

Online Appendix

Entrepôt: Hubs, Scale, and Trade Costs

Sharat Ganapati, Woan Foong Wong, and Oren Ziv

Appendix A Data Construction

A.1 Shipment Microdata

We compile and combine two proprietary microdata sets in this project: global ports of call data for all containerships, which allows us to reconstruct the routes taken by specific ships, and United States bill of lading data for containerized imports, which gives us shipment-level data on imports into the United States. Independently, each of these datasets allows us to partially describe the global shipping network. By merging them, we are able to reconstruct nearly the entire journey most shipments entering the United States take, from their initial origin point or place of receipt to the port of entry into the United States. To our knowledge, we provide the most comprehensive reconstruction of the global shipping network and routes undertaken by individual shipments into the United States (Panjiva, 2014; Astra Paging, 2014; CEPII, 2017; KGM Associates, 2014).

Port of call data We partner with Astra Paging, which provides us with the port of call data for containerships. Astra Paging’s data captures vessel movements using the transponders on these ships (known as the automatic identification system, AIS). A network of receivers at ports collects and shares AIS transponder information (including ship name, speed, height in water, latitude, and longitude). Using the geographic variables in the AIS data, Astra Paging marks entry and exit into a number of ports all over the world and provides us with a dataset of ships’ entry and exit from ports of call, timestamps, and ships’ height in the water, or draft. Using these data elements, we are able to calculate an estimated shipment volume between each port pair by taking the observed draft relative to the maximum observed draft and multiplying by total ship capacity.

Our sample covers a six-month period, from April to October 2014. Over this period, we have information on 4,986 unique container ships with a combined capacity of 30.6 million TEU. This represents over 90% of the global container shipping fleet. Ports with no AIS receivers or where information is not shared do not show up in our data. In addition, if transponders are turned off or transmissions are not recorded, ports of call can be missed. However, transponders are required to be operational by the International

Maritime Organization on ships engaging in international voyages 300 gross tons, applying to all containerships in our sample (International Maritime Authority, 2003).

Bill of lading data We partner with Panjiva Inc. (Now a division of Standard and Poor’s) to acquire bill of lading information for all seaborne US imports from April to October 2014. Panjiva cleans this data to standardize the names of the ports, ships, companies, and container volumes. We subset this data to only consider goods that arrive on seaborne container ships.

International shipping relies on an industry-standardized system of bills of lading, which act as receipts of shipment, recording all information on the shipment and all the parties involved in the shipping process. The US Customs and Border Patrol (CBP) agency collects these bills in addition to customs information at all ports of entry into the US and this data is obtained from the agency by Panjiva.¹

Over six months of US imports from April to October 2014, we see a total of 14.8 million TEUs weighing 106 million tons were imported into the US from 221 shipment origin countries and 144 countries with ports of lading. This accounts for about three quarters of the 2014 TEU and tonnage imports, 77 percent and 74 percent respectively (Maritime Administration, US Department of Transportation, 2014).² The countries in our data are categorized using the three-digit alphabetical codes assigned by the International Organization for Standardization (ISO) by the Statistics Division of the United Nations Secretariat. Accordingly, Hong Kong, Macau, Taiwan, as well as dependencies and areas of special sovereignty like Guam have their own designated codes. Non-containerized goods, including goods on roll-ons (vehicle carriers), bulk cargo liners (for commodities), and non-containerized cargo ships are not observed in our data.

Our data captures the following location information for each shipment into the US: the foreign location where the shipment originated from (*shipment origin*), the foreign port where it was loaded on the containership which brings it into the US (*port of lading*), and the US port where it was unloaded from the containership (*port of unloading*). In addition, we know the name and identification number of the containership (IMOs) which transported the shipment as well as the shipment’s weight, number of containers (TEUs), and product information.

This data set allows us to start tracing the journey of a shipment from its origin to its

¹US Bill of Lading data is immediately available for direct purchase from the Department of Homeland Security or through a lag using a Freedom of Information Act. However, this raw data requires substantial computing resources for processing and needs to be standardized over time.

²In particular, we miss containers that arrive on trucks and trains from either Mexico or Canada. Our estimation strategy explicitly accounts for this unobserved data.

destination US port. In particular, we can determine whether this shipment was loaded at its origin location onto the vessel that brings it directly to its final US destination, or if it went through at least one other location during its journey. When matched with the port of call data, we can reconstruct most of its remaining journey after the port where it was loaded onto a US-bound vessel (from its port of lading).

Reconstructing shipment routes Using the containership information, port of arrival information, timing of unloading and ports of call at US ports, and port of lading information, we are able to match the bills of lading to the journeys of specific containerships, then use the ports of call between lading and unloading to reconstruct each shipment’s path from its foreign origin to US destination.

First, we identify containerships using Vessel IMOs. Vessel IMOs are identifiers unique to containership vessels and stay constant for the lifetime of their operation. By IMO, we identify about 4000 ships in the Bills of lading data. An additional (roughly) 2,000 ships are matched to IMOs using a fuzzy string match, after which matches are made by hand with the help of undergraduate research assistants.

Second, we match the port calls that the containerships make with the ports of arrival of shipments. Ports of arrival are recorded using UNLOCODEs in the AIS port of call data and US Census Schedule D codes in the Bills of Lading data. We construct a crosswalk to match these ports with the help of undergraduate research assistants.

Third, we match the port calls that the containerships make with the ports of lading of shipments. Ports of lading are recorded using UNLOCODES in the AIS port of call data and the US Customs and Border Protection’s listing of foreign ports (Schedule K) in the Bills of Lading data. We construct a crosswalk to match these ports with the help of undergraduate research assistants.

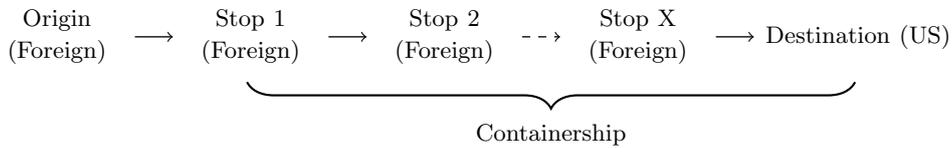
What remains unobserved is the shipment’s journey between its Origin and its first stop (port of lading location). This portion of the shipment’s journey takes place in a container and could be transported overland (by trucks or rail) or by sea on another containership. While this information is not recorded by both our datasets and therefore unobserved, the amount of indirectness that we establish in our stylized facts is a lower bound since we assume that this portion of the journey is direct. The amount of transshipment that we establish in our stylized facts is also a lower bound since at most we observe one transshipment port. To the best of our knowledge, we capture the most detailed information on shipments’ journeys by merging these two datasets.

For each bill of lading, we match ship, date of unloading, and port of unloading to

the AIS data on ships’ port of call. Once we match shipments to ships, we record each port of call in the AIS data before the port of unloading as a stop the shipment makes, then remove all stops observed before the ship stopped at the port of lading. If the port of lading is not observed, the route is discarded and the shipment remains unmatched. Furthermore, any routes that include the port of unloading before the date of unloading are discarded, as they represent loops where the port of call for the port of lading is missing.

Over 90% of containerized TEUs entering the US in the bills of lading data can be matched to routes using this method.³ Appendix Figure A.1 visualizes this merge.

Figure A.1: Combined Dataset: Routes Undertaken by Shipments into the US



Notes: Origin is the foreign location where the shipment originated from, Stop 1 is the location where the shipment was loaded on its US-bound containership (also known as the port/location of lading), Stop 2 to Stop X are the subsequent stops that the US-bound containership made while the shipment remains on the ship, and Destination is the US port where the shipment was unloaded from containership.

As an example, Figure A.2 plots for all containerized trade from the United Arab Emirates (UAE), the proportion that stops in each country. This illustrates the paths shipments take when being transported from the UAE on to the US. Shipments from the UAE collectively stop in many countries before continuing onto the US. Many of the most popular are regional neighbor hubs, including Egypt, Pakistan, but Spain and China also facilitate UAE-US trade.

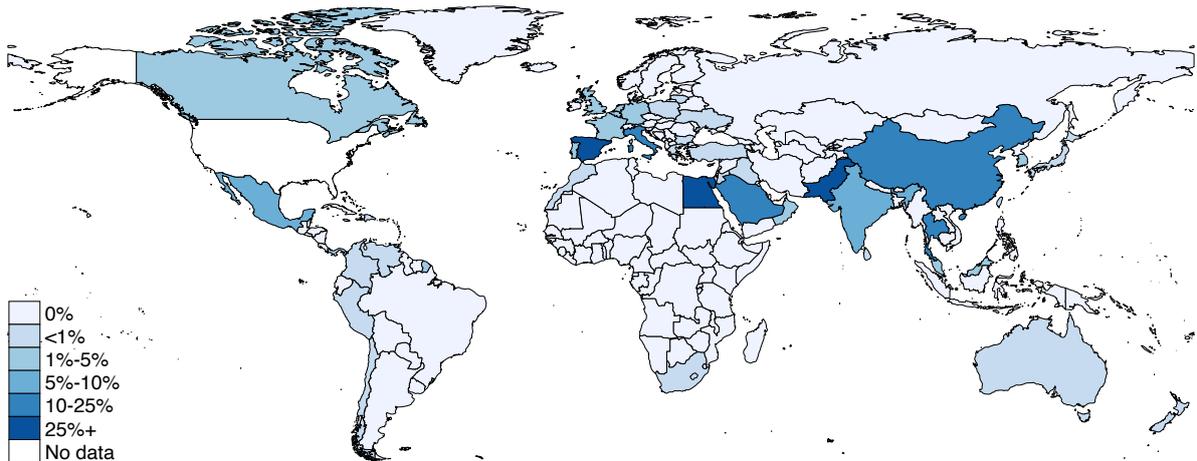
A.2 Geographic Distance Data

Geographic distance data is computed using two rasterized (with pixels) world maps. One map consists of all the navigable oceans and large seas, with a polar ice cap, as well as the Suez and Panama canals. The second map assumes that the Arctic ice sheet melts away due to anthropogenic climate change. In both maps, we compute the sea distance between ports of call, and aggregate to the national level using port-to-port container flows. We do this computation in R using Dijkstra’s algorithm on a world map with and without Arctic ice caps.⁴ We argue this with distance data from Bertoli, Goujon

³Unmatched shipments may have missing and unrecoverable ship information, or ports of call that do not match lading and unloading records on bills of lading. In addition, a small number of reconstructed routes have implicit voyage speeds above 50KPH, and are discarded.

⁴For more information, see the ‘gdistance’ package and mapping files from Kelso and Patterson (2010).

Figure A.2: Percent of UAE-US trade that stops in each country



Notes: Each country’s color represents the share of shipments from the UAE to the United States that stop in that country. Stops computed at the country level and weighted by total container volume (TEU). The United States and the UAE are denoted in white.

and Santoni (2016) and Conte et al. (2021), as well a data on landlocked countries from Encyclopedia Britannica (2022).

A.3 Aggregate Economic and Trade Statistics

For our main estimation, we also require data on the value of containerized trade between countries. We use aggregate trade data from Centre d’études Prospectives et d’Informations Internationales (CEPII) and their BACI international database for 2014. This database aggregates data from the UN Comtrade Database, aligning data from origin and destination countries. This provides us data on trade volumes from origin to destinations by industry using Harmonized System (HS) codes.

To aggregate industry trade to industries that use container shipments versus trade that does not, we use aggregate data from 2014 from the United State Customs, as disseminated by Schott (2008).⁵ This data reports the share of shipments by HS Codes that arrive by containerships. We consider 4-digit HS Codes as a consistent level of aggregation. The distribution of containership share by HS code is bi-modal, with one peak around 0% and another around 100%. We use a cutoff of 80%. So HS codes that are shipped by containership to the US over 80% of the time are classified as “containerizable” trade.

For aggregate trade and economic statistics for using in the counterfactual, we use the Eora global supply chain database with a multi-region input-output table (EORA-

⁵This data has been continually updated by the author following the initial publication

MRIO).⁶ We collapse all world trade into three categories; those that are non-tradable, those that are typically traded over oceans by containerized vessels, and those that are not typically traded over oceans by containerized vessels.⁷ We again classify industries using the methods of Schott (2008). We augment this with GDP data from Feenstra, Inklaar and Timmer (2019); World Bank (2018); OECD (2018).

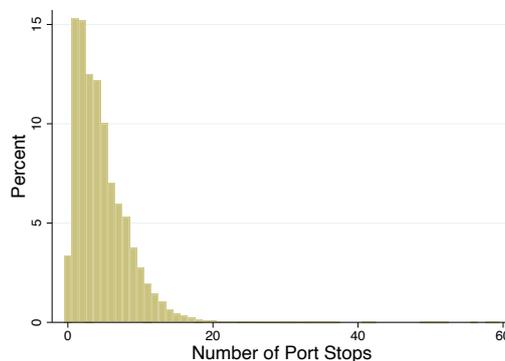
Appendix B Additional Descriptive Results

In this section we report additional results and robustness checks related to the analysis in Section 3.

B.1 Additional Indirectness Results

Figure A.3 reports the histogram of number of port stops minus the port of lading if the port of lading is in the country of origin, and the port of unloading. We exclude landlocked countries. The mean number of third-port stops is 4.6 and fewer than 5% of shipments do not stop at additional ports.

Figure A.3: Distribution of Port Stops per Container (TEU)



Notes: This figure reports the distribution at the shipment level of the number of unique port stops minus the port of lading if the port of lading is in the country of origin, and the port of unloading, weighted by shipment TEU. Shipments from landlocked countries are excluded.

Next, in Figure A.4, we rerun the analysis in Panel (A) of Figure 2 weighting by Tons in Panel (A) and USD in Panel (B). For the latter, a minority of shipment data report dollar values. Overall, the results are similar to our main results using TEU.

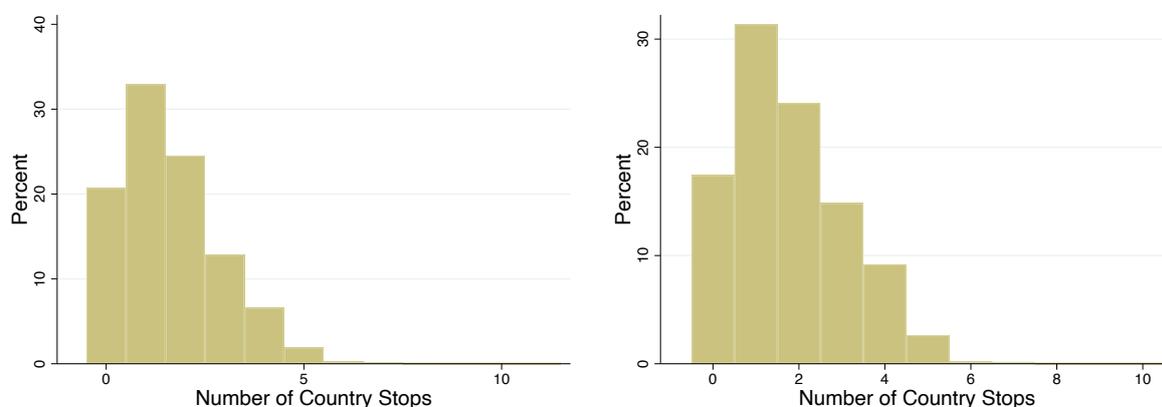
Figure A.5 reports the percent of shipments loaded onto a US-bound ship in a third-party country by country of origin. Countries that are closer and trade more with the US are less likely to transship goods at third-party countries—a fact we explore in more detail in Appendix B.4.

⁶Freely available for academic use from <https://worldmrio.com/>.

⁷This includes bulk shipping, roll-on roll-off ships, as well as air freight.

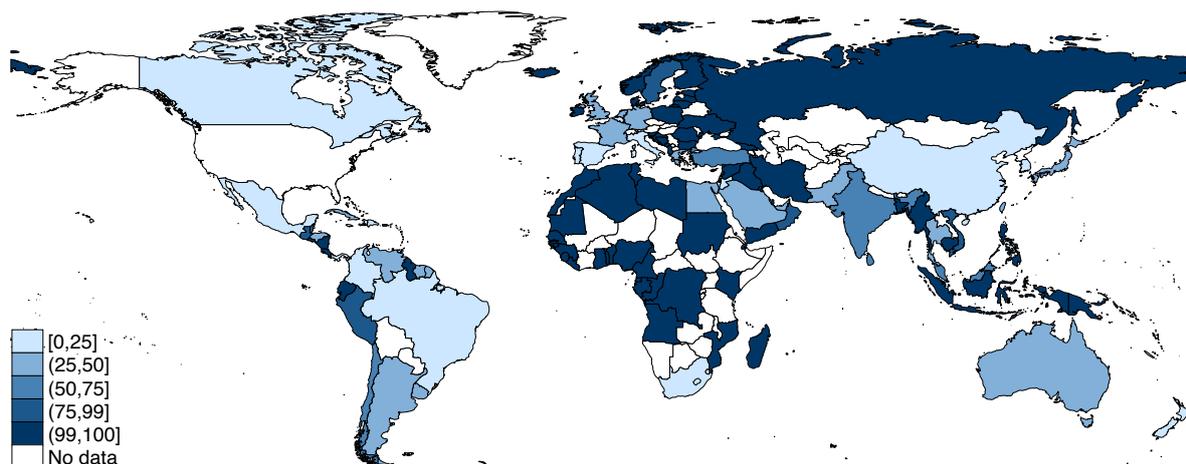
Figure A.4: Distribution of Third-Party Countries Involved in Bilateral Trade by Weight and Value

Panel (A) Number of Countries per Ton Panel (B) Number of Countries per USD Value



Notes: Panel (A) reports the distribution of the total number of unique third-party country stops made by shipments entering the US, weighted by shipment tons (kg). Panel (B) reports the same but weights by value for the portion of shipments for which value measures are reported.

Figure A.5: Transshipped Trade Share between Origin and US Destination



Notes: This figure plots for each country the share of its originated shipments transshipped in a third-party country, weighted by TEU. Lighter colors indicated lower levels of transshipped trade share (ie. more direct trade). The US is not included since it is the destination country. Landlocked countries are also excluded. 34 of the shipment origin countries are landlocked accounting for 1.6 percent of total TEUs. The missing remaining countries are either due to lack of overall trade with the US (e.g. Somalia) or due to the merge process (e.g. Namibia).

Finally, we further explore the result that additional stops increase the distance and time costs of trade. In Table A.1, we regress, at the shipment level, log of observed distance (Columns (1)-(4)) and time (Column (5)), on the number of country stops made by a shipment. All port distances are computed using Dijkstra’s algorithm, and time is computed by the difference in AIS logs for port of lading and unloading. Results are clustered two ways by port of unloading and port of lading.

Column (1) reports the baseline relationship: an elasticity of 0.112 (SE 0.022) on

stops, controlling for the computed direct sea distance between the port of lading and port of unloading. Adding port of lading fixed effects (Column (2)) or port of unloading fixed effects (Column (3)) does not significantly change the result. In Column (4), we add port of lading-by-unloading fixed effects. Here identification comes from variation between routes where goods come on and off boats at exactly the same ports, but where different ships take different routes (the existence of this variation is explored further in Appendix B.4). The elasticity here remains stable as 0.104 (SE 0.03). Column (5) repeats our most heavily controlled-for exercise in Column (4) but with time traveled as the variable of interest. We find an elasticity of 0.333 (SE 0.0819) which implies that for shipments loaded and unloaded at the same ports, routes with double the stops along the way increase journey time by 33%.

Table A.1: The Relationship Between Indirectness, Distance, and Time

	(1)	(2)	(3)	(4)	(5)
	ln Observed Dist	ln Observed Dist	ln Observed Dist	ln Observed Dist	ln Time Travelled
ln Country Stops	0.112 (0.0223)	0.109 (0.0237)	0.101 (0.0270)	0.104 (0.0300)	0.333 (0.0819)
ln Direct Dist	0.881 (0.0276)	0.918 (0.0347)	0.896 (0.0282)		
Lading Port FE		Y			
Unlading Port FE			Y		
Lading-Unlading Ports FE				Y	Y
Observations	215,655	215,655	215,655	215,656	215,656
R^2	.942	.954	.945	.966	.774
F-stat	1360.62	1818.20	1242.46	12.11	16.49

Notes: This table presents regression coefficients for regression of ln Observed Distance, the natural log of sea distance traveled between all reported ports of call, or ln Time Travelled, the natural log of time between port of lading and port of unloading, and ln Country Stops, the natural log of unique third-country stops, as well as ln Direct Distance, the natural log of the sea distance between the port of lading and unloading. Distances are calculated using Dijkstra’s algorithm and measured in kilometers while time is measured in hours. Observations are shipment level and weighted by TEU. Shipments originating in landlocked countries are omitted. Standard errors in parentheses are clustered two ways by the port of lading and port of unloading.

B.2 Additional Concentration Results

List Countries by Entrepôts Activity. Table A.2 reports our index of entrepôt activity for all countries in our data, using data on trade and transportation that have been adjusted as in Section 5 and normalized so that the lowest value (for the US) is zero.

Countries towards the top of the list have more third-country activity, with the 15 countries at the top of this list defined as entrepôts for the purposes of our counterfactual analyses. These include Egypt (Suez Canal) and Singapore. Countries in the middle of the list neither differentially depend on nor are used as third countries. Countries that

are small and/or less open dominate this section like Papua New Guinea or North Korea. Countries towards the end of the list differentially depend on others as third countries (like Ireland and Malaysia), with the bottom of the list dominated by the largest economies, who account for large portions of global trade but not large portions of global traffic (like China and Germany).

Table A.2: Entrepôt Index by Country

Country Name	Index Value	Country Name	Index Value	Country Name	Index Value
Egypt	11.42	Congo	5.77	Ecuador	5.73
Singapore	10.39	Barbados	5.77	Bangladesh	5.72
Netherlands	10.22	Suriname	5.77	Tunisia	5.71
China Hong Kong SAR	8.65	Aruba	5.77	Angola	5.70
Belgium-Luxembourg	7.78	Guinea	5.77	Iraq	5.70
Taiwan	7.26	New Caledonia	5.77	Croatia	5.70
Spain	6.99	Lao Peoples Dem. Rep.	5.77	Qatar	5.69
Saudi Arabia	6.72	Mauritania	5.77	Peru	5.69
Rep. of Korea	6.68	Ghana	5.77	Bulgaria	5.68
United Arab Emirates	6.63	Cyprus	5.77	Viet Nam	5.65
Morocco	6.47	Nicaragua	5.77	Nigeria	5.64
Panama	6.44	Georgia	5.77	Chile	5.62
Malta	6.30	Dem. Peoples Rep. of Korea	5.77	New Zealand	5.60
Portugal	6.17	Madagascar	5.76	Kazakhstan	5.60
United Kingdom	6.09	Albania	5.76	Algeria	5.60
Greece	5.99	Honduras	5.76	Venezuela	5.56
Bahamas	5.94	Lithuania	5.76	Kuwait	5.56
Pakistan	5.90	United Rep. of Tanzania	5.76	Romania	5.55
Israel	5.88	Mauritius	5.76	Malaysia	5.54
Lebanon	5.87	Papua New Guinea	5.76	Finland	5.50
Russian Federation	5.85	Mongolia	5.76	So. African Customs Union	5.50
Jamaica	5.83	Cambodia	5.76	Ukraine	5.49
Uruguay	5.83	Slovenia	5.76	Iran	5.49
Dominican Rep.	5.82	Cameroon	5.76	Poland	5.46
Sri Lanka	5.81	Gabon	5.76	Philippines	5.45
Djibouti	5.79	Brunei Darussalam	5.75	Australia	5.45
Benin	5.78	Côte d'Ivoire	5.75	Argentina	5.44
Senegal	5.78	Guyana	5.75	Indonesia	5.30
Togo	5.78	Trinidad and Tobago	5.75	Brazil	5.30
Colombia	5.77	Belarus	5.75	Denmark	5.24
Gambia	5.77	Yemen	5.75	Ireland	5.23
Liberia	5.77	Iceland	5.75	Thailand	5.23
Somalia	5.77	Latvia	5.75	Norway	5.21
Eritrea	5.77	Paraguay	5.75	Czech Rep.	5.12
Antigua and Barbuda	5.77	Kenya	5.75	Mexico	5.02
Cabo Verde	5.77	Turkey	5.75	Sweden	4.99
Greenland	5.77	Cuba	5.75	Switzerland	4.93
Cayman Isds	5.77	Libya	5.74	France	4.90
Belize	5.77	Guatemala	5.74	India	4.80
Sierra Leone	5.77	Bolivia Plurinational State of	5.74	Austria	4.70
Montenegro	5.77	China Macao SAR	5.74	Italy	4.63
Mozambique	5.77	Syria	5.74	Canada	4.45
Maldives	5.77	Estonia	5.73	China	4.10
Haiti	5.77	Costa Rica	5.73	Japan	4.09
Bahrain	5.77	Oman	5.73	Germany	3.38
				USA	0.0

Notes: Table presents measure of entrepôt activity, calculated, as defined in Section 3, as the percent of global trade minus the percent of global traffic, with adjustments made for overland traffic, with the US normalized to zero.

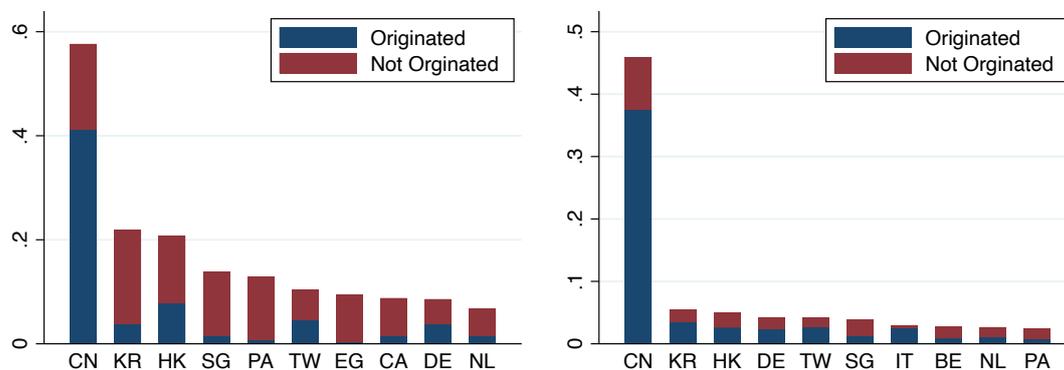
Concentration of US-Bound Shipments Panel (A) of Figure A.6 tabulates, for each of the top ten countries, the percent of all goods entering the US stopping in that

country. The share of shipments accounted for by shipment origination is in blue while shipments observed stopping in the country but not originating in the country is in red. Unsurprisingly, many recognizable entrepôts are listed, including Korea, Panama, Singapore, and Egypt. Perhaps more surprisingly, more than 50% of the containers entering into the US stop in China. While this panel sums to over 1, since each container stops in more than one country, over 80% of shipments to the US stop in at least 1 of 5 countries: China, Panama Singapore, Korea, or Egypt.⁸

Panel (B) replicates Panel (A) but for the country of lading. Here the total of all bars (including those not graphed) sum to 1, and China again dominates as a source of lading. A few of these top countries, like Germany in (A) and Italy in (B) are majority blue, implying they are important to the US because of their role as an origination country. Other countries, like Singapore, are differentially red, and appear important as entrepôt rather than as countries of origin.

Figure A.6: Roles of Countries in Bilateral Trade: Origin vs Entrepôts

(A) Share of Shipments Stopping in Country, for Top Ten Countries (B) Share of Shipments Laded in Country, for Top Ten Countries



Notes: The blue portion in Panel (A) highlights the share of all incoming US shipments that originate in the indicated country while the red accounts for the percent of all incoming US shipments stopping in that country (not originated), weighted by TEU.

Panel (B) replicates Panel (A) but for country of lading.

B.3 Spokes Disproportionately Use Entrepôts

Conceiving the shipping network as a hub and spoke system implies that spokes largely access their trading markets using hubs. While in Section 3 we find the network is characterized by having hubs, we clarify here that the excess concentration of shipments at entrepôts are in part due to their disproportionate use by smaller, less well-connected

⁸Of course, the sum of these five bars is greater than 80% because the average shipment makes multiple stops).

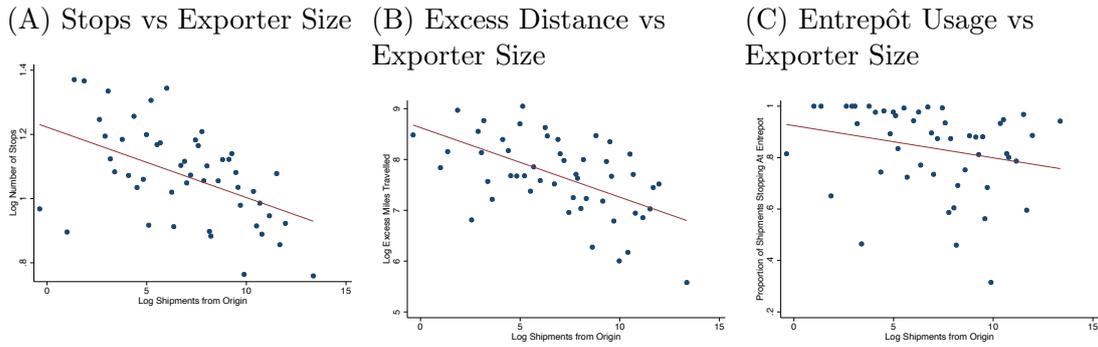
Table A.3: Concentration Ratios

	Third-Party Stops	Transshipment	Trade
Max/50	429	476	400
99/50	390	476	76
95/50	215	135	27
90/50	120	91	15

Notes: Data present concentration ratios across countries in our data. Third-party Stops are the sum total TEU-weighted shipments that use a country as a third-party country. Transshipments are the TEU-weighted sum total of shipments transhipped at a country, and Trade is the total volume of trade from a country. Countries are ranked and percentile ratios are presented. For example, the country used the most (by TEU shipments) as a third-country stop acts as such for 429 times the number of shipments stopping at the median (50th-percentile) country.

origins, or, in other words, the spokes of the network.

Figure A.7: Smaller Exporters Are Disproportionately Indirect and More Likely to Use Entrepôts



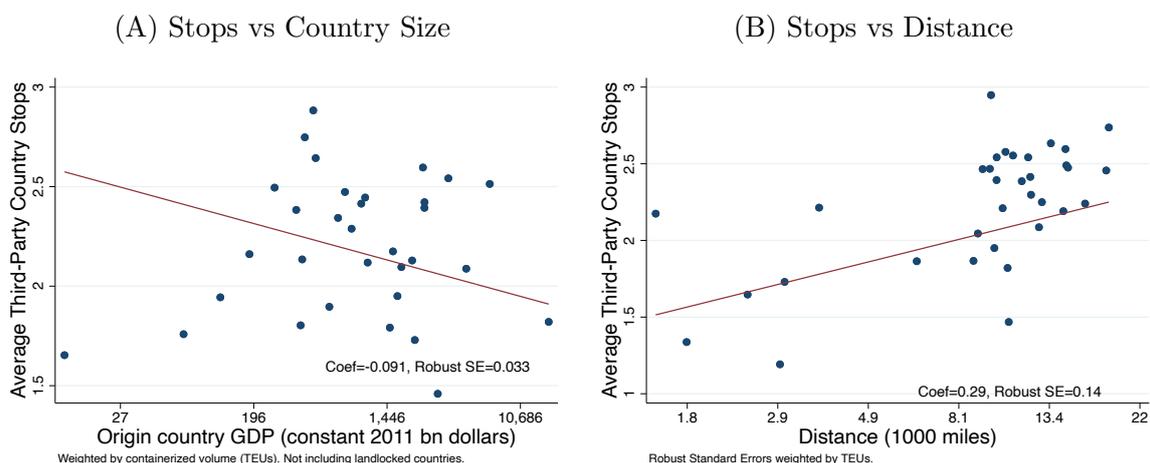
Notes: Binned scatter plots with observation at the origin level weighted by total TEU with 50 bins. The x-axis for each of the three panels features the size of each origin’s exports to the US. Panel (A) shows the relationship between the origin’s size and its average number of stops before its US destination. Panel (B) shows the relationship between the origin’s size and the average excess distance traveled by its exports before its US destination. Panel (C) shows the relationship between the origin’s size and the share of its exports which stopped at an entrepôt before its destination.

In Figure A.7, we zoom in on the set of origins that are simultaneously the most indirect and most likely to send goods through hubs. The three panels are binned scatterplot with 50 bins of origin-level measures of average TEU (A) number of stops, (B) excess distance, and (C) likelihood of passing through an entrepôt. Panel (A) of Figure A.7 confirms that smaller origins are more indirectly connected to the US, and Panel (B) confirms that shipments from smaller origins move further distances to get to their destination. Panel (C) shows that shipments from smaller origins are more likely to use entrepôts. These relationships are echoed in shipment-level regressions which add additional controls and cluster by origin. In sum, the smallest origins constitute the spokes of the hub-and-spoke network.

B.4 Variation in Connectivity

There is a high degree of variance in indirectness across countries, as shown in Figures 2 and A.5. This variation is reasonable explained by traditional gravity variables. In Panel (A) of Figure A.8, we find that countries with higher GDPs are more likely to have less stops on their journeys to the US. In Panel (B), we find that countries which are closer to the US are more likely to have less stops on their journeys (i.e. have more direct trade with the US). These results are robust to using port stops instead of country stops (Table A.4) as well as to weighting by containers, tons, and value. One natural interpretation of this would be the endogenous response of shippers to the scale of shipments from these countries. Of course, the availability of direct trade to the US could in principle reverse the causality.

Figure A.8: Larger and Closer Countries Have Lower Number of Average Stops



Notes: Binned scatter plots with observation at the origin level weighted by total TEU. Landlocked countries are excluded.

Table A.4: Relationship Between Stops and Country Size as Well as Distance

	(1)	(2)	(3)	(4)	(5)	(6)
	ln Ctry Stops	ln Ctry Stops	ln Ctry Stops	ln Port Stops	ln Port Stops	ln Port Stops
ln GDP	-0.0371 (0.0187)		-0.0488 (0.0140)	-0.00226 (0.00966)		-0.00935 (0.00719)
ln Distance		0.166 (0.0851)	0.212 (0.0934)		0.119 (0.0352)	0.128 (0.0384)
Observations	133	133	133	133	133	133
F-stat	3.933	3.795	8.878	0.0546	11.50	5.644
R ²	0.120	0.142	0.339	0.00185	0.305	0.335

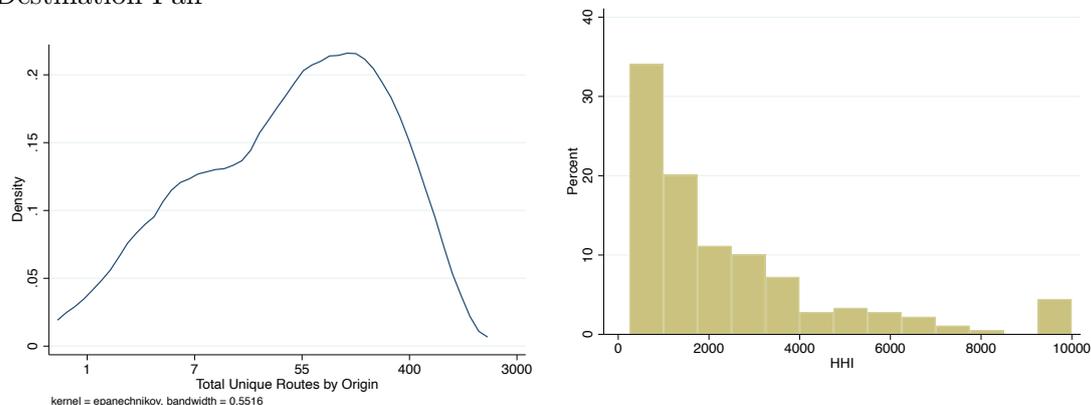
Notes: This table presents coefficients from country-level regression of ln Ctry Stops, the natural log of the TEU-weighted average number of third-party country stops made for shipments from country against ln GDP, the natural log of the country's GDP, and ln Distance, the natural log of the sea distance between the countries. Observations are weighted by total TEU. Landlocked countries are excluded. Robust standard errors in parentheses.

Do shipments from a given origin follow a unique path to the US? Panel (A) in Figure A.9 shows the distribution in the number of unique routes to the US by origin country. With an average of about 397 routes with wide variation (sd 681), observed routes from a single origin are indeed varied. The countries with the highest number of unique routes are big trading partners like China, the United Kingdom, Germany, and well-established entrepôts like Hong Kong. Countries with the lowest unique routes are smaller trading partners like American Samoa, Nauru, Tonga, and Montserrat. The existence of this within-origin route variation will be a particularly important assumption in our model and external validity checks.

We can measure the concentration of these unique routes by constructing a Herfindahl-Hirschman Index (HHI) for each origin country using the container shares of each route. Panel (B) in Figure A.9 shows that almost 70 percent of origin countries have fairly low concentrations of routes (HHI less than 1500). The average HHI overall is 1475 (sd 1974). Examples of countries with high levels of concentration are like Vanuatu, Cuba, and Liberia while countries with low levels of concentration are Macau, Hong Kong, and Belgium-Luxembourg.

Figure A.9: Variation in Trade Indirectness

(A) Number of Unique Routes by Origin-Destination Pair (B) Distribution of Route Concentration



Notes: Panel (A) plots the kernel density plot for the total number of unique routes from a given origin country. Panel (B) plots the distribution of the HHI index for routes from each country.

Appendix C Additional Theoretical Results

C.1 Definition of Entrepot

The share of imports from origin i to destination j in industry n which passes through leg kl is:

$$\pi_{ij}^{kl} = \left[(c_i \kappa_{ij})^{-\theta} \cdot b_{ik} a_{kl} b_{lj} \right] \cdot \Phi_j^{-1}, \quad (14)$$

where $a_{ij} = t_{ij}^{-\theta}$ and $b_{ij} = \tau_{ij}^{-\theta}$.

Summing over shipment origins, we write the share of global shipping to a destination j that goes through kl as follows:

$$\begin{aligned} \pi_j^{kl} &= \sum_i \left[(c_i \kappa_{ij})^{-\theta} \cdot b_{ik} a_{kl} b_{lj} \right] \cdot \Phi_j^{-1} \\ &= \sum_i \left[(c_i \kappa_{ij})^{-\theta} \cdot b_{ik} \right] a_{kl} b_{lj} \cdot \Phi_j^{-1} \\ &= \Phi_k a_{kl} b_{lj} \cdot \Phi_j^{-1} \end{aligned}$$

Summing over all k 's generates the share of traffic to j which flows through node l :

$$\begin{aligned} \pi_j^l &= \sum_k \Phi_k a_{kl} b_{lj} \cdot \Phi_j^{-1} \\ &= b_{lj} \cdot \Phi_j^{-1} \cdot \sum_k \Phi_k a_{kl} \end{aligned}$$

Now, the total share of shipments originating at l and going to j is:

$$\pi_{lj} = (c_l \kappa_{lj})^{-\theta} b_{lj} \Phi_j^{-1}.$$

We define node l 's measure of Entrepôt with respect to destination j as

$$Entrepôt_{lj} \equiv \pi_j^l - \pi_{lj} = \sum_k \Phi_k a_{kl} b_{lj} - (c_l \kappa_{lj})^{-\theta} b_{lj} \quad (15)$$

which is the difference between node l 's weighted network position with respect to destination j – how close to j other locations k in the network are when moving through l , where weights are multilateral resistance at k – and the marginal cost of production and transport from l . The former predicts transit, as higher values here mean l is in a more important position in the network to move goods to j , while the latter predicts exports from l to j .

We further note that, holding constant all other leg level costs $a_{k'l'}$ for $k' \neq k \vee l' \neq l$, a reduction in leg-level trade cost to a node l –i.e. an increase in a_{kl} , increases this measure

for l . In particular,

$$\frac{d\pi_j^l - \pi_{lj}}{da_{kl}} > 0 \quad (16)$$

noting that $d\Phi_j/da_{kl} > 0$ and $db_{lj}/da_{kl} > 0$.

This functional form has the convenient property that it aggregates from micro- to macro-data – as we see below – and thus allows for consistency in measure between our micro-level approach in Figure 4 and our macro-level approach when estimating the model and counterfactuals.

Similarly we can write share of global trade moving through kl as

$$\pi^{kl} = a_{kl} \cdot \sum_j b_{lj} \Theta_j \frac{\Phi_k}{\Phi_j}. \quad (17)$$

where Θ_j is country j share of global GDP. And, summing across locations k ,

$$\pi^l = \sum_k \Phi_k a_{kl} \cdot \sum_j \Theta_j \frac{b_{lj}}{\Phi_j}. \quad (18)$$

The global share of trade from l as

$$\pi_l = \sum_j \Theta_j \frac{(c_l \kappa_{lj} \tau_{lj})^{-\theta}}{\Phi_j} \quad (19)$$

The difference between the two is

$$\begin{aligned} \pi^l - \pi_l &= \sum_k \Phi_k a_{kl} \cdot \sum_j \Theta_j \frac{b_{lj}}{\Phi_j} - \sum_j \Theta_j \frac{(c_l \kappa_{lj} \tau_{lj})^{-\theta}}{\Phi_j} \\ &= \sum_j \frac{\Theta_j}{\Phi_j} \left[\sum_k \Phi_k a_{kl} \cdot b_{lj} - (c_l \kappa_{lj} \tau_{lj})^{-\theta} \right] \\ &= \sum_j \frac{\Theta_j}{\Phi_j} \left[\pi_j^l - \pi_{lj} \right] \end{aligned}$$

which is a weighted average of our individual country j measure, where each destination country j measure is weighted by its share of total global trade and network proximity, respectively Θ_j and Φ_j . That is, we can either take (on the left hand side) measure of global shares of traffic and trade, or (on the right hand side) an average of the same difference for each country j derived from micro-data.

C.2 The Network Effect of Adjustments on Trade

A change in the leg cost between k and l (t_{kl}) can affect trade volumes between an origin i and destination j through the trade network. However, Ricardian competition can interact with the trade network to generate unexpected effects. For any change to the

cost t_{kl} , trade volumes between i and j will adjust according to the following equation:

$$\frac{dX_{ijn}}{dt_{kl}} = \frac{\partial X_{jn}}{\partial t_{kl}} \cdot \pi_{ijn} + X_{jn} \cdot \left[\frac{\partial c_{in}^{-\theta}}{\partial t_{kl}} \cdot \frac{\pi_{ijn}}{c_{in}^{-\theta}} + \frac{\partial \tau_{ijn}^{-\theta}}{\partial t_{kl}} \cdot \frac{\pi_{ijn}}{\tau_{ijn}^{-\theta}} + \frac{\partial \Phi_{jn}^{-\theta}}{\partial t_{kl}} \cdot \frac{\pi_{ijn}}{\Phi_{jn}^{-\theta}} \right].$$

The first term on the right is the effect of t_{kl} on trade with i through a change in the volume consumed at j in industry n . In square parentheses, the first term is the effect through any changes to the production costs at i , which can happen if the price of inputs changes or through a change in wages. The second term is the effect through trade costs between i and j in industry n , and the final term is the effect through multilateral resistance.

What can we say about the signs on these terms? As the trade cost matrix is endogenous to trade volumes, these terms are ambiguous, as a change in t_{kl} , by changing trade volumes, changes traffic volumes at each leg, and therefore equilibrium effects on the full matrix of trade costs.

However, if we consider a change in t_{kl} which holds fixed all other leg costs $t_{k'l'}$ for $k' \neq k \vee l' \neq l$, only the final term can be negative. Intuitively, a reduction in trade costs between k and l can increase consumption at j , reduce expected trade costs between i and j , and reduce production costs at i , all of which result in an increase in trade volumes between i and j . However, a reduction in trade costs between k and l also stiffens competitions at j . If this last effect is large enough, it can overturn the sign of the first three.

In the scale-free case, the total effect is positive if and only if the elasticities of consumption at j ($\epsilon_{X_{jn}, t_{kl}}$), production costs at i ($\epsilon_{c_{in}, t_{kl}}$), and trade costs between i and j ($\epsilon_{\tau_{in}, t_{kl}}$) with respect to t_{kl} are larger than the elasticity of multilateral resistance at j with respect to t_{kl} ($\epsilon_{\Phi_j, t_{kl}}$). Furthermore, $\frac{\partial X_{ijn}}{\partial t_{kl}} > 0$ if and only if:

$$\epsilon_{X_{jn}, t_{kl}} + [\epsilon_{c_{in}, t_{kl}} + \epsilon_{\tau_{in}, t_{kl}}] (1 - \pi_{ijn}) > \sum_{i' \neq i} (\epsilon_{c_{i'n}, t_{kl}} + \epsilon_{\tau_{i'jn}, t_{kl}}) \pi_{i'jn}. \quad (20)$$

The sum of the effects on production and transport costs between all other countries i' (other than i) and j has to be less than a function of the effects on production and transport cost at i and the overall propensity of consumption at j to grow. This last expression shows most clearly that the effect of a decline in trade costs between k and l has the potential to negatively affect trade flows between i and j if it differentially lowers trade and production costs from i 's competitors.

C.3 Extension: Market Power

This section addresses the question of firm behavior and how does it fit in with the estimation of scale economies. We investigate this issue using a simple adaptation of the Cournot framework with endogenous entry (Sutton, 1991). Suppose we have origin and destination countries denoted by k and l respectively. Demand for shipping on this route is $\Xi_{kl}(t_{kl})$, where t_{kl} is the equilibrium cost of shipping on that leg, determined by the shippers on that route.

Consider a game with two stages. First, shippers with constant marginal costs c decide to enter after paying cost ϵ . Second, shippers play a nash-in-prices entry game to determine the shipping price t . A particular shipper's i market share on route kl is $s_{i,kl} = \frac{\exp(at_{i,kl})}{\sum \exp(at_{i,kl})}$ where $i \in 1 \dots N_{kl}$. Total demand for shipping on the route is $\Xi_{kl} = \delta \times \left(\sum_{i=1}^{N_{kl}} \exp(at_{i,kl}) \right)^\gamma$, where a , δ , and γ are constants that governs consumer sensitivity to shipping prices.

Starting with backward induction and the first stage, each symmetric shipper i on route kl will charge a shipping cost

$$t_{i,kl} = \frac{\delta}{a [1 - (1 - \gamma) s_{i,kl}]} + c.$$

In the first stage, we then determine the number of shippers N_{kl} who are willing to pay entry cost ϵ . This is pinned down by the equation:

$$\frac{\Xi_{kl}}{N_{kl}} = \left[\frac{\delta \Xi_{kl}}{a\epsilon} + (1 - \gamma) \right]^{-1}.$$

So in equilibrium,

$$\ln(t_{kl} - c) = -\ln(a) - \ln \left(1 - (1 - \gamma) \left[\frac{\delta \Xi_{kl}}{a\epsilon} + (1 - \gamma) \right]^{-1} \right)$$

This relationship is sensitive to competition and market share forms. Additionally we write trade costs in iceberg form and if marginal costs are low, then $c \equiv 1$. A rough approximation is close to our main estimating relation, for some constant ϕ :

$$\ln(t_{kl} - 1) \approx \ln(\phi) + \alpha \ln \left(\Xi_{kl}^{data} \right).$$

Effectively one source for scale economies comes from an increase market size that increases entry and thus drives down prices.

Appendix D Estimation

This section reports additional details, results, and robustness checks from our estimation strategy, as well as discusses the potential threats to identification.

D.1 Recovery of Predicted Trade Costs

Table A.5 shows the results of our estimation that predicts leg-level trade costs. Positive values for β indicate increases in trade costs and negative values indicate decreases in trade cost.

However, these estimates are not causal, and cannot be used for either inference or counterfactuals. They represent the power of various (including highly endogenous) variables in predicting a trade cost matrix that rationalizes leg-level containerized traffic flow. We find high correlations between observed and model-predicted shares, including for shares that we do not target (Figure 8). We find a correlation between trade shares of 0.7 (which we do not target) and traffic shares of 0.9 (which we target). If we had more possible useful predictive variables, we could use a machine learning technique to tease out the best basis of variables to predict model-consistent trade costs.

Table A.5: Predictive Trade Cost Estimates

Coefficient	Estimate
β_0 (intercept)	7.968
β_1 (log distance)	-0.006
β_2 (log route traffic)	-1.033
β_3 (log outgoing port traffic)	0.273
β_4 (log incoming port traffic)	0.275
β_5 (land borders)	-0.386
β_6 (trade volume)	-0.000

Notes: Results presented here are the moments from the GMM estimation in Section 5. These results are not causal, and cannot be used for either inference or counterfactuals. They represent the predictive power of various (possibly endogenous) variables in predicting a trade cost matrix that rationalizes leg-level containerized traffic flow.

This analysis reflects the spirit of pure prediction and cannot satisfy the “Lucas Critique” as they are purely observational and do not reflect fundamental economic parameters, forces, or relationships. In Section 5 we address endogeneity and causality, using an instrument to find the relationship between route-level volume and trade costs.

D.2 Scale Elasticity and Trade Cost Estimation

In this Appendix we explore the potential for a mechanical relationship between traffic volumes and costs in our model, which is also present in the Allen and Arkolakis (2019)

framework. We first show how such a potential correlation is a form of omitted variable bias and conditions under which an instrument corrects it. We then run Monte Carlo simulations confirming the existence of the bias in the model and showing how our instrument can remove it. We proxy for the bias using the difference between model-generated costs and a small set of external cost estimates, and show both the existence of the bias in the OLS relationship between volumes and cost and that the bias is eliminated by the instrument. Finally, we use our external cost estimates in a parallel estimation and find a similar scale economy.

D.2.1 Identification Strategy

As mentioned in the main text, we recognize that there is traffic volumes and trade costs are endogenous. As a result, we introduce a demand shifter as our instrument to recover the causal impact of traffic on trade costs. In this section, we show how a potential mechanical relationship between traffic volumes and costs in our model can be a form of omitted variable bias and conditions under which an instrument can correct for it. This issue is also present in the Allen and Arkolakis (2019) framework.

Suppose we observe traffic volumes with measurement error ($\Xi_{kl}^{data} = \Xi_{kl} + \chi_{kl}$) and \hat{t}_{kl} is our estimated trade cost as part of the estimation procedure. Our OLS specification for our scale elasticity from Equation (11) would be slightly modified to the following:

$$\ln(\hat{t}_{kl}^\theta - 1) = \alpha_0 + \alpha_1 \cdot \ln \Xi_{kl}^{data} + \alpha_2 \cdot \ln d_{kl} + \varepsilon'_{kl}, \quad (21)$$

Suppose the error on trade costs as part of the estimation process, due to mismeasured traffic volumes, is as follows:

$$\ln(\hat{t}_{kl}^\theta + 1) = \ln(t_{kl}^\theta + 1) + \nu_{kl}, \quad (22)$$

where t_{kl}^θ is the true cost and ν_{kl} is some error in the estimation. The measurement error from ν_{kl} can create a mechanical correlation between \hat{t}_{kl}^θ and Ξ_{kl}^{data} if $Cov(\nu_{kl}, \Xi_{kl}^{data}) \neq 0$, i.e. when the error between the true and estimated costs are correlated with observed traffic flows.

Using Equation (22) in order to recover the true trade costs t_{kl}^θ from the OLS specification in Equation (21):

$$\begin{aligned} \ln(t_{kl}^\theta + 1) + \nu_{kl} &= \alpha_0 + \alpha_1 \cdot \ln \Xi_{kl}^{data} + \alpha_2 \cdot \ln d_{kl} + \varepsilon'_{kl} \\ \ln(t_{kl}^\theta - 1) &= \bar{\alpha}_0 + \bar{\alpha}_1 \cdot \ln \Xi_{kl}^{data} + \bar{\alpha}_2 \cdot \ln d_{kl} + \nu'_{kl} \end{aligned} \quad (23)$$

where $\nu'_{kl} = \varepsilon'_{kl} - \nu_{kl}$, and the mechanical correlation can be interpreted as a stan-

standard concern that the error is correlated with the regressor, in this case through ν'_{kl} if $Cov(\nu_{kl}, \Xi_{kl}^{data}) \neq 0$.

While measurement error on the dependent variable is not unique to this setting, the specific concern here is the measurement error is correlated with traffic. For example, measurement error in observed traffic flows will show up both in Ξ_{kl}^{data} as well as in estimated costs \hat{t}_{kl} , and using the estimated costs may recover a mechanical correlation which could bias our scale economy estimates. However, even if this is the case, the instrument can recover the true scale parameter so long as the instrument is correlated with traffic but uncorrelated with the measurement error – i.e. under a specific version of an exclusion restriction.

Using an instrument z_{kl} , the estimate of α_1 from Equation (21) is

$$\alpha_{1,IV} = \frac{Cov(z_{kl}, \hat{t}_{kl})}{Cov(z_{kl}, \Xi_{kl}^{data})}$$

whereas the IV estimate of $\bar{\alpha}_1$ from Equation (23) is

$$\bar{\alpha}_{1,IV} = \frac{Cov(z_{kl}, t_{kl})}{Cov(z_{kl}, \Xi_{kl}^{data})}$$

Crucially, these two estimates are identical if and only if

$$Cov(z_{kl}, t_{kl}) = Cov(z_{kl}, \hat{t}_{kl}) = Cov(z_{kl}, t_{kl} + \nu_{kl})$$

where all variables are residualized for distance. This condition holds if our instrument is uncorrelated with the error in our estimation of leg costs, $Cov(z_{kl}, \nu_{kl}) = 0$. As such, in order to recover the correct coefficient on scale, we have a second restriction, $Cov(z_{kl}, \nu_{kl}) = 0$ in addition to the standard exclusion restriction, $Cov(z_{kl}, \varepsilon_{kl}) = 0$.

D.2.2 Monte Carlo Simulations

Here we run Monte Carlo simulations to demonstrate the potential sources of bias in the OLS estimate of the scale elasticity, and how an instrument satisfying the conditions outlined in the previous section recovers an unbiased estimate. The results are shown in both the main text in Figure 6 and here in Table A.6.

Simulation Procedure We run the following simulation procedure:

1. Generate distances between 15 countries from a unit uniform distribution.
2. Generate a graph with links between countries, where 1/3 of pairs have a bilateral transport link kl .

3. Generate the invariant part of trade costs on link between k and l . This is a linear transformation of the invariant part of trade cost in the main text (Equation (2)):

$$a_{0,kl} = \frac{1}{300} Distance_{kl}.$$

4. Generate origin-destination trade values between countries i and j :

$$X_{ij} = 3 \times Distance_{ij} + \nu_{ij},$$

where ν_{ij} is drawn from a normal distribution with mean 1 and standard deviation 1 ($N(1, 1)$).

5. Generate model consistent trade costs that satisfy the following relationship in Equation (10) using matrix notation:⁹

$$\Xi = A \otimes (B' (X \otimes B) B') \quad (24)$$

where Ξ is the true matrix of traffic volumes, lower case a_{kl} denotes the k, l element of the matrix A such that $a_{kl} = \exp(a_{0,kl} + \alpha_1 \ln(\Xi_{kl}))$, lower case b_{kl} denotes the k, l element of the matrix B such that $b_{kl} = \tau_{ij}^\theta$, and X is the matrix of trade values (each element is X_{ij}). To translate this back to Equation (10), note that $a_{kl} = t_{ij}^{-\theta}$ and $b_{kl} = \tau_{ij}^\theta$ (Equations (2) and (3)). We use $\alpha_1 = 0.01$ for this simulation (True Value, first row, Table A.6).

6. Assume that the econometrician observes mis-measured traffic $\ln(\Xi^{data}) = \ln(\Xi) + \epsilon$, where ϵ is drawn from $N(1, 1)$.
7. Generate an instrumental variable Z such that $E(Z\epsilon) = 0$, but $E(Z\Xi) \neq 0$.
8. Use our routine from the main text to recover $\hat{A}(\tilde{\Xi}, X, Distance)$, based on the mismeasured $\tilde{\Xi}$ from step 6. Each element in this \hat{A} matrix is denoted as \hat{a}_{kl} .

Specifications We run the four specifications below, 500 times each, and report the median estimate of α_1 and its standard deviation in Table A.6. The true value of α_1 is 0.01 for this simulation (first row, Table A.6). The distributions of each specification are plotted in Figure 6 in the main text.

1. **No Errors Scenario** If we perfectly observe Ξ_{kl} without measurement error, we would be able to generate the true model-consistent trade costs a_{kl} (Equations (10))

⁹We use the notation of Allen and Arkolakis (2019) (See Corollary 1, Equation (22)).

and (24)) and run the following OLS specification:

$$\ln(a_{kl}) = \alpha_0 + \hat{\alpha}_{1,OLS} \ln(\Xi_{kl}) + \alpha_2 \ln d_{kl} + \psi_{kl}$$

where Ξ_{kl} denotes the k, l element of the matrix Ξ and $\ln d_{kl}$ is the log of distance. Note that even though the estimate of a_{kl} is generated from Ξ , there is no bias in our estimates and recovers the true value (second row, Table A.6).

2. **Circularity Bias Scenario** If we do not perfectly observe traffic and instead observe traffic with error $\tilde{\Xi}$ (Step 6 above), we will generate trade costs with measurement error (\hat{a}_{kl}) per Step 8. This will lead to the following OLS specification:

$$\ln(\hat{a}_{kl}) = \alpha_0 + \hat{\alpha}_{1,OLS,noise} \ln(\tilde{\Xi}_{kl}^{data}) + \alpha_2 \ln d_{kl} + \psi_{kl}.$$

The measurement error here biases our OLS estimates upwards, due to the mechanical relationship of our traffic to implied trade costs (third row, Table A.6).

3. **Independent Variable Error Scenario** Here we consider classic measurement error in observed traffic volumes: $\ln(\tilde{\Xi}_{kl}) = \ln(\Xi_{kl}) + N(0, 1)$. Assuming that our trade costs are estimated correctly, this will result in the following OLS specification:

$$\ln(a_{kl}) = \alpha_0 + \hat{\alpha}_{1,OLS,measurement} \ln(\tilde{\Xi}_{kl}) + \alpha_2 \ln d_{kl} + \psi_{kl}$$

This classic measurement error will lead to a classic attenuation bias in the results (third row, Table A.6).

4. **IV with Circularity Bias Scenario** With mismeasured traffic volumes that generate mismeasured trade costs, we run two-stage least squares using the simulated instrument Z from Step 7. The first and second stages of our specification are as follows:

$$\begin{aligned} \ln(\tilde{\Xi}_{kl}^{data}) &= \beta_0 + \hat{\beta}_{1,IV,noise} \ln(Z_{kl}) + \beta_2 \ln d_{kl} + \psi'_{kl} \\ \ln(\hat{a}_{kl}) &= \alpha_0 + \hat{\alpha}_{1,IV,noise} \ln(\tilde{\Xi}_{kl}^{data}) + \alpha_2 \ln d_{kl} + \psi_{kl} \end{aligned}$$

The instrumental variable approach restores the upward bias of the measurement error (last row, Table A.6).

D.2.3 Circularity Bias and Geographic Instrument

As is generally the case with an exclusion restriction, we cannot directly test the condition $Cov(z_{kl}, \nu_{kl}) = 0$ (Section D.2.1). However, we can proxy for ν_{kl} by comparing our model's

Table A.6: Monte Carlo Estimates - With Scale

	Estimate	Median Estimate	Standard Deviation
α_1	True Value	0.10	
$\hat{\alpha}_{1,OLS}$	No Errors	0.10	0.00
$\hat{\alpha}_{1,OLS,noise}$	Circularity Bias	0.13	0.04
$\hat{\alpha}_{1,OLS,measurement}$	Independent Var Error	0.08	0.01
$\hat{\alpha}_{1,IV,noise}$	IV with Circularity Bias	0.10	0.02
	N		500

Notes: The distribution of these results is shown in both the main text in Figure 6. The estimate for $\hat{\alpha}_{1,OLS}$ in the No Errors specification (second row) shows no bias and recovers the true value. The estimate for $\hat{\alpha}_{1,OLS,noise}$ in the third row illustrates our upward circularity biases and shows a larger scale economy than the true scale economy. The estimate for $\hat{\alpha}_{1,OLS,measurement}$ shows how different our upward bias is from classic attenuation bias (the fourth row). Lastly, our estimate for $\hat{\alpha}_{1,IV,noise}$ shows how our instrument corrects for this upward attenuation bias and recovers the true value for our scale economy a_1 (last row).

estimates of leg costs t_{kl} with external estimates of pecuniary shipping costs from Wong (2022).

We calculate an estimate of our model's mismeasurement $\hat{\nu}_{kl}$ as follows:

$$\hat{\nu}_{kl} = t_{kl} - t_{Wong,kl}.$$

for the 209 links for which we have external freight cost estimates from Wong (2022), residualizing both for distance.

In Panels (A) and (B) of Figure 7, we plot a scatter plot of this estimated mismeasurement against traffic and our estimated costs, controlling for sea distance. The existence of a correlation between the three is consistent with a potential circularity bias. In the context of our Monte Carlo Simulations, this correlation opens the door to an upward bias in an OLS estimate of the scale elasticity.

In Panel (C), we plot the same estimated mismeasurement against our geography-based instrument z_{kl} . Here the correlation in Panel (A) vanishes. While we caution that this lack of correlation is not evidence that the exclusion restriction is met, as that condition is inherently unknowable, this exercise can be thought of as a balancing test, where the observed correlation between proxies for the unobserved error and the endogenous variable is not present with the instrument.

D.2.4 Estimating Scale Elasticities without Imputed Costs

In order to test the robustness of our identification strategy, we find a similar scale elasticity using observed freight rates for only a subset of routes from Wong (2022). First, we find suggestive evidence for potential scale economies for this subset of routes

where we directly observe freight rates for. These findings further support our findings in Section 7.1 on the presence of scale economies in our context. This result holds even when we include origin and destination-level port fees, which would be correlated with port-level congestion. Second, we apply our instrument to this subset of routes, but due to a small number of observations and a relatively weak first stage, cannot reject the null hypothesis.

We note that estimating our scale elasticity using this approach has two main drawbacks. First, external pecuniary freight rates such as those in Wong (2022) do not include all possible elements of network leg costs that are consistent with our model. Second, our goal is to estimate a global set of leg-level trade costs and a global dataset on observed freight rates does not exist. Estimating a scale elasticity within this context creates both a power issue (as we show below) and an external validity issue. Nevertheless, these estimates provide a measure of scale elasticity that is free of any potential bias from our trade cost estimation and has the potential to indirectly confirm the results from our instrumented estimation.

First, we find a statistically significant and negative correlation between freight rates and traffic in Table A.7. In Column (1), we find this negative correlation using all-in freight rates which is the sum of base freight rates of the route, origin port fees, destination port fees, and bunker fuel. Distance between routes is included as a control and is positively correlated with freight rates. Using an even smaller set of routes for which we observe base freight rates directly, we find that the coefficient between freight rates and traffic in Column (2) retains the same sign and is within one confidence interval of the results in Column (1). Given that we observe origin and destination port fees for this even smaller subset of routes, we can include these fees in Column (3). These fees are potentially correlated with congestion at the origin and destination ports. Including these proxies for origin and destination port congestion, the coefficient between base freight rates and traffic retains the same sign and is within one standard error of the results in Column (2). We conclude that this is further suggestive evidence for the presence of scale economies in this context.

Second, we apply our instrument to this subset of routes and find a scale elasticity that is within a standard error of our results in Table 1. Due to the small number of observations, however, this estimate is noisy and our first stage is weak.

Table A.7: Correlation between Freight Rates and Traffic for Subset of Routes

	(1)	(2)	(3)
	All-in FR	Base FR	Base FR
Traffic	-0.0478 (0.0241)	-0.116 (0.0337)	-0.108 (0.0256)
Distance	0.404 (0.0848)	0.555 (0.137)	0.516 (0.104)
Origin Fees			-0.164 (0.241)
Dest Fees			0.557 (0.308)
Specification	OLS	OLS	OLS
Observations	142	142	142
R^2	.38	.24	.29

Notes: All variables are in logs. Robust standard errors clustered by origin and destination ports in parentheses. weighted by route trade values. The all-in freight rates used in Column (1) are the costs paid by firms, in dollar terms, to transport a standard full container load between port pairs. These all-in rates include the base ocean rate, fuel surcharge, as well as port handling fees at both origin and destination. For a much smaller subset of routes, we observe the direct origin and destination fees breakdown of these freight rates. The base rate is used in Column (2) while the base rate and port fees are used in Column (3). In order to make this comparison directly, the observations are restricted to routes where the direct breakdown is observed.

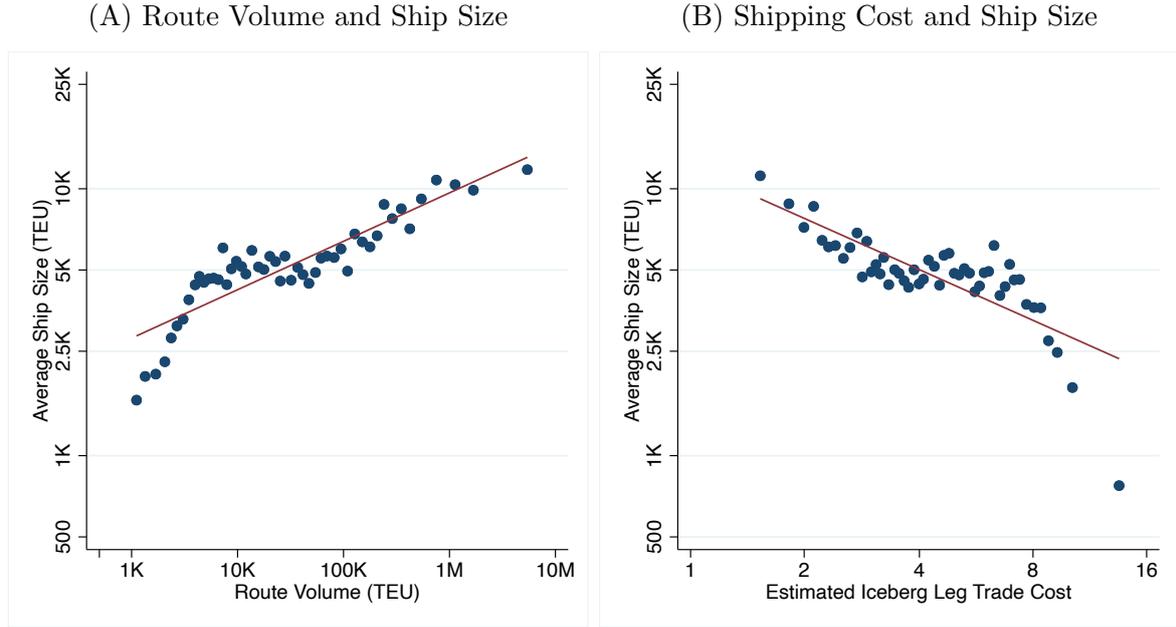
D.3 Ship Sizes, Trade Volumes, and Recovered Trade Costs Robustness

Figure A.10 replicates main text Figure 9 without distance controls. Results are broadly similar. Tables A.8 and A.9 replicate Panels (A) and (B) in Figure 9 in regression form respectively. Columns (2)-(4) sequentially add controls for route distance, origin fixed effects, and destination fixed effects. Results are broadly consistent with baseline results. Table A.8 shows a 10% increase in trade volumes is correlated with a 2.2-2.5% in average ship sizes. Table A.9 shows a 10% decrease in estimated trade costs corresponds to 6.3-10.6% increase in average ship sizes.

D.4 Shipment-Level Data: Ship Size

We pair the visual analysis in Figure 10 with Table A.10, which displays shipment-level regressions. Column (1) regresses, for our sample of shipments, the log of ship size against the log of total origin country volumes shipped (TEUs), confirming a positive relationship. Column 2 adds the log of quantity loaded at each shipment's port of lading—the port where the shipments are loaded onto a US-bound ship (Stop 1 in Figure A.1). Both coefficients are positive but the coefficient on origin volume is almost halved (0.084 in Column (1) compared to 0.043 in Column (2)), indicating that much of the correlation

Figure A.10: Link Between Recovered Trade Costs and Ship Size - No Distance Controls



Notes: These figures are bin-scatter plots over all observed containership routes, with 100 bins. (A) plots the relationship between the total containers on a route and the average containership’s size on that route. (B) plots the relationship between the estimated trade cost t_{kl} with $\theta = 4$ and the average containership’s size on that route. Containership size reflects the size of the ship for the average container on that route.

Table A.8: Correlation Between Route Volume and Ship Size

	(1)	(2)	(3)	(4)
	log(Ship Size)	log(Ship Size)	log(Ship Size)	log(Ship Size)
log(Trade Volume)	0.223 (0.00671)	0.245 (0.00609)	0.238 (0.00607)	0.216 (0.00727)
FE Origin	0	1	1	1
FE Good	0	0	1	1
FE Origin-Good	0	0	1	1
R^2	0.315	0.515	0.627	0.703
N	2304	2304	2304	2304

Notes: We consider the relationship between the total containers on a route and the average containership’s size on that route. Containership size reflects the size of the ship for the average container on that route. We use robust standard errors. Column (2), controls for logarithm of shipping distance. Column (3), adds controls for the origin port. Column (4) adds fixed effects for the destination port.

between origin volume and ship size acts through the size of the lading port. Column (3) fully interacts the variables in Column (2) with an indicator variable for shipments that are laded in their origin countries. As suggested by the figure, for shipments whose origin country differs from lading country—an indicator value of 0—the correlation between ship size and lading volume is considerably higher (0.130), and shipments’ ship sizes are not strongly correlated with origin country volumes when they lade in third countries (0.009).

Finally, stopping at larger ports matters, even when goods remain on board: goods

Table A.9: Correlation Between Shipping Costs and Ship Size

	(1)	(2)	(3)	(4)
	log(Ship Size)	log(Ship Size)	log(Ship Size)	log(Ship Size)
log(Trade Cost)	-0.632 (0.0353)	-1.060 (0.0310)	-0.986 (0.0279)	-0.842 (0.0282)
FE Origin	0	1	1	1
FE Good	0	0	1	1
FE Origin-Good	0	0	1	1
R^2	0.136	0.457	0.596	0.703
N	2304	2304	2304	2304

Notes: We consider the relationship between the estimated trade cost t_{kl} with $\theta = 4$ and the average containership's size on that route. Containership size reflects the size of the ship for the average container on that route. We use robust standard errors. Column (2), controls for logarithm of shipping distance. Column (3), adds controls for the origin port. Column (4) adds fixed effects for the destination port.

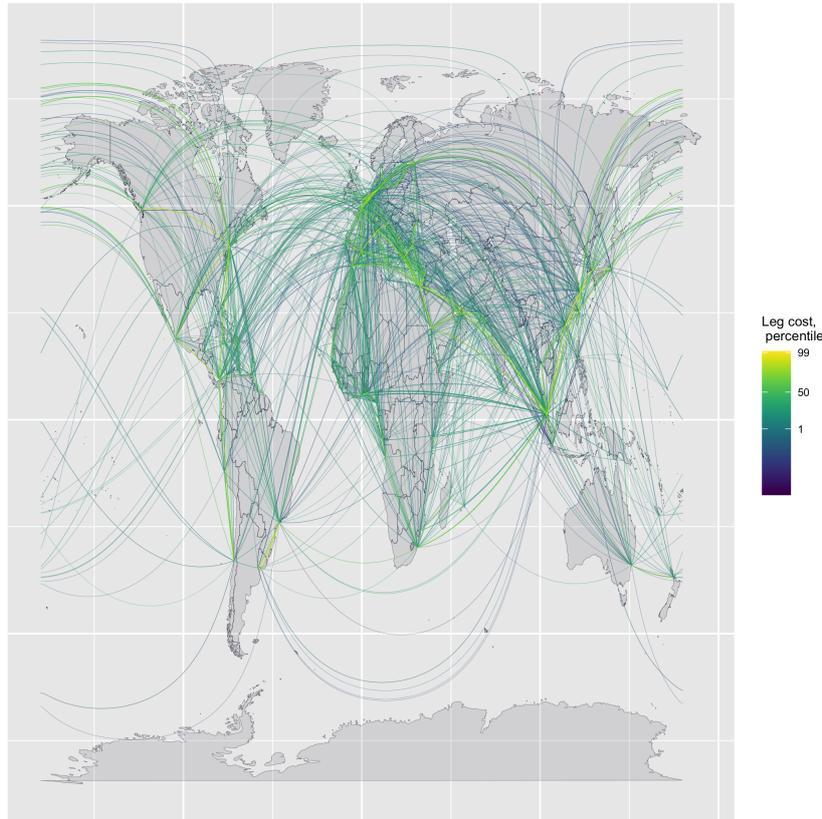
Table A.10: Determinants of Ship Size

	(1)	(2)	(3)	(4)
	ln Ship Size	ln Ship Size	ln Ship Size	ln Ship Size
ln Volume at Origin	0.0843 (0.0163)	0.0432 (0.0179)	0.00925 (0.0121)	
ln Volume at Lading		0.0803 (0.0202)	0.127 (0.0230)	0.0282 (0.0182)
$\mathbb{1}(\text{Lading is Origin})=1$			-0.0220 (0.300)	
$\mathbb{1}(\text{Lading is Origin})=1 \times \ln \text{Volume at Lading}$			-0.0937 (0.0295)	
$\mathbb{1}(\text{Lading is Origin})=1 \times \ln \text{Volume at Origin}$			0.0861 (0.0220)	
ln Largest Port Stop				0.121 (0.0250)
Observations	215,656	215,656	215,656	215,656
R^2	.124	.174	.199	.21
F-stat	26.82	14.66	13.51	26.73

Notes: Observations are at the shipment level, weighted by TEU, representing all matched imported containers to the United States. ln Ship Size is the natural log of maximum ship capacity in TEU. ln Volume at Origin is the natural log of the sum of all shipments' TEU by shipment origin country. ln Volume at Lading is the sum of all shipments' TEU by shipment lading country. The indicator takes a value of 1 if the shipment is laded at the country of origin. ln Largest Port Stop is the maximum of the natural log of the volume of lading at all ports visited between the port of lading and unloading. Standard errors are clustered two ways by lading and destination ports.

lading at smaller transshipment points that travel along major routes are also on larger ships. Column (4) of Table A.10 regresses shipments' log ship size against the log volume laded at their port of lading and the log volume laded at the largest port at which we observe the shipment making a port call. The effect of the max-port-size variable is large, positive, and overall stronger than the effect of lading port volumes alone. Additional

Figure A.11: Trade Cost Estimates, All Legs



Notes: This map displays the recovered trade cost between all origins and destinations for containership legs in the AIS data. Lighter colors indicate lower trade costs.

stops that move through entrepôts allow shipments laded in smaller ports to travel on larger ships. Indirectness facilitates larger ship sizes beyond transshipment alone.

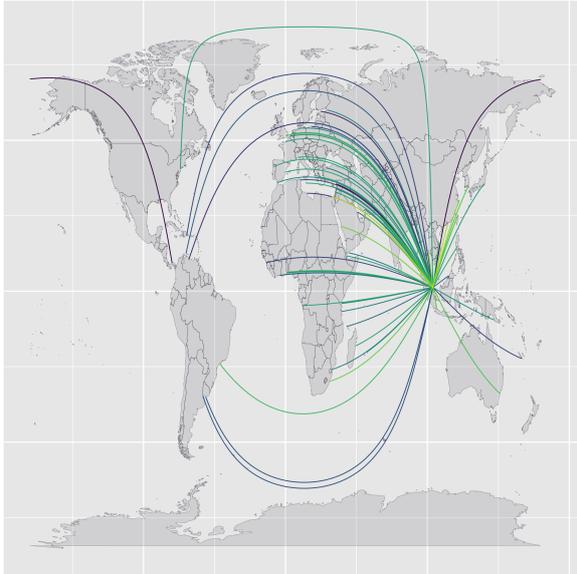
D.5 Additional Estimation Results

We plot our estimated route costs in Figure A.11. Thicker and lighter colors indicate lower-cost routes. Shorter and more heavily trafficked routes are the cheapest. The effect of scale is observable here: Syria to France is one of the highest cost legs, significantly higher than Singapore to Gibraltar, a much longer distance. Even among the subset of bilateral pairs for which we observe traffic, the triangle inequality is violated 280 times.

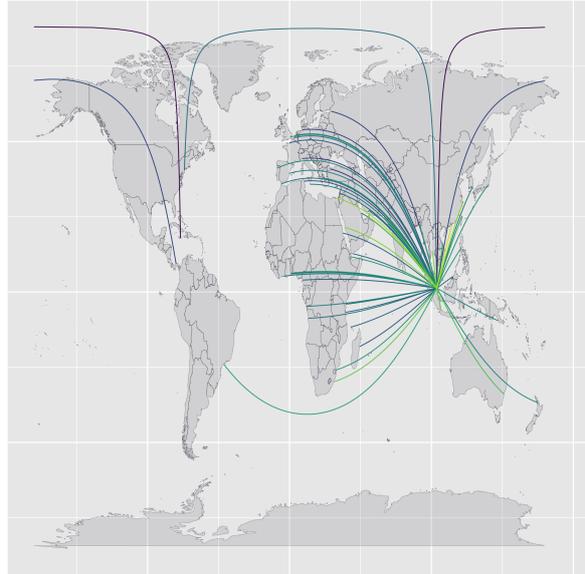
Figure A.12 plots bilateral incoming and outgoing trade costs for Singapore and Lebanon separately. Singapore is not only well-connected both as an origin and destination, but also has some of the cheapest legs. Lebanon, on the other hand, has both fewer and shorter connections.

Figure A.12: Trade Costs by Country

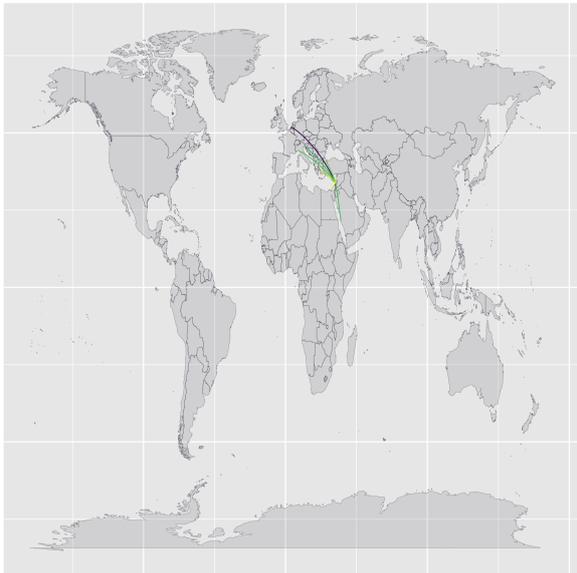
(A): Singapore, Origin



(B): Singapore, Destination



(C): Lebanon, Origin



(D): Lebanon, Destination



Notes: This map plots estimated link costs from Singapore in Panel (A) to Singapore in Panel (B), from Lebanon in Panel (C), and to Lebanon in Panel (D). Lighter colors indicate lower trade costs.

D.6 Analysis of Trade costs

Figure A.13 plots country-level market access for producers and consumers, which are averages of the expected trade cost (from the B-matrix) weighted by the GDP of origins and destinations, respectively. Entrepôts such as Egypt, Panama, and (not visible) Singapore and Gibraltar have generally cheaper trade costs, as does China, due to the scale of shipping as well as access to nearby low-cost entrepôt (Korea, Singapore, and Japan).

Table A.11 reflects the log-linear relationship between our estimated trade cost τ , aggregate bilateral trade values, and distance. These results highlight the reduced form relationships between these three variables, as well as the predictive power of our computed trade costs. Without origin or destination fixed effects, our trade costs alone can explain 29% variation of global trade. The logarithm of distance can account for less than 3%. We do not take this as a horse race, but rather indication that these two measures are distinct: Our cost estimates τ measure network proximity and real shipping network relationships. Distance is a proxy for other orthogonal variables which impact trade volumes as well.

Table A.11: The Relationship between Trade Volumes and Network-Consistent Trade Costs and Distance

	(1)	(2)	(3)	(4)	(5)	(6)
	Log trade values					
Log $\tau_{ij}^{-\theta}$	0.462*** (0.0297)		0.444*** (0.0306)	0.756*** (0.0310)		0.516*** (0.0278)
Log Dist		-0.755*** (0.0993)	-0.393*** (0.0937)		-1.372*** (0.0626)	-0.673*** (0.0597)
Constant	12.71*** (0.354)	14.82*** (0.909)	16.04*** (0.772)	15.67*** (0.311)	20.36*** (0.561)	19.29*** (0.439)
Orig, Dest FEs	No	No	No	Yes	Yes	Yes
Observations	23,344	22,985	22,985	23,344	22,985	22,985
R-squared	0.290	0.028	0.292	0.762	0.753	0.771

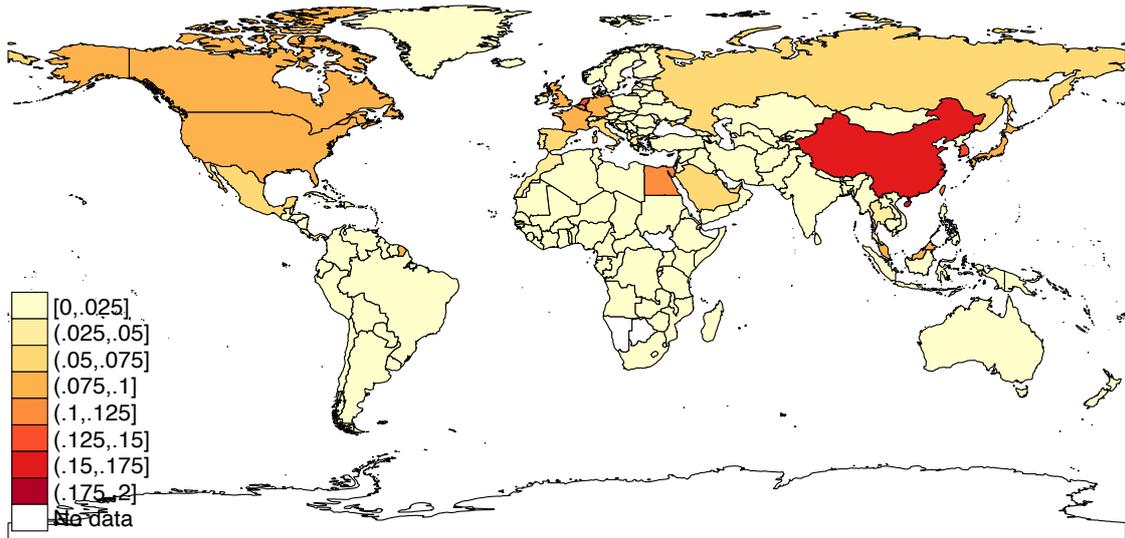
Notes: This table presents regression coefficients from the regression of the natural log of trade volumes on the natural log of $\tau_{ij}^{-\theta}$, the natural log of model-estimated origin-destination trade costs raised to the trade elasticity, and the natural log of distance, and the sea distance between the origin and destination measured in kilometers. Column (1)-(3) report results for cost and distance independently, then combined. Columns (4)-(6) rerun regressions in (1)-(3), respectively, adding origin and destination fixed effects.

Appendix E General Equilibrium Model in Changes

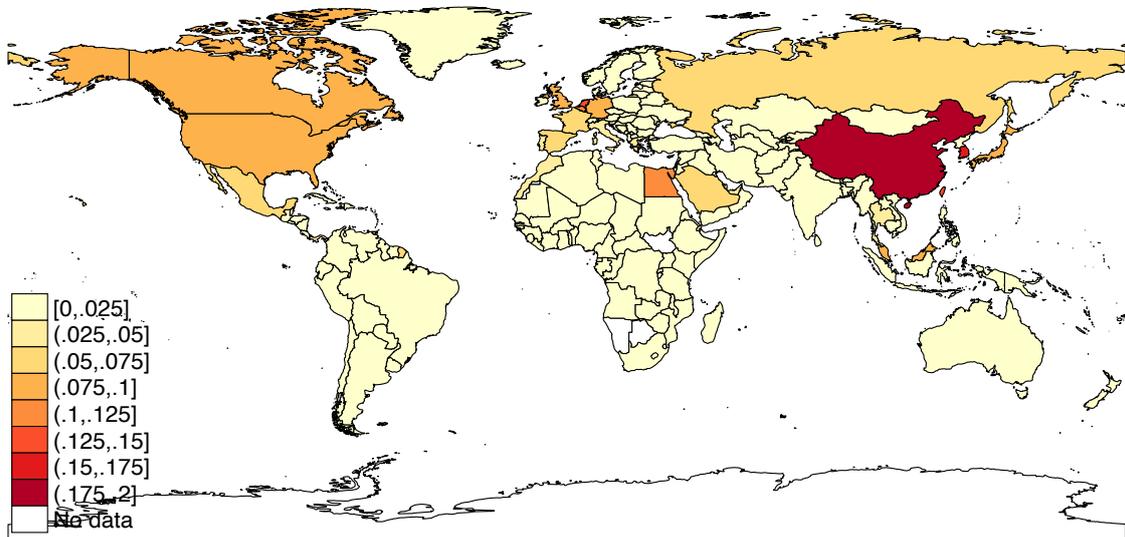
To close our model, we adopt the Caliendo and Parro (2015) framework. A continuum of intermediate goods ω_n are used in the production of composite goods that are in turn used domestically both as final goods and as materials for intermediate production by firms in

Figure A.13: Market Access

(A) Consumer Market Access



(B) Producer Market Access



Notes: Figure plots producer and consumer market access for each country according to the transportation costs estimated in Section 5. Countries with higher market access are in darker reds, while countries with lower market access are in lighter yellows. Countries with missing data are in white.

each industry n . We assume there are three sectors ($N = 3$): containerized tradables c , non-containerized tradables nc , and nontradables nt ($n \in [c, nc, nt]$). Intermediates in the nt sector are only sourced domestically while ω_{nc} and ω_c goods are sourced internationally. Trade routes are modeled for all three sectors but we only consider transportation cost changes for the containerized sector ω_c .

Consumption In each country i , consumers consume composite goods m_{in} from each sector n , maximizing Cobb-Douglas utility.

$$U_i = \prod_n m_{in}^{\eta_n} \text{ where } \sum_n \eta_n = 1,$$

where η_n is the Cobb-Douglas industry share, $\sum_n \eta_n = 1$.

Intermediate goods production The traded goods are intermediates, which are used in each country as building blocks for the industry composite goods. In each country i and industry n , firms produce a continuum of intermediate goods, indexed in each industry by $\omega_n \in \Omega_n$. There are two types of input required for the production of ω : labor and composite goods. The production of intermediate goods across countries differs in their efficiency by a country-industry specific constant z_{in} , a Ricardian technology. The production technology for intermediate ω is

$$q_{in}(\omega) = z_{in} [l_{in}]^{\gamma_{in}} \prod_{n'} [m_{in}^{n'}]^{\gamma_{in}^{n'}},$$

where l_{in} is labor. $\gamma_{in}^{n'}$ is share of materials from sector n' used in production of intermediate good ω , γ_{in} is share of value added, with $\sum_{n'} \gamma_{in}^{n'} = 1 - \gamma_{in}$. The marginal cost of production for firms is

$$c_{in} \equiv \frac{\Upsilon_{in} w_i^{\gamma_{in}} \prod_{n'} P_{in'}^{\gamma_{in}^{n'}}}{z_{in}}, \quad (25)$$

where w_i is the wage in country i , $P_{in'}$ is the price of a composite good from sector n' , and constant $\Upsilon_{in} = \prod_{n'} (\gamma_{in}^{n'})^{\gamma_{in}^{n'}} (\gamma_{in})^{\gamma_{in}}$.

Composite goods production In each country i , composite goods in industry n are produced using a CES aggregate of intermediates Ω_n , purchased and sold domestically at marginal cost. In traded industries, intermediates are sourced internationally from lowest-cost suppliers. Using the standard aggregation, the resulting price at j of the

composite in industry n is expected to be the following (where A_n is a constant):

$$P_{jn} = A_n \left[\sum_{i=1}^I c_i^{-\theta_n} \kappa_{ijn}^{-\theta_n} \tilde{\tau}_{ijn}^{-\theta_n} \right]. \quad (26)$$

The production costs in country i and industry n respond to a shock to a given t_{kl} according to the equation:

$$\dot{c}_{in} = \dot{w}_i^{\gamma_{in}} \prod_{k=1}^N \dot{P}_{ik}^{\gamma_{ink}}. \quad (27)$$

The change in the price of the composite intermediate good in country i and industry n relative to shock to t_{kl} is:

$$\dot{P}_{in} = \left[\sum_{i=1}^J \pi_{ijn} [\dot{\tau}_{ijn} \dot{c}_{in}]^{-\theta_n} \right]^{-1/\theta_n}. \quad (28)$$

Bilateral trade shares between i and j in industry n will change according to standard changes through production and transport costs:

$$\dot{\pi}_{ijn} = \left[\frac{\dot{c}_{in} \dot{\tau}_{ijn}}{\dot{P}_{in}} \right]^{-\theta_n}. \quad (29)$$

Trade volumes similarly adjust:

$$X'_{in} = \sum_{k=1}^N \gamma_{ink} \sum_{j=1}^I \frac{\pi'_{ijn}}{1 + \kappa_{ijn}} X'_{jk} + \alpha_{in} I'_i. \quad (30)$$

Lastly, trade is balanced to a deficit shifter such that:

$$\sum_{n=1}^N \sum_{i=1}^I \frac{\pi'_{ijn}}{1 + \kappa_{ijn}} X'_{in} - D_i = \sum_{n=1}^N \sum_{i=1}^I \frac{\pi'_{jin}}{1 + \kappa_{jin}} X'_{jn}, \quad (31)$$

where $I'_i = \dot{w}_i w_i L_i + \sum_{n=1}^N \sum_{i=1}^I \tau'_{ijn} \frac{\pi'_{ijn}}{1 + \kappa_{ijn}} X'_{in} + D_i$.

Appendix F Counterfactual Results

F.1 Counterfactual Procedure

Algorithm 1 describes the algorithm for finding the new equilibrium after an adjustment that induces an endogenous scale response in the network. We limit the counterfactual trade cost on any route to be no lower than the minimum observed initial trade cost.

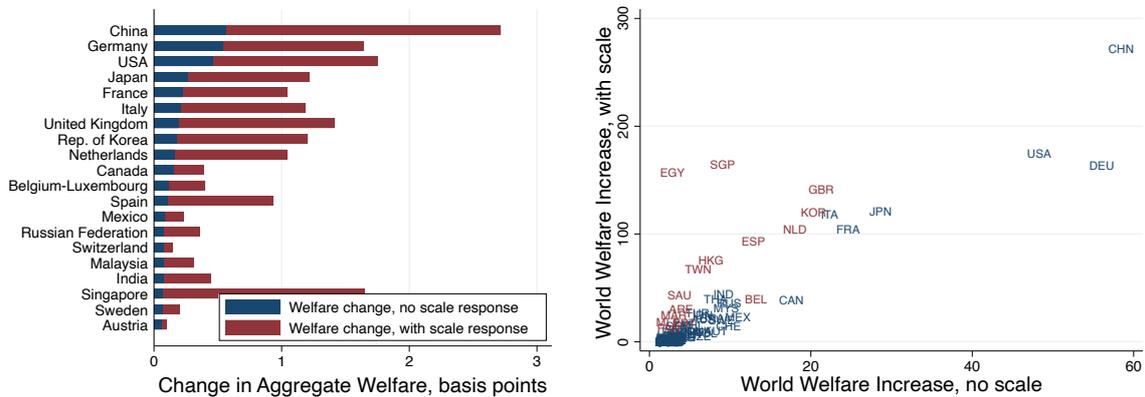
F.2 Counterfactuals: Additional Figures and Tables

The Global Impact of Local Infrastructure Improvements Table A.13 reports, for each case (denoted by column), the mean global welfare impact (first row) and

Algorithm 1 Scale Counterfactual Algorithm

- 1: **procedure** WELFARE CHANGE(X_0, Ξ_0, \hat{t}) ▷ Find a new equilibrium
 - 2: Initialize current trade flows X_0 and traffic Ξ_0
 - 3: Initialize changes in cost fundamentals \hat{t} ▷ Example: shipping distances changes
 - 4: Compute $A_0 = A(\Xi_0; \hat{t})$ ▷ Following equation 11
 - 5: Compute $B_0 = (I - A_0)^{-1}$
 - 6: Initialize difference = ∞ , tolerance = ϵ
 - 7: **while** *difference* < *tolerance* **do**
 - 8: Update trade flows $X_1 = X(B_0)$ ▷ Solving 8.1
 - 9: Update traffic $\Xi_1 = \Xi(X_1, A_0, B_0)$ ▷ Following equation 10
 - 10: Update leg costs $A_1 = A(\Xi_1)$
 - 11: Update trade costs $B_1 = (I - A_1)^{-1}$
 - 12: Compute *difference* = $\sum_{ij} (B_1 - B_0)^2$
 - 13: Update $A_0 = A_1$ and $B_0 = B_1$
 - 14: Return final trade flows X_1
 - 15: Compare welfare and price index changes between X_1 and X_0 ▷ Solving 8.1
-

Figure A.14: Most Pivotal Nodes: Change in Welfare Excluding Own

 (A) Non-Transportation Cost Reductions: Highest Global Welfare Changes
 (B) Non-Transportation Costs Reductions with vs without scale


Notes: Panel (A) shows absolute values for aggregate net change in global welfare after non-infrastructure cost reductions in the listed country, excluding the country’s own, for the 20 countries with the largest global impact calculated without scale economies. Overlaid grey bars represent welfare changes allowing for the network’s endogenous response to scale economies. Panel (B) compares, for each country, the change in world welfare, excluding the country’s own welfare, from a 1% decrease in non-transportation costs excluding the endogenous scale response (X-axis) vs a the same including the scale response (Y-axis). Markers are ISO Country codes. Entrepôts are in red.

standard deviation (second row) across all 136 targeted countries. Rows three through six consider results separately for counterfactuals where targeted countries are entrepôts and non-entrepôts. Column (1) reports welfare changes from non-transportation cost reductions without scale responses. Raw effects from counterfactuals targeting entrepôts are roughly twice as large, reflecting entrepôts’ greater global integration—a difference eliminated below in Table A.14. In Column (2), the scale response, which incorporates the

Table A.12: Welfare and Trade Outcomes from Improvements in Transportation and Non-Transportation Costs, Basis Points

	Non-Transportation Improvement		Transportation Improvement	
	Baseline Effect	Total Effect (Network & Scale)	Network Effect	Total Effect (Network & Scale)
	(1)	(2)	(3)	(4)
Δ Average Global Welfare				
Mean	0.08%	0.26%	0.18%	0.54%
Standard Deviation	(0.20)	(0.59)	(0.41)	(1.33)
Δ Container Trade Volumes				
Mean	0.87%	2.89%	2.02%	6.11%
Standard Deviation	(2.22)	(6.65)	(4.67)	(14.98)

Notes: This table reports results for our first counterfactual, transportation and non-transportation cost declines for each of 136 countries. Columns (1) and (2) present results for cases where non-transportation trade costs are reduced. Columns (3) and (4) present results for cases where transportation costs are reduced (infrastructure improvements). The top panel presents aggregate welfare changes. The bottom panel presents changes to aggregate container trade. Columns (1) and (3) correspond to cases where no scale economy feedback loops are allowed. Columns (2) and (4) present results allowing for scale economy feedback.

Table A.13: Counterfactual Reductions in Local Trade Costs, by Targeted Country Entrepôt Status

	Welfare Change from Cost Reduction				Trade Change from Cost Reduction			
	Non-Transportation		Transportation		Non-Transportation		Transportation	
	$\Delta\kappa_{kl}$	$\Delta\kappa_{kl}$ with Scale	Δt_{kl}	Δt_{kl} with Scale	$\Delta\kappa_{kl}$	$\Delta\kappa_{kl}$ with Scale	Δt_{kl}	Δt_{kl} with Scale
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Global Changes								
Mean	0.08%	0.26%	0.18%	0.54%	0.87%	2.89%	2.02%	6.11%
Standard Deviation	(0.20)	(0.59)	(0.41)	(1.33)	(2.22)	(6.65)	(4.67)	(14.98)
Reductions at Entrepôts								
Mean	0.16%	0.88%	0.87%	2.91%	1.56%	9.64%	9.73%	32.77%
Standard Deviation	(0.16)	(0.64)	(0.67)	(2.41)	(1.64)	(7.03)	(7.43)	(27.03)
Reductions at Non-Entrepôts								
Mean	0.07%	0.18%	0.09%	0.24%	0.78%	2.05%	1.07%	2.81%
Standard Deviation	(0.20)	(0.54)	(0.27)	(0.71)	(2.28)	(6.13)	(3.13)	(8.21)

Notes: This Table replicates the results for our first counterfactual, transportation and non-transportation cost declines for each of 136 countries (Table A.12), breaking out the mean results from 136 targeted countries (rows one and two) into those 15 targeting entrepôts (rows three and four) and all others (rows five and six). Columns (1)-(4) present aggregate welfare changes. Columns (5)-(8) present changes to aggregate container trade. Columns (1), (2), (5), and (6) present results for cases where non-transportation trade costs are reduced. Columns (3), (4), (7), and (8) present results for cases where transportation costs are reduced (infrastructure improvements). Odd columns correspond to cases where no scale economy feedback loops are allowed. Even columns present results allowing for scale economy feedback. In each case, we report the mean impact and its standard deviation in parentheses.

effects of each shock on the transportation network, augments this to 5-fold. In Columns (3) and (4), infrastructure investments at entrepôts generate on average 10 times the global welfare impact relative to investment elsewhere.

Table A.14 compares the welfare impact of 136 counterfactuals in each of the four cases, the relationship between a welfare increase from a given reduction in trade costs and the entrepôt status of the targeted location, controlling for GDP at the targeted location, the distance between targeted and impacted countries, and a fixed effect for impacted country. The latter controls for whether a specific impacted country is particularly sensitive to trade cost reductions. There are 18,340 bilateral pairs of targeted and impacted countries. Regressions are weighted by impacted country’s GDP.

The strong controls in these regressions reduce the differential impact of entrepôts: Column (1) controls for gravity variables and impacted fixed effects, fully accounting for entrepôt countries’ raw positive impact (double non-entrepôts’ in Table A.12). However, once scale economies’ impact on the transportation network are accounted for (in Column (2)), the impact from counterfactuals targeting entrepôt countries are an order of magnitude larger. In Column (3), when the transportation network is directly impacted by infrastructure investment, entrepôts are at baseline more than two-thirds more impactful (52 log points) and over 200% more impactful when scale economies are allowed.

Note that because some welfare effects are negative, we add a constant to all results before taking logs. This makes the indicator variable not directly comparable to the raw numbers in Table A.12. However, the relative size and direction of results are robust to using raw percent changes on the left-hand side.

Figure A.14 repeats the exercise in Figure 12 for non-transportation cost reductions with and without the endogenous response of costs throughout the network when accounting for scale economies. The black bars in Panel (A) underscore that without transportation network impacts, smaller entrepôts are generally not pivotal. The grey bars show results accounting for the endogenous response of the transportation network. The dramatic difference for Singapore in particular underscores that conflating network changes with non-network adjustments such as tariff changes can bias results.

In Panel (B), we plot the results for each country with and without scale. Here the average relationship as well as the average error is nearly identical as in Figure 12, as is the bias at entrepôts.

Brexit Figure A.15 shows the impact of our two counterfactual cases on the UK’s 20 largest trading partners in welfare percent changes. Black bars show the impact of

Table A.14: Bilateral Welfare Impacts, by Entrepôt status

	ln % Δ Welfare from Non-Transport Cost Reduction		ln % Δ Welfare from Transport Cost Reduction	
	$\Delta\kappa$ (1)	Δt with Scale (2)	Δt (3)	Δt with Scale (4)
$\mathbb{1}_{entrep\hat{o}t} = 1$	-0.0428 (0.0393)	0.655 (0.117)	0.517 (0.0870)	1.145 (0.181)
ln GDP, targeted 0.220	0.364 (0.0130)	0.495 (0.0219)	0.568 (0.0335)	(0.0255)
ln Distance	-0.352 (0.0540)	-0.551 (0.0531)	-0.353 (0.0455)	-0.435 (0.0491)
Obs	18340	18340	18340	18340
R^2	0.575	0.729	0.817	0.809

Notes: Results weighted by impacted country GDP. Outcome values are shifted by a constant in order to include negative values. Standard errors in parentheses are clustered two ways by targeted and impacted countries. Columns (1) and (3) correspond to cases where no scale economy feedback loops are allowed. Columns (2) and (4) present results allowing for scale economy feedback.

increased non-transportation trade friction with the UK. Grey bars show the impact with scale effects changing transportation costs through the UK. All partners experience outsized losses due to scale economies. Most of these losses come through increased trade costs in the Netherlands and Belgium, which far from benefiting from our counterfactual, lose because of decreased volumes as well. Ireland in particular, which our microdata tells us sends 50% of goods to the US through the UK, experiences large additional losses.

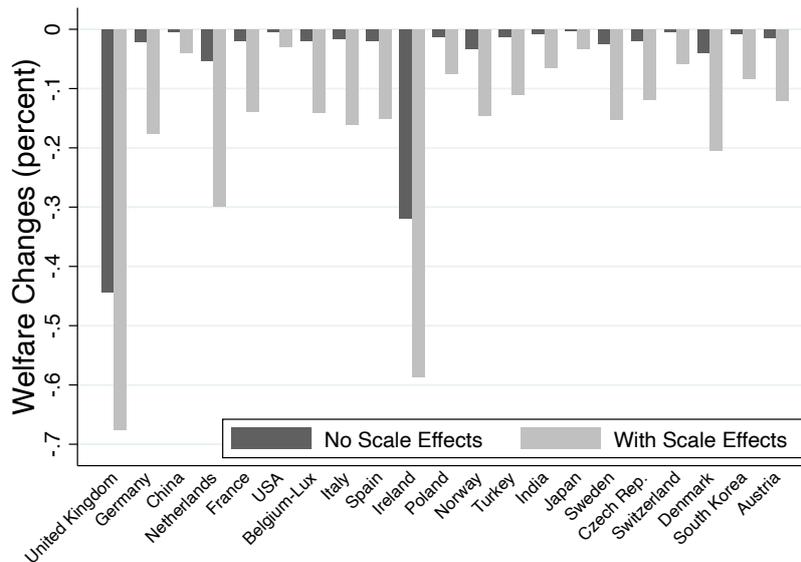
Global trade volume changes under these two cases are reported in Figure A.16. These results largely mirror our welfare results in the main text.

The Opening of the Arctic Passage Figure A.17 shows changes in the relative wage-adjusted price index (interpreted as national welfare, if we omit the costs of climate change) across the three cases.¹⁰ In the baseline scenario in Panel (A), we see increases in trade between countries that are along the Northeast passage, and small spillover impacts at countries not directly impacted—reflecting classic multilateral resistance and cascading effects from value chains. Figure A.17 Panel (B) shows how, through indirect trade, the benefits of the passage pass on to nearby countries not directly impacted. In Panel (C), scale economies amplify these effects.

Since some of the Asian entrepôts are smaller and harder to see on a global map,

¹⁰Appendix Figure A.19 shows related changes in country-by-country containerized exports.

Figure A.15: Welfare Changes - Brexit - Largest Trading Partners



Notes: Bars show the percent change in welfare (the relative price index) of a simulated 5% increase in trading costs with the United Kingdom the largest 15 trading partners. The first bar reflects changes if shipping costs remain constant, reflecting only welfare changes due to changes in prices. The second bar allows for endogenous network adjustment to scale economies.

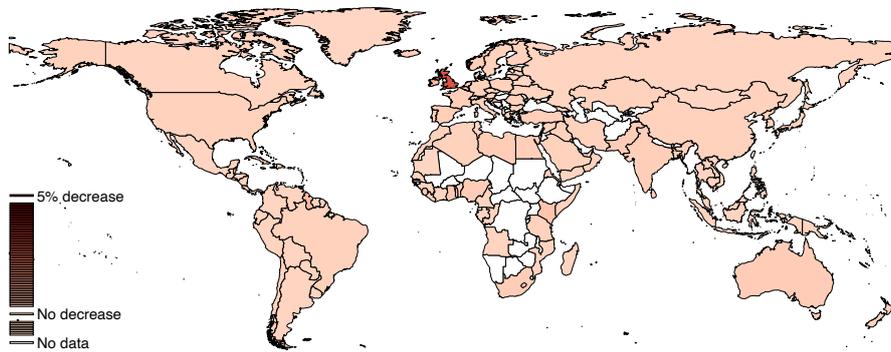
Figure A.18 zooms in on the welfare changes of Singapore, Hong Kong, and Taiwan as well as their surrounding countries as a result of the opening of the Arctic Passage. In the baseline scenario in Panel (A), we see that these entrepôts have a direct welfare increase from the passage opening since they have direct routes to Northern European countries and North America. When allowing for indirect trade in Panel (B), the neighboring countries of these entrepôts see an increase in welfare because they are now able to benefit from using these entrepôts to trade with the Northern European countries and North America. When allowing for scale economies to amplify effects in Panel (C), the entrepôts and their neighboring countries are going to benefit even further as a result of this indirect trade.

The concentration of welfare gains in entrepôts from this counterfactual highlights a novel source of agglomeration—scale economies in transportation and transport networks can help contribute to and shape entrepôts. This is further explored in our first counterfactual in Subsection 8.2.

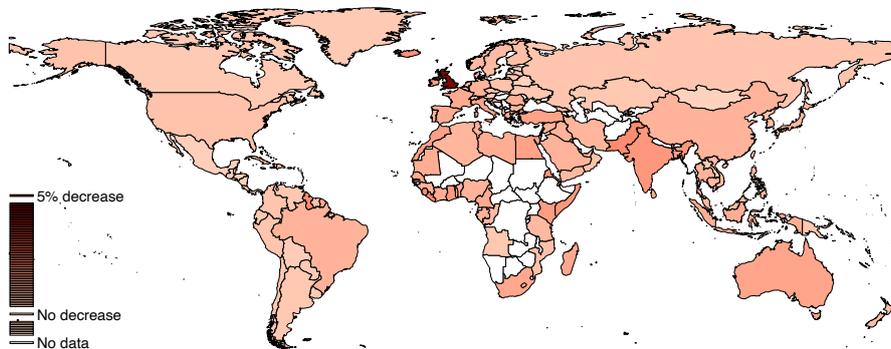
Figure A.19 reports global trade volume changes under the three cases. These results highlight the significant heterogeneity in trade changes across countries and largely mirror our welfare results in the main text.

Figure A.16: Export Volume Changes - Brexit

(A) Trade Cost Change, No Network Scale Effects



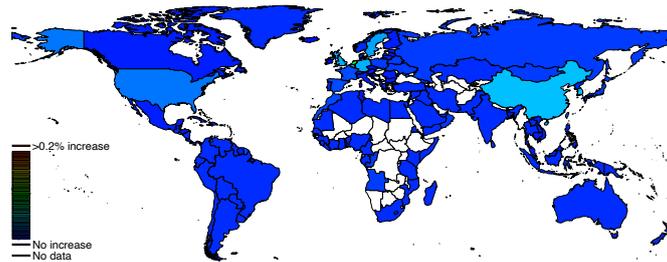
(B) Full Trade Network Effects and Scale Economies



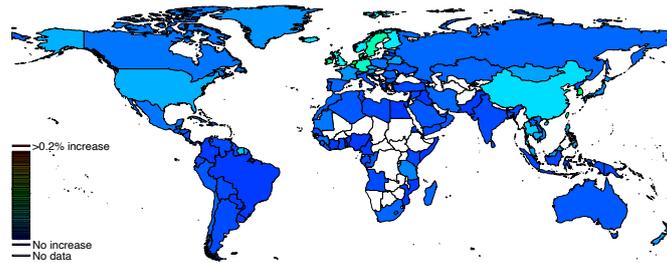
Notes: These two plots show the percent change in exports of a simulated 5% increase in trading costs with the United Kingdom for all countries in our dataset. Darker reds reflect a greater increase. White represents omitted countries. Panel (A) reflects changes if shipping costs remain constant, reflecting only trade changes due to changes in prices. Panel (B) allows for endogenous network adjustment to scale economies.

Figure A.17: Welfare Changes - Arctic Passage

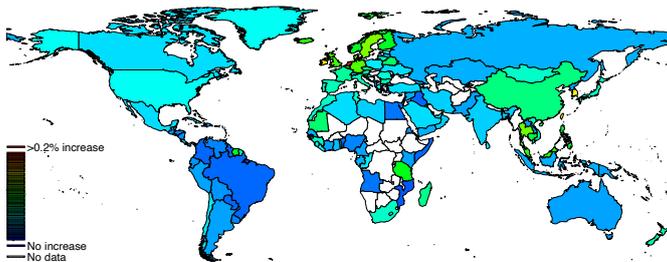
(A) Only Directly Affected Routes (Exogenous Trade Costs)



(B) Full Trade Network Effects



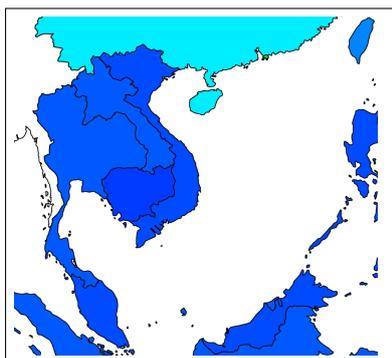
(C) Full Trade Network Effects and Scale Economies



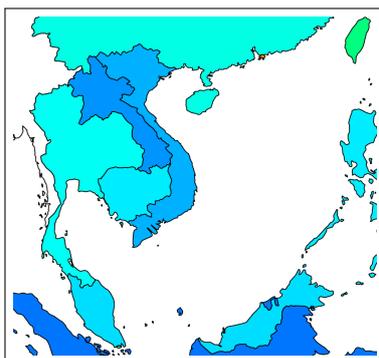
Notes: Plots show the percent change in welfare (the relative price index). Darker reds reflect a greater increase and blue represents no change. Omitted countries are white. Panel (A) reflects changes only allowing trade costs to decrease on routes whose distance is directly reduced to the Arctic Passage. Panel (B) reflects changes allowing all countries to indirectly access the Arctic Passage through the trade network. Panel (C) allows for the network's endogenous response to scale economies.

Figure A.18: Welfare Changes on Asian Entrepôts - Arctic Passage

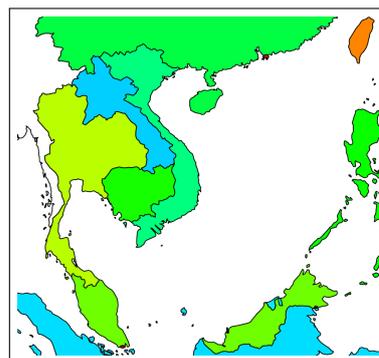
(A) Only Directly Affected Routes



(B) Full Trade Network Effects



(C) Full Trade Network Effects & Scale Economies



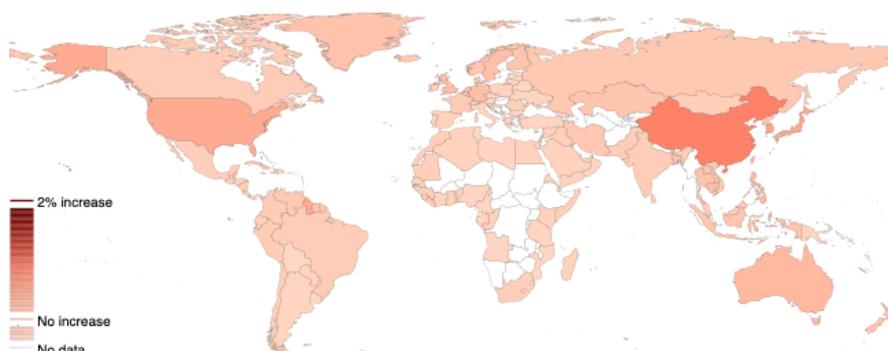
Notes: These three plots are a magnified part of figure A.17 to show the percent change in welfare (the relative price index) for a subset of Asian Entrepôts in our dataset. Darker reds reflects a greater increase and blue represents no change. White represents omitted countries. Panel (A) reflects changes if we only allow trade costs to decrease on routes whose distance is directly reduced to the Arctic Passage. Panel (B) reflects changes if we allow all countries to indirectly access the Arctic Passage through the trade network. Panel (C) allows for the endogenous network response to scale economies.

Figure A.19: Export Volume Changes - Arctic Passage

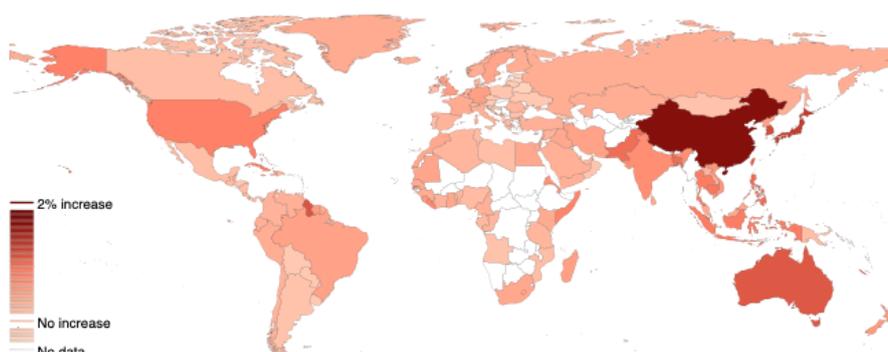
(A) Only Directly Affected Routes



(B) Full Trade Network Effects



(C) Full Trade Network Effects and Scale Economies



Notes: These three plots show the percent change in exports from all countries in our dataset. Darker reds reflect a greater increase in exports. White represents omitted countries. Panel (A) reflects changes if we only allow trade costs to decrease on routes whose distance is directly reduced to the Arctic Passage. Panel (B) reflects changes if we allow all countries to indirectly access the Arctic Passage through the trade network. Panel (C) allows for the endogenous network response to scale economies.