $X_{i,j,s,t}$.²³ These controls are intended to account for possible confounding supply effects that affect production (the logarithm of total exports excluding the United States), legal or other regulatory issues specific to the route (the logarithm of imports by air to that port), or financial issues (the logarithm of the exchange rate). In addition, imports by air capture any substitution of trade from ship to air transportation.²⁴ The results and the figure show that hurricanes decrease trade on the route by a semi-elasticity of -0.054 in the month following the hurricane, and that there is no significant effect in other months.

The delayed effect on trade is likely because the majority of hurricanes in our data are near Asia—in the Sea of Japan or the South China Sea—and it takes most ships about 18 days to cross the Pacific Ocean. Thus, a hurricane that delays a ship from departing by a few days is most likely affecting imports in the following month if the hurricane is after the first week of the month. One, perhaps surprising, result is that there is no evidence that the lost trade is made up in subsequent months, which implies that the lost trade is a permanent effect in levels. This result suggests that there may be capacity constraints, either at the dock or in terms of available fleet capacity, that prevent exporters from sending additional ships in subsequent months to catch up. Alternatively, it could represent demand-side factors—for example, not receiving a loaf of bread in one month does not necessarily imply that a consumer will want 2 loaves in the following month.

Columns (3) to (4) present slight variations on our main specification to show the robustness of our results to different specifications of $X_{i,j,s,t}$.²⁵ Columns (1) and (2) show the sensitivity of our results to the inclusion or not of the $\zeta_{i,j,s,t \pm \tau}$ terms. The estimated effect of a hurricane crossing a route is consistent across all specifications which suggests

²²We cluster the standard errors by route, $\{i, j, s\}$, to account for serial correlation within route and also correlation by type of ship, see, for example, Larch, Wanner, Yotov, and Zylkin (2019) for a discussion. As we discuss in our local projection results below, we find some evidence of correlation by ship type, which suggests that shipping constraints may be salient.

²³Table 12 in the Supplementary Appendix presents estimates from a simple PPML model and regresses the hurricane indicator on the value of US vessel imports without any control variables or fixed effects. These estimates show the effect of seasonal correlation between hurricanes and trade values, column (1), and the dynamic effects of hurricanes on the control group, column (2). We provide these estimates only for transparency in our empirical analysis and do not refer to them further.

²⁴For example, if within the same month a ship did not leave the port because of an expected hurricane and part of the load was rerouted via air, then we may observe higher import values of trade by air in the month of the hurricane.

²⁵In column (3) we include the logarithm of total exports net of exports to the US for country j in month t and the logarithm of the average monthly exchange rate with the USD; in column (4) we include the logarithm of the total value of trade by air with the US port and the logarithm of the average monthly exchange rate with the USD.

that the specification of $X_{i,j,s,t}$ in our baseline sample is not driving our results. These estimates are significantly different from zero at the 1% level of significance. There is also no evidence of anticipatory effects, as none of the coefficient estimates prior to the hurricane are significantly different from zero. There is some evidence of a contemporaneous effect of hurricane crossings in column (1). This specification does not include the $\zeta_{i,j,s,t\pm\tau}$ which control for hurricanes in adjacent periods. Since our remaining estimates suggest that hurricanes impact trade with a one-month lag, the specification in column (1) would appear to include hurricane effects in the implied counterfactual that are correlated with the contemporaneous hurricane dummy (which is likely the case, since hurricane seasons are spatially correlated). This suggests that including the $\zeta_{i,j,s,t\pm\tau}$ terms is appropriate in our baseline specification.

4.1.1 The presence of zeros

As we have discussed, the monthly imports data feature many zeros, both within and across routes, which poses a potential concern because the zero value could be attributed to either the presence or absence of a hurricane on the route. Zero values themselves are not problematic for estimation because we use PPML. However, PPML's functional form implies that a hurricane that crosses a route that has zero trade will have an estimated effect, since PPML does not fit precise zeros. We consider different sample restrictions to test the robustness of our results to zeros. Our first additional sample restriction is to include observations only if there was any positive trade in that calendar year. This restriction reduces the proportion of zeros to 30%. Our second sample restriction is to remove all observations in which there was zero trade in any of the 3 months prior to the hurricane event. This further reduces the share of zero-trade flows to 8% in our estimation sample.

Table 3 and Figure 9 present the event study estimates for the sample restrictions described in the previous paragraph. There is almost no difference qualitatively or quantitatively from the additional sample restrictions on zero-trade observations. The point estimates are essentially identical to the baseline estimates, which suggests that the zero-trade routes are not driving the estimated coefficients on the hurricanes. This may, in part, reflect our careful attention to excluding temporally adjacent routes that are also affected by hurricanes, as many of these routes may, in fact, have zero trade and would otherwise affect our counterfactual estimates through their impact on the fixed effects.

4.1.2 Route distance

A second concern is that our estimated hurricane effects are based on routes of potentially different distances. It is possible that the estimated effect of hurricanes in Tables 2 and 3, and in particular the timing of the effects, reflect the impact of hurricanes on longer routes.

If, for example, a route takes more than a month to traverse and a hurricane impacts that route near its port of departure, then it is logical that the effect of that hurricane would not be observed in the import data until the following calendar month. Conversely, for a short trade route, it is possible that a hurricane would have very little effect, because the volume of trade that can be traded is based on the number of voyages a ship can make. If a ship is delayed for a few days by a hurricane, but then can increase its speed to traverse the route faster, then there may be, in fact, no change in the total volume of cargo it can carry in that month.²⁶

For each exporting port in our data, we calculate the Haversine distance in km between that port and the US destination port. Then, for each exporter, we take the average distance of its ports to the US destination port as its distance. The Haversine formula is an approximation to the port-to-port distance because it is based on spherical geometry and the Earth is not an exact sphere. However, it is sufficiently accurate to discriminate broadly between routes in our data. The average logarithm of route distance is 8.97 log points, and the median is 9.17, which suggests that skewness is not of great concern. The minimum distance is 5.28, and the maximum is 11.08. We interact the logarithm of the Haversine distance with the hurricane indicator $h_{i,j,s,t}$ and re-estimate the baseline event study regression.

Table 4 and Figure 10 report the estimated effects for hurricane crossings by distance. The point estimate of -0.007 implies an effect of $-0.007 \times 8.97 \approx -0.063$ that is consistent with our baseline event study estimates. The distance estimates also suggest that the effects of hurricanes are heterogeneous by distance, which is illustrated by Figure 10. Using the minimum and maximum log distances, back-of-the-envelope calculations suggest that the effects of hurricanes range between -0.037 and -0.078 . One additional difference is that the contemporaneous effect of hurricanes on trade is marginally significant, with an estimated effect of -0.004 (pvalue 0.085). This is suggestive of an interaction between distance, timing of trade interruption, and the location of the hurricane.

4.1.3 US exports

As discussed above, the estimated costs of hurricanes that we find could reflect ships taking longer routes around storms or choosing not to leave port. While we cannot directly address this issue using the import data, we can examine the response of US exports to hurricanes. One key difference between the export and the import data is that the import data is registered in trade statistics after ship arrival at the port, while exports are registered in trade statistics when they clear customs and before they are loaded on ships and leave the port. This timing difference implies that US exports should not be affected

²⁶It is likely that the shipping costs would be higher, but we do not observe the shipping costs; we observe simply the value of total trade imported.

by a hurricane intersecting a route if a ship did not leave port because of the hurricane. To test this prediction, we estimate our baseline event study replacing the imports data with export data.

The results in Table 10 show that hurricanes do not have a significant impact on export flows in all specifications with different control variables. Figure 11 shows the associated event graph with the pre- and post- hurricane dummies. All dummies are not statistically different from zero. These findings suggest that, conditional on the controls for expectations, hurricanes are unanticipated events; otherwise, we may observe ships not being loaded, leading to a lower export value in the month of the hurricane.

4.2 Local projection

The local projection specification, Equation (2), estimates the effect of hurricanes on trade by leveraging the time series structure of the data. Because trade is often seasonal, we include 12 lags of the dependent variable and 12 lags of the control variables G . In our baseline specification, G includes the logarithm of total exports and the logarithm of the exchange rate. The import data is also characterized by periods of intermittent trade and we define a dummy variable equal to 1 if a route has had positive trade in that month and then include this dummy variable in G . The route-month dummy variables for positive trade help to flexibly control for the seasonal patterns of trade on the route, which may reflect expectations of hurricane activity. Finally, we consider a second specification, where we include the logarithm of imports by air in G to examine the sensitivity of the estimates of β_q to the variables included in G .

Our baseline local projection sample is identical to our baseline event study sample and includes observations once there has been positive trade observed on that route. Because different types of ships appear to have different seasonal patterns of operation, e.g., bulk carrier routes often have non-zero trade after harvest seasons, and because the estimation residuals are likely to negatively covary across routes by ship type if there are capacity constraints, we cluster the standard errors for our regressions by type of ship.²⁷ Since we have only 5 types of ship, typical cluster-robust standard errors may be biased, and we follow MacKinnon, Nielsen, and Webb (forthcoming) and report a jackknife cluster-robust standard error (their CV3). Table 5 and Figure 12 present the estimates from the baseline local projection for the two specifications of G . Similar to the event study results, the estimated effect of the hurricane on imports is observed in the month following the storm. The point estimate of -0.16 is larger than the point estimate in the event study, but

²⁷Indeed, we find some evidence of a negative covariance by type of ship, which we interpret as an indication that ships themselves are a constraint on, at least some types of, trade. For instance, larger than usual demand for bulk carriers by one exporter may reduce the supply of bulk carriers available for a different exporting country. An alternative approach would be to cluster the standard errors by exporting country; however, we do not find much qualitative difference to the results.

in general the responses appear qualitatively similar. There is no evidence that trade is ‘caught up’ in the months following the hurricane event. As we have previously discussed, it is not particularly concerning that the point estimate is larger using the local projection strategy since some part of the hurricane effect is soaked up by the US port, type of ship, and date fixed effect. Indeed, the estimates presented in Table 5b, which include the inverse hyperbolic sine of air imports, are nearly identical to those using the logarithm of total exports in G , which was also the case in the event study results. This observation suggests that there is some consistency between the empirical specifications.

4.2.1 The presence of zeros and route distance

Although we include lags of a extensive margin dummy variable for trade to account for seasonal or intermittent trade, including zero-trade months, the local projection regression typically does not predict counterfactual zero trade, and may over-estimate the effect of hurricanes for zero-trade routes (like in the event study). Similarly, the baseline local projection results are based on routes of differing distances, and this finding may impact the timing of the hurricane effect on trade.

We consider two robustness exercises. In the first, we restrict our sample to routes that had positive calendar-year trade for the month of the hurricane. In general, this restriction allows us to keep the time-series specification of the local projection regression unchanged, as there are typically at least several months of data prior to the hurricane to inform the estimates. This specification provides some illustration of how zero-trade observations may be affecting the estimated effect of hurricanes. We note, however, that this sample restriction does make interpreting the coefficient estimates a little challenging as it may potentially change the time series behaviour of the data. However, by keeping all observations with at least some positive calendar-year trade, we are not directly selecting on the short-run dynamics of the trade process that is used to estimate the local projection coefficients. Table 6 and Figure 13 present the estimates and the impulse responses for the effect of a hurricane in the sample restricted to positive calendar-year trade. The point estimate is larger in absolute terms at 0.225, and it remains significantly different from zero. Perhaps more importantly, the timing of the hurricane impact remains unchanged from the baseline local projection specification.

The second robustness exercise we consider is to interact the logarithm of the trade route’s distance with the hurricane crossing indicator. Table 7 reports the point estimates for the heterogeneous impulse response for the hurricane crossing by distance. The hurricane impacts trade one month after the hurricane crossing by -0.018 log points, which implies an average effect of $-0.018 \times 8.97 \approx -0.162$, which is almost exactly identical to the baseline local projection results. The point estimate also suggests an effect range of -0.095 to -0.199 , depending on the route distance (see Figure 14). One difference between the local projection and the event study estimates is, however, that the contemporaneous

hurricane effect is insignificantly different from zero and no longer marginally significant.

4.2.2 Short-run trade adjustment

One advantage of the local projection approach is that we can examine whether exporters reroute trade because of hurricanes.²⁸ There are two possibilities. The first is that an exporting country facing a hurricane disruption along a route to Los Angeles, for example, could reroute the intended shipment to arrive at a different US port, e.g., Seattle, that was anticipated to be unaffected by the hurricane. The second possibility is that an exporting country with more than one port may reroute shipments within its borders if routes from only one port are affected. Both options are likely to be costly. Rerouting to a different US port suggests that there may be additional logistical costs to deliver the imports to their final destination, perhaps especially for those importers that are geographically local to the original port. It may also be the case that changing the destination port would entail additional and prohibitive administrative costs, or that changing ports is simply impossible because there is no available berth at the new US port. Similarly, changing the port of export is likely to lead to additional logistics costs for exporters, who must now deliver their product to a different location for shipment. It is also likely to increase costs for shippers, who may be required to relocate a ship between ports.

To address the first possibility, we construct a new ‘spillover’ dummy variable that equals 1 for any route that is not directly affected by a hurricane but for which a route to a different US port is affected, for each type of ship by exporter country. If exporters adjust to a hurricane by shifting trade to other US ports, then the coefficient estimate on this indicator variable should be positive. We also restrict the sample to exclude routes that were directly hit by a hurricane and exporting countries that have only one route to the US. Table 8 reports the point estimates for the spillover dummy variable with the standard errors again clustered using the CV3 jackknife cluster-robust standard error recommended by MacKinnon, Nielsen, and Webb (forthcoming). Figure 15 presents the estimated impulse response. As the table and figure show, we find no evidence that exporters reroute shipments to other US ports.

We next investigate whether there is any evidence that exporting countries adjust to hurricanes by shipping from different ports within their country. We define a ‘changeport’

²⁸Recall that the local projection regressions include monthly route-specific fixed effects, which control for hurricane expectations, and lagged dependent and control variables to control for route-specific dynamic adjustments. Thus, the identification of hurricane effects relies on a comparison of affected routes versus unaffected routes, conditional on route-specific monthly factors and dynamic trade patterns, across exporting countries and over time. We identify short-run adjustments through these conditional comparisons. In contrast, the event study specification includes, in addition to monthly route-specific fixed effects, time-varying port and time-varying ship-type specific fixed effects which capture most, if not all, variation due to short-run adjustments. Identification relies on within period comparison of affected and unaffected routes and precludes estimating separate treatment effects for each of those routes.

dummy variable equal to 1 for an exporter that has more routes to a given US port than routes directly affected by a hurricane for each ship type. For example, if there is a hurricane on a RoRo route from Osaka to Los Angeles, but no hurricane on a RoRo route from Yokohama to Los Angeles, then the changeport dummy variable is equal to 1, because Japanese automobile exporters could switch the port they use for export. We re-estimate our baseline local projection including the changeport dummy in addition to the usual hurricane indicator variable. Our interest is on the coefficient estimates for the changeport dummy. Table 9 and Figure 16 present the estimated impulse response differences (the coefficients on the changeport dummy) for our baseline sample. The standard errors are again clustered using the CV3 jackknife cluster-robust standard error recommended by MacKinnon, Nielsen, and Webb (forthcoming). As the table and figure illustrate, there is no evidence that exporters adjust to hurricanes by rerouting within their borders.

4.3 US Exports

In the event-study analysis, we examined US export data to determine whether there was any evidence that hurricanes were anticipated events. Our identification strategy exploits hurricanes as conditionally random events on trade routes—controlling for seasonal expectations using route-month fixed effects. While we found no evidence of anticipatory effects using the event-study specification, it is possible that the control variables in that specification absorb variation that would reveal such effects.

We re-examine the effects of hurricanes on US exports using a local projection specification that has the same basic structure as Equation (2). We include the following control variables in the matrix analogous to G : the logarithm of the exchange rate; the logarithm of total US west coast imports excluding that route (to control for supply conditions); the inverse hyperbolic sine of exports by air for that route; and a dummy variable for positive trade for that route by month. We also specify the same lag structure and include route-month fixed effects as Equation (2). The results in Table 11 and Figure 17 are qualitatively identical to those for the event study specification and show no evidence of a hurricane impact on export flows.

5 Hurricanes and aggregate trade

The estimates and robustness exercises from the event study and local projection specifications all imply that hurricanes disrupt trade with a one-month delay, and that the losses are permanent, in the sense that we find no evidence of higher-than-expected trade in subsequent months. The estimated costs are between 5.4% and 16.0% of route-level trade depending on the exact specification. These costs appear economically salient, at least at the route level. In this section, we discuss their implications for aggregate trade and economic activity.

As we have previously noted, hurricane-route intersections are both spatially and temporally correlated. Hurricanes at sea are not localized events, and hurricane eyes often cross many optimal trade routes. There is also variation in the value of trade across routes. Thus, the aggregate trade costs of hurricanes almost certainly vary across hurricane events.

5.1 The aggregate cost of a hurricane season

To calculate the aggregate costs of hurricanes for US west coast imports, we simulate tropical cyclone seasons and combine these with our route-level cost estimates. The simulations are based on historical storm activity from 2003-2019 in our sample; therefore, the results of these simulations are representative of recent trade costs.

We begin by calculating a vector, P , whose cells contain the probability of a hurricane hitting a route (which is ship-type specific) between an exporting country and a US port, in a given month of the calendar year. The probability of a cyclone hitting a country-port pair in a month is calculated as the total number of storms that hit the route in a calendar month over 2003-2019, divided by the total number of years (17). In our data, we have a total of 180 exporting country-US port-ship-type routes, and with twelve months in the calendar year, P has 2160 rows, with each cell representing the probability that a country-port pair has been hit by a hurricane in, say, January, calculated as the average across all January-country-port observations in the sample.

For the baseline simulations—which calculate the total cost of an average storm season over the 2003-2019 period—we draw from a continuous uniform distribution, with values between 0 and 1, for each cell in P . If the cell value in P exceeds the value of the uniform draw, then we record that a hurricane has hit that route in the simulation. Once a full storm season has been simulated, for each route that has been exposed to a storm, the month-specific sample average trade share of that route $\omega_{i,j,s,m}$ is multiplied by $(1 - |\hat{\beta}|)$ to create a simulated trade share, $\omega_{i,j,s,m}^S$. For country-port pairs not hit by a major storm in a simulation, $\omega_{i,j,s,m}^S = \omega_{i,j,s,m}$, and for those hit by a storm, $\omega_{i,j,s,m}^S < \omega_{i,j,s,m}$. The country-port-month trade shares are calculated as the mean share for each country-port pair over all years in the sample, such that $\omega = \sum_{i,j,s,m} \omega_{i,j,s,m} \approx 1$ and $\omega^S < \omega$.²⁹ The impact of a storm season on total trade in percentage terms is then calculated as $\Gamma^S = (1 - \frac{\omega^S}{\omega}) \cdot 100$.

We simulate the model 1000 times and take the mean value of Γ^S as the baseline simulated storm season (call it Γ^B). This methodology suggests that an average storm season based on historical storm patterns over the 2003-2019 period cost approximately 1.15% of US west coast imports when $\hat{\beta} = -0.054$. The loss is 3.42% when $\hat{\beta} = -0.16$. One advantage of the simulations is that they allow us to explore “good” and “bad” storm

²⁹By averaging shares across years, the sum of the shares is slightly above 1, but this has little bearing on our simulation calculations. In addition, the trade shares will include cyclone effects; however, by averaging over the whole sample, these likely reflect expectations regarding major storms rather than the exogenous, unexpected impact of storms, which is what we are capturing here.

years. For example, the 10th percentile of the baseline simulation (a “good” year) results in a trade cost of about 2.43% for the case where $\hat{\beta} = -0.16$, whereas the 90th percentile (a “bad” year) has a cost 4.43% of US west coast imports.

To provide some back-of-the-envelope evidence on the magnitude of the hurricanes effect, following Arkolakis, Costinot, and Rodríguez-Clare (2012), we can calculate the hurricane impact on consumer welfare using the percentage change in the domestic expenditure share ($d \ln \lambda_{ii}$) divided by the trade elasticity (ϵ). The percentage change in the domestic expenditure share is given by the negative value of the percentage point change in the import expenditure share divided by the domestic expenditure share.³⁰ Using our estimated loss of 3.42% of west coast imports by ship, the fact that these imports account, on average, for 19.7% of total imports and Costinot and Rodríguez-Clare (2014)’s estimate of 0.9147 for the domestic expenditure share, the implied change in welfare is 0.016%.³¹ This welfare loss is likely a lower bound, because it implicitly assumes some expenditure switching between imports and domestically produced goods. If switching to domestic goods is not possible within a short period or if missing foreign intermediate inputs cannot be replaced, the implied welfare losses are likely higher. Given that we examine losses from unanticipated hurricanes, the feasibility of expenditure switching seems unclear. We can also express the trade loss in dollar terms. Overall, our estimates suggest that in 2019 the projected loss from a hurricane season is about \$17.1 billion while total US imports were \$2.5 trillion. For comparison, (Deryugina, 2017) estimates that the total yearly physical damage of hurricanes hitting land are, on average, \$8.1 billion.

5.2 Discussion of the implications for economic activity

Our estimates of the aggregate costs of hurricanes for trade suggest that they lower annual US west coast imports by up to 3.42% on average. In this section, we relate the fall in imports to Gross Domestic Product (GDP) and briefly discuss the channels through which these losses may impact aggregate economic activity.

Possibly the most direct approach is to note the expenditure approach to GDP (Y) is:

$$Y = C + I + G + E - M,$$

where C is consumption, I is investment, G is government spending, E is exports and M is imports. Since the expenditure calculation of GDP is an identity, any change in M must necessarily impact one of the other terms in the equation. For simplicity, we

³⁰We use the fact that the import expenditure share m_{ii} is 1 minus the domestic expenditure share. This implies that the percentage change in the domestic expenditure share equals $\ln \lambda_{ii} = \frac{d\lambda_{ii}}{\lambda_{ii}} = \frac{-dm_{ii}}{\lambda_{ii}}$.

³¹The percentage point change in the import share is the import share (1-0.9147) multiplied by the west coast import share of total imports 0.197 multiplied by the percentage change in trade loss 0.0342. We then divide this percentage point change by the domestic expenditure share and the trade elasticity (which we assume to be 4).

assume that government spending and exports are unaffected. Thus, the effect of a fall in M must be (1) to increase Y , which would imply substitution to domestically produced goods; (2) to lower C and/or I , which would suggest that imports are final goods; (3) some combination of (1) and (2); or (4) a decrease in Y and C or I , which could occur if imports are intermediate goods in a domestic supply chain. While the analysis in this paper does not differentiate between intermediate and final goods, it would seem probable that the imports we examine are a mix of intermediate and final goods. The exact distribution of the trade losses will also depend on several factors, such as whether shipments are insured, pre-paid or priced using within firm transfer prices. The distribution also likely depends on the type of good being traded or whether prices adjust—for example, if US inventories are used to smooth trade shocks.³² Because our estimated effects are the route-specific percentage loss from a hurricane intersection, they identify the value of missing imports without attributing which entity bears that loss. In addition, the contracting environment and the presence of insurance may matter. For example, if goods are paid for in advance, then the importer may bear the losses, whereas if goods are paid for upon delivery, the exporter may incur the losses.

To further consider how trade losses resulting from unanticipated hurricanes impact economic activity, imagine a hypothetical firm in LA that imports and sells 5 red bikes from a supplier in China every May. Suppose that when a storm interrupts this route, only 3 bikes are delivered, and that the firm believes there is a 50% chance that a storm intersects this route in May. Thus, the firm expects to receive (import) 4 bikes each May. To insure against the anticipated import shortfall, assume that the firm holds 1 bike in inventory. In our data, we would observe that, on average, only 4 bikes are expected to be delivered on this route each May, and this expected difference between 4 and 5 would be the route-month fixed effect (5 bikes being the no-storm counterfactual identified from the cross section and/or time series, depending on our identification scheme). When a storm occurs, we observe that only 3 bikes are delivered, and the estimated hurricane impact is the difference between 4 and 3 since the expected shortfall in the route-month fixed effect already accounts for 1 of the missing bikes. From the importer perspective, they would have 1 bike in inventory and 3 from imports, for a total of 4 bikes—1 short of the 5 they want to sell. Relating this thought experiment to national accounts, a hurricane disruption would have $C = 4$, $I = -1$ and $M = 3$ which would imply $Y = 0$ —that there is no effect on GDP from the storm. In this simple example, although consumption of red bikes is lower (4 instead of 5) and imports are lower (3 instead of 5), there is no effect on GDP because this example does not account for retail profits or other value-chain effects. Understanding multiplier effects from trade or the effects of import variability for price indices are, however, interesting topics that we defer to future research.

³²See, for example, estimates of the inventory response to oil supply shocks in Brannlund, Dunbar, and Ellwanger (2022).

6 The cost of climate variability for trade

Hurricanes are naturally occurring climate phenomena, so it is perhaps natural to wonder what our estimates imply for the costs of climate change for trade. The IPCC reports that climate change in the coming decades is expected to affect storm activity along three dimensions: the frequency of storms, the intensity of storms, and poleward migration of the average latitude at which major storms reach their maximum intensity (Pörtner, Roberts, Masson-Delmotte, Zhai, Tignor, Poloczanska, and Weyer, 2019). We combine our estimates of the effects of hurricanes on trade with potential outcomes for frequency, intensity and poleward migration to calculate the costs of climate variability for trade. This approach assumes no adaptation by exporters or trade networks. By implication, the costs presented in this section should be interpreted as indicative of either the medium-term costs of climate change on trade (before any adaptations occur) or as the costs of non-adaptation. We are agnostic regarding which of the scenarios we construct is most probable. Instead, we calculate costs for a range of outcomes so that readers can evaluate the expected costs of future climate variability based on their own subjective expectations of climate change.

For clarity in what follows, we refer to the 3.42% cost estimate—derived using the local projection estimate of β —as our baseline cost estimate. Only the climate-change scenario results based on the local projection estimates of β will be presented in the text, while the simulation results based on the event study estimate of β are available in the figures only.

6.1 Climate change simulations

To explore the effects of climate change, we adjust P to reflect projected changes in the frequency, intensity and poleward migration of storms. We detail the steps taken to simulate outcomes associated with climate change along these three dimensions individually and present stand-alone estimates for the effect of each. We then combine them to provide an overall picture of the climate change on international trade, or, more specifically, US west coast imports.

6.1.1 Changes in the frequency of storms

Knutson, Camargo, Chan, Emanuel, Ho, Kossin, Mohapatra, Satoh, Sugi, Walsh, and Wu (2020) summarize and assess the existing literature on the possible effects of global warming on tropical storm patterns, and present multistudy aggregate estimates and ranges scaled to a 2°C increase in mean surface air temperature. The majority of the studies surveyed by Knutson, Camargo, Chan, Emanuel, Ho, Kossin, Mohapatra, Satoh, Sugi, Walsh, and Wu (2020) project that the total number of global major storms will decrease with global warming, with a median change across the studies of -14%. However, since our focus is on imports to the US west coast, the majority of which pass through the north Pacific Ocean,

we draw on the range of estimates provided for the northeast Pacific. For this zone, the median estimate is for a 0% change in the frequency of tropical cyclones, with a range from -40% to +40% for the 5th and 95th percentile ranges across the studies, and it is this range that we use to simulate the effects of changes in the frequency of storms.³³

To simulate increased frequency of storm activity, we multiply P by a scalar representing the desired change in frequency. For example, to simulate a 10% increase in the frequency of storm activity, we multiply each cell in P by 1.1. For those route-months where a hurricane has never hit, there is no change to that cell in P . For those route months that have a positive probability of being hit by a hurricane, that probability increases by 10%.³⁴ We do this for each decile within the -40 to +40% range. Call the new frequency P vectors, P^{freq} where $freq \in \{-40, -30, -20, -10, 0, 10, 20, 30, 40\}$.³⁵

The frequency simulations are then conducted in the same fashion as the baseline estimate, except that the P^{freq} vectors are used instead of P , and cost estimates are derived for each value of $freq$. As shown in Figure 18, a 40% decrease in the frequency of storm activity would reduce the average cost of a storm season to 2.05%—from our baseline estimate of 3.42%—whereas a 40% increase in the frequency of storms would increase costs to 4.50%. Finally, in an environment with a 40% increase in storm activity, a “bad” year (90th percentile of the simulations) would result in a cost of 5.28% of US west coast imports.

6.1.2 Changes in the intensity of storms

For the intensity simulations, we again rely on the estimates provide in Knutson, Camargo, Chan, Emanuel, Ho, Kossin, Mohapatra, Satoh, Sugi, Walsh, and Wu (2020), which are scaled to reflect a 2°C increase in global mean surface air temperature. Our definition of storm intensity is the maximum wind speed reached by a storm, which is broadly consistent with the studies surveyed by Knutson, Camargo, Chan, Emanuel, Ho, Kossin, Mohapatra, Satoh, Sugi, Walsh, and Wu (2020), where they find a general consensus that, worldwide, there will be an increase in the intensity of storms, with a median estimate of about a 5% increase in storm wind speed. This is the same median change projected throughout the Pacific. The ranges around this estimate are rather tight—in the northwest Pacific, the 10th percentile of the estimate range is -2.5% and the 90th is 10%, and there is a similar range for the northeast Pacific.

³³Our frequency simulations are based on the range provided in Knutson, Camargo, Chan, Emanuel, Ho, Kossin, Mohapatra, Satoh, Sugi, Walsh, and Wu (2020) for the northeast Pacific. They also provide estimates for the northwest Pacific, which are relevant to our trade routes. The median estimate is similar for the two zones, but the range of estimates is greater for northeast Pacific.

³⁴The fact that the frequency simulations have no effect on routes that have never been hit by a storm is by design. Both the intensity and poleward migration simulations detailed below will affect some routes that have never crossed a hurricane, and the frequency simulations were designed to be able to separate these effects.

³⁵The case where $freq = 0$ is equivalent to the baseline simulation in the previous subsection.

The first step of our intensity simulations will be to calculate a series of intensity P vectors associated with different projected changes in the wind speed of storms. Because our estimates do not vary with the strength of hurricanes, increasing the wind speeds of existing hurricanes will have no effect on our cost estimates. Where increases in intensity will result in higher costs is via adjustments in the extensive margin of hurricanes—that is, increasing the wind speeds of near hurricane-strength observations, thereby turning them into hurricanes. As a result, the costs associated with intensity adjustments are relatively small compared to the frequency and poleward migration simulations. For this reason, we choose a broad range of intensity adjustments—beyond the ranges in Knutson, Camargo, Chan, Emanuel, Ho, Kossin, Mohapatra, Satoh, Sugi, Walsh, and Wu (2020)—to provide a better sense of the relationship between the intensity shocks and trade costs.

We define P^{int} as the probability vectors associated with adjustment in the intensity of storms, where the range of intensity adjustments (in terms of percentage adjustment in maximum wind speed) is $int \in \{-10, -5, 0, 5, 10, 15, 20, 25, 30\}$. To generate the P^{int} vectors, we take the maximum wind speed within a given route-month pair in the observed data and adjust it by $int\%$. We then re-calculate whether a hurricane has hit a route based on the adjusted wind speed using the Saffir-Simpson wind scale—where maximum wind speeds 64 knots or above constitute a hurricane. Finally, with a new history of hurricanes, we calculate P^{int} using the same methodology used to construct P .

The simulations then proceed as in the baseline and frequency simulations. Figure 19 presents the outcome of the simulations, and we see that a 5% increase in the intensity of storms raises the cost of an average storm season to 3.45%. A decrease in storm intensity of 10% reduces the cost to 3.20%, while a 30% increase raises costs to 4.18%. Finally, a “bad” year (90th percentile of the simulations) in an environment associated with a 30% increase in intensity could cost up to 5.16% of US west coast imports.

6.1.3 Poleward migration

The third and final dimension of climate-change induced changes in storm behaviour that we explore is poleward migration. The IPCC has highlighted that climate change is likely leading to upward shifts in the latitudes at which storms reach their maximum wind speed. Trends in poleward migration are identified in Kossin and Vecchi (2014), who estimate a global-average migration of major storm activity away from the tropics at a rate of roughly 1° of latitude per decade. Given this finding, we attempt to model a 1° latitude shift in poleward migration, which is analogous to a ten-year horizon.

Simulating scenarios that reflect a poleward migration of hurricanes requires yet a different approach than our simulations for frequency and intensity. Conceptually, the simulations of changes to frequency and intensity held constant the geospatial location of storm tracks. With poleward migration, storm tracks themselves are likely to be affected. A second issue is that poleward migration of hurricane locations implies that the range of

latitudes where hurricanes may be found is increasing, but this does not imply that all hurricanes are likely to move poleward.³⁶ If we hold constant the frequency and intensity of hurricanes, then the choice of which hurricanes shift poleward is the main determinant of the simulated costs.

We proceed as follows. For each hurricane observed in the data, we take a draw a from a continuous uniform distribution with values between and including 0 and 2— $a \sim U[0, 2]$ —and then adjust the entire path of the hurricane northward by a° latitude.³⁷ This will result in an average northward shift of 1° latitude of hurricane paths, which is consistent with a ten-year horizon in Kossin and Vecchi (2014). We next replace the original data for these hurricanes with the adjusted locations so that we have a new sample of “historical” hurricane activity. We then recompute hurricane and route crossings for this new sample and use these new data to recompute a probability vector reflecting the new hurricane data, P^{pm} . Finally, to average out the random effects from the individual draws of a , we generate 10 probability vectors, P^{pm} where $pm \in \{1, \dots, 10\}$, each associated with a unique random number generator for the draws of a . Simulations for each P^{pm} then proceed as in the baseline simulations and the mean cost across the 10 simulations is our overall cost of poleward migration.

Our results suggest that poleward migration will increase the cost of an average hurricane season to 4.05%. The lowest and highest estimates across the 10 simulations can also provide a sense of the variation of our estimates: the lowest estimate is 3.99% and the highest estimate is 4.18%. At least for the average effect, there is little variation across the 10 random samples pm .

6.1.4 Combined shifts in frequency, intensity, and poleward migration

To provide a composite picture of the costs of climate change, we combine the effects of changes in frequency, intensity, and poleward migration. We begin with the adjustments for poleward migration, which identify new routes being hit by hurricanes, and then apply the adjustments for changes in intensity using the newly created poleward migration storm data. With these adjustments, we can create a series of hurricane-probability vectors associated with the different values of pm and int . Next, for each of these new probability

³⁶There remains a significant amount of uncertainty regarding how poleward migration will affect the overall paths of major storms (Kossin and Vecchi, 2014). For example (as it applies to our particular analysis), poleward migration may shift the entire storm path toward higher latitudes (thereby increasing storm intersections with routes at higher latitudes and decreasing intersections with routes nearer the equator) or the poleward shift may stretch out the maximum intensity of storms paths (thereby increasing the number of routes affected by a single storm). A lack of reliable historical data on storm formation and duration makes it difficult to analyze the relationship between poleward migration and adjustment in storms’ full paths. Without clear guidance on this, we assume that the entire path of a selected storm shifts in our simulations.

³⁷Because the supermajority of routes and hurricanes in our data are from the northern hemisphere, we consider only one direction of poleward migration—northward shifts.

vectors, we make frequency adjustments, *freq*. We then have 810 probability vectors on which to run 1000 simulations and for each value of *freq* and *int* we will average across the 10 *pm* draws to get a single value associated with poleward migration. For comparison, we also compute probability vectors for the interaction of frequency and intensity without poleward migration to provide some sense of the relative importance of route locations (recall the frequency and intensity simulations hold the routes fixed).

We begin by presenting the results for the interaction between the frequency and intensity adjustments with no poleward migration in Figure 20. The median estimates presented above from Knutson, Camargo, Chan, Emanuel, Ho, Kossin, Mohapatra, Satoh, Sugi, Walsh, and Wu (2020) are a 0% change in frequency and a 5% increase in storm intensity, which would result in a trade cost of 3.45%. At the upper end of the ranges in Knutson, Camargo, Chan, Emanuel, Ho, Kossin, Mohapatra, Satoh, Sugi, Walsh, and Wu (2020)—a 40% increase in frequency and a 10% increase in intensity—the cost would rise to 4.69%. And, if the intensity of storms were to increase by 30% (the top end of our estimates), the trade cost would be 5.33%.

In Figure 21, we present the interaction of the frequency and intensity adjustments along with the average latitude increase—about 1° —for hurricanes. Here we see that the median estimates for frequency and intensity are associated with a 4.08% trade cost. For upper-end frequency and intensity estimates in Knutson, Camargo, Chan, Emanuel, Ho, Kossin, Mohapatra, Satoh, Sugi, Walsh, and Wu (2020)—40 and 10%, respectively—the cost is 5.40%, while for the lower end of the frequency and intensity estimates the associated cost is 2.22%. Finally, if the frequency of storms were to increase by 40% and the intensity by 30%, the trade loss due to major storms would rise to 6.10%.

We acknowledge that the costs presented in this section are mechanical transformations and omit behavioural responses along trade networks. Our defense is that to model endogenous changes in trade networks or hurricane expectations that result from climate change is a daunting prospect. We view the results in this section, however, as an indication that such changes are likely to occur and deserve study. Certainly, in the absence of adaptation, our calculations suggest that the costs of hurricanes for trade could substantially increase. We also note that our scenarios follow Knutson, Camargo, Chan, Emanuel, Ho, Kossin, Mohapatra, Satoh, Sugi, Walsh, and Wu (2020) and assume a 2° change in global average temperatures. Higher or lower average temperature scenarios would likewise affect our simulated outcomes.

7 Conclusion

In this paper we have identified and estimated the effect of extreme weather and climate change for international trade. Our preferred estimates suggest that unanticipated hurricanes reduce monthly trade flows by 5.4% to 16% on affected trade routes. We

find no evidence that these trade losses are recouped by trade diversion either intra- or inter-temporally. We use our estimates and scenario projections for climate change from the IPCC to quantify the potential costs of climate change for trade. The scenario analysis does not focus on a single future outcome and we present ranges of possible future costs. On average these costs are expected to rise.

In terms of policy recommendations from our study, beyond the obvious statement that efforts to reduce climate change will likely lower future trade losses from extreme weather, our estimates may suggest that trade networks are likely to adapt endogenously to mitigate these costs in coming decades. For example, existing hub and spoke trade networks may no longer remain optimal configurations if the poleward migration of hurricanes accelerates. Future research on trade adaptation and configuration and the optimal policy response appear to be required.

8 Bibliography

- ALESSANDRIA, G., S. YAR KHAN, A. KHEDERLARIAN, C. MIX, AND K. J. RUHL (2022): “The aggregate effects of supply-chain delays,” *mimeo*.
- ARKOLAKIS, C., A. COSTINOT, AND A. RODRÍGUEZ-CLARE (2012): “New trade models, same old gains?,” *American Economic Review*, 102(1), 94–130.
- BARATTIERI, A., AND M. CACCIATORE (2020): “Self-harming trade policy? Protectionism and production networks,” Working Paper 27630, National Bureau of Economic Research.
- BELASEN, A. R., AND S. W. POLACHEK (2008): “How hurricanes affect wages and employment in local labor markets,” *American Economic Review*, 98(2), 49–53.
- BELLEMARE, M. F., AND C. J. WICHMAN (2020): “Elasticities and the inverse hyperbolic sine transformation,” *Oxford Bulletin of Economics and Statistics*, 82(1), 50–61.
- BENTLEY, J. L., AND T. A. OTTMANN (1979): “Algorithms for reporting and counting geometric intersections,” *IEEE Transactions on Computers*, C-28 (9), 643–647.
- BRANCACCIO, G., M. KALOUPSTIDI, AND T. PAPAGEORGIOU (2020): “Geography, transportation, and endogenous trade costs,” *Econometrica*, 88(2), 657–691.
- BRANNLUND, J., G. DUNBAR, AND R. ELLWANGER (2022): “Are temporary oil supply shocks real?,” Working Paper 2022-52, Bank of Canada/Banque du Canada.
- COSTINOT, A., AND A. RODRÍGUEZ-CLARE (2014): “Trade theory with numbers: Quantifying the consequences of globalization,” in *Handbook of International Economics*, vol. 4, pp. 197–261. Elsevier.
- CRUZ, J.-L., AND E. ROSSI-HANSBERG (2021): “The economic geography of global warming,” Working Paper 28466, National Bureau of Economic Research.
- DE CHAISEMARTIN, C., AND X. D’HAULTFOEUILLE (2022): “Difference-in-differences estimators of intertemporal treatment effects,” Working Paper 29873, National Bureau of Economic Research.
- DERYUGINA, T. (2017): “The fiscal cost of hurricanes: Disaster aid versus social insurance,” *American Economic Journal: Economic Policy*, 9(3), 168–98.
- DESCHÊNES, O., AND M. GREENSTONE (2007): “The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather,” *American Economic Review*, 97(1), 354–385.
- DINGEL, J. I., K. C. MENG, AND S. M. HSIANG (2019): “Spatial correlation, trade, and inequality: Evidence from the global climate,” Working Paper 25447, National Bureau of Economic Research.
- DUBE, A., D. GIRARDI, Ò. JORDÀ, AND A. M. TAYLOR (2022): “A local projections approach to difference-in-differences event studies,” *mimeo*.
- GALLAGHER, J., AND D. HARTLEY (2017): “Household finance after a natural disaster: The case of hurricane Katrina,” *American Economic Journal: Economic Policy*, 9(3), 199–228.

- GANAPATI, S., W. F. WONG, AND O. ZIV (2021): “Entrepôt: Hubs, Scale, and Trade Costs,” Working Paper 29015, National Bureau of Economic Research.
- HEILAND, I., A. MOXNES, K. H. ULLTVEIT-MOE, AND Y. ZI (2019): “Trade from space: Shipping networks and the global implications of local shocks,” *CEPR Discussion Paper*, DP14193.
- HELPMAN, E., M. MELITZ, AND Y. RUBINSTEIN (2008): “Estimating trade flows: Trading partners and trading volumes,” *The Quarterly Journal of Economics*, 123(2), 441–487.
- JACKS, D. S., AND M. STUERMER (2021): “Dry bulk shipping and the evolution of maritime transport costs, 1850-2020,” Working Paper 28627, National Bureau of Economic Research.
- JORDÀ, Ò. (2005): “Estimation and inference of impulse responses by local projections,” *American Economic Review*, 95(1), 161–182.
- KNAPP, K. R., H. J. DIAMOND, J. P. KOSSIN, M. C. KRUK, AND C. J. SCHRECK (2018): “International best track archive for climate stewardship (IBTrACS) project, version 4 (all),” Discussion paper.
- KNAPP, K. R., M. C. KRUK, D. H. LEVINSON, H. J. DIAMOND, AND C. J. NEUMANN (2010): “The international best track archive for climate stewardship (IBTrACS): Unifying tropical cyclone best track data,” *Bulletin of the American Meteorological Society*, 91, 363–376.
- KNUTSON, T., S. J. CAMARGO, J. C. L. CHAN, K. EMANUEL, C.-H. HO, J. KOSSIN, M. MOHAPATRA, M. SATOH, M. SUGI, K. WALSH, AND L. WU (2020): “Tropical cyclones and climate change assessment: Part II: Projected response to anthropogenic warming,” *Bulletin of the American Meteorological Society*, 101(3), E303–E322.
- KORINEK, J. (2011): “Clarifying trade costs in maritime transport,” Discussion paper, OECD: Working Party of the Trade Committee.
- KOSSIN, E., AND G. VECCHI (2014): “The poleward migration of the location of tropical cyclone maximum intensity,” *Nature*, 509, 349–352.
- LARCH, M., J. WANNER, Y. V. YOTOV, AND T. ZYLKIN (2019): “Currency unions and trade: A PPML re-assessment with high-dimensional fixed effects,” *Oxford Bulletin of Economics and Statistics*, 81(3), 487–510.
- MACKINNON, J. G., M. Ø. NIELSEN, AND M. D. WEBB (forthcoming): “Leverage, influence, and the jackknife in clustered regression models: Reliable inference using summclust,” *Journal of Econometrics*.
- MENDELSON, R., W. D. NORDHAUS, AND D. SHAW (1994): “The impact of global warming on agriculture: A ricardian analysis,” *The American Economic Review*, 84(4), 753–771.
- MONTIEL OLEA, J. L., AND M. PLAGBORG-MØLLER (2021): “Local projection inference is simpler and more robust than you think,” *Econometrica*, 89(4), 1789–1823.
- PÖRTNER, H.-O., D. ROBERTS, E. POLOCZANSKA, K. MINTENBECK, M. TIGNOR, A. ALEGRÍA, M. CRAIG, S. LANGSDORF, S. LÖSCHKE, V. MÖLLER, AND A. OKEM

- (2022): “IPCC, 2022: Summary for policymakers,” in *Climate Change 2022: Impacts, Adaptation, and Vulnerability. Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, ed. by H.-O. Pörtner, D. Roberts, E. Poloczanska, K. Mintenbeck, M. Tignor, A. Alegría, M. Craig, S. Langsdorf, S. Löschke, V. Möller, A. Okem, and B. Rama, vol. In Press. Cambridge University Press.
- PÖRTNER, H.-O., D. C. ROBERTS, V. MASSON-DELMOTTE, P. ZHAI, M. TIGNOR, E. POLOCZANSKA, AND N. WEYER (2019): “The ocean and cryosphere in a changing climate,” *IPCC Special Report on the Ocean and Cryosphere in a Changing Climate*.
- ROBERTS, M. J., AND W. SCHLENKER (2010): “Is agricultural production becoming more or less sensitive to extreme heat? Evidence from U.S. corn and soybean yields,” Working Paper 16308, National Bureau of Economic Research.
- SCHLENKER, W., W. M. HANEMANN, AND A. C. FISHER (2005): “Will U.S. agriculture really benefit from global warming? Accounting for irrigation in the hedonic approach,” *The American Economic Review*, 95(1), 395–406.
- SILVA, J. S., AND S. TENREYRO (2006): “The log of gravity,” *The Review of Economics and Statistics*, 88(4), 641–658.
- STROBL, E. (2011): “The economic growth impact of hurricanes: Evidence from US coastal counties,” *Review of Economics and Statistics*, 93(2), 575–589.

9 Appendix

9.1 Tables

Table 1: Summary statistics

	Mean	Min	Max	SD
Number of routes	180			
Number of exporting countries	39			
Routes per country	9.39	1	18	5.82
Route type per country	3.32	1	5	1.39
Hurricane frequency per route by month	0.098	0	1	0.298
Monthly imports by US port (Million \$)				
Los Angeles	524	0	34,500	2960
Portland	42.5	0	1,650	185
San Diego	94	0	922	153
San Francisco/Oakland	73	0	2,710	283
Seattle	230	0	6,940	839

Notes: SD refers to the standard deviation and the total number of monthly observations is 36,715.

Table 2: Event study: Baseline regression

Dependent variable	Value of imports				
	(1)	(2)	(3)	(4)	(5)
Hurricane	-0.040* [0.016]	0.025 [0.019]	0.004 [0.017]	0.007 [0.017]	-0.002 [0.016]
Hurricane t+1	-0.054** [0.018]	-0.060*** [0.015]	-0.058*** [0.014]	-0.059** [0.021]	-0.054*** [0.014]
Hurricane t+2	-0.020 [0.019]	-0.018 [0.016]	-0.031 [0.016]	-0.018 [0.025]	-0.023 [0.017]
Hurricane t+3	0.020 [0.019]	0.020 [0.018]	-0.009 [0.017]	0.015 [0.027]	-0.002 [0.018]
Hurricane t+4	0.039 [0.021]	0.039 [0.021]	0.007 [0.022]	0.030 [0.023]	0.007 [0.022]
Hurricane t+5	0.029 [0.016]	0.028 [0.016]	-0.001 [0.017]	0.026* [0.012]	0.001 [0.017]
Hurricane t-1	0.000 [.]	0.000 [.]	0.000 [.]	0.000 [.]	0.000 [.]
Hurricane t-2	0.004 [0.014]	-0.006 [0.015]	-0.005 [0.013]	-0.013 [0.016]	-0.012 [0.012]
Hurricane t-3	0.003 [0.018]	-0.007 [0.020]	-0.016 [0.018]	-0.015 [0.023]	-0.016 [0.018]
Hurricane t-4	-0.004 [0.021]	-0.012 [0.024]	-0.033 [0.022]	-0.017 [0.027]	-0.033 [0.022]
Hurricane t-5	0.005 [0.015]	0.003 [0.016]	-0.012 [0.014]	0.006 [0.018]	-0.008 [0.015]
ln(exchange rate)			-0.043 [0.161]	-0.301 [0.163]	-0.025 [0.170]
ln(total exports)			0.619*** [0.126]		0.583*** [0.126]
ln(imports by air)				0.086** [0.031]	0.019 [0.024]
Lead and lags	Yes	Yes	Yes	Yes	Yes
Exporter, port, ship-type, month FE	Yes	Yes	Yes	Yes	Yes
Port, ship-type, time FE	Yes	Yes	Yes	Yes	Yes
Dummy for adjacent hurricanes	No	Yes	Yes	Yes	Yes
Observations	26563	26563	26563	24438	24438

Notes: Estimation method is Poisson Pseudo Maximum Likelihood (PPML). All regressions include lead and lag dummy variables for the 5 months prior and post a hurricane event. Robust standard errors in parentheses: ***, **, * indicate the statistically significant difference from zero at the 0.1, 1 and 5 percent levels respectively.

Table 3: Event study: Baseline regression with sample restrictions

Dependent variable	Value of imports	
	(1)	(2)
Hurricane	-0.001 [0.016]	-0.003 [0.014]
Hurricane t+1	-0.056*** [0.014]	-0.062*** [0.013]
Hurricane t+2	-0.026 [0.017]	-0.025 [0.015]
Hurricane t+3	-0.005 [0.018]	-0.005 [0.016]
Hurricane t+4	0.003 [0.022]	0.005 [0.021]
Hurricane t+5	-0.003 [0.017]	-0.004 [0.016]
Hurricane t-1	0.000 [.]	0.000 [.]
Hurricane t-2	-0.012 [0.012]	-0.016 [0.011]
Hurricane t-3	-0.018 [0.018]	-0.025 [0.016]
Hurricane t-4	-0.034 [0.022]	-0.042* [0.021]
Hurricane t-5	-0.009 [0.015]	-0.017 [0.013]
ln(exchange rate)	-0.035 [0.169]	-0.010 [0.163]
ln(total exports)	0.588*** [0.124]	0.564*** [0.117]
ln(imports by air)	0.017 [0.024]	0.029 [0.022]
Lead and lags	Yes	Yes
Exporter, port, ship-type, month FE	Yes	Yes
Port, ship-type, time FE	Yes	Yes
Dummy for adjacent hurricanes	Yes	Yes
Sample	Non-zero annual trade	Non-zero trade past 3 months
Observations	21679	12743

Notes: Estimation method is Poisson Pseudo Maximum Likelihood (PPML). All regressions include lead and lag dummy variables for the 5 months prior and post a hurricane event. Robust standard errors in parentheses: ***, **, * indicate the statistically significant difference from zero at the 0.1, 1 and 5 percent levels respectively.

Table 4: Event study: Distance regression

Dependent variable	Value of imports
	(1)
Hurricane	-0.001 [0.002]
Hurricane t+1	-0.007*** [0.002]
Hurricane t+2	-0.004 [0.002]
Hurricane t+3	-0.001 [0.002]
Hurricane t+4	0.000 [0.002]
Hurricane t+5	-0.000 [0.002]
Hurricane t-1	0.000 [.]
Hurricane t-2	-0.000 [0.001]
Hurricane t-3	-0.001 [0.002]
Hurricane t-4	-0.003 [0.003]
Hurricane t-5	-0.001 [0.002]
ln(exchange rate)	-0.033 [0.175]
ln(total exports)	0.599*** [0.135]
ln(imports by air)	0.008 [0.024]
Lead and lags	Yes
Exporter, port, ship-type, month FE	Yes
Port, ship-type, time FE	Yes
Observations	23156

Notes: Estimation method is Poisson Pseudo Maximum Likelihood (PPML). All regressions include lead and lag dummy variables for the 5 months prior and post a hurricane event. Robust standard errors in parentheses: ***, **, * indicate the statistically significant difference from zero at the 0.1, 1 and 5 percent levels respectively.

Table 5: Local projection: Baseline regression

	(a) with $\ln(\text{total exports})$					
	(0)	(1)	(2)	(3)	(4)	(5)
β_{-q}	0.033	-0.160***	-0.121	-0.027	-0.071	0.099
SE	0.403	0.031	0.089	0.054	0.066	0.064

	(b) including imports by air					
	(0)	(1)	(2)	(3)	(4)	(5)
β_{-q}	0.033	-0.163***	-0.119	-0.025	-0.070	0.102
SE	0.404	0.030	0.092	0.054	0.064	0.063

Notes: Estimation method is local projection with 12 lags of the dependent variable (DV) and the control variables. The control variables are logarithm of total exports, logarithm of the bilateral exchange rate and inverse hyperbolic sine of the value of imports by air. All regressions include monthly fixed effects that vary by US port, exporting country and ship type. Robust standard errors clustered by type of ship in parentheses: ***, **, * indicate the statistically significant difference from zero at the 0.1, 1 and 5 percent levels respectively.

Table 6: Local projection: Baseline regression with positive annual trade

	(0)	(1)	(2)	(3)	(4)	(5)
β_{-q}	0.015	-0.225***	-0.203	-0.018	-0.131	0.205
SE	0.518	0.056	0.140	0.119	0.097	0.082

Notes: Estimation method is local projection with 12 lags of the dependent variable (DV) and the control variables. The control variables are logarithm of total exports and logarithm of the bilateral exchange rate. All regressions include monthly fixed effects that vary by US port, exporting country and ship type. Robust standard errors clustered by type of ship in parentheses: ***, **, * indicate the statistically significant difference from zero at the 0.1, 1 and 5 percent levels respectively.

Table 7: Local projection: Distance regression

	(0)	(1)	(2)	(3)	(4)	(5)
β_{-q}	0.004	-0.018***	-0.012	-0.002	-0.006	0.012
SE	0.045	0.004	0.010	0.005	0.007	0.007

Notes: Estimation method is local projection with 12 lags of the dependent variable (DV) and the control variables. The control variables are logarithm of total exports and logarithm of the bilateral exchange rate. All regressions include monthly fixed effects that vary by US port, exporting country and ship type. Robust standard errors clustered by type of ship in parentheses: ***, **, * indicate the statistically significant difference from zero at the 0.1, 1 and 5 percent levels respectively.

Table 8: Local projection: Rerouting to different US port

	(0)	(1)	(2)	(3)	(4)	(5)
β_{-q}	-0.008	-0.019	0.141	0.236	0.149	0.021
SE	0.159	0.036	0.179	0.174	0.130	0.226

Notes: Estimation method is local projection with 12 lags of the dependent variable (DV) and the control variables. The control variables are logarithm of total exports and logarithm of the bilateral exchange rate. All regressions include monthly fixed effects that vary by US port, exporting country and ship type. Robust standard errors clustered by type of ship in parentheses: ***, **, * indicate the statistically significant difference from zero at the 0.1, 1 and 5 percent levels respectively.

Table 9: Local projection: Rerouting to different exporting port

	(0)	(1)	(2)	(3)	(4)	(5)
β_q	-0.303	0.088	-0.048	0.363	-0.398	0.217
SE	0.164	0.073	0.068	0.187	0.158	0.153

Notes: Estimation method is local projection with 12 lags of the dependent variable (DV) and the control variables. The control variables are logarithm of total exports and logarithm of the bilateral exchange rate. All regressions include monthly fixed effects that vary by US port, exporting country and ship type. Robust standard errors clustered by type of ship in parentheses: ***, **, * indicate the statistically significant difference from zero at the 0.1, 1 and 5 percent levels respectively.

Table 10: Exports: Event study

Dependent variable	Value of exports					
	(1)	(2)	(3)	(4)	(5)	(6)
Hurricane	-0.006 [0.020]	-0.009 [0.021]	0.013 [0.023]	-0.008 [0.021]	0.010 [0.022]	0.011 [0.022]
ln(exports by air)	0.011 [0.013]			0.011 [0.013]		0.015 [0.015]
Hurricane t+1	-0.004 [0.017]	-0.007 [0.022]	0.015 [0.012]	-0.007 [0.023]	0.011 [0.017]	0.012 [0.017]
Hurricane t+2	-0.022 [0.020]	-0.025 [0.023]	-0.010 [0.025]	-0.025 [0.023]	-0.014 [0.027]	-0.013 [0.027]
Hurricane t+3	-0.025 [0.019]	-0.027 [0.022]	-0.014 [0.021]	-0.027 [0.022]	-0.018 [0.024]	-0.018 [0.024]
Hurricane t+4	-0.007 [0.012]	-0.009 [0.014]	0.003 [0.011]	-0.010 [0.014]	-0.000 [0.012]	-0.000 [0.012]
Hurricane t+5	0.023 [0.025]	0.021 [0.025]	0.030 [0.025]	0.021 [0.025]	0.027 [0.025]	0.028 [0.025]
Hurricane t-1	0.000 [.]	0.000 [.]	0.000 [.]	0.000 [.]	0.000 [.]	0.000 [.]
Hurricane t-2	-0.027 [0.021]	-0.029 [0.022]	-0.013 [0.018]	-0.029 [0.023]	-0.015 [0.019]	-0.015 [0.019]
Hurricane t-3	-0.026* [0.015]	-0.028* [0.017]	-0.011 [0.015]	-0.028 [0.017]	-0.014 [0.016]	-0.014 [0.016]
Hurricane t-4	-0.009 [0.018]	-0.011 [0.019]	0.006 [0.016]	-0.011 [0.019]	0.004 [0.017]	0.004 [0.017]
Hurricane t-5	-0.024 [0.017]	-0.026 [0.018]	-0.012 [0.017]	-0.026 [0.018]	-0.014 [0.017]	-0.014 [0.017]
ln(exchange rate)		0.058 [0.188]		0.056 [0.187]	0.081 [0.179]	0.078 [0.178]
ln(imports ROW)			0.363** [0.171]		0.365** [0.166]	0.367** [0.166]
Lead and lags	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Ctry., port, ship, month FE	Yes	Yes	Yes	Yes	Yes	Yes
Port, ship, time FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32140	32140	32140	32140	32140	32140
Pseudo R^2	0.95	0.95	0.95	0.95	0.95	0.95

Notes: Estimation method is Poisson Pseudo Maximum Likelihood (PPML). All regressions include lead and lag dummy variables for the 5 month prior and post a hurricane event. Robust standard errors clustered at the port-country level in parentheses: ***, **, * indicate the statistically significant difference from zero at the 0.1, 1 and 5 percent levels respectively.


Table 11: Exports: Local projection

	(0)	(1)	(2)	(3)	(4)	(5)
β_q	0.164	-0.010	0.146	-0.128	0.032	-0.188
SE	0.091	0.304	0.092	0.056	0.081	0.145

Notes: Estimation method is local projection with 12 lags of the dependent variable (DV) and the control variables. The control variables are the logarithm of total imports excluding imports on that route, logarithm of the bilateral exchange rate, the inverse hyperbolic sine of exports by air for that route and a dummy for positive route-month trade. All regressions include monthly fixed effects that vary by US port, importing country and ship type. Robust standard errors clustered by type of ship in parentheses: ***, **, * indicate the statistically significant difference from zero at the 0.1, 1 and 5 percent levels respectively.

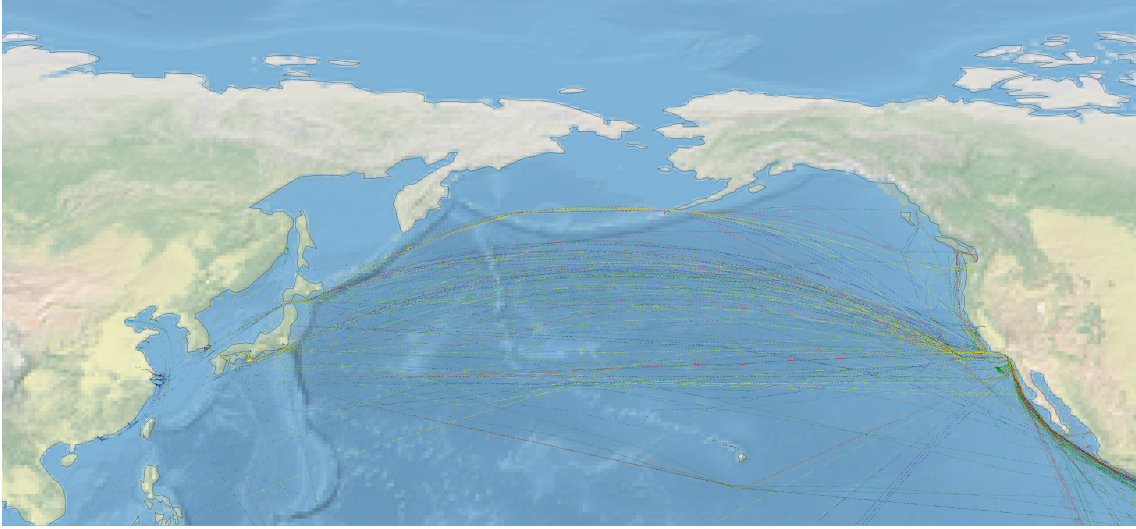
9.2 Figures

Figure 1: All ship crossings for the port of LA



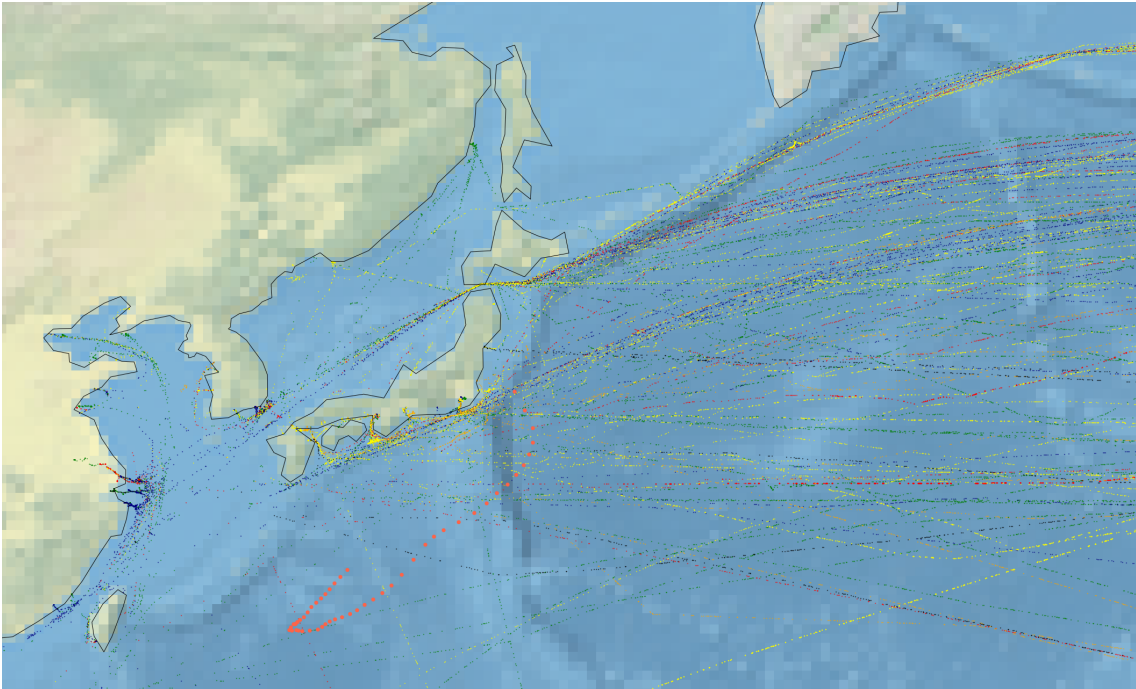
Note: The line colours represent different ship types. Green is oil and chemical tanker; red is general cargo; blue is container; yellow is bulk carrier; orange is roll-on roll-off, and black is gas tanker.

Figure 2: Optimal routes for the port of LA



Note: The lines represent the optimal routes, defined as the 95th percentile of the route-time distribution. The line colours represent different ship types. Green is oil and chemical tanker; red is general cargo; blue is container; yellow is bulk carrier; orange is roll-on roll-off, and black is gas tanker.

Figure 3: Typhoon Lionrock, late August, 2016



Note: The path of Typhoon Lionrock is represented by the thin bright red dots located south of Japan. Typhoon Lionrock occurred between August 17th and 30th in 2016 and was a category 4 storm on the Saffir-Simpson scale. The remaining lines are the optimal ship routes from Figure 2. The line colours represent different ship types. Green is oil and chemical tanker; red is general cargo; blue is container; yellow is bulk carrier; orange is roll-on roll-off, and black is gas tanker.

Figure 4: Seasonal pattern of hurricanes

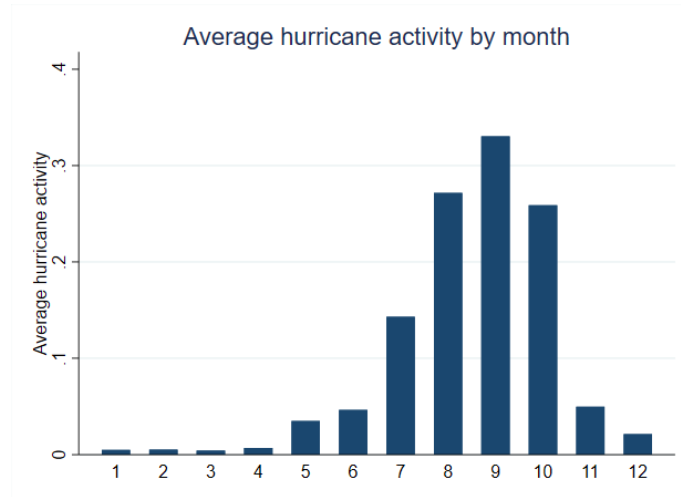


Figure 5: Average number of hurricanes per route by country in a given month




Figure 6: Seasonal pattern of trade

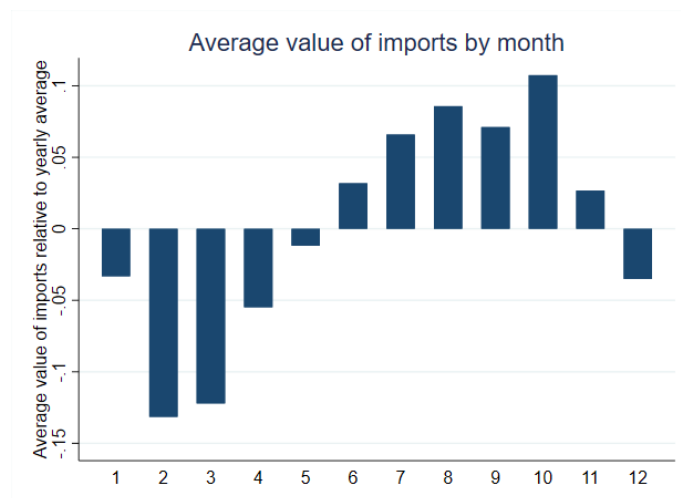
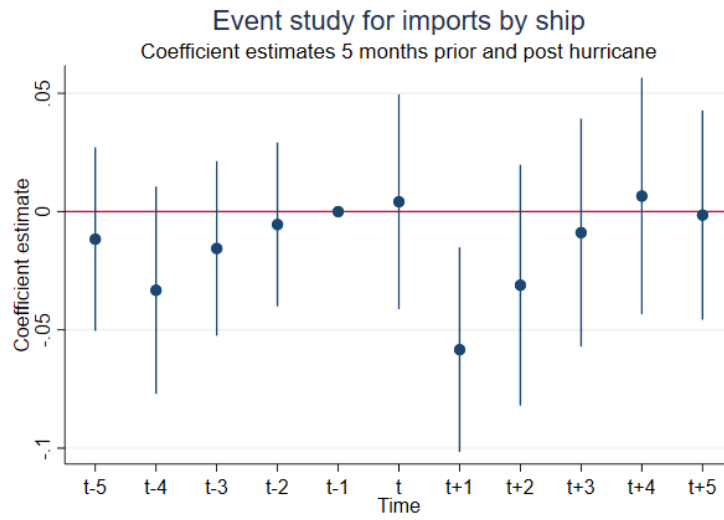


Figure 8: Event study: Baseline

(a) Baseline: Total imports (without controls)



(b) Baseline: Total imports (with controls)

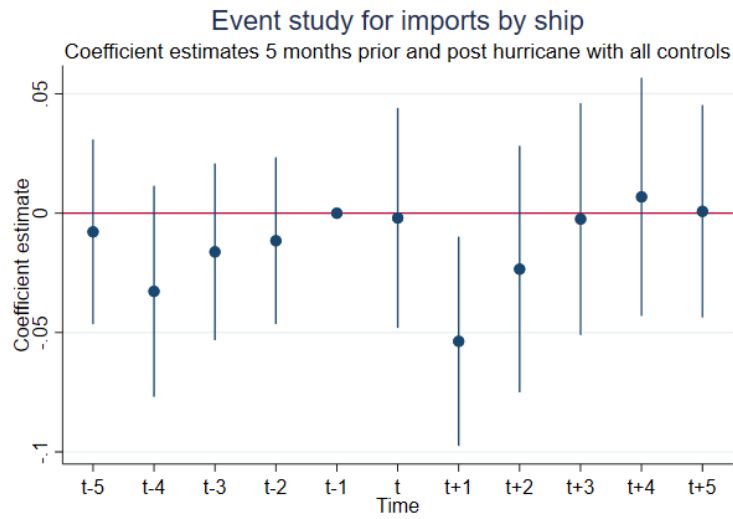



Figure 9: Event study: Zeros robustness

(a) Positive trade in calendar year



(b) Positive trade in past 3 months

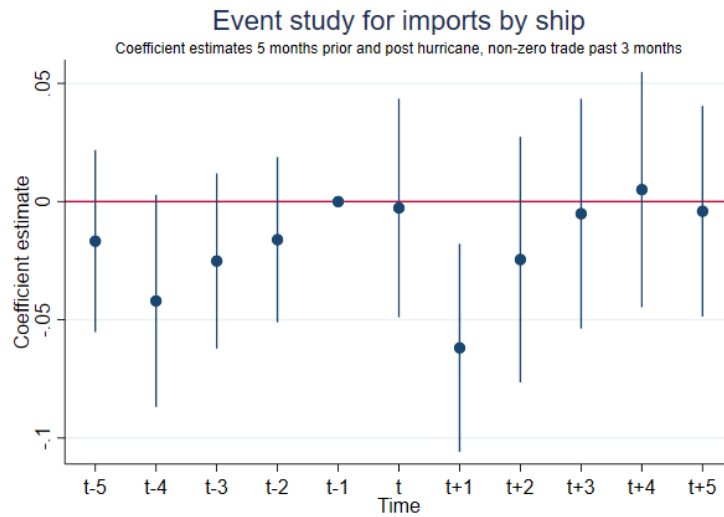


Figure 10: Event study: Heterogeneous effects of route distance

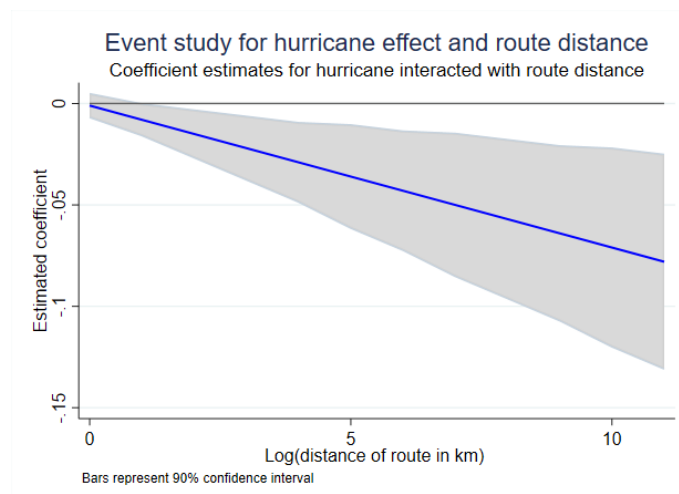


Figure 11: Exports

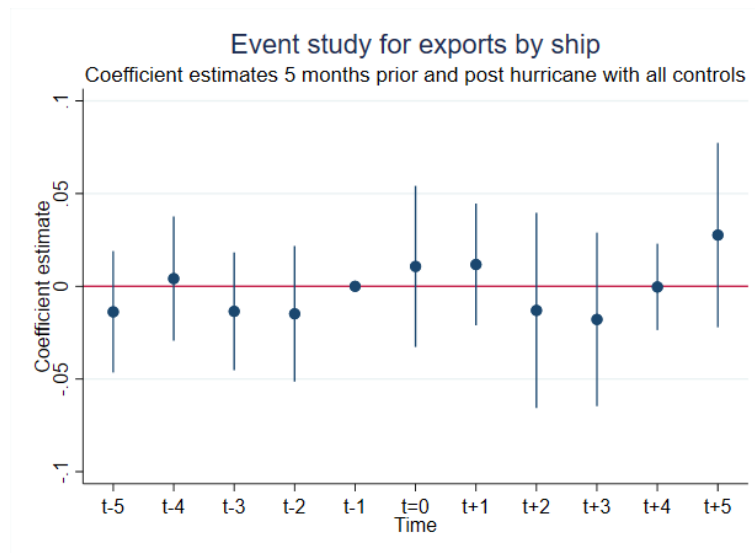



Figure 12: Local projection: Baseline

(a) Baseline: Total exports



(b) Baseline: Imports by air

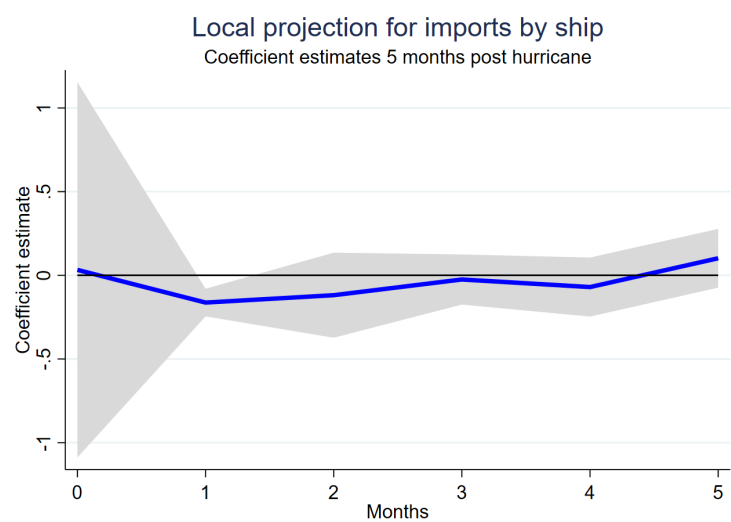


Figure 13: Local projection: Zeros




Figure 14: Local projection: Heterogeneous affects of route distance




Figure 15: Local projection: Rerouting across US ports




Figure 16: Local projection: Rerouting across exporter country ports

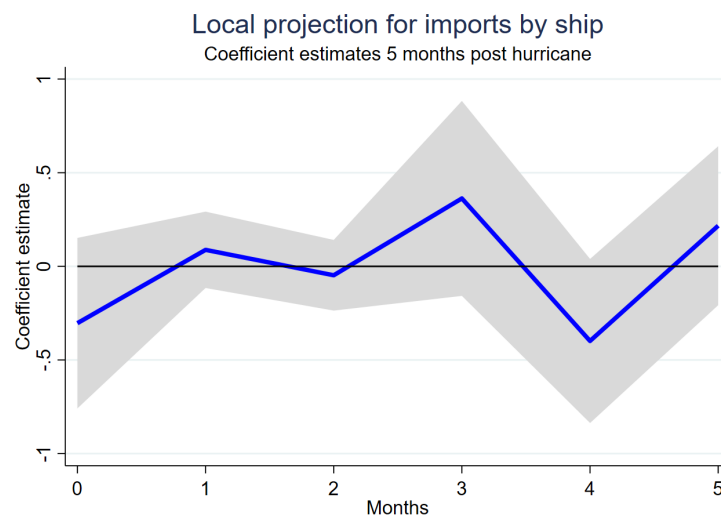


Figure 17: Local projection: Exports


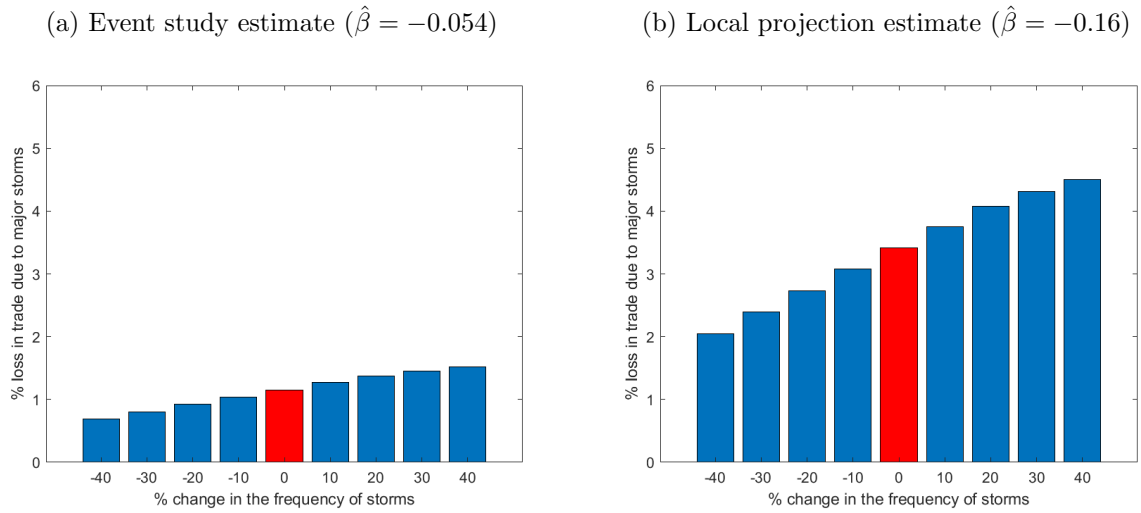
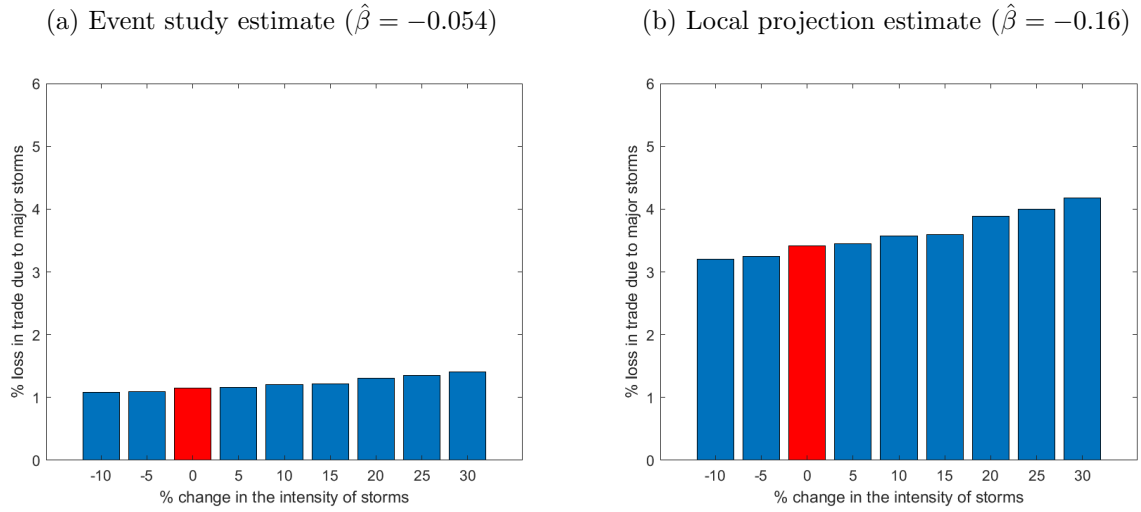


Figure 18: Independent frequency simulations




Note: The red bar indicates the baseline simulation with no frequency adjustment.

Figure 19: Independent intensity simulations




Note: The red bar indicates the baseline simulation with no intensity adjustment.

Figure 20: Joint frequency and intensity interactions (no poleward migration)



Note: The flat blue plane represents the baseline cost estimate without any frequency, intensity or poleward migration adjustments (i.e. no climate change effects).

Figure 21: Joint frequency and intensity interactions with 1° poleward migration



Note: The flat blue plane represents the baseline cost estimate without any frequency, intensity or poleward migration adjustments (i.e. no climate change effects). The bottom coloured plane is the frequency and intensity interactions *without* the poleward migration adjustment.

10 Supplementary Appendix

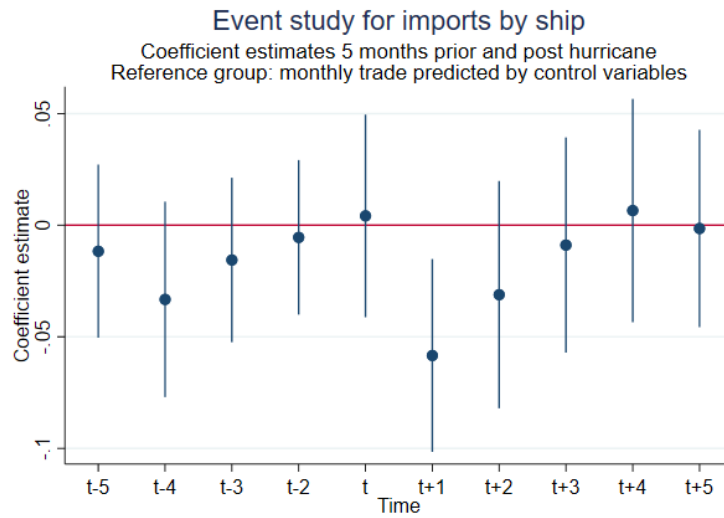
Table 12: Event study additional regressions

Dependent variable	Value of imports			
	(1)	(2)	(3)	(4)
Hurricane	0.801*** [0.110]	-0.056 [0.034]	-0.078** [0.036]	-0.040* [0.021]
Hurricane t+1			-0.054 [0.040]	-0.054** [0.025]
Hurricane t+2			-0.007 [0.041]	-0.020 [0.030]
Hurricane t+3			0.040 [0.052]	0.020 [0.030]
Hurricane t+4			0.085 [0.056]	0.039 [0.031]
Hurricane t+5			0.088 [0.046]	0.029 [0.028]
Hurricane t-1			0.000 [.]	0.000 [.]
Hurricane t-2			-0.004 [0.042]	0.004 [0.022]
Hurricane t-3			-0.009 [0.043]	0.003 [0.022]
Hurricane t-4			-0.033 [0.047]	-0.004 [0.025]
Hurricane t-5			-0.048 [0.047]	0.005 [0.022]
Lead and lags	No	No	Yes	Yes
Exporter, port, ship-type, month FE	No	Yes	Yes	Yes
Port, ship-type, time FE	No	No	No	Yes
Dummy for adjacent hurricanes	No	No	No	No
Control variables, $X_{i,j,s,t}$	No	No	No	No
Observations	31323	27145	27145	26563

Notes: Estimation method is Poisson Pseudo Maximum Likelihood (PPML). Robust standard errors in parentheses: ***, **, * indicate the statistically significant difference from zero at the 0.1, 1 and 5 percent levels respectively.

Figure 22: Event study: Redefining control group to routes that never experience a hurricane crossing

(a) Baseline: Total imports (without controls)



(b) Baseline: Total imports (with controls)

