American Economic Journal: Macroeconomics 2024, 16(4): 239–278
https://doi.org/10.1257/mac.20220250

Entrepôt: Hubs, Scale, and Trade Costs†

By Sharat Ganapati, Woan Foong Wong, and Oren Ziv*

We study the global trade network and quantify its trade and welfare impact. We document that the trade network is a hub-and-spoke system where 80 percent of trade is shipped indirectly and largely via entrepôts—major hubs that facilitate trade between many origins and destinations. We estimate indirect-shipping-consistent trade costs using a model where shipments can be sent indirectly through an endogenous transport network and develop a geography-based instrument to estimate scale economies in shipping. Network and scale effects propagate local trade cost changes globally. Counterfactual infrastructure improvements at entrepôts generate ten times the global welfare impact relative to nonentrepôts. (*JEL* F12, F14, L92)

Exchanging goods over borders involves more than production and consumption: shipping, transshipping, and distribution can include multiple agents and additional countries beyond producers and consumers. These activities are concentrated at entrepôts—trading hubs that goods travel through and from other origins and are bound for other destinations. The idea that entrepôts are integral to the trade network and are engines of growth has been the impetus behind many policies aimed at attaining or maintaining entrepôt status (Smith 2015; Nidhi and Das 2016; Paris 2021).

This paper studies entrepôts, the trade network they form, and their impact on international trade. Using novel data on the trade network and developing a quantitative general equilibrium spatial trade model, we answer the following questions: (i) How do goods move from their origins to their destinations and

* Ganapati is Assistant Professor of International Economics at Georgetown University and Faculty Research Fellow at the National Bureau of Economic Research (email: sg1390@georgetown.edu); Wong is Assistant Professor of Economics at the University of Oregon and Research Afilliate at the Centre for Economic Policy and Research (email: wfwong@uoregon.edu); and Ziv is Assistant Professor of Economics at Michigan State University (email: orenziv@msu.edu). Simon Gihchrist was coeditor for this article. We thank Treb Allen, Costas Arkolakis, Panle Jia Barwick, Bruce Blonigen, Johannes Boehm, Mark Colas, Kerem Coşar, Anca Cristea, Meredith Crowley, Carsten Eckel, Stefania Garetto, Matthew Grant, Keith Head, Reka Juhasz, Myrto Kalouptsidi, Steve Matusz, Ezra Oberfield, Nina Pavcnik, Andrés Rodríguez-Clare, Robert Staiger, and Meredith Startz; seminar participants at Stanford University, Princeton University, University of Michigan, University of Virginia, Boston College, the Federal Reserve Board, Ludwig-Maximilians-Universität Munich, Indiana University, University of Michigan, University of Notre Dame, Syracuse University, University of Cambridge, University of Mannheim, and University of Warwick; and participants at the NBER Conference on Cities, Labor Markets, and the Global Economy, 2019 Mid-Atlantic International Trade Workshop (Duke University), 2020 CeMENT Mentoring Workshop, 20th Annual Nordic International Trade Seminars, 2021 NBER Summer Institute and other conferences for helpful comments. Andrew Castro, Ray McCormick, Giacomo Romanini, and Philip Valtadoros provided outstanding research assistance.

† Go to https://doi.org/10.1257/mac.20220250 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

what role do entrepôts play in facilitating this process? (ii) What trade costs and scale economies can explain the observed routes that goods take and the existence of entrepôts? (iii) How does this pattern of trade through entrepôts impact global and regional trade as well as welfare?

We start by constructing a new dataset mapping the journeys that containerized shipments take through the global trading network. This microdata allows us to observe indirect trade, which we define as trade journeys that make stops with the shipment either on-board or transshipped—transferred onto a ship—at additional countries beyond the shipment's origin and destination.

Our first contribution is to establish two stylized facts about the global trade network. Our first stylized fact is that the majority of trade—80 percent—is shipped indirectly. The median shipment stops at two additional countries before reaching its destination. The majority of trade is also transshipped via an additional country before its destination. This indirectness is not incidental and significantly increases shipping times and distances.

Our second stylized fact is that indirectness is incredibly concentrated, with over 90 percent of indirect trade channelled through a small number of entrepôts, establishing a hub-and-spoke network. These facts highlight a trade-off and trace the existence of a potential scale-cost relationship: indirect trade concentrated through entrepôts increases the observable distance and time costs of trade, but by revealed preference, it implies lower trade costs, especially for the spokes of the network that disproportionately choose to ship via entrepôts.

In order to rationalize the documented direct and indirect trade through the global trading network, we build a general equilibrium model of trade with entrepôts and endogenous trade costs that flexibly accommodates input-output linkages. Producers choose shipping routes and compete for foreign consumers in a generalized Ricardian setting. Low-cost routes can involve indirect shipping through additional countries, and entrepôts endogenously arise where trade costs are lowest. We allow for both scale economies and diseconomies to govern shipping costs on these network links.

Our second contribution is to use our model to estimate a global set of indirect-shipping consistent trade costs and the economies of scale in shipping. Expanding from our microdata to global seaborne container shipping and trade data, our estimation yields trade costs for each link of the global shipping network and a global set of model-consistent origin-destination trade costs that are distinct from typical distance-based costs. We establish the validity of both our estimates and modeling approach by finding a tight match between our estimated trade costs and external freight rate data, as well as between our model-predicted network flows and microdata on shipment journeys. Our trade cost estimates are publicly available online.

We use a geography-based instrument to identify the causal effect of increasing shipping volumes on decreasing trade cost using an instrumental variable approach. Embedded in our model is the intuition that some links have inherently higher traffic because of their geographic position in the network. For example, links that include Singapore are close to the lowest-distance route between many European and Asian countries due to Singapore's location in the Strait of Malacca. For each link, we compute the distance to and from the link relative to the shortest distance between each origin and destination, recovering a weighted average of each link's proximity to global trade. Increasing traffic volume on a link by 1 percent reduces costs by 0.06 percent. As the typical journey in our microdata has 2.5 links, a 10 percent increase in overall origin-destination trade translates into a 0.17 percent decrease in trade costs.

Our third contribution uses our estimates and model to quantify the impact of the trade network on global trade and welfare, highlighting how trade cost changes at node countries—entrepôts and nonentrepôts—as well as links can have widespread impacts through the network that are subsequently magnified due to scale economies. Our main counterfactual quantifies the trade and welfare benefits of transport infrastructure improvements for each country in our sample. Entrepôts are pivotal to the global trade network: welfare impacts of infrastructure investment are, on average, 10 times higher at entrepôts than nonentrepôts. Conflating transport and nontransport trade costs impacts estimated welfare effects by an order of magnitude. This is especially true at entrepôts, which differentially concentrate infrastructure improvement benefits locally relative to nonentrepôts. Scale economies in transportation further concentrate these gains locally at and around entrepôts, highlighting that scale economies in transportation act as a source of agglomeration. We establish that Singapore and Egypt (the Suez Canal) are the top two most pivotal locations in the trade network, as reflected by the strain in global supply chains when Egypt was blocked in March 2021 (Paris and Malsin 2021; Sheppard, Dempsey, and Saleh 2021; Gambrell and Magdy 2021).

Our second counterfactual investigates how nontransportation cost changes at an entrepôt can have widespread impacts beyond the countries that are directly impacted through endogenous adjustments in trade network. We illustrate this by studying the ramifications of worsening trade relations between one hub—the United Kingdom—and its trading partners—Brexit. When only considering the direct impact of increased nontransportation trade costs, Brexit's consequences are largely proportional to a country's direct trade exposure with the United Kingdom. When our analysis accounts for the impact of scale economies on the trade network, we find that smaller countries like Ireland and Iceland that use the United Kingdom as an entrepôt to access all other trading partners are disproportionately hurt (as recognized in *Financial Times* 2020). This illustrates how trade network and scale interactions can lead to distinct distributional outcomes in welfare even when the initial changes are unrelated to transport.

Our last counterfactual evaluates the importance of endogenous trade costs by demonstrating the welfare and trade impacts from the two endogenous mechanisms in our model: (i) network effects—allowing countries to ship indirectly—and (ii) scale effects—allowing countries to ship indirectly and take advantage of scale economies. To illustrate this, we study the effects of opening up the Arctic Ocean to regular year-round shipping, connecting countries in East Asia and Europe. Allowing for network effects doubles the welfare relative to a naïve exogenous trade cost case with no network effects, and allowing for scale economies triples the welfare relative to the network effects case.

This paper ties two broad literatures together, combining detailed microdata on the flow of goods through the trade network with a structural model of trade and transportation. The first dives deeply into the technology underpinning the fundamentals of international trade, such as container shipping and infrastructure investment (Coşar and Demir 2018). The second considers the geography and cost structures of transportation networks within a class of gravity models (Head and Mayer 2014; Allen and Arkolakis 2022).

With regard to the technologies underpinning trade, we make two contributions. First, a wide literature shows how both containerization and infrastructure investments have local outcomes (Heiland et al. 2019; Ducruet et al. 2019; Wong 2022; Coşar and Demir 2018; Bernhofen, El-Sahli, and Kneller 2016; Rua 2014).[1] We demonstrate the global welfare impacts of the container shipping network, which accounts for two-thirds of annual trade moved by sea (World Shipping Council 2018). Using our general equilibrium spatial trade framework, our counterfactuals show how endogenous changes in trade costs propagate via the network and through entrepôts as well as quantify their trade and welfare impacts. Allowing for network effects double the welfare relative to a baseline case with no network effects and allowing for the effect of scale economies further triples welfare impacts.[2]

Second, we explore the general equilibrium effects of scale economies in shipping. For the median route into the United States, our leg-level scale economy implies that a 10 percent increase in volume leads to a 1.7 percent decrease in costs.[3] The role of localized scale economies in production is well known in general (Allen and Arkolakis 2014; Allen and Donaldson 2018) and in the context of trade in particular (Lashkaripour and Lugovskyy 2023; Bartelme et al. 2019; Kucheryavyy, Lyn, and Rodríguez-Clare 2019). In these settings, scale economies typically generate agglomerations by acting on local productivity. By contrast, in our setting, scale economies generate agglomerations by affecting trade costs. Our counterfactuals find that, by acting on endogenous transport costs over the network, scale economies further concentrate transportation,trade, and welfare gains at entrepôts.

With respect to the geography and structure of the trade network, we make two contributions. First, we provide empirical evidence for a growing quantitative literature investigating the role of trade networks (Allen and Arkolakis 2022; Fajgelbaum and Schaal 2020; Redding and Turner 2015). We provide the first systematic documentation of indirect trade through the containerized shipping network and the pivotal role

[1] Hummels, Lugovskyy, and Skiba (2009); Grant and Startz (2022); and Asturias (2020) study transport costs in the context of market power. While container shipping firms may hold market power, we generalize away from the profits of the shipping companies. Models allowing for leg-level oligopoly, fixed costs, and endogenous entry competition fit within our framework (Sutton 1991), but we leave the study of how market power works through the hub-and-spoke network for future study.

[2] Allen and Arkolakis (2022) studies the endogeneity of trade costs to traffic congestion on highways. We find the presence of scale economies in shipping. Brancaccio, Kalouptsidi, and Papageorgiou (2020) studies two aspects of trade cost endogeneity for the network of dry bulk ships carrying homogeneous commodities where all trade is direct: the loading opportunities of dry bulk ships after delivering their cargo relative to the country's trade balance (the equilibrium bargaining position of these ships) and the trade balance of neighboring countries (the network effects). Wong (2022) focuses on the round trip effect from container shipping: a bilateral trade cost endogeneity.

[3] Our estimate is about three-quarters of the estimates in Asturias (2020) and Skiba (2017). Asturias (2020) reports an origin-destination country trade-volume trade-cost elasticity of 0.23 while Skiba (2017) reports an elasticity of 0.26 using product-level import data from Latin America. See also Alder (2015); Holmes and Singer (2018); and Anderson, Vesselovsky, and Yotov (2016).

that entrepôts play within this network.[4] Our microdata on the movement of shipments through the trade network document the widespread nature of indirect trade and its concentration. In contemporaneous work, Heiland et al. (2019) studies the impact of the Panama Canal expansion on global ship movements and uses model-based imputations to estimate the physical movement of goods. We further estimate a set of network-consistent trade costs, distinct from and more predictive of trade than distance. Finally, our counterfactuals demonstrate how transport costs behave differently from nontransport costs, particularly at entrepôts. For example, Egypt ranks top two in terms of global welfare impacts from infrastructure improvements, while it is not among the top 20 in terms of the welfare impacts from nontransportation trade cost reductions.

Second, our model embeds transportation networks within a class of gravity models (Head and Mayer 2014). We extend the Armington framework in Allen and Arkolakis (2022)—where route cost shocks are born by consumers—to a general Ricardian setting—where traffic volumes reflect both route choice and head-to-head competition on prices at destinations and demonstrate how to estimate the model in a multi-industry setting with nontransport barriers to trade and in the presence of unobserved traffic flows. Methodologically, we adopt an approach from the literature on marginal cost estimation (Ackerberg et al. 2007), combining market level data and exogenous instruments with equilibrium assumptions—the indirect routing of trade in our case, or market conduct in the Industrial Organization literature's case—to recover unobserved costs. We establish that our estimates reflect actual costs and indirect flows by comparing our model predictions to external cost estimates, ship sizes, and observed trade routes in our microdata. These results serve as a check to the validity of the Allen and Arkolakis (2022) framework within the international trade setting.

I. Data

Our paper uses two distinct sets of data. To establish the stylized facts of the international trade network (Section II), we use a microdata on the detailed journey of US-bound shipments. To estimate global trade costs that are network consistent (Section IV), we use global data on trade and shipping traffic.

To construct the microdata on US shipments, we merge two proprietary datasets: global ports of call data for containerships, which allow us to reconstruct the routes taken by specific ships, and the US bill of lading data for containerized imports, which gives us shipment-level information on US imports. Independently, these datasets partially describe the global shipping network. Merged, they reconstruct the journey of individual shipments as they navigate the trade network from their origin to their US port of entry. To our knowledge, we provide the most comprehensive

[4] The emergence of entrepôts as hubs in geographically advantageous locations is consistent with the findings of Barjamovic et al. (2019). This is related to studies of airlines hub-and-spoke networks. This literature takes large parts of the network as fixed (Berry 1992) or is restricted to simple entry games (Ciliberto and Tamer 2009).

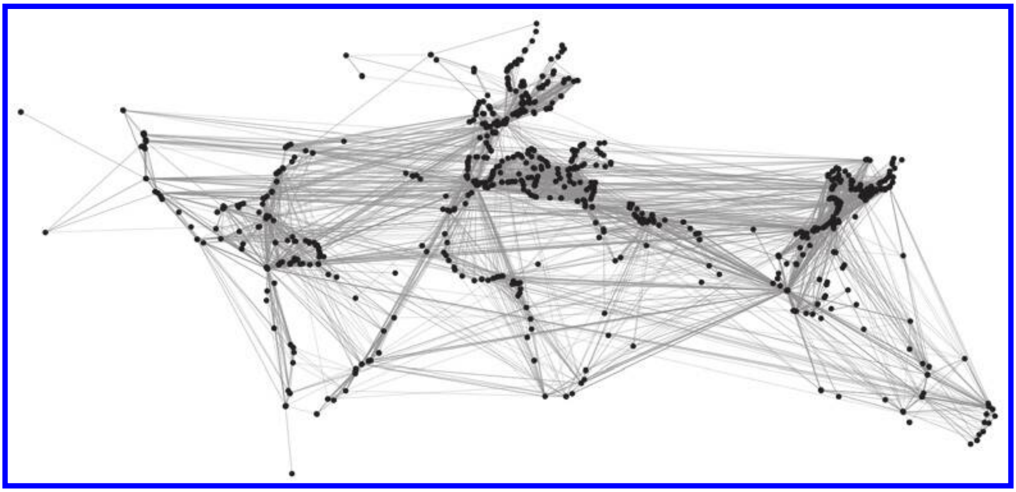

Figure 1. Global Network of Ships and Ports of Call Data

Notes: Dots represent 1,230 ports. Lines represent journeys between port pairs undertaken by a containership (total of 4,986 ships). We show direct distances here. Analysis uses sea-route distance.

reconstruction of the global trading network and routes undertaken by individual shipments into the United States.[5]

Our ports of call data capture vessel movements using automatic identification system (AIS) transponders.[6] For each vessel, these data capture the vessel's characteristics, time-stamped ports of call, capacity, and height in the water before and after stopping at each port. The latter two pieces of information indicate the vessel's load at these ports, allowing us to observe volumes shipped between port pairs. We measure volume in twenty-foot equivalent container units (TEUs).

Our sample covers 4,986 unique container ships with a combined capacity of 30.6 million TEUs—over 90 percent of the global container shipping fleet—making 397,625 calls at 1,230 ports from April to October 2014. Figure 1 shows the coverage of the shipping network in our port of call data. Each line represents a containership journey. We use these global data along with Centre d'études Prospectives et d'Informations Internationales (CEPII) global trade data—aggregated into containerized and noncontainerized industries according to the procedure outlined in online Appendix A.3—to estimate our model in Section IV.

With these port of call data alone, shipment journeys within the trading network remain unobserved. We do not observe containers being loaded or unloaded. To remedy this, we merge the port of call data with US bills of lading data, which capture shipment-level information for all containerized imports. We observe each shipment's origin country, the port where they are loaded onto containerships (also known as port of lading), and the US port where they are unloaded (port of unlading). We observe the name and identification number of the containership

[5] Online Appendix A.1 explains both datasets and their merge procedure in detail.

[6] Port receivers collect and share AIS transponder information (including ship name, speed, height in water, latitude, and longitude). Using Astra Paging (2014) data, we track global port entry and exit data.

that transported the shipment as well as the shipment's weight, number of TEUs, and product information. Over the same six-month period, we see a total of 14.8 million TEUs weighting 106 million tons were imported into the United States from 227 origin countries and loaded onto US-bound containerships (laded) in 144 countries. This accounts for about three quarters of the 2014 TEU and tonnage imports—77 percent and 74 percent respectively (Maritime Administration 2018).

Using details on containerships, ports, and arrival times, we reconstruct each shipment's journey from its foreign origin to its US destination by matching each shipment to the containership that it was transported on (online Appendix Figure A.1 visualizes this merge). While the shipments' exact journey between origin and the first stop (the port where they are loaded onto containerships) remain unobserved, this initial portion can either take place overland (by trucks or rail) or by sea on another containership because they are containerized. Not observing this portion, in fact, leads us to undercount the overall level of indirectness. We empirically deal with unobserved transit in Section IV.

II. Stylized Facts

We analyze the international trade network and the routes taken by goods entering the United States along that network. We find that the majority of trade takes place indirectly in a manner that is costly—increasing both shipping time and distance traveled. We further show that the global trade network is a hub-and-spoke system, concentrating a large number of shipments through a small number of entrepôts.

A. *The Majority of Trade Is Indirect*

Panel A in Figure 2 reports the distribution of the number of observed country stops made by each shipment, weighted by TEU containers. Only 20 percent of containers are exported to the United States directly from their origin countries, making no stops in between. The average container entering the United States stops at around two third-party countries that are neither the origin nor destination.[7] The map in Figure 2, panel B shows that this is also true at the country level: the majority of US trading partners export to it indirectly. Only shipments from nine countries typically enter the United States directly.[8] Similarly, the average shipment from a majority of US trading partners is transshipped in a third-party country—60 percent of US trading partners transship more than 90 percent of their US-bound goods.[9] Figure A.5 reports the percent of goods transshipped at third-party countries.

We explore the high degree of variation in connectivity in online Appendix B.4, showing that this variation is, in part, explained by traditional gravity variables. We

[7] Mean of 1.5 and standard deviation of 1.3. Landlocked countries are excluded. The average number of port stops is higher (Figure A.3, mean of 4.6 and standard deviation of 3.5). This result is robust for shipment weight and value (Figure A.4). Multiple stops at the same third-party country are not counted.

[8] These countries are Canada, Mexico, Panama, Japan, South Korea, Spain, Portugal, South Africa, and New Zealand. We treat mainland China, Hong Kong, Taiwan, and Macau as separate locations.

[9] Both on-board stops and transshipment are important measures of indirect trade. For completeness, all results are broken out here or in the online appendix using transshipment only. Examples of countries transshipping more than 90 percent of goods include Denmark, Bangladesh, Cambodia, and Ecuador.

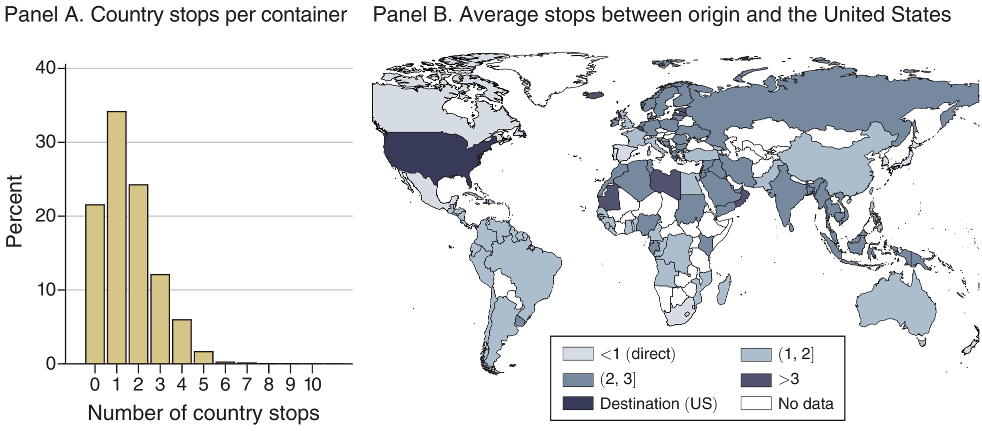

Figure 2. Indirect Trade Distributions by Container and Country

Notes: Panel A shows the distribution of containers by number of unique third-party countries visited. In panel B, for each origin country, we calculate the average number of third-party country. The destination country (United States) is excluded (in white). Plots are at the shipment level and weighted by the aggregate exported containers (TEU). Landlocked countries are also excluded (in white), since they would mechanically need to stop at a coastal country. The missing remaining countries are excluded either due to lack of overall trade with the United States (e.g., Somalia) or due to the merge process (e.g., Namibia).

show that there is substantial variation in routes from unique origins into the United States, which is an important assumption in our model and is used in our validity checks (online Appendix Figure A.9, panel B).

Indirect Trade Increases Shipping Distances and Time.—Are the additional country stops simply incidental stops along the way, or do they constitute a trip that is distinct from a direct path? One possibility is that the observed indirectness is optimal but only incidental—perhaps additional stops only have small effects on costs, and therefore may be optimal, even if the benefit of indirectness is small. As an example, goods transiting the Strait of Malacca can perhaps stop at Singapore since it is on the way. However, the significant additional distance and time incurred by indirect travel relative to the direct path, documented here, implies this is unlikely to be the case.

On average, the actual traveled distance between a shipment's origin and its US destination is 31 percent more than its direct ocean distance (panel A in Figure 3). Panel B shows the actual traveled distance between the location where the shipment was last loaded onto a ship and its final destination. Here, the remaining gap is still substantial at 23 percent. Online Appendix Table A.1 further evaluates the relationship between indirectness and journey length. Controlling for direct journey length or origin-by-destination fixed effects, doubling the number of stops adds 10 percent to distance traveled and 33 percent to time traveled (columns 2 and 5 in online Appendix Table A.1, respectively). These distance and time costs do not include pecuniary costs of transshipment. Consequently, this indirectness is meaningful in the sense that it is costly. These longer shipping routes imply a cost reduction from

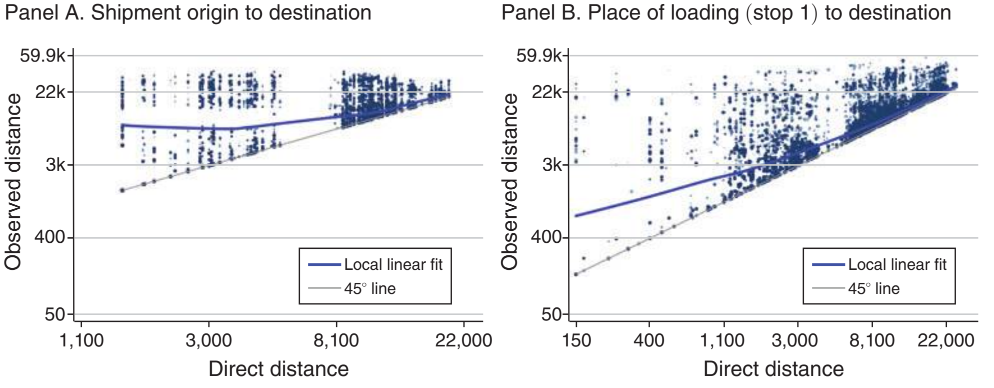

Figure 3. Difference between Observed Distance and Direct Distance

Notes: These figures show only indirect shipments with different direct and observed distances. Dots are shipments, shaded by TEU. Panel A compares the direct shipping distance from the shipment's origin country to the United States with the actual route traveled. Panel B compares the direct distance from the place a shipment was last loaded onto a US-bound ship (Stop 1 in online Appendix Figure A.1) with the actual route traveled. Sea distances for observed and direct routes are calculated using Dijkstra's algorithm. The local linear fit line is a locally weighted regression of the observed on direct pair-wise distance.

indirectness that is over and above the additional time and distance costs. From these results, we can summarize our first stylized fact:

Stylized Fact 1: The majority of containerized trade into the United States is indirect and results in a significant increase in shipping distance and time.

B. *Indirect Trade Is Routed through Entrepôts*

When shipments stop in third-party countries, how are they routed? We show that the stops along indirect shipping routes are not arbitrarily distributed throughout the world. Instead, they are channelled through a small number of hubs, which disproportionately service shipments originating in other countries.

Panel A of Figure 4 plots each country's share of total third-party country stops against its share of total US trade. Some locations are both popular stopping points and major countries of origin for goods like China, Germany, and Japan. Key countries like Korea, Singapore, Panama, and Egypt disproportionately participate as third-party countries in US-bound shipments.[10] This leads to our measure of entrepôt activity:

$$Entrep\hat{o}t_{l,j} \equiv \pi_j^l - \pi_{l,j} \tag{1}$$

[10] Figure A.6 tabulates the percent of all goods entering the United States stopping in that country, broken into goods originated there and elsewhere.

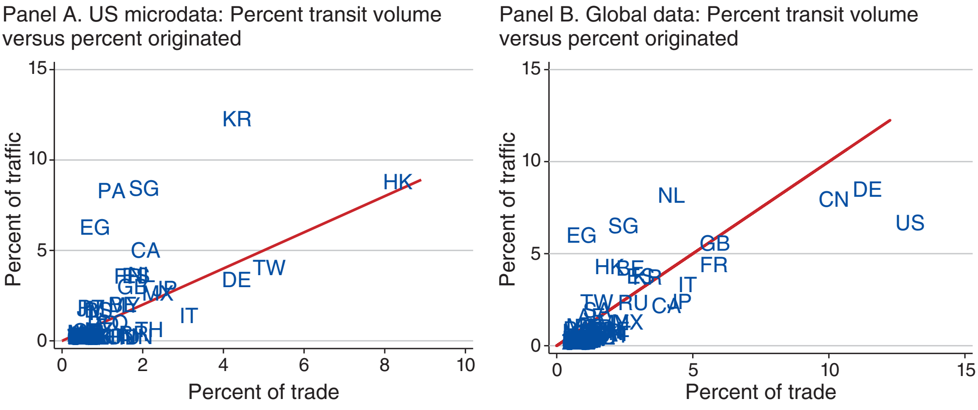

Figure 4. Concentration of Indirect Shipments

Notes: Panel A uses US microdata to compare, for each country, the share of US imports that originated in a country (x-axis) to the the share that passed through that country (y-axis), weighted by TEU. For readability, China is omitted in panel A. Panel B replicates panel A using global port of call and trade data with adjustments made for unobserved overland traffic as discussed in Section IV.

where country j's usage of entrepôt l for its imports is the difference between π_j^l—the share of j's imports flowing through l—and $\pi_{l,j}$—the share of j's imports originating at l. This captures the use of location l above and beyond its role as an exporter to j.[11]

Panel B of Figure 4 repeats the exercise in panel A using global traffic minus trade shares.[12] While the results are broadly consistent with the microdata in panel A, some countries such as Canada and Panama that are specifically integral to the US network are now below or closer to the 45-degree line. In both panels, third-party country stops (the y-axes) are significantly more concentrated than trade (the x-axes).[13] Our measure of entrepôt activity in equation (1) is the distance to the 45-degree line. Online Appendix Table A.2 lists our measure for all the countries and territories in our data, normalized by the value of the country with the lowest measure: the United States.

Definition of Entrepôts.—We define the top 15 countries using this metric as our set of global entrepôts, a natural break after which the measure rapidly flattens (online Appendix Table A.2). This list of 15 includes several well-known global hubs, but our results are robust to changes in this threshold as well as to using a continuous measure.[14] This threshold and definition will be used again in counterfactual

[11] $Entrepôt_{l,j}$ is directly proportional to the total volume of goods moving through l that do not originate at l. Online Appendix C shows how this measure arises from our model as the difference between l's on-board marginal cost selling to j and its network relation to j, and that lowering location l's leg-level transport costs to other origins increases $Entrepôt_{l,j}$. Our results here and throughout are robust to other functional forms—for example log differences.

[12] We subtract country l's share of observed global containerized trade π_l from its observed share of global container traffic π^l, with an adjustment for unobserved overland traffic as described in Section IV. Online Appendix C clarifies how this is a consistent aggregation of the country-level measure in equation (1).

[13] Table A.3 reports the concentration ratios for trade, transshipment, and third-party-country stops.

[14] Our set of global entrepôts are: Egypt, Singapore, Netherlands, Hong Kong, Belgium, Taiwan, Spain, Saudi Arabia, South Korea, the United Arab Emirates, Morocco, Panama, Malta, Portugal, and the United Kingdom.

analyses, where we explore the impacts of cost changes at these hubs. For US shipments, we see 73 percent of all shipments pass through at least one entrepôt. Of indirect shipments, 92 percent pass through an entrepôt.

Additionally, we find that smaller origin countries disproportionately use entrepôts. They are simultaneously more likely to ship their goods indirectly and more likely to use entrepôts (see online Appendix B.3 and Figure A.7 for further details). Jointly, this confirms that smaller countries are spokes that disproportionately use entrepôts for their trade.[15] These relationships can be summarized in our second stylized fact:

Stylized Fact 2: Indirect shipping routes are concentrated through entrepôts. International trade occurs over a hub-and-spoke network.

Our two facts outline an inherent trade-off: indirectness increases observable distance and time costs of trade, but by revealed preference, implies lower costs, especially for the spokes of the network that disproportionately choose to send goods indirectly through entrepôts.[16] The goal of our empirical estimation is to measure this trade-off within the context of the full global trading network by finding a set of node-to-node costs that describes the shipping network and is consistent with the indirect trade we observe.

These facts also trace the existence of a size-cost relationship: shipment along high-concentration entrepôts routes appears, by revealed preference, to be cost reducing. As with any scale-cost relationship, both directions of causation may be operational. We model the shipping decision in a way that allows for but does not impose a reduced-form scale economy, and in our estimation, identify the causal impact of scale on costs.

III. Theoretical Framework

We present a model of global trade where shipments are sent indirectly through an endogenously formed transport network. We embed the Allen and Arkolakis (2022) route selection model in a generalized Eaton and Kortum (2002) framework where production technologies in each industry and country are nonstochastic, but idiosyncratic variation in the products' optimal route generates random variation in product-origin pair prices.

Entrepôts emerge as locations where goods pass through but are neither the goods' origin nor their destination. We maintain a production and consumption setting that is as general as possible, allowing for any number of goods, industries, and input-output linkages. This model is agnostic to scale economies or

[15] Section IV addresses the extent to which exogenous characteristics like geography are responsible for lower costs at—and hence, higher concentration of shipments through—entrepôts.

[16] While some entrepôts lie along lowest-cost routes, routes stopping at entrepôts are 3–9 percent longer. This is true even when comparing shipments sent from the same origin, to the same destination, and using the same total number of stops, and when comparing total distance traveled as well as distance from port of lading to US destination.

diseconomies in transportation costs, which could work to either amplify or attenuate shipments through entrepôts.

A. *Setup*

Consumption and Production.—In each country j, consumers consume goods $\omega_n \in \Omega_n$ from each n of N industries according to function $U_j = U_j(C_j)$, where $U_j(\cdot)$ is a continuous, twice differentiable function and C_j is a matrix of quantities of an arbitrarily large number of goods ω_n in industry $n \in N$ in country j. Within each industry and product category, goods are homogeneous and normal.[17]

Goods are produced using a variety of traded and nontraded inputs, including labor, capital, and traded and nontraded varieties from any industry. The production technology for good ω is common for all goods in the same industry n and includes a vector of factor inputs L, as well as inputs of other goods.[18] Production functions can vary across industries and countries. Cost minimization results in identical production costs among competitive firms within an industry in each country. The marginal cost of a good ω is

$$c_{in} \equiv c_{in}(z_{in}, W_i, P_i),$$

where P_i is the matrix of prices of all goods ω in industries n in country i, and where W_i is the vector of factor prices in country i. Because producers in the same industry and country share the same input prices and production function, costs are shared within country industries. These costs correspond to the classic Ricardian comparative advantage.

Pricing.—To sell goods abroad at any destination $j \in J$, a firm producing product ω in industry n must pay nontransport trade costs κ_{ijn} and iceberg transport costs $\tau_{ijnr}(\omega)$ after optimally choosing the route r between i and j to minimize the shipping costs incurred. Competitive firms in i selling to j price their goods at marginal cost. The observed prices for these products at j are

$$p_{ijn}(\omega) = c_{in}\kappa_{ijn}\tau_{ijnr}(\omega),$$

where purchasers of good ω in industry n at j source the lowest cost supplier globally.

Shipping.—Producers seek to minimize shipping costs, choosing the lowest cost shipping route available. Shipping route r is comprised of K_r legs of a journey with $K_r - 1$ stops along the way between the origin, i (or $k = 1$), and destination, j (or $k = K_r$).

[17] The model and empirics can accommodate arbitrarily fine industry classifications in order to ensure this assumption holds.

[18] The production function is given by $q_{in}(\omega) = f_{in}(z_{in}, L_{in}, Q_{in})$, where $f_{in}(\cdot)$ is a continuous and twice differentiable country-industry specific production function, z_{in} is the production technology common to industry n and country i, L_{in} is a vector of nontradable factor inputs, and Q_{in} is a country-industry specific matrix of inputs of other goods ω from all industries. All inputs are treated as homogeneous.

Following Allen and Arkolakis (2022), moving stop to stop involves iceberg transport costs as well as product- and route-level idiosyncratic cost shocks $\epsilon_{ijnr}(\omega)$.[19] We place minimal structure on these direct leg-level costs $t_{r_{k-1},r_k}(\cdot)$ between locations r_{k-1} and r_k on route r, allowing them to be a function of exogenous and endogenous variables:

$$t_{r_{k-1},r_k} = f(\Xi, \varepsilon_{r_{k-1},r_k}), \tag{2}$$

where Ξ is a matrix of endogenous containerized traffic over the entire network, and $\varepsilon_{r_{k-1},r_k}$ reflects exogenous transportation cost elements such as distance.

Route-specific idiosyncratic shocks are drawn from the Fréchet distribution such that $F_{ijn}(\epsilon)$, the cumulative distribution function of the idiosyncratic draws, is as follows:[20]

$$F_{ijn}(\epsilon) \equiv \Pr(\epsilon_{ijnr}(\omega) \leq \epsilon) = \exp(-\epsilon^{-\theta}),$$

where shape parameter $\theta > 0$ captures the randomness or dispersion in the choice of routes from i to j.[21] Higher $\epsilon_{ijnr}(\omega)$ draws mean industry n has lower costs for route r.

Accordingly, product ω's shipping cost along route r from country i to country j is

$$\tau_{ijnr}(\omega) = \frac{1}{\epsilon_{ijnr}(\omega)} \prod_{k=1}^{K_r} t_{r_{k-1},r_k}(\Xi, \varepsilon_{r_{k-1},r_k}) \equiv \frac{1}{\epsilon_{ijnr}(\omega)} \tilde{\tau}_{ijr}, \tag{3}$$

where $\tilde{\tau}_{ijr}$ is the product of all leg-specific costs $t_{r_{k-1},r_k}(\Xi, \varepsilon_{r_{k-1},r_k})$ and is common to all products taking route r. Product ω in industry n's realized shipping cost from i to j is that of the transport-cost minimizing route from the set of all routes from i to j.[22] We treat t_{k_{r-1},k_r} in equation (3) as ad valorem, corresponding to the iceberg costs typically considered in the literature (Allen and Arkolakis 2022; Fajgelbaum and Schaal 2020). To test the validity of this modeling approach, we consider the fit between our cost estimates with two sets of external data and find significant correlations (Section VI).[23]

This structure is consistent with a host of mechanisms, including but not limited to port-level effects and leg-level scale economies.[24] With regard to market power, we do not directly model the decision of shipping firms. Instead, our equilibrium can

[19] Because of the max-stable property of the Frechét distribution, an isomorphic specification would have firm-specific cost shocks with a finite mass of potential competitive firms in each country. This would affect the interpretation of the source of idiosyncratic variation (firm variation or product variation) and of shape parameter θ.

[20] This distribution is identical across industries, so product-industry subscript n is dropped.

[21] This dispersion assumption is reflected in our microdata (panel B in Figure A.9 of online Appendix B.4) Almost 70 percent of origin countries have a fairly low concentration of routes (Herfindahl-Hirschman Index less than 1,500).

[22] The price of a product ω in industry n from i to j conditional on route r is $p_{ijnr}(\omega) = c_{in} \kappa_{ijn} \tau_{ijnr}(\omega)$.

[23] Using an additive cost assumption through the network, Allen and Arkolakis (2022) derives a similar expression for the iceberg cost structure (online Appendix D.1; Allen and Arkolakis (2022)).

[24] It also allows for spatial correlation in link costs, say, between t_{kl} and t_{lm}.

be considered as an overall industry equilibrium within a Sutton (1991) framework, where larger markets induce more entrants and lower marginal costs, with profits being absorbed by fixed costs.[25] Differences between these mechanisms will not impact the model estimation but will manifest in the interpretation of scale economies and for counterfactual predictions.

B. *Equilibrium*

Route Volume.—Firms from origin i select the lowest-cost route before consumers in j select the lowest-cost intermediate good supplier across all the origins countries. We observe good ω being shipped on route r from i to j only if the final price of ω, which includes both the marginal cost of production and shipping cost on route r from i to j $(p_{ijnr}(\omega))$, is lower than all other prices of good ω from all other origin country-route combinations.

We can define the joint probability that a route r is the lowest-cost route from i to j for good ω and that country i is the lowest-cost supplier of good ω to j as

$$\text{(4)} \qquad \pi_{ijnr\omega} \equiv \Pr\Big(p_{ijnr\omega} \leq \min_{i' \in I \setminus i,\, r' \in R_{ij} \setminus r} p_{i'jnr'\omega}\Big) = \frac{\left[c_{in}\kappa_{ijn} \cdot \tilde{\tau}_{ijr}\right]^{-\theta}}{\sum_{i' \in I}\left[\left(c_{i'n}\kappa_{i'jn}\right)^{-\theta} \cdot \sum_{r' \in R_{i'j}} \tilde{\tau}_{i'jr'}^{-\theta}\right]}.$$

By the law of large numbers, this is also the share of goods sold in j in industry n, coming from i, and taking route r. Introducing auxiliary matrix $A_n = \left[t_{ijn}^{-\theta}(\Xi, \varepsilon_{ij})\right]$, where each element is a function of the leg-specific transport cost, we define the expected transport cost matrix as

$$\text{(5)} \qquad \left[\tau_{ijn}\right] \equiv \left[\left(I - A_n(\Xi, \varepsilon)\right)^{-1}\right]^{\circ(-\theta)},$$

where $\circ$ is the element-by-element Hadamard power.[26] Substituting the definition of $\tilde{\tau}_{ijr}$ (equation (3)) into equation (4) and summing across routes r that pass between leg k to l, we can express the share of imports in industry n in destination j that come from origin i that passes through leg kl as

$$\text{(6)} \qquad \pi_{ijn}^{kl} = \left[c_{in}\kappa_{ijn} \cdot \tau_{ikn}(\Xi, \varepsilon) \cdot t_{kln}(\Xi, \varepsilon) \cdot \tau_{ljn}(\Xi, \varepsilon)\right]^{-\theta} \Phi_{jn}^{-1},$$

where $\Phi_{jn} = \sum_{i'}\left[c_{i'n}\kappa_{i'jn} \cdot \tau_{i'jn}(\Xi, \varepsilon)\right]^{-\theta}$ is the key distinction from Allen and Arkolakis (2022)—a multilateral resistance term that accounts for average costs, openness, and connectivity of competitors from all other countries i'. With optimal

[25] We omit discussion of the optimal shipping network from the perspective of a firm with market power and focus on leg-level scale instead. In our time period (2014), we do not find diseconomies of scale using nonlinear least squares. See Section V for further discussion.

[26] The expected transport cost from i to destination j is also $\tau_{ijn} = \gamma^{-1/\theta}\left(\sum_{r \in R_{ij}} \tilde{\tau}_{ijr}^{-\theta}\right)^{-1/\theta}$, where γ is the function $\Gamma(t) = \int_0^\infty x^{t-1} e^{-x}\, dx$ evaluated at $\left[(1+\theta)/\theta\right]^{-\theta}$.

route selection and competition on price both accounted for, equation (6) is the realized and observable share of traffic that flows through leg kl from i to j.

Next, the model yields a gravity equation. The sum of products sold in j in industry n from country i equals the share of products sold in j in industry n coming from i and taking route r, summed across all r routes:

$$(7) \qquad \pi_{ijn} \equiv \sum_r \frac{\left[c_{in}\kappa_{ijn} \cdot \tilde{\tau}_{ijr}\right]^{-\theta}}{\sum_{i' \in I}\left[\left(c_{i'n}\kappa_{i'jn}\right)^{-\theta} \cdot \sum_{r' \in R_{i'j}} \tilde{\tau}_{i'jr'}^{-\theta}\right]} = \frac{\left[c_{in}\kappa_{ijn} \cdot \tau_{ij}(\Xi, \epsilon)\right]^{-\theta}}{\Phi_{jn}}.$$

Equations (6) and (7) will jointly generate our estimation equation in Section IV.

Finally, we derive an expression for the share of global shipping passing through kl:

$$(8) \qquad \pi^{kl} = \sum_n \sum_j \sum_i \pi_{ijn}^{kl} = \sum_n t_{kln}(\Xi, \varepsilon)^{-\theta} \cdot \sum_j \Theta_{jn} \tau_{ljn}(\Xi, \varepsilon)^{-\theta} \cdot \frac{\Phi_{kn}}{\Phi_{jn}},$$

where Θ_{jn} is j's global consumption share of industry n. Because optimal route selection and competition on price are both accounted for, equation (8) corresponds to the observable shares of all goods passing through leg kl, including shipments bound for l and those continuing onward to other destinations. In Section VI, we compare our model-implied leg-level trade flows to those observed in the US microdata. We find high correlations that also hold true for higher levels of aggregation across origins and levels. In online Appendix C.2, we show how a change in the leg cost between k and l $(t_{kl}(\Xi, \varepsilon_{kl}))$ can affect trade volumes between an origin i and destination j through the trade network.

Closing the Model.—In order to close the model, we require markets to clear for factors and goods as well as the balanced trade condition. Unnecessary for estimation, we defer them to Section VII where we conduct counterfactuals.

IV. Estimation

We now show how to link our model to real world data, use the model to recover the trade costs underlying the global trade network, and estimate a scale elasticity in shipping.

A. *Taking the Model to Data*

Using equations (6) and (7), we can calculate the probability of any good traveling through link kl conditional on being sold from origin i to destination j. With the total value of trade between origin i and destination j in industry n (X_{ijn}), we can express the total volume of traffic between k and l in a given industry n (Ξ_{kln}) as

$$(9) \qquad \Xi_{kln} \equiv \sum_i \sum_j X_{ijn} \cdot \left(\tau_{ikn} t_{kln} \tau_{ljn} \tau_{ijn}{}^{-1}\right)^{-\theta}.$$

In our setting, expensive trade routes suffer from Ricardian selection at destination markets—the route's impact on prices make them less competitive relative to other routes. Yet, this does not impact the trade cost estimation as seen in equation (9), which is identical to Allen and Arkolakis (2022), despite differences in framework. While Ricardian selection, nontransportation trade costs such as tariffs, and multilateral resistance all reduce total trade, they do not differentially favor one route from an origin i to a destination j. Instead, they reduce traffic flows proportionally along all links kl.

Mapping our model into the data requires that for a set of industries $\bar{N}$, trade costs are identical and all origin-destination trade $\left(X_{\bar{N}} \equiv \sum_{n \in \bar{N}} X_n\right)$ and link-level traffic $\left(\Xi_{\bar{N}} \equiv \sum_{n \in \bar{N}} \Xi_{kln}\right)$ are observable. Summing equation (9) over industries $n \in \bar{N}$ yields

$$\Xi_{kl\bar{N}} \equiv \sum_i \sum_j X_{ij\bar{N}} \cdot \left(\tau_{ik\bar{N}} t_{kln} \tau_{lj\bar{N}} \tau_{ij\bar{N}}{}^{-1}\right)^{-\theta}. \tag{10}$$

Equation (9) tells us that to accurately measure transport costs, we only need data on origin-destination trade and link-level traffic for all goods in an industry. Equation (10) tells us that we can use traffic across multiple industries so long as we have the correct trade aggregate, we see all traffic for those industries, and we can assume transport costs are identical in those industries. We implement equation (10) using observed total containerized traffic and trade in containerized industries, where transportation costs are likely similar, and apply it in estimation only to legs where all traffic is observed.

B. *Recovering Scale Elasticities*

The Cost–Scale Relationship.—The existence of a scale economy in shipping implies that perturbations to the global shipping network that affect traffic volumes will, in turn, impact the link cost matrix estimated in the next section. Such effects must be accounted for in order to correctly estimate counterfactual adjustments.

Using leg-level trade costs from equations (5) and (10), we consider the regression

$$\ln\left(\hat{t}_{kl}^{-\theta} - 1\right) = \alpha_0 + \alpha_1 \cdot \ln\left(\Xi_{kl}^{data}\right) + \alpha_2 \cdot \ln\left(d_{kl}\right) + \varepsilon_{kl}, \tag{11}$$

where α_0 is a constant, Ξ_{kl}^{data} is the traffic volume between link kl that we observe in the ports call data, α_1 is the relationship between price and quantity (traffic volumes), and $\alpha_2 \cdot \ln\left(d_{kl}\right)$ is the coefficient and measure of log sea distance from k to l respectively. $\left(\hat{t}_{kl}^{-\theta} - 1\right)$ allows us to interpret α_1 as the elasticity between cost and traffic volumes to a trade elasticity θ.[27] That is, to interpret results from equation (11)

[27] In our model, θ serves as both the route dispersion parameter and trade elasticity. As an alternative, we can model a nested elasticity and decompose the total trade elasticity into a transportation route elasticity of substitution and nontransportation component, estimating the former using the observed dispersion of routes in the US microdata.

as elasticities, they are deflated by θ. The functional form in equation (11) presumes scale economies exist at the leg level. In online Appendix Section D.1, we discuss alternative specifications.

Of course, this relationship cannot be taken as causal. Lower cost legs may face larger demand precisely because unobserved cost-reducers induce higher levels of demand on those legs. Essentially, we wish to observe the supply elasticity, but we have only market-clearing prices and quantities. We therefore need a demand shifter.

Geography-Based Instrument.—We use the intuition of our model to construct a geography-based instrument for demand. Demand for a given leg will be higher, all else equal, if the leg lies along the most direct route between an origin and a destination. For example, consider routes from origin South Korea to destination the Netherlands. Routes that include a China-Singapore link are closer to the direct Korea-Netherlands route compared to routes that include the China-Australia link. As such, more Korea-Netherlands trade should flow through the China-Singapore leg than the China-Australia leg, which would involve a longer detour. Links that are effectively out of the way for most journeys should, all else equal, face lower demand, such as Australia on routes between East Asia and Europe compared to Singapore.

Operationalizing this intuition, we relate the direct sea distance between an origin and a destination to the distance of two legs as part of a three-leg journey, where the omitted middle leg is the object of interest. We calculate the instrument z_{kl} as

$$z_{kl} = \sum_{i \backslash k,l} Pop_{i,1960} \sum_{j \backslash \{k,l\}} Pop_{j,1960} \frac{d_{ij}^2}{\left(d_{ik} + d_{lj}\right)^2}, \tag{12}$$

where d_{ij} is the sea distance between origin i and destination j, and the square of the relative excess distance between links ik and $lj\left(d_{ik} + d_{lj}\right)$ is weighted by the year 1960 population at each origin i and destination j, $Pop_{i,1960}$ and $Pop_{j,1960}$.[28] Figure 5 shows the robust first-stage relationship between our instrument and traffic.

For plausible identification, our demand shifter instrument has to be generally uncorrelated with unobserved cost determinants for a particular leg controlling for its sea distance $(\operatorname{corr}(\varepsilon_{kl}, \ln z_{kl}) = 0)$. Locations that are close in sea distance are also close in land distance and may have easier access to other modes of transportation like road or rail. As a robustness check, we recalculate our instrument in equation (12) in a simplified setting by omitting the shortest 10 percentile distances for each origin i and destination j respectively and find similar results.

As previously noted, the observed scale economy in our setting can be generated by a number of mechanism, including but not limited to internal or external scale

[28] The population in 1960 here stands in place of GDP, which may be endogenous to the trade costs in our model. The year is chosen because immigration and populations prior to 1960 could not plausibly be impacted by 2014 containerized shipping costs. While this squared deviation functional form has an intuitive interpretation, our analysis is robust to other functional forms.

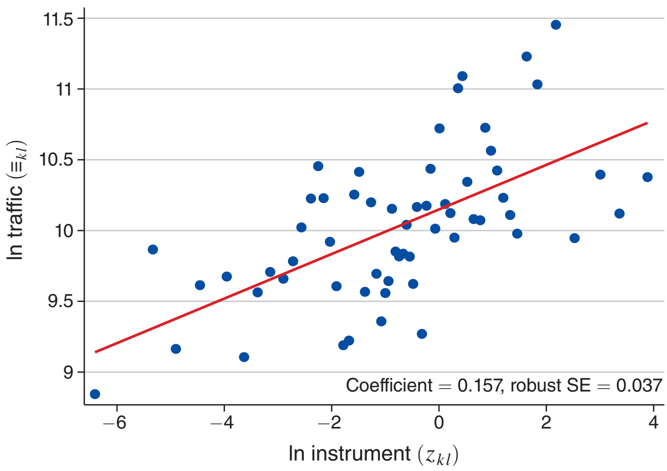

Figure 5. Residualized Plot of Correlation between Instrument and Traffic

Notes: The figure shows a binned scatter plot of 1,946 observations of link kl with the logarithm of sea distance between k and l as a control. The x-axis is the logarithm of the instrument z_{kl}. The y-axis is the natural log of traffic on leg kl. Standard errors are clustered two ways by nodes k and l.

economies and market power. These mechanisms may generate different out-of-sample results; further work should be done to isolate and test for these. In order to accommodate this multitude of mechanisms simultaneously, we implement a model-consistent and agnostic approach in our estimation of scale. Formally, we construct moments $m_1(\alpha, \beta) = Z\varepsilon(\alpha, \hat{\mathbf{t}})$ based on equation (11) with vector α and matrix of trade costs $\hat{\mathbf{t}}$. First, however, we need to recover leg-level trade costs $\hat{t}_{kl}$.

C. *Recovering Trade Costs*

We require two observable objects in order to recover a global set of trade costs: origin-destination trade values and link-level traffic volumes (equation (10)).[29] Our traffic data come from our global port of call AIS shipping data.[30] We use aggregate origin-destination trade data from CEPII and their BACI international database for 2014, segregating containerized and noncontainerized commodities.[31] Note that we do not rely on the merged US microdata in our estimation.

[29] This procedure is agnostic to the exact specification of any particular trade model that generates trade value flows X. We control for all origin, destination, and origin-destination factors by conditioning our estimation on trade flows X. In particular, items such as all origin-destination tariffs and nontariff barriers are accounted for. This does not mean that we can disentangle the two—rather, we can directly account for these factors collectively.

[30] Units for traffic are in TEU. Recall we estimate ship-by-leg TEUs by combining reported ship draught and maximum TEU. This process does not rely on the merged US customs data.

[31] We use 2014 US customs data on containerized and noncontainerized shipments to construct the share of each harmonized system 4-digit commodity code that is transported by container. All commodities with a containerized share above 80 percent are labeled as containerized. This procedure shuts down the substitution between containerized and noncontainerized transport. In practice we find a bimodal distribution, with some commodities being never containerized (e.g., oil and iron ore) and others always containerized (e.g., washing machines and children's toys). This process is documented in online Appendix A.3.

In an ideal world, estimation would recover the trade costs that directly rationalize observed bilateral containerized traffic flows—a just-identified case. While we directly observe ocean containerized traffic, our data omit movement of containers overland, across, and within borders. We overcome this limitation by assuming a functional form that allows for estimation without requiring the direct observation of overland links. We consider the mapping[32]

$$\hat{t}_{ij}^{-\theta} = \frac{1}{1 + \exp(\mathbf{Y}\beta)} \in [0,1],$$

where the matrix $\mathbf{Y}$ is a vector defined as

$$\mathbf{Y}\beta = \beta_0 + \beta_1 \log(SeaDistance_{ij}) + \beta_2 \log(traffic_{ij}) + \beta_3 \log(traffic_i) + \beta_4 \log(traffic_j) + \beta_5 \log(trade_{ij}) + \beta_6 \mathbf{1}\{i,j \in LandBorders\},$$

where β_0 is an intercept, β_1 considers the sea distance between the nearest principal ports,[33] and β_2 considers port-to-port traffic. β_3 and β_4 consider the total incoming and outgoing traffic at ports i and j, respectively. β_5 considers trade flows from ports in i to j. Finally, β_6 is an indicator for a shared land border.[34]

It is crucial to note two things. First, while the equations above posit relationships between observables, our objective at this stage is not the vector β of coefficients—which may reflect endogenous variables—but the resulting predictions for $\hat{t}_{ij}$. Instead, we seek to fully saturate the variation in the data in order to generate the closest empirical prediction for the matrix of trade costs relative to the just-identified case, which yields the model-perfect estimates of trade costs for each link. This allows us to recover the trade costs while remaining agnostic to their underlying determinants, including potential economies of scale as well as possible geographic indicators. Secondly, while the parameters for β yield estimates of every trade cost $\hat{t}_{ij}$, we need not discipline β by comparing traffic on every link. This allows us to still recover estimates of $\hat{t}_{ij}$ although we do not observe within-country traffic and between-countries traffic that share overland routes.

We create a moment m_2 that finds the vector β that minimizes the difference between the matrix of expected traffic, $\hat{\Xi}(\beta | \mathbf{X}, \mathbf{Y}, \theta)$, and observed traffic, Ξ^{data}, for countries that do not share a land border:

$$m_2(\beta) = \left(\hat{\Xi}(\beta | \mathbf{X}, \mathbf{Y}, \theta)\right) - \left(\Xi^{data}\right),$$

where expected traffic is a function of β, trade elasticity θ, and observed trade values $\mathbf{X}$.

As noted, we do not fully observe the traffic flows of containerized goods on geographically contiguous legs, and we do not perform our estimation procedure

[32] This functional form maps from the real numbers to the unit interval, as is required by our theory.
[33] For each country pair, we calculate the volume-weighted mean sea distance across all port pairs.
[34] We do not estimate within-country trade costs directly due to data constraints and assume that they do not change in the counterfactual.

using traffic data from these legs. Instead, our trade cost estimates, even for overland links, are disciplined by the observed traffic flows of sea-only legs that do not share a land border.

D. *Joint Estimation*

We combine our scale estimation and recovery of trade costs into a single stage:

$$
\begin{aligned}
m_1(\alpha,\beta) &= Z\varepsilon\left(\alpha,\hat{\mathbf{t}}(\beta)\right) \\
m_2(\beta) &= \left(\hat{\Xi}(\beta|\mathbf{X},\mathbf{Y},\theta)\right) - \left(\Xi^{data}\right).
\end{aligned}
$$

We conduct a two-stage generalized method of moments (GMM) procedure, using optimal instrumental variable weights estimation for the first set of moments m_1, which accounts for our causal estimates of scale, and trade volumes on the second set of moments m_2, which rationalizes a global set of link-level trade costs t_{kl} conditional on observable origin-destination trade values $\mathbf{X}$ and link-level traffic flows Ξ^{data}. We reiterate that inference can only be conducted on α. β contains incidental parameters—important for estimation, but not inference.[35]

E. *Simultaneous Identification of Scale and Trade Costs*

Our approach parallels the industrial organization literature, which seeks to recover unobserved cost structures, and identification depends both on instrumental variables and behavioral assumption. For example, Ackerberg et al. (2007) take market level data and instruments to recover demand and then use equilibrium assumptions on behavior to recover marginal costs, which are then projected on product attributes. Similarly, we rely jointly on the structure of equilibrium shipping flows embedded in the Allen and Arkolakis (2022) framework and our demand-shifting instrument.

However, this approach opens the door for a mechanically driven result. Specifically, we are concerned with estimating the causal scale impact of traffic volumes on trade cost (equation (11)) while, at the same time, our cost estimates are themselves recovered from our model prediction, which is a function of traffic volumes (equation (10)). This circularity can introduce a mechanical correlation if, for example, measurement error in traffic feeds both into trade cost estimates and traffic.

We approach this problem through multiple methods. First, we establish how this issue can arise due to measurement error in our context. We show how this error can be considered a form of omitted variable bias and the conditions under which an instrumental variable can correct for this bias. Second, we run Monte Carlo simulations that confirm the existence of this bias in the presence of measurement error and show how our instrument eliminates it. Third, we use external data on freight costs to estimate potential traffic-correlated errors, both to illustrate the potential

[35] The second stage computes an optimal weighting matrix using the first stage results.

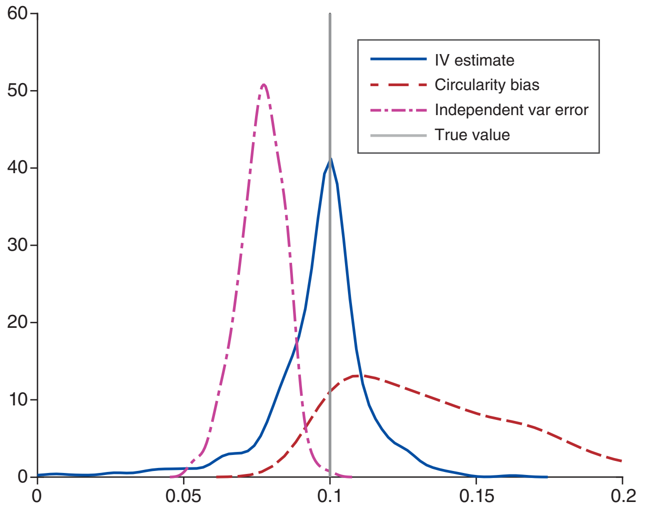

Figure 6. Monte Carlo Simulations Illustrating Estimation Biases

Notes: The figure shows 500 simulated estimates. The blue solid line is our preferred instrumental variable estimator. Our instrument is correlated with the true shipping traffic on a particular route. The purple dot-dash line illustrates classic measurement error in the independent variable (shipping traffic on a route), leading to classic attenuation bias in ordinary least squares (OLS). The red dash line illustrates our principle worry—an upward bias in OLS, due to our recovered trade costs being a function of observed shipping traffic that could be measured with error. A valid IV can correct for this bias (blue solid line). See online Appendix D.2 for full details.

bias in the OLS and show how our instrument removes this bias. Finally, we run a parallel scale estimation purely on our external freight costs and find similar results. See online Appendix D.2.

Figure 6 summarizes our findings using Monte Carlo simulations. First, we show that with true trade costs, typical measurement error in traffic volumes would bias OLS estimates downward (purple dot-dash line). If measurement error in traffic affects the trade cost estimates, the OLS estimates would bias upward (red dash line), since the dependent variable (trade costs) is partially derived from the independent variable (traffic). However, a valid instrumental variable can correct for this bias (blue solid line). Online Appendix D.2 further elaborates on the simulation procedure.

We show the lack of correlation between our instrument and an approximation of the error, estimated as the difference between our measured costs and external measures of Drewry maritime research freight costs from Wong (2022). Details for this exercise are found in online Appendix D.2. Panels A and B of Figure 7 show a positive and negative correlation between this approximation of the error and estimates link costs and link traffic, respectively, controlling for distance, consistent with the circularity bias in online Appendix D.2 and the Monte Carlo. Panel C shows a weak and insignificant correlation between this residualized approximation of the error and our instrument, again controlling for sea distance. The lack of correlation is consistent with an instrument that is uncorrelated with the true error. While this is insufficient to validate our instrument, it performs the same role as a balancing test, showing an absence of evidence of exclusion restriction violations.

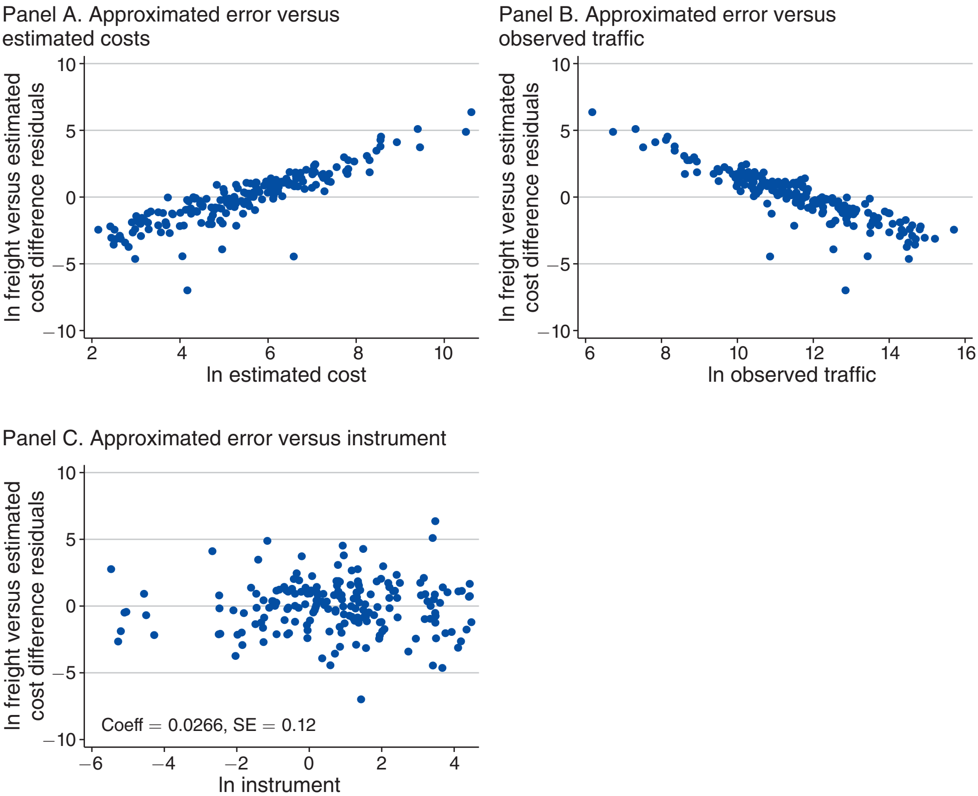

Figure 7. Balancing Test

Notes: Figures are scatter plots of, on the x-axis, the natural log of the estimated leg costs in Section IV, the observed traffic, and the geography-based instrument used in Section IV (panels A, B, and C, respectively) against, on the y-axis, the difference between the natural logs of the estimated leg costs in Section IV and from Wong (2022), residualized after controlling for sea distance for 209 legs for which both costs exist. Standard errors are clustered two ways by the nodes on each link. See online Appendix D.2 for full details.

V. Results

Scale Economy.—Table 1 reports our instrumented scale elasticity from our scale moments (equation (11)). For the widely used trade elasticity value of $\theta = 4.4$ (Simonovska and Waugh 2014), the interpretation of our causal estimate is that increasing traffic volume on a link by 1 percent would reduce costs by 0.06 percent. As the typical journey observed in our microdata has 2.5 links, this translates into a 0.17 percent decrease in overall origin-destination trade costs. Our estimate is within one standard error of Hummels and Skiba (2004), who estimate an elasticity of freight to quantity of 0.18 using an IV and trade data from six importers, and Asturias (2020), who reports an elasticity of 0.23 using US port data.[36] Additionally, Skiba (2017) reports an elasticity of 0.26 using product-level import

[36]The six importers in Hummels and Skiba (2004) are Argentina, Brazil, Chile, Paraguay, Uruguay, and the United States. We compare our estimates to theirs for all countries since our scale estimate is based on global data.

TABLE 1—GMM ESTIMATION RESULTS

	$\ln(c_{kl})$ (1)
$\ln\left(\Xi_{kl}^{data}\right)$	−0.28 (0.02)
$\ln(d_{kl})$	0.59 (0.03)
Constant	4.06 (0.35)

Notes: We conduct a two-stage GMM procedure, first using optimal instrumental variable weights estimation on the first set of moments and the inverse of trade volumes on the second set of moments. The second stage computes an optimal weighting matrix W using the first-stage results. $\ln(c_{kl})$ is the natural log of transportation trade cost on link kl. $\ln\left(\Xi_{kl}^{data}\right)$ is the natural log of traffic volume on link kl. $\ln(d_{kl})$ is the natural log of sea distance between k and l, computed using Dijkstra's algorithm.

data from Latin America. Our estimate is also broadly consistent with the literature on scale in production. Bartelme et al. (2019) estimates a sector-level scale elasticity of 0.21, while Lashkaripour and Lugovskyy (2023) finds an elasticity of 0.20 after jointly estimating both scale and trade elasticities.

Link and Average Bilateral Trade Costs.—Online Appendix Figure A.11 graphs our resulting matrix of pairwise trade costs. We present the vector β estimates in online Appendix Table A.5 as purely predictive parameters, not fundamentals that we can alter in the counterfactuals (see online Appendix D.1 for further details). Instead, we simply need to know if our β estimates can predict containerized traffic that reflects the actual observed traffic volumes. With a full link-level trade cost matrix $[t_{kl}]$, we also can generate an average bilateral transport cost between locations $[\tau_{ij}]$. We provide our network-consistent trade-link and origin-destination cost estimates to researchers, and they are available for download on our websites. Online Appendix Table A.11 compares these network-consistent bilateral trade costs to more commonly used distance measures. Our cost measures have more predictive power than distance alone and both are significant in a combined specification, implying that both measures have distinct predictive power for trade.

Robustness and Alternative Specifications.—First, to mitigate the risk of model misspecification (in equation (11))—for example, port-level economies or diseconomies of scale—we explore alternative specifications. Adding origin or destination fixed effects increases the magnitude of leg-level scale economy. We choose our current specification that yields a more conservative scale measure. We also search for, but do not find, nonlinearities in our estimated scale economy indicative of port congestion or scale economies that would result in altered counterfactual costs. Additionally, as an alternative estimation approach for equation (11), we use external cost measures—freight rates from 140 bilateral pairs from Wong (2022)—and find similar, larger, and noisier point estimates (online Appendix Section D.2.4). These pecuniary freight rates are available for just a subset of routes compared to

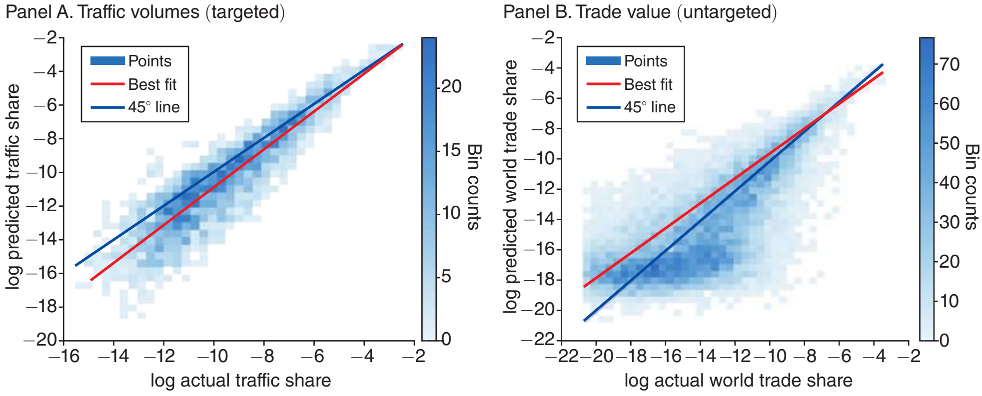

FIGURE 8. MODEL FIT COMPARISONS

Notes: Panel A compares our targeted moment: predicted container traffic volumes from any two ports (y-axis) to the actual container traffic volumes (x-axis, normalized as a share to total world container traffic). Panel B compares untargeted aggregate trade shares (x-axis) versus predicted trade shares for containerized traffic (y-axis), where predicted trade shares are computed using the full model described in Section VII.

our setting and do not include all possible elements of link trade costs that are consistent with our model.

Finally, locations that are strategically close to each other in sea distance are close in land distance and potentially have easier access to alternative modes of transportation like road or rail. We recalculate our instrument omitting the shortest 10 percentile distances for each origin-destination pair and find that our results retain the same signs and stay within a standard error of our baseline estimates.

Model Fit.—Figure 8 compares our model-predicted traffic and trade values against their observed counterparts in the data. In panel A, we compare actual observed global container traffic shares with the our model-predicted shares using our estimated trade costs. We include both a best fit line and a 45-degree line. We fit the data extremely well, with a correlation between the observed and predicted shares (in logs) of 0.9. Panel B compares our estimated trade shares to actual observed trade shares, which we do not target.[37] We fit the data well here as well with a correlation (in logs) of 0.7.

VI. Comparison of Model-Predicted Estimates to Data

We compare our model's results with three separate sets of external data. First, we link our results to ship size estimates to highlight one possible scale-economy mechanism. Second, we compare our trade cost estimates with freight rates. Third, we compare our model-predicted traffic flows for US-bound shipments to our US microdata.

[37] To generate trade flows, we close the model using the full setup in Section VII.

A. *Symptoms of Scale Economies: Ship Size*

Using our model, we estimate leg-level shipping scale economies. A number of mechanisms can generate the cost reductions that coincide with these scale economies. Internal or external scale economies in shipping and competition among shippers could all generate a negative relationship between volume and costs, as could factors such as port infrastructure.[38] Lacking data to directly test these mechanisms, we turn to one symptom of a scale economy observable in our US microdata that lends further credibility to our results: ship size. Relying on the idea that larger ships enable lower shipping costs (Cullinane and Khanna 2000), we consider the correlations between ship sizes, trade volumes, and our recovered leg-level trade costs and then investigate the relationship between indirect shipping and ship size.

Ship Sizes, Traffic Volumes, and Recovered Trade Costs.—In panel A of Figure 9, we show the positive relationship between the average containership size on a route and the traffic volume on that route, controlling for the distance between origin and destination. In panel B, using the route-level containership size measure, we show the positive link between ship size and our corresponding recovered trade costs. Routes with more container traffic use larger ships: a 10 percent increase in route volumes correspond to a 2 percent increase in ship size (column 1, Table A.8 in the online Appendix). Routes with lower trade costs use larger ships: a 10 percent decrease in our estimated iceberg trade costs corresponds to 6 percent increase in ship sizes (column 1, Table A.9 in the online Appendix).[39]

Ship Size and Indirect Trade.—Figure 10 further investigates the relationship between entrepôt usage and ship size, plotting ship size (x-axis) against US-bound traffic volume (y-axis) by country of origin—separately for traffic that is routed through an entrepôt and traffic that is not—such that each origin country is associated with two data points. Larger origins transport goods to the United States on larger ships. However, shipments from smaller origins routed through entrepôts also arrive on large ships, such that indirect shipping through entrepôts appears to close the ship-size gap for smaller origins.[40]

B. *Cost Estimates with Freight Rates Data*

Next, we compare our expected trade cost estimates τ_{ij} at the origin-destination level with container freight rates from Wong (2022). These rates are the costs paid

[38] High-traffic routes are served by many carriers, using ships capable of carrying 25,000 containers with automated loading and unloading.

[39] Online Appendix D.3 reports shipment-level regressions controlling for origins, destinations, and without route distance controls. Results are similar.

[40] For shipments with the same origin, US destination, and controlling for the total number of stops, shipments stopping at entrepôts arrive on ships that are on average 15 percent larger. For shipments with the same origin and US destination, shipments sent directly arrive on ships that are, on average, 8 percent smaller. Further shipment level analysis in online Appendix D.4 confirms the positive relationships between shipment volume and ship size and robustness to different notions of origin, lading, and transshipment.

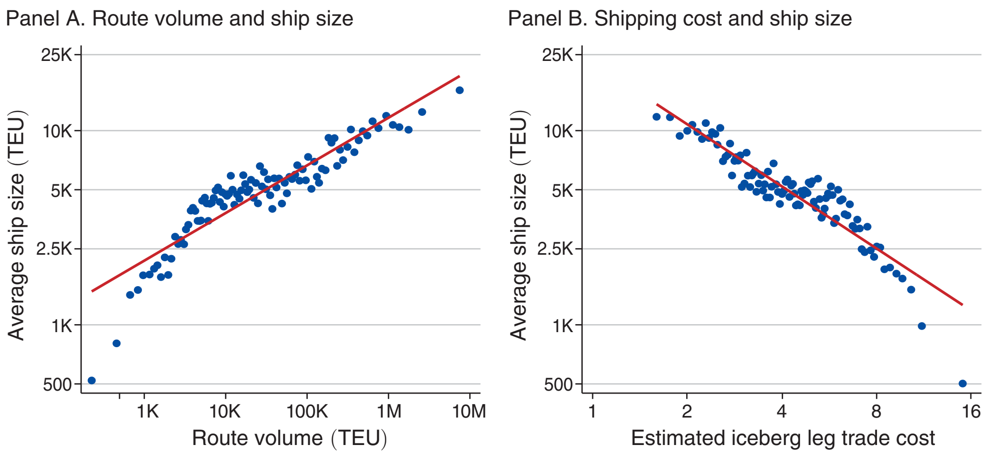

Figure 9. Link between Recovered Trade Costs and Ship Size

Notes: Figures are bin-scatter plots over all observed containership routes with 100 bins. We control for the $\log(SeaDistance)$ between origin and destination ports but add variable means back for the plots. Panel A plots the relationship between the total containers on a route and the average containership's size on that route (weighted by utilized capacity). Panel 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.

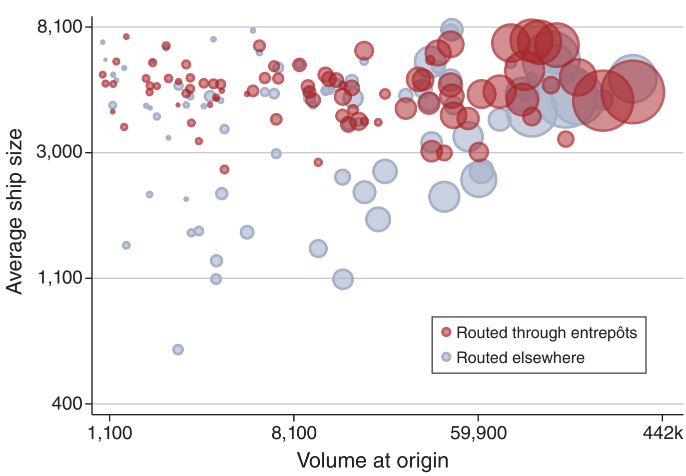

Figure 10. Link between Indirect Trade and Ship Size

Notes: The *x*-axis shows the total export volume in TEUs from an origin country to the United States. The *y*-axis shows the average ship size that arrives from an origin country to the United States. Each country is represented by two data points: a blue and a red circle. The red circle indicates the corresponding information for trade from an origin that is routed through an entrepôt while the blue circle is for trade that is not. Circle size denotes shipping volume specific to the route (either through an entrepôt or not). Note that trade that is not routed through an entrepôt (blue circle) could either be shipped directly to the United States or shipped via a nonentrepôt.

by firms to transport a standard full container load between port pairs and include the base ocean rate, fuel surcharge, and terminal handling charges at both origin and destination. They are for the largest ports globally that handle more than 1 million

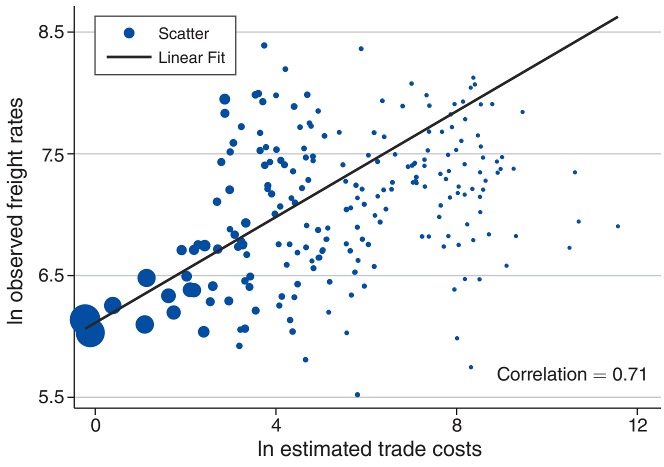

Figure 11. Correlation between Cost Estimates with Actual Freight Rates

Notes: Data points compare origin-destination predicted costs τ_{ij} to average freight rates from Wong (2022) and Dewry Maritime Research (2014). Circle size are weights for container volumes (TEU).

containers annually and account for about 73 percent of global container volumes during this time period (World Bank 2018). While we are only comparing a subset of the cost estimates from our entire sample with these freight rates, we find a correlation of 0.7 (Figure 11).

C. *Traffic Estimates with US Microdata*

In order to assess our model's ability to capture actual shipment journeys and trade indirectness, we compare our model predictions for the paths of US-bound shipment traffic to the actual observed paths in our US microdata. Our estimation, which uses global traffic data rather than the US microdata, delivers predictions for how US-bound shipments travel through the shipping network. Equations (6) and (7) imply

$$\hat{\pi}_{iUS}^{kl} = \left[\tau_{ik} t_{kl} \tau_{lj} \tau_{ij}^{-1}\right]^{-\theta}, \tag{13}$$

as the ratio of all shipments from i to the United States that are observed flowing through leg k, l.

We compare our model-predicted value of equation (13) to the proportion of goods coming into the United States from any origin i on leg kl, which we call $\pi_{iUS,Data}^{kl}$, by aggregating shipments using link kl in our microdata. Note that while our microdata is described in Section I and used to generate our stylized facts in Section II, it is not used to estimate our trade costs in Section IV. Column 1 of Table 2 reports the univariate regression outcome between these two measures, weighted by total origin TEU. We find a significantly positive relationship with a coefficient of 1 in the confidence interval. Over half of the variation in the observed distribution can be explained using the predicted probabilities.

Table 2—Correlation between Traffic Estimates with Microdata

	$\hat{\pi}^{kl}_{iUS}$ (1)	$\hat{\Xi}^{kl}$ (2)	$\hat{\pi}^{l}_{US} - \hat{\pi}_{l,US}$ (3)	$\hat{\pi}^{kl}_{iUS}$ (4)	$\hat{\Xi}^{kl}$ (5)	$\hat{\pi}^{l}_{US} - \hat{\pi}_{l,US}$ (6)
$\pi^{kl}_{iUS,Data}$	0.844			0.870		
	(0.117)			(0.119)		
Ξ^{kl}_{Data}		1.217			1.233	
		(0.123)			(0.121)	
$\pi^{l}_{US,Data} - \pi_{l,US,Data}$			0.942			0.963
			(0.224)			(0.220)
Observations	13,763	650	95	365,330	2,149	186
Data				All	All	All
R^2	0.518	0.666	0.420	0.517	0.675	0.425
F	51.88	97.98	17.64	53.15	103.9	19.15

Notes: $\hat{\pi}^{kl}_{iUS}$ is the model-predicted share of goods from origin i to US destination flowing through leg k,l, $\hat{\Xi}_{kl}$ is the model-predicted total US-bound traffic on a given leg k,l, and $\hat{\pi}^{l}_{US} - \hat{\pi}_{l,US}$ is the model-predicted total excess US-bound traffic through node l. Their corresponding variables observed in the compiled microdata are indicated with subscript *Data*: $\pi^{kl}_{iUS,Data}$, $\Xi_{kl,Data}$, and $\pi^{l}_{US,Data} - \pi_{l,US,Data}$. Columns 1 to 3 are restricted to nonzero traffic volumes in the US microdata, while columns 4 to 6 include journeys with zero traffic volumes in the US microdata (all data). Columns 1 and 4 results are robust to tobit specifications, which allow for lower and upper censoring limits. Standard errors clustered by origin and destination countries.

Next, summing the predicted probabilities in equation (13) across all origins i, the model delivers a prediction for the total amount of US-bound traffic on a given leg kl:

$$\hat{\Xi}^{kl} = \sum_i X_{iUS} \cdot \hat{\pi}^{kl}_{iUS},$$

where X_{iUS} is the total trade flow from origin i to the United States. Column 2 compares this to the total volume of shipments moving between a given leg in the microdata (Ξ^{kl}_{Data}), again finding a positive and significant coefficient with 1 in the confidence interval.

Finally, summing probabilities in equation (13) across origins i and nodes k, we obtain the total traffic through node l. Subtracting volume of exports from l, we obtain the entrepôt usage of l for US-bound shipments:

$$\hat{\pi}^{l}_{US} - \hat{\pi}_{l,US} \propto \sum_k \hat{\Xi}^{kl} - X_{l,US} = \sum_k \sum_i X_{iUS} \cdot \hat{\pi}^{kl}_{iUS} - X_{l,US},$$

Column 3 compares this to its counterpart in the microdata ($\pi^{l}_{US,Data} - \pi_{l,US,Data}$), finding a positive and significant result with 1 within the confidence interval.

In the microdata, a number of legs have zero traffic volumes. However, our model predicts some small amount of traffic on every leg. In columns 4 through 6, we rerun the regressions for each corresponding predicted traffic estimate including legs with zero observed volumes (increasing our observation count). Our results do not significantly change because our model predicts extremely low volumes on these legs.

Our paper provides a new set of global trade costs that accounts for the trade network. The tight matches between our estimates—trade costs and traffic—and separate sets of observed data external to our estimation demonstrate that our estimates reflect actual costs and indirect traffic flows in the trade network. Additionally, these results serve as a check to the validity of our modeling approach and the Allen and Arkolakis (2022) framework. Allen and Arkolakis (2022) impute traffic and trade flows within the US highway system for their estimation.[41] Despite the strong structural assumptions made and the limited data requirements, our checks curtail the risk that our estimates are wildly off the mark. In addition to our leg and origin-destination cost estimates, we provide model-implied indirectness measures for ocean shipping as well as resulting market access measures to researchers on our websites.

VII. Counterfactuals

We quantify the welfare importance of the trade network and the specific role entrepôts play within that network in three counterfactual exercises. In our first counterfactual, we demonstrate that (i) transportation improvements at entrepôts have significant global welfare impacts (not including their own gains), as well as localized benefits for nearby neighboring countries as a result of the trade network; (ii) the global impact of transportation improvements differs meaningfully from nontransportation improvements for all countries—not just, but especially for especially entrepôts—due to the network structure of trade; and (iii) scale economies in transportation further magnifies these impacts.

In our second counterfactual, we illustrate how nontransportation cost changes at an entrepôt generate widespread impacts through the trade network—beyond directly impacted countries—by considering the impact of a negative trade shock on an entrepôt node country in the form of the United Kingdom leaving the European Union. Changes to the trade network due to scale economies generate different consequences for Brexit, both in effects' magnitudes and their distributions.

Our third counterfactual evaluates the welfare and trade impacts of the two endogenous mechanisms in our model: (i) network effects—allowing countries to ship indirectly—and (ii) scale effects—allowing countries to ship indirectly and take advantage of scale economies. To illustrate this, we study the effects of the Arctic opening up to trade between the Pacific and Atlantic Oceans, bypassing the Suez and Panama canals.

A. *Counterfactual Methodology*

To estimate these counterfactuals, we first introduce structural assumptions into our general framework as well as factor and goods market clearing and balanced trade conditions in order to deliver a quantifiable general equilibrium model.

[41] They assume that the observed traffic for a link is proportional to the underlying value of trade on that link. This assumption is later on verified by comparing their predicted trade flows to actual flows from the commodity flow survey.

Closing the Model.—We adopt the Caliendo and Parro (2015) framework and assume there are three sectors ($N = 3$): containerized tradables c, noncontainerized tradables nc, and nontradables nt ($n \in [c, nc, nt]$), all three of which are used as final goods and intermediates in roundabout production. See online Appendix E for full details.

Equilibrium in Changes.—Defining the general equilibrium using hat algebra, we consider two sets of changes: (i) link-level transport costs $\dot{t}_{kl} = t'_{kl}/t_{kl}$, which change expected trade costs $\dot{\tau}_{ijn} = \tau'_{ijn}/\tau_{ijn}$, and (ii) changes in nontransportation trade costs $\dot{\kappa}_{kl} = \kappa'_{kl}/\kappa_{kl}$. Both alter the endogenous costs of production, price indices, wage levels, trade flows, and welfare. We solve for how wages and prices change $\{\dot{w}_i, \dot{P}_i\}$ as a function of changes to model primitives, $\{\dot{\tau}_{ijn}, \dot{z}_{in}, \dot{\kappa}_{ijn}\}$, and compute changes in marginal costs $\dot{c}_{in}$ and trade volumes $\dot{X}_{ij}$.

Additional Data.—We combine our trade volume data with country-level input-output data from the Eora Global Supply Chain database, aggregating to three sectors: nontraded, containerized traded, and noncontainerized traded goods. We use country-level consumption and production data to compute Cobb-Douglas shares η and γ. This gives us a sample size of 136 countries. We conservatively set $\theta = 4$ (Simonovska and Waugh 2014).[42]

Procedure.—Changes to transport costs are implemented as changes to link costs $\dot{t}_{kl}$, which, translated through the model, generate changes in the expected trade cost between every bilateral trading pair in our data—even those that are not directly connected with each other. Once calculated, these bilateral changes enter isometrically to changes in bilateral nontransportation costs. For analysis that includes the impact of scale, we model a new equilibrium in the short-to-medium run by following an iterated procedure in algorithm 1 in online Appendix F.1. In this procedure, we start at today's equilibrium and allow all shippers to optimize their transportation patterns. We then recalculate trade costs at new volumes according to equation (11). We iterate, allowing reoptimization until a new stable equilibrium is reached. Our model theoretically admits multiple equillibria (as in Brancaccio, Kalouptsidi, and Papageorgiou 2020); however, we focus on the unique equilibrium from our current starting point—the world today.[43]

B. *Importance of Entrepôts in the Trade Network*

Overview.—We consider the role of the shipping network in international trade and the specific importance of entrepôts in that network. We run two types of

[42] An alternative approach decomposes the total trade elasticity into a transportation route elasticity of substitution and nontransportation component, estimating the former using the observed dispersion of routes in the US microdata.

[43] Kucheryavyy, Lyn, and Rodríguez-Clare (2019) establishes a common mathematical structure that characterizes the unique equilibrium in multi-industry gravity trade models with industry-level external economies of scale. Their structure requires that the product of the trade and scale elasticities to be not higher than one, which is satisfied in our case.

counterfactuals. For all countries, we consider the impact of transportation infrastructure investment in the form of a 1 percent reduction in transportation costs (t_{kl}) to and from a targeted country. We contrast this with a 1 percent reduction in nontransportation trade costs (κ_{ij}) to and from the targeted country, such as a unilateral tariff reduction or reduction in information frictions. For each type of counterfactual, we evaluate two cases—equilibrium changes with and without accounting for the endogenous impact of scale economies on transport costs throughout the shipping network. Reductions in κ_{ij} without scale effects consider changes in a manner that ignores the shipping network, while the other three cases involve exogenous and/or endogenous changes to the shipping network. In each of these four cases, we consider welfare and bilateral trade changes to the targeted country as well as to all other impacted countries, and focus specifically on differences between entrepôts and nonentrepôts. With 136 targeted countries and four cases, we have 544 counterfactuals.

Which Countries Are Pivotal to the Trade Network.—Our general equilibrium model yields a convenient metric for how pivotal a country or node is within the trade network: the impact of changes at the country on global welfare excluding a country's own. Pivotal locations are those that generate the largest adjustments throughout the network. Panel A in Figure 12 lists the global welfare impact of infrastructure improvements at the 20 most pivotal nodes in the network excluding countries' own welfare change, for cases both with and without scale responses. Our 15 entrepôts dominate this list. Singapore and Egypt are top two, evocative of the strain in global supply chains when the Suez Canal was blocked in March 2021 (Paris and Malsin 2021; Sheppard, Dempsey, and Saleh 2021; Gambrell and Magdy 2021).[44] Scale economies' impact on the transportation network (overlaid gray bars) further augment the differential impact of entrepôts.[45] Infrastructure investments at entrepôts generate, on average, 10 times the global welfare impact relative to investment elsewhere (columns 3 and 4, Appendix Table A.13).[46]

When Does Accounting for the Trade Network Matter.—Panel B of Figure 12 plots the average welfare impact, excluding the targeted country's own welfare change, of a transportation cost reduction (as in panel A) against the same for nontransportation trade costs. While driven by gravity, there is a strong overall relationship between the two counterfactuals, the average difference is roughly an order of magnitude: the effects of one type of counterfactuals will be a poor predictor of the other for any given country. For entrepôts, (red in panel B), the one-to-one relationship is violated. For example, Egypt ranks top two in terms of global impact from infrastructure improvements, while it is not among the top 20 in terms of nontransportation trade cost reductions. While the effect of nontransportation cost reductions in Egypt has a similar global welfare effect to that of Colombia, Egypt's impact is larger than

[44] Within drybulk shipping, Brancaccio, Kalouptsidi, and Papageorgiou (2020) finds that removing the Suez Canal has a higher welfare impact than the Panama Canal and Strait of Gibraltar.

[45] Panel A of online Appendix Figure A.14 repeats the exercise for cases with nontransportation cost reductions, finding that the top 20 list is dominated by the largest economies instead.

[46] Online Appendix Tables A.13 and A.14 examine the differential impact of targeting entrepôts.

Panel A. Transportation improvements: Highest global welfare changes

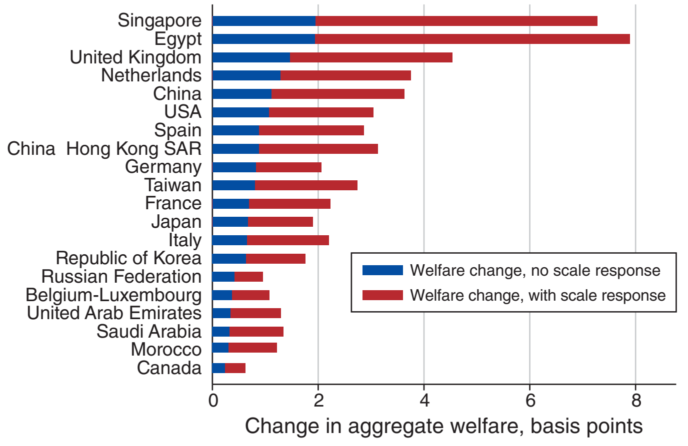

Panel B. Transportation versus nontransportation improvements

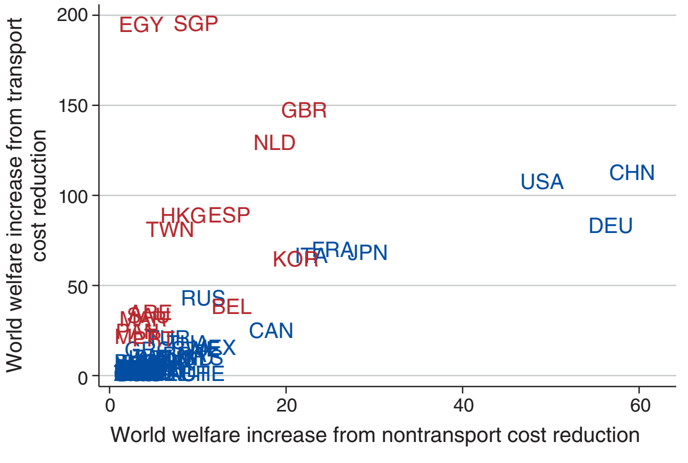

Figure 12. Most Pivotal Countries in the Network: Change in Global Welfare

Notes: Panel A shows aggregate net change in global welfare after infrastructure investment in the targeted country, excluding the country's own welfare change, for the countries with the largest global impact calculated without scale economies. Panel B compares, by country, the change in world welfare, excluding the country's own welfare, from a 1 percent decrease in nontransportation costs (x-axis) versus a 1 percent decrease in transportation costs (y-axis). Entrepôts are labelled in red.

that of the United States in the transportation cost reduction exercise.[47] The pivotal nature of the entrepôts are specific to their role in the trade network.

Ignoring the trade network impacts of policy rolls the quantitatively large network impacts into the effects of nontransportation cost changes. On the one hand, the impact from any one individual trade cost change will be highly nonpredictive. On the other hand, this may not qualitatively impact analysis at the spokes of the

[47] Panel B online Appendix Figure A.14 finds similar results comparing nontransportation cost reductions with and without an endogenous scale response. Country-pair bilateral trade results are similar.

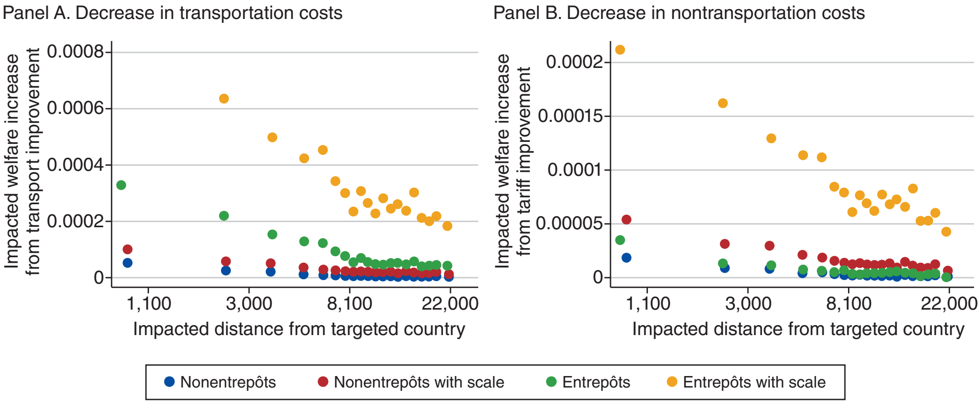

Figure 13. Spatial Decay of Benefits by Entrepôt Status

Notes: Panel A is a binned scatter of welfare effects of transportation infrastructure on impacted countries versus distance between targeted and impacted countries. Targeted countries receive the cost reduction and impacted countries trade with them. Blue and red dots are the no-scale and scale cases for counterfactuals where targeted countries are not entrepôts, respectively. Green and orange dots are no-scale and scale cases, respectively, for counterfactuals where targeted countries are entrepôts. Panel B presents the same for reductions in nontransportation trade costs.

network—those origins or destinations that do not significantly participate in trade as third countries—but substantially obfuscates the role of entrepôts in trade.

The Impact of Entrepôts Are Localized.—To account for the differential impacts of entrepôts, we drill down to one particular margin at which the impact appears most distinct: locally. Figure 13 is a binned scatter plot considering the welfare effects on the impacted country (y-axis) relative to its distance from the targeted country (x-axis), adjusting for the impacted country fixed effects. Nearly overlapping blue and green dots in Figure 13 panel B show a nearly identical distance gradient for nonentrepôts and entrepôts respectively for counterfactual nontransportation cost reductions without scale economies. The blue and green dots in Figure 13 panel A show the overall larger impact of infrastructure investments at entrepôts is relatively more localized—decaying at five times the rate. Scale economies amplify the localization, with orange dots decaying at seven times the rate compared to red.[48]

Scale Economies Concentrate Gains to Entrepôts.—Finally, we turn our attention to how these cost reductions differentially affect the impacted countries when they are entrepôts versus nonentrepôts. Figure 14 plots the differential welfare gains to entrepôts relative to nonentrepôts, as impacted countries, controlling for impacted country size, distance between targeted and impacted countries, and targeted country fixed effects. Without scale economies, we find that the welfare gains for both entrepôts and nonentrepôts are not significantly different (in blue). However, the differential benefits to entrepôts is significant and large when allowing for scale

[48] The orange dots in panel B, which include the endogenous scale response through the transportation network, echo these results.

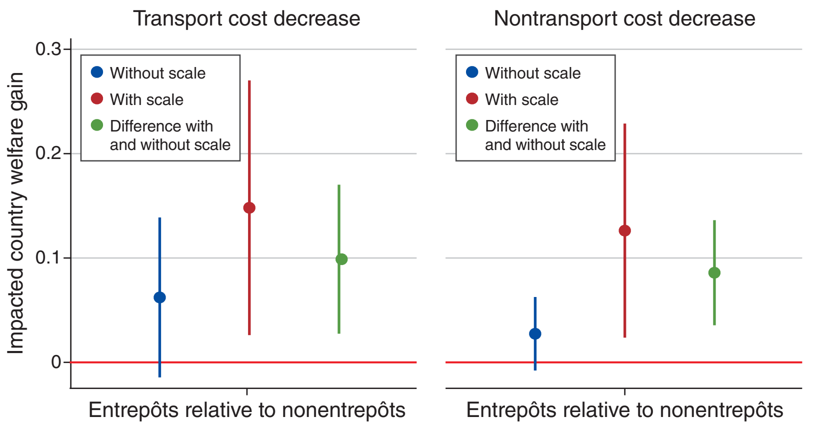

Figure 14. Differential Welfare Gains of Impacted Countries by Entrepôts Status

Notes: Figure plots the coefficients (dots) and confidence intervals (lines) for indicators for entrepôt status from a country-pair level regression of impacted countries' log percent welfare gains from a transportation cost reduction or an infrastructure improvement (left panel) or nontransportation trade cost reduction (right panel) at targeted countries, controlling for impacted country GDP, bilateral distance, and targeted country fixed effects. Targeted countries receive the cost reduction and impacted countries trade with them. Standard errors are clustered by targeted country.

economies (in red). Scale economies disproportionately accrue gains to entrepôts as impacted countries. The coefficient on the entrepôt dummy is 0.15 (SE of 0.06) and 0.13 (SE of 0.05) for transportation and nontransportation counterfactuals, respectively. The pairwise difference between the two cases (in green) is statistically significant. These results that scale economies in transportation concentrate gains locally at and around hubs highlight scale economies in transportation as a source of agglomeration.

C. *Impact of Nontransport Trade Costs on the Network*

In order to illustrate the trade network consequences of nontransportation trade cost changes on a node, we study the effects of Brexit—a 5 percent increase in nontransportation trade costs for goods that originate or are destined for the United States. We assume these increases will not be charged to goods that temporarily stop or are transshipped at British ports.

We model two cases: first without, then with the impact of scale on the trade network. In our first case, as in a traditional model, outcomes are only affected through changes in trade with the United Kingdom or multilateral resistance. However, with scale economies, the decrease in UK trade will raise trade costs of neighboring countries through the trade network. Lower trade volumes lead to increased transport costs, not only for the United Kingdom, but also countries that use the United Kingdom as an entrepôt. Irish exports to the United States will now be more costly, as they will either pay the increased costs of travelling through Britain, use an alternative entrepôt, or take a low-volume, more costly direct trip. Panel A of Table 3 reports aggregate effects. The direct effect decreases global welfare by 2.3 basis

Table 3—Welfare and Trade Impact of Brexit and Arctic Passage Opening, Basis Points

	Direct effect (1)	Network effect (2)	Total effect (network and scale) (3)
Panel A. Brexit: impact of nontransport trade costs			
Δ Average global welfare	−2.3		−9.2
Δ Container trade volumes	−19.5		−100.7
Panel B. Arctic Passage: impact of endogenous trade costs			
Δ Average global welfare	1.4	2.9	6.4
Δ Container trade volumes	23.2	44.0	94.7

Notes: Panel A presents results for Brexit, a 5 percent increase in nontransportation trade costs κ_{ij} between the United Kingdom and its trading partners. The direct effect of Brexit only accounts for changes in direct trade with the United Kingdom or multilateral resistance. The total effect allows for the change in direct UK trade to impact trade costs with neighboring countries through the trade network. Panel B presents results for the Arctic Passage counterfactual. The direct effect of the passage opening only accounts for direct changes in physical distance between countries. The network effect results allow for indirect shipping through the trade network as a result of the passage opening. The total effect adds in the scale impact.

points (column 1). The introduction of scale economies leads to a decrease of 9 basis points. Trade volumes follow a similar pattern. Figure 15 highlights the distributional effects in terms of welfare (see online Appendix Figure A.16 for trade volumes). Scale economies amplify the Brexit impact, especially for European countries. Notably, the impact of scale is not well predicted by the nonscale case (panel B, Figure 15). We document significant negative welfare impacts on Ireland, Iceland, and other Nordic countries that rely on UK feeder routes to get their goods to large vessels that ply transoceanic trade (online Appendix Table A.15).

D. *Impact of Endogenous Trade Costs on the Network*

We evaluate the importance of endogenous trade costs by demonstrating the welfare and trade impacts from the two endogenous mechanisms in our model: (i) network effects—allowing countries to ship indirectly—and (ii) scale effects—allowing countries to ship indirectly and take advantage of scale economies. We achieve this by studying the physical trade route changes due to the opening of the once-fabled Northeast and Northwest Passages through the Arctic Ocean between North America, northern Europe, and East Asia as a viable shipping route due to global warming. For example, a ship traveling from South Korea to Germany would take roughly 34 days via the Suez Canal but only 23 days via the Northeast and Northwest Passages (*Economist* 2018). For every link within the network, we compute the difference in sea distance using Dijkstra's algorithm between world maps with and without arctic ice caps (online Appendix A.2). Panel A of Figure 16 compares existing shipping routes today and shortest ocean-going distance of these routes after the Arctic Passage is viable.

We compare three different cases. First, we consider a network-naive exogenous trade cost case where we only allow for changes in origin-destination trade costs between country pairs for which the direct bilateral distance decreased. Second, for all observed links with positive traffic, we recalculate t_{kl} using new distances with

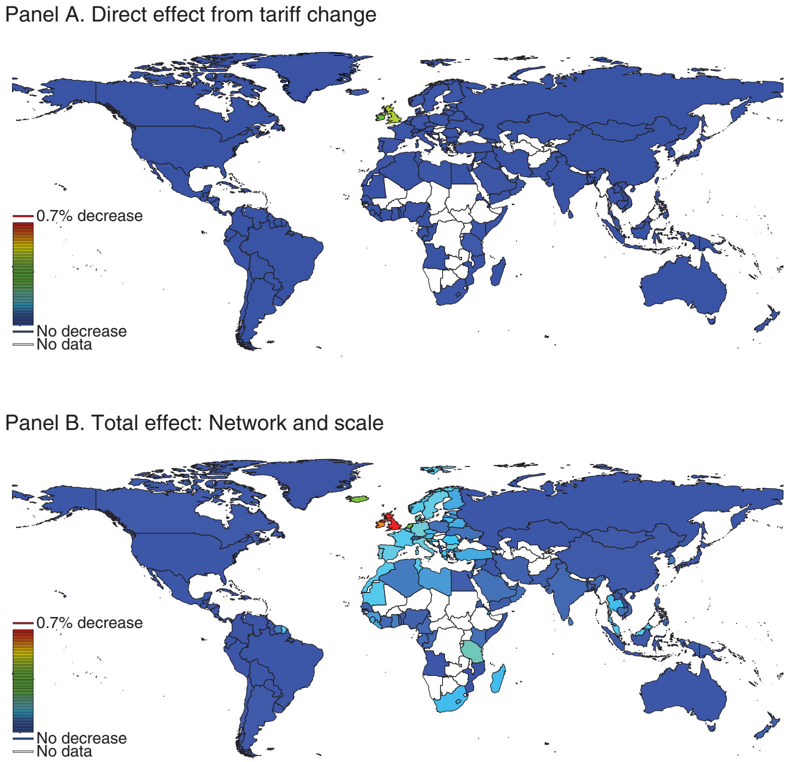

Figure 15. Welfare Changes—Brexit

Notes: These two plots show the percent change in welfare (the relative price index) of a simulated 5 percent increase in trading costs with the United Kingdom. Darker reds reflects a greater increase and blue represents no change. Omitted countries are white. Panel A reflects changes if shipping costs remain constant, reflecting only welfare changes due to changes in prices. Panel B allows for a scale economy feedback loop on transportation costs for all countries.

the option of traveling through the Arctic Passage and α_2 in equation (11). Here, countries without direct connections through the passage—for example, China and Ukraine—experience trade cost changes due to the trade network effects. Third, we repeat the second case accounting for the impact of scale: as trade costs change, trade volumes change, reducing trade costs further.

Assuming exogenous trade costs with our input-output structure, column 1, panel B of Table 3 shows that the network-naive and direct effects of the Arctic Passage are positive, with aggregate welfare increasing 1.4 basis points, and container trade volumes increasing 23 basis points. Endogenizing trade costs to allow for the trade network impact of the passage—including indirect shipping—doubles the aggregate welfare effect to 2.9 basis points and increases worldwide container volumes by 44 basis points (column 2, panel B of Table 3). Allowing for both scale and network effects triples and doubles the welfare and trade impact relative to the network results.

Panel B in Figure 16 plots the top 20 most impacted countries, showing gains are particularly pronounced in East Asian entrepôts like Hong Kong and Singapore that disproportionately benefit from the scale economy. Scandinavian countries also

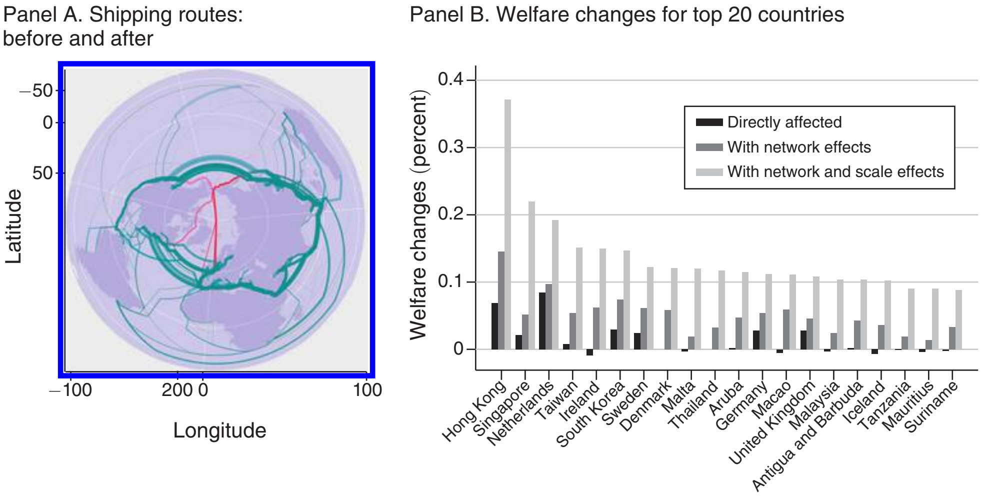

Figure 16. The Opening of the Arctic Passage

Notes: The red lines in panel A indicate counterfactual shipping. Blue lines indicate existing shipping. Their overlap is brown. Most global shipping utilize similar routes, which results in many overlapping brown lines. Route width reflects the number of containers (TEU). Panel B shows the percent change in welfare of the simulated opening of the Arctic Passage for the 20 countries with the largest welfare changes. The first bar reflects only the trade cost changes on routes that are directly affected from the opening. The second bar allows for the trade costs to affect indirect trade with network effects while the third bar allows for the endogenous response to scale economies.

gain due to their geography. Denmark and Finland, which in the baseline first case have zero or a small trade diversion impact, gain due to the trade network and scale response. Online Appendix Figure A.17 shows changes in the relative wage-adjusted price index.

VIII. Conclusion

This paper studies entrepôts, the trade network they form, and their impact on international trade. We characterize the global container shipping network as a hub-and-spoke system by documenting that the majority of trade is indirect and flows from origins to destinations through entrepôts (hubs). To rationalize these stylized facts, we develop a general equilibrium model of world trade with endogenous trade costs and entrepôts, estimating both the underlying trade costs on all routes, and scale economies. We quantify the impact of the trade network on global trade and welfare, highlighting how changes at nodes operate through the network, entrepôts, and scale economies to create widespread impacts. We find that infrastructure investments at entrepôts generate on average ten times the global welfare impact relative to investment elsewhere.

While we are singularly focused on containerized shipping because containerized trade accounts for the majority of global seaborne trade, the hub-and-spoke network is not specific to just containerized trade (Rodrigue, Comtois, and Slack 2013). Such networks are also prevalent in freight services like UPS or DHL in addition to air transport. And while our estimates of scale economies are agnostic to underlying mechanisms, future work should consider the roles of fixed costs in enabling the scale economies in containerized shipping, especially the costs incurred by potential oligopolies in setting shipping networks and the endogenous creation of firm-specific hub-and-spoke networks. In particular, we can account for leg-level monopolies and variable markups but not within-firm spillovers in sea route selection. While sector-specific research has been done on these networks, future work should consider a tractable general equilibrium framework able to quantify welfare effects.

REFERENCES

Ackerberg, Daniel, C. Lanier Benkard, Steven Berry, and Ariel Pakes. 2007. "Econometric Tools for Analyzing Market Outcomes." In *Handbook of Econometrics*, Vol. 6A, edited by James J. Heckman and Edward E. Leamer, 4171–276. Amsterdam: North-Holland.

Alder, Simon. 2015. "Chinese Roads in India: The Effect of Transport Infrastructure on Economic Development." Unpublished.

Allen, Treb, and Costas Arkolakis. 2014. "Trade and the Topography of the Spatial Economy." *Quarterly Journal of Economics* 129 (3): 1085–140.

Allen, Treb, and Costas Arkolakis. 2022. "The Welfare Effects of Transportation Infrastructure Improvements." *Review of Economic Studies* 89 (6): 2911–57.

Allen, Treb, and Dave Donaldson. 2018. "The Geography of Path Dependence." Unpublished.

Anderson, James E., Mykyta Vesselovsky, and Yoto V. Yotov. 2016. "Gravity with Scale Effects." *Journal of International Economics* 100: 174–93.

Astra Paging. 2014. *Automatic Identification System Data*. Historical AIS Data.. https://www.astrapaging.com/data-services (accessed January 25, 2015).

Asturias, Jose. 2020. "Endogenous Transportation Costs." *European Economic Review* 123: 103366.

Barjamovic, Gojko, Thomas Chaney, Kerem Coşar, and Ali Hortaçsu. 2019. "Trade, Merchants, and the Lost Cities of the Bronze Age." *Quarterly Journal of Economics* 134 (3): 1455–503.

Bartelme, Dominick G., Arnaud Costinot, Dave Donaldson, and Andrés Rodríguez-Clare. 2019. "The Textbook Case for Industrial Policy: Theory Meets Data." NBER Working Paper 26193.

Bernhofen, Daniel M., Zouheir El-Sahli, and Richard Kneller. 2016. "Estimating the Effects of the Container Revolution on World Trade." *Journal of International Economics* 98: 36–50.

Berry, Steven T. 1992. "Estimation of a Model of Entry in the Airline Industry." *Econometrica* 60 (4): 889–917.

Bertoli, Simone, Michaël Goujon, and Olivier Santoni. 2016. "The CERDI-Seadistance Database." HAL Working Paper halshs-01288748.

Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papageorgiou. 2020. "Geography, Transportation, and Endogenous Trade Costs." *Econometrica* 88 (2): 657–91.

Caliendo, Lorenzo, and Fernando Parro. 2015. "Estimates of the Trade and Welfare Effects of NAFTA." *Review of Economic Studies* 82 (1): 1–44.

CEPII. 2017. "BACI: International Trade Database at the Product-Level." http://www.cepii.fr/cepii/en/bdd_modele/bdd_modele_item.asp?id=37 (accessed January 19, 2020).

Ciliberto, Federico, and Elie Tamer. 2009. "Market Structure and Multiple Equilibria in Airline Markets." *Econometrica* 77 (6): 1791–828.

Conte, Maddalena, Pierre Cotterlaz, and Thierry Mayer. 2022. "The CEPII Gravity Database." CEPII Working Paper 2022-05.

Coşar, A. Kerem, and Banu Demir. 2018. "Shipping Inside the Box: Containerization and Trade." *Journal of International Economics* 114: 331–45.

Costa, Daniel. 2023. "Landlocked Country." *Encyclopedia Britannica*, September 14, 2023. https://www.britannica.com/topic/landlocked-country.

Cullinane, Kevin, and Mahim Khanna. 2000. "Economies of Scale in Large Containerships: Optimal Size and Geographical Implications." *Journal of Transport Geography* 8 (3): 181–95.

Drewry Maritime Research. 2011–2016. *Container Freight Rate Insight*. London, UK: Drewry Maritime Research. https://www.drewrycfri.co.uk/ (accessed February 1, 2017).

Ducruet, César, Réka Juhász, Dávid Krisztián Nagy, and Claudia Steinwender. 2019. "All Aboard: The Aggregate Effects of Port Development." Unpublished.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography, and Trade." *Econometrica* 70 (5): 1741–79.

Fajgelbaum, Pablo D., and Edouard Schaal. 2020. "Optimal Transport Networks in Spatial Equilibrium." *Econometrica* 88 (4): 1411–52.

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer. 2022. *Penn World Table Version 9.1*. Groningen Growth and Development Centre. https://doi.org/10.15141/S50T0R.

Gambrell, Jon and Samy Magdy. 2021. "Shipping Losses Mount from Cargo Vessel Stuck in Suez Canal." *Associated Press*, March 24, 2021. https://apnews.com/article/cargo-ship-stuck-suez-canal-egypt-d8b6696716ab45794cba583ecf0e4001 (accessed July 7, 2021).

Ganapati, Sharat, Woan Foong Wong, and Oren Ziv. 2024. Data and Code for: "*Entrepôt: Hubs, Scale, and Trade Costs*." Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. https://doi.org/10.3886/E188242V1.

Gaulier, Guillaume, and Soledad Zignago. 2010. "BACI: International Trade Database at the Product-Level: The 1994–2007 Version." CEPII Working Paper 2010-23.

Grant, Matthew, and Meredith Startz. 2022. "Cutting Out the Middleman: The Structure of Chains of Intermediation." NBER Working Paper 30109.

Head, Keith, and Thierry Mayer. 2014. "Gravity Equations: Workhorse, Toolkit, and Cookbook." In *Handbook of International Economics*, Vol. 4, edited by Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, 131–95. Amsterdam: North-Holland.

Heiland, Inga, Andreas Moxnes, Karen Helene Ulltveit-Moe, and Yuan Zi. 2019. "Trade From Space: Shipping Networks and the Global Implications of Local Shocks." CEPR Discussion Paper 14193.

Holmes, Thomas J., and Ethan Singer. 2018. "Indivisibilities in Distribution." NBER Working Paper 24525.

Hummels, David, Volodymyr Lugovskyy, and Alexandre Skiba. 2009. "The Trade Reducing Effects of Market Power in International Shipping." *Journal of Development Economics* 89 (1): 84–97.

Hummels, David, and Alexandre Skiba. 2004. "Shipping the Good Apples Out? An Empirical Confirmation of the Alchian-Allen Conjecture." *Journal of Political Economy* 112 (6): 1384–402.

International Maritime Organization. 2003. *Guidelines for the Installation of a Shipborne Automatic Identification System* (*AIS*). London: International Maritime Organization.

ISO (**International Organization for Standardization**)**.** 2023. *Country Codes, Officially Assigned Codes*. Online Browsing Platform (OBP). https://www.iso.org/obp/ui/#search (accessed June 20, 2018).

Kelso, Nathaniel Vaughn, and Tom Patterson. 2010. "Introducing Natural Earth Data — naturalearthdata.com." *Geographia Technica* 5 (S): 82–89.

KGM & Associates. 2014. *Eora Global Supply Chain Database*. https://worldmrio.com/eora/ (accessed April 1, 2018).

Kucheryavyy, Konstantin, Gary Lyn, and Andrés Rodríguez-Clare. 2019. "Grounded by Gravity: A Well-Behaved Trade Model with External Economies." Unpublished.

Lashkaripour, Ahmad, and Vova Lugovskyy. 2023. "Scale Economies and the Structure of Trade and Industrial Policy." *American Economic Review* 113 (10): 2759–2808.

Maritime Administration. 2018. *US Waterborne Foreign Container Trade by US Customs Ports 2000–2017*. Washington, DC: US Department of Transportation. https://www.maritime.dot.gov/data-reports/data-statistics/us-waterborne-foreign-container-trade-us-customs-ports-2000-%E2%80%93-2017 (accessed April 2023).

OECD. 2018. *Annex Table 55: Import Penetration*. Paris: OECD Publishing; distributed by OECD Economic Outlook, Volume 2018 Issue 2. https://doi.org/10.1787/eco_outlook-v2018-2-en.

Panjiva. 2014. *US Bills of Lading*. Panjiva Supply Chain Intelligence. (accessed June 21, 2019).

Paris, Costas. 2021. "Singapore Looks to Tech to Bolster Its Role in Shipping." *Wall Street Journal*, April 8, 2021. https://www.wsj.com/articles/singapore-looks-to-tech-to-bolster-its-role-in-shipping-11617909788 (accessed July 7, 2021).

Paris, Costas, and Jared Malsin. 2021. "Suez Canal Is Blocked by Container Ship Causing Huge Traffig Jam." *Wall Street Journal*, March 24, 2021. https://www.wsj.com/articles/the-suez-canal-is-blocked-by-a-giant-container-ship-11616560437 (accessed July 7, 2021).

Redding, Stephen J., and Matthew A. Turner. 2015. "Transportation Costs and the Spatial Organization of Economic Activity." In *Handbook of Regional and Urban Economics*, Vol. 5, edited by Gilles Duranton, J. Vernon Henderson, and William C. Strange, 1339–98. Amsterdam: North-Holland.

Rodrigue, Jean-Paul, Claude Comtois, and Brian Slack. 2013. *The Geography of Transport Systems*. 3rd ed. London, UK: Routledge.

Rua, Gisela. 2014. "Diffusion of Containerization." FEDS Discussion Paper 2014-88.

Schott, Peter K. 2008. "The Relative Sophistication of Chinese Exports." *Economic Policy* 23 (53): 6–49.

Sheppard, David, Harry Dempsey, and Heba Saleh. 2021. "Suez Canal Blocked after Huge Container Ship Runs Aground." *Financial Times*, March 24, 2021. https://www.ft.com/content/eec9f3a6-2817-45f5-b007-a290f3e530c6 (accessed July 7, 2021).

Simonovska, Ina, and Michael E. Waugh. 2014. "The Elasticity of Trade: Estimates and Evidence." *Journal of International Economics* 92 (1): 34–50.

Skiba, Alexandre. Forthcoming. "Regional Economies of Scale and Regional Welfare." *Review of International Economics*.

Smith, Shaw. 2015. "Saudi Port Initiatives sat Forefront of Push to Diversify." *Financial Times*, June 28, 2015. https://www.ft.com/content/2913363e-09e6-11e5-a6a8-00144feabdc0 (accessed October 23, 2019).

Sutton, John. 1991. *Sunk Costs and Market Structure: Price Competition, Advertising, and the Evolution of Concentration*. Cambridge, MA: MIT Press.

US Customs and Border Protection. 2017. *Automated Manifest Interface Requirements – Ocean ACE M1; Appendix F; Schedule K*. Washington, DC: US Customs and Border Protection. https://www.cbp.gov/sites/default/files/assets/documents/2017-Feb/appendix_f_0.pdf (accessed November 2, 2018).

US Customs and Border Protection. 2020. *CBP Automated Mianifest Interface Requirements; Appendix E; Schedule D; US CBP Port Codes*. Washington, DC: US Customs and Border Protection. https://www.cbp.gov/sites/default/files/assets/documents/2020-Jan/ACE%20Appendix%20E%20Schedule%20D%20January%202020%20%281%29.pdf (accessed June 20, 2018).

Nidhi, Verma, and Krishna Das. 2016. "With Eye on China, India Doubles Down on Container Hub Ports." *Reuters*, July 28, 2016. https://www.reuters.com/article/world/with-eye-on-china-india-doubles-down-on-container-hub-ports-idUSKCN1080NT/ (accessed April 1, 2018).

Wong, Woan Foong. 2022. "The Round Trip Effect: Endogenous Transport Costs and International Trade." *American Economic Journal: Applied Economics* 14 (4): 127–66.

World Bank. 2018. *Development Indicators*. Databank. https://databank.worldbank.org/source/world-development-indicators (accessed January 8, 2018).

World Shipping Council. 2018. "About the Industry." https://www.worldshipping.org/about-the-industry/global-trade/trade-statistics (accessed February 3, 2018).