Entrepôt*

Sharat Ganapati — Georgetown University
Woan Foong Wong — University of Oregon
Oren Ziv — Michigan State University

Preliminary and incomplete. Please do not circulate or post.

August 2019

Abstract

We study the effects of transportation networks in international trade and the potential for local and regional spillovers from the concentration of commerce at entrepôts, trading hubs where goods travel through—from other origins, and bound for other destination. Using novel data, we retrace the path of each containerized good entering the US and document the ubiquitous use of entrepôts in global shipping, the significant heterogeneity in connectivity, and the concentration of shipments through a small number of entrepôts. We build a model of endogenous entrepôt formation which incorporates route choice by exporters within a Ricardian setting and economies of scale in shipping, allowing for entrepôts to emerge endogenously. We use the model to estimate trade costs on each leg and the effect of scale on shipping costs. Finally, we embed our estimation in a quantifiable general equilibrium model to run counterfactuals which show that (1) changes in trade costs of a node in the transportation network results in regional trade and welfare spillovers and (2) the magnitude of this spillover is larger when the node is an entrepôt.

*Contact: sganapati@georgetown.edu, wfwong@uoregon.edu, and orenziv@msu.edu. We thank Meredith Crowley, Carsten Eckel, and seminar participants at the Federal Reserve Board, Ludwig-Maximilians-Universität Munich, University of Cambridge, Georgetown University, as well as University of Mannheim for very helpful comments. We are also very grateful for the valuable feedback from participants at the 2019 Society for Advancement of Economic Theory Conference. Andrew Castro provided outstanding research assistance.

1. Introduction

International trade is generally thought of as a bilateral arrangement between an exporter and an importer. However in reality, the act of exchanging goods over borders involves more than the production and consumption of these goods: shipping, transshipping and distributing these goods can include multiple agents and additional countries beyond the producers and consumers.

This activity is often concentrated at entrepôts, which are trading hubs where goods travel through—from other origins, and bound for other destination. The notion that countries stand to gain by becoming hubs for these commercial activities has a long history and continues to be a powerful narrative. Governments and local port authorities have and continue to invest billions of dollars with the specific aim of becoming or maintaining their role as entrepôts.[1]

This paper studies the role of transportation networks and entrepôts in international trade, incorporating trade spillovers due to network linkages and indirect trade. We seek to answer three questions in this paper: (1) How indirect is trade? (2) What are the international trade implications from indirect trade (i.e. the transportation networks)? (3) What are the positive (or negative) regional spillovers of entrepôts? To address these questions, we begin by documenting the indirectness of and the use of entrepôts in US containerized trade, constructing and using two new datasets to map individual journeys at the container level. We use this data to describe the role of entrepôts in international trade. We then build a general model of international trade where route choice endogenously gives rise to hubs, and use the resulting estimation equations to estimate leg-specific trade costs and estimate the role of scale economies and the concentration of commerce at entrepôts in reducing shipping costs. Finally, we embed our results in a quantitative general equilibrium model to quantify local and regional effects of en-

[1] Saudi Arabia is implementing a $7 billion project to expand the container capacity of King Abdullah port with explicit plans to "be a major east-west marine transshipment location."(Financial Times, 2015) while India is spending $4 billion to construct its first container hub port to rival Chinese facilities in the region (Reuters, 2016). On the other hand, established entrepôts are investing similar amounts to maintain their position. Singapore is investing $1.1 billion to boost its capacity by 21 million TEUs to "stay ahead of the curve as a world-class hub port" (Port Technology, 2018) following a $3 billion project to construct an automated container yard (Ship and Bunker, 2012).

trepôts on trade volumes and welfare.

We compile two novel data sets which allow us to uniquely characterize the global trading network: a shipment-level data set of bills of lading for the universe of US container imports, and a global data set on container ships' ports of call. By combining these two data sets, we are able to reconstruct the origin to destination journeys taken by individual shipments for the majority of in-bound US shipments.

We use our data to generate the first comprehensive look at the nature of the global shipping network and shipments into the United States. We find indirectness is both ubiquitous and heterogeneous, and the global trade network is incredibly concentrated in a small number of entrepôts. We summarize our findings in three stylized facts: First, the average shipment into the US is indirect, stopping at on average two countries; and the average shipment from 90% of countries makes at least one stop at another country.[2] Second, the majority of countries export to the US indirectly with significant country-level heterogeneity. Figure 1 shows the number of average stops for a container from a foreign origin country to its US destination. Distance and size matter: average shipments from countries close to the US and bigger countries have relatively fewer stops. Third, within country-pairs, significant heterogeneity exists: the average country has 151 unique shipping routes to the US in our data. Countries with more US trade exhibit more route variety.

To explain the data, we build a model of global trade with entrepôts, where individual firms choose shipping routes and compete for consumers in destination countries in a generalized Ricardian setting which is flexible enough to accommodate a broad range of input-output linkages. Lowest-cost routes can involve shipments through third-party countries, and entrepôts endogenously arise at ports through which shipping costs are lowest. Crucially, scale economies may work to reduce shipping costs on individual links as traffic grows at individual ports or links.

Using data on shipment flows through entrepôts, we use the model to estimate trade costs for each country pair in a subset of our data. A key intuition for the disproportion-

[2]This is true for 132 countries out of a total of 147. This number does not include landlocked countries because landlocked countries by definition would need to stop at a coastal country in order to ship to the US.

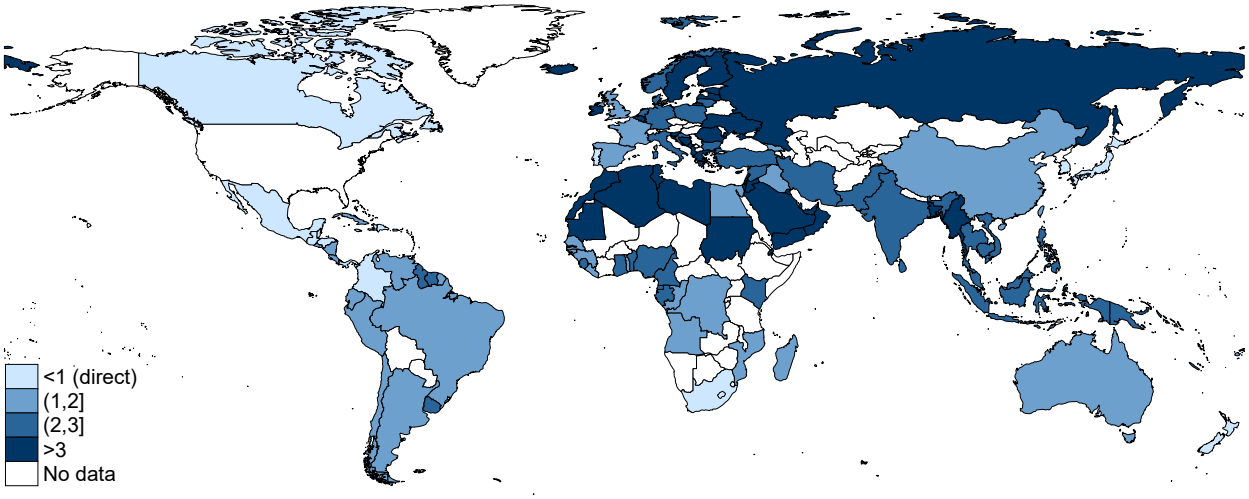


Figure 1: Number of Stops between Origin and US Destination

Stops are by country and weighted by container volume (TEU). The US is not included since it is the destination country. Landlocked countries are also not included, since by definition they would need to stop at a coastal country. 34 of the shipment origin countries are landlocked accounting for 1.6 percent of total TEUs. The missing remaining countries are either due to lack of overall trade with the US (e.g. Somalia) or due to the merge process (e.g. Namibia).

ate use of a few ports, as well a potential mechanism through which capturing indirect trade may benefit locations and induce competition for hub status is the existence of scale economies. We estimate a strong scale elasticity: all else equal, a 1% increase in traffic reduces costs on the same leg by nearly 0.35%.

Finally, to estimate the impact of the trade network on global shipping as well as the localized effects and spillovers from hubs, we adopt further structure to our general model. We use the resulting quantified general equilibrium model to run counterfactual predictions. Our counterfactuals drive home the point that small and highly connected countries have out-sized effects on world trade. Decreasing trade costs with Singapore has a similar effect as decreasing trade costs with China, even though Singapore is an order of magnitude smaller. Even though the UK is a large exporter, Belgium trade costs are more important to world trade flows.

This paper makes theoretical and empirical contributions to the literature. Our first empirical contribution is the datasets we compile. With notable exceptions,[3] these data

[3]Holmes and Singer (2018) use bill of lading data for Walmart containers, and Brancaccio, Kalouptsidi and Papageorgiou (2017) use port call data for to study bulk carriers.

are rare in the literature and to our knowledge never used in this comprehensive fashion. We believe our data are singular in their ability to reconstruct both the global shipping network and the routes taken by individual shipments.[4]

Our model is most closely related to Allen and Arkolakis (2019). The theoretical contribution of this paper is to extend their Armington framework where route cost shocks are born by consumers to a general Ricardian setting, where link-specific trade volumes reflect not only choice of routes but head-to-head competition and selection on prices at destinations. We show that their basic estimating equations generalize to this setting.

Both Allen and Arkolakis (2019) and Fajgelbaum and Schaal (2017) consider diseconomies of scale in transportation in an intra-national context. Our focus is on international container shipping network where infrastructure investment is localized at nodes which includes entrepôts and where spillovers between nodes may even be negative because of scale economies. In that respect, this paper is related to a literature on economies of scale in trade. Holmes and Singer (2018) investigates indivisibilities in shipping and how it can create scale economies in distribution using bill of lading data. Utilizing Canada-US trade data, Anderson, Vesselovsky and Yotov (2016) identify economies of scale using a structural gravity framework. We contribute to this literature by introducing another aspect of scale, in shipping, and estimating a scale economy with respect to volume of traffic.

We are also related to a literature on endogenous transport costs (Hummels (2007), Limao and Venables (2001), and Wong (2019)). Most importantly, Brancaccio, Kalouptsidi and Papageorgiou (2017) estimate a model of endogenous trade costs and search frictions between exporters and ships. Their port of call data for bulk carriers is most similar to ours, but their focus is on dry bulk shipping and homogeneous commodities.[5]

[4] Lazarou (2016) is a complementary study which utilizes network theory to study transshipment hubs within a monopolistic competition framework. As mentioned in his paper, the lack of actual observed indirect route data in his empirical exercise leads to "ambiguity as to the existence, followed by the identification, and therefore selection, of good-specific routes and their respective measurement of distance. Our data set is a contribution to the literature in that we observe exactly this.

[5] The main difference between dry bulk shipping and container shipping is similar to the distinction between taxis and buses. Dry bulk ships likely have to depart from their destination without cargo which gives rise to the search for cargo, whereas this would be less true for containerships by virtue of containers being able to accommodate a lot of different products. Containerships are more like buses in the regard that they travel on fixed schedules between a few locations (the "round trip effect", Wong (2019)).Wong

Because we model endogenous transport costs in a general equilibrium framework, our counterfactuals are able to contribute to this literature by addressing the overall welfare impacts of these endogenous transport costs and demonstrating the spillover effects on trade and welfare.

Last but not least, this paper is connected to the literature on containerization and trade. Coşar and Demir (2018) investigates the impact of containerization on transport by studying the choice between containerization and breakbulk shipping. Bernhofen, El-Sahli and Kneller (2016) studies the impact of containerization on world trade flows while Rua (2014) investigates the diffusion of containerization. Blonigen and Wilson (2008) estimates the impact of port efficiency on trade.

The rest of this paper is organized as follows: Section 2 describes both our novel datasets and how we recover the journeys of individual shipments by combining them. Section 3 uses the dataset to characterize the indirectness of containerized trade into the US. Section 4 lays out our general model which shows how the trade network can be analyzed under a wide possible set of functional forms. Section 5 uses the model to estimate trade costs and scale economies. Section 6 estimates a quantitative general equilibrium model as a special case of the general model and uses it to generate counterfactuals to illustrate the differential effects of changes to the trade network. Section 7 concludes.

2. Data

We compile and combine two proprietary data sets in this project: global ports of call data for containerships, which allows us to reconstruct the routes taken by specific ships, and US bill of lading data for containerized imports, which gives us shipment-level data on imports into the United States. Independently, each of these datasets allow us to partially describe the global shipping network. By merging them, we are able to reconstruct nearly the entire journey most shipments entering the United States take,

(2019) studies the trade implications of endogenous transport costs within the context of joint transportation costs between hub locations. As a result, dry bulk shipping does not have the same hub-and-spoke network structure which would not give rise to systematic trade spillovers due to network linkages like our container shipping framework would.

from their initial origin point or place of receipt, to the port of entry into the United States. To our knowledge, we provide the most comprehensive reconstruction of the global shipping network and routes undertaken by individual shipments into the US.

2.1. Port of call data

Our proprietary ports of call data from Astra Paging captures vessel movements using the transponders on these ships (known as the automatic identification system, AIS). A network of receivers at ports collects and shares AIS transponder information (including ship name, speed, height in water, latitude and longitude). Using the geographic variables in the AIS data, our data marks entry and exit into a number of ports all over the world.

For each vessel, this data captures its name and identification information, the port it enters or exits, the time stamp of this record, the vessel's maximum capacity and the vessel's height in the water before and after stopping in the port, which indicates the vessel's load. Using these data elements, we are able to calculate an estimated shipment volume between each port pair.

Figure 2 shows the coverage of the shipping network in our port of call data. Each line represents a journey between port pairs, the dots, undertaken by a containership. The data in the figure is unweighted by volume or number of journeys.[6]

Our preliminary sample covers the same six months period, from April to October 2014. Over this period, we have information on 4986 unique container ships with a combined capacity of 18.13 million TEU. This is over 90% of the global container shipping fleet. These ships make 429,868 calls at 1203 ports.[7]

[6]We can further identify the communities—groups of nodes that are more densely connected to one another than to other nodes—within the shipping network by using the Louvain method, a widely used community detection algorithm (figure A.1).

[7]The limitation of this data corresponds to the limitations of the AIS receiver network. Ports where receivers are not located or where information is not shared will not show up in our data. In addition, if transponders are turned off (sometimes for security reasons) or transmissions are not recorded (sometimes for weather-related reasons), ports of call can be missed.



Figure 2: Map of Global Port of Call Network

The dots represent the ports in our data set. Each line represents a journey between port pairs undertaken by a containership.
Source: Authors' calculations (April to October 2014).

2.2. Bill of lading data

We put together proprietary bills of lading data which captures shipment-level information for all imports into the United States. International shipping relies on an industry-standardized system of bills of lading. The US Customs and Border Patrol agency collects these bills in addition to customs information at all ports of entry into the US and this data is obtained from the agency via the Freedom of Information Act from Panjiva.

Bills of Lading act as receipts of shipment, recording all information on the shipment, all the parties involved in the shipping process, the shipment's origin and destination, and some customs information. Specifically, for the purposes of this study, our data captures the foreign location where the shipment originated from, the foreign port where it was loaded on the containership which brings it into the US, and the US port where it was unloaded from the containership. In addition, we know the name and identification number of the containership which transported the shipment as well as the shipment's weight, number of containers (TEUs), and product information. For a smaller subset of the shipments, we observe value information.

Our preliminary Bill of Lading dataset covers six months of US imports from April

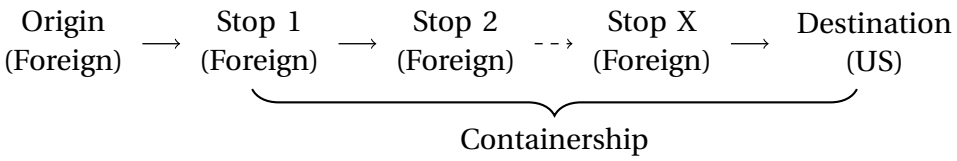


Figure 3: Combined Dataset: Routes Undertaken by Shipments into the US

to October 2014. Over this period, a total of 14.8 million TEUS weighting 106 million tons were imported into the US. This accounts for about three quarters of the 2014 TEU and tonnage imports (77 percent and 74 percent respectively, Maritime Administration, US Department of Transportation). There are 227 shipment origin countries, 225 place of receipt countries, and 194 ports of lading countries.[8] Noncontainerized goods, including goods on roll-ons (vehicle carriers), bulk cargo liners (for commodities), and non-containerized cargo ships are not observed in our data.

2.3. Reconstructing shipment routes

Using the containership information,[9] port of arrival information,[10] timing of unlading and ports of call at US ports, and port of lading information,[11] we are able to match the bills of lading to the journeys of specific containerships, then use the ports of call between lading and unlading to reconstruct each shipment's path from its foreign origin to US destination.

Figure 3 outlines the information our data captures for the purposes of this study. We see the following location information for each shipment into the US: the foreign location where the shipment originated from (*Origin*), the foreign location where it was

[8]The United Nation recognizes 233 countries.

[9]We use Vessel IMOs which are identifiers that are unique to containership vessels and stay with vessel hulls for the lifetime of their operation. Only about 4000 ships are identified in Bills of lading by IMO. An additional (roughly) 2,000 ships are matched to IMOs using a fuzzy string match, after which matches are made with the help of excellent undergraduate research assistants.

[10]Ports of arrival are recording using UNLOCODEs in the port of call data and US Census Schedule D codes in the Bill of Lading data. We construct a crosswalk with the excellent help of undergraduate research assistants.

[11]Port of lading are recorded using CBP's Schedule K Foreign ports on Bill of Lading and UNLOCODES in the port of call data. We construct a crosswalk between these with the help of our excellent undergraduate research assistants.

loaded on the containership which brings it into the US (*Stop 1* or port of lading), subsequent stops which the containership makes (*Stop 2*, ... , *Stop X*) and the US destination where it was unloaded from the containership (*Destination*).

What remains not directly observable is the shipment's journey between its Origin and the first stop we observe in our data. In particular, this initial portion of the shipment's journey could take place overland (by trucks or rail) or by sea on another ship. This information is not recorded by both our datasets and therefore is impossible for us to observe. We take special care in our estimation to reduce bias from this omission where possible.

Over two-thirds of all bills of lading can be matched to routes using this method. Our matched bill-of-lading and port-of-call dataset includes over 1 million bills of lading.[12]

3. Stylized Facts

In this section, we use the datasets above to explore the nature of international shipping network and the routes goods entering the US take along that network. Our analysis generates four stylized facts about trade and the international shipping network. (1) Trade is generally indirect through entrepôt; (2) The extent of indirectness varies by country, with larger countries more likely to send goods direct to the US; (3) Conditional on a given origin, a high degree of variation exists in the routes goods take to the US; (4) Shipping to the US is highly concentrated in a small number of entreôt.

3.1. Indirect Trade by Shipment

We start by describing how realistic the traditional bilateral view of trade is. Figure 4 reports the distribution of the number of observed country stops made by each shipment, weighted by TEU.

Only about 30 percent of containers are exported to the US directly, making no stops along the way. The majority of containers are indirect (about 70 percent), making at least one stop before reaching its final destination. The average number of TEU-weighted

[12] Unmatched shipments may have missing and unrecoverable ship information, or ports of call that do not match lading and unlading records on bills of lading.

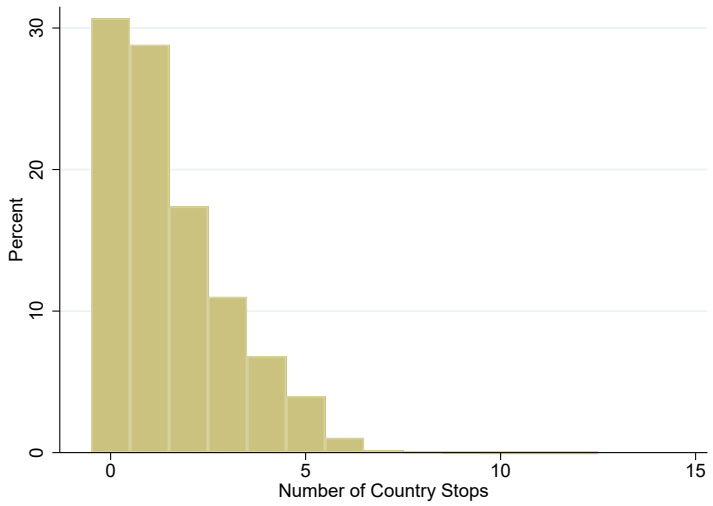


Figure 4: Distribution of Third-Party Countries Involved in Bilateral Trade
Number of Country Stops per Container (TEU)

Note: The US is not included since it is the destination country. Landlocked countries are also not included, since by definition they would need to stop at a coastal country. 34 of the shipment origin countries are landlocked accounting for 1.6 percent of total TEUs. The missing remaining countries are either due to lack of overall trade with the US (e.g. Somalia) or due to the merge process (e.g. Namibia).

stops is about 2, which implies that a container stops at an average of around 2 countries that are not involved as either a producer or a consumer. The average number of port stops is higher at 4 stops (figure A.2).

Panels (A) and (B) of Appendix Figure A.3 show that this pattern is robust to using shipment weight and value. Less than 30 percent by weight measured in tons and less than 20 percent by value are transported between an exporting origin and the US directly.

From these results, we can summarize the how indirect trade into the US is:

Stylized Fact 1 *The majority of containerized trade into the US is indirect, involving an average of two countries that are not a part of the bilateral trade.*

3.2. Variation in Directness by Country

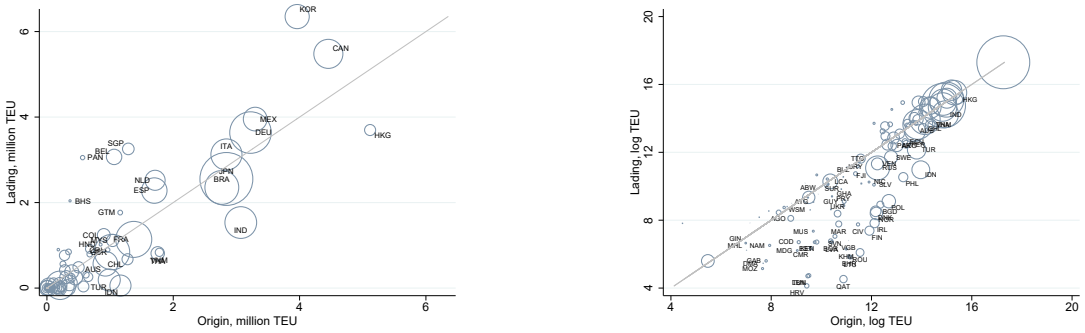
If indirectness is common, to what extent does indirectness vary at the country level? for the average country, the average container will stop at 2 additional countries, confirming that fact 1 is not driven by outlier countries. However, as is evident in Figure 1, there is a high degree of variance at the country level. The average shipments from Mexico, Canada, Panama, Haiti, and South Africa are all direct to the US, while the average shipments from 22 countries stop in four or more additional countries.

Another way to consider directness of trade is to focus our attention on the countries that are the first observed stop in our data. As mentioned earlier, this first observed stop is unique in that it is the port the shipment is loaded onto the final ship in its journey to the US. One reason entrepôts are so important may be related to their role as distribution hubs for countries whose shipments are not loaded onto ships bound directly for the United States.

In Panel (A) of Figure 5, we plot each country's total volume as a first stop (lading volume) against their origin volume. Those countries above the 45° line have higher lading volumes compared to origination, implying they are more involved with US trade as third-party countries than as exporters. This metric captures countries long-associated with entrepôt, including Singapore (SGP), Belgium (BEL), Netherlands (NLD), Korea (KOR), and Panama (PAN), as well as others. Since this dataset is only for US imports, this relative ranking may miss entrepôts that deal mainly in non-US-bound traffic or may overemphasize the importance of US-proximate locations like Jamaica and the Bahamas.

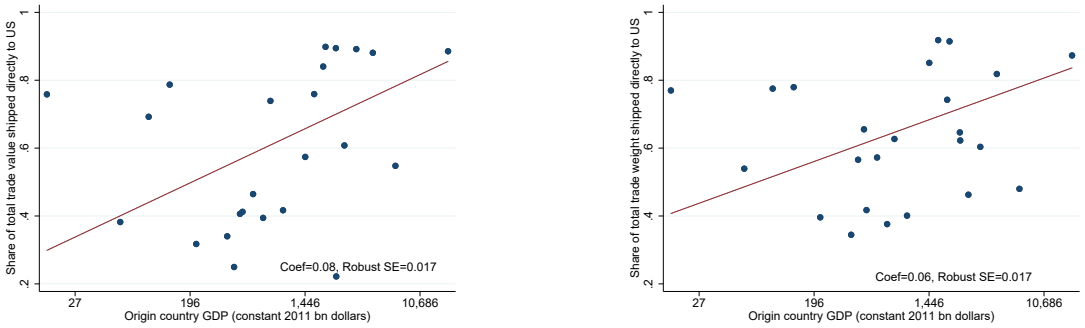
Panel (B) of Figure 5 repeats Panel (A) in logs, which shows the extent to which countries are below 45°, this implies that the goods originating from these countries make their way to the US via a third country and that they are more involved with US trade as exporters than as third-party countries.

What if anything makes sense of the variation in directness? In Figure 6, we plot GDP against trade share at the first stop. By value and by TEU, shipments are more likely to be direct from bigger countries. We find positive and significant correlations between origin country GDP and indirect trade share, by value, weight, and TEU.



(A) Levels (B) Logs

Figure 5: Origin vs First Stop Volume



(A) Share of total value (B) Share of total weight

Figure 6: Positive correlation between origin country GDP and the share of trade shipped directly to the US

Note: Binned scatter plot with observation at the origin level. Robust standard errors clustered by origin. Weighted by total container TEU trade by origin. Source: World Bank WITS and author's calculations (April to October 2014).

Higher GDP countries are more likely to ship to the US, and their shipments are also more likely to go direct. One natural interpretation of this would be the endogenous response of shippers to the scale of shipments from these countries. Of course, the availablility of direct trade to the US could in principle reverse the causality.

We can summarize the following based on our results above:

Stylized Fact 2 *The majority of countries export to the US indirectly. There is significant variation in directness, and larger countries are more likely to ship directly.*

3.3. Variation in Connectivity within Country Pairs

Given that the majority of trade is indirect by shipment and also by country, a natural question to ask is how much variation is there in the indirectness of trade with the US for any particular country. Figure 7 shows the distribution in the number of unique routes to the US by origin country. The average is about 151 routes with wide variation (about 193). The countries with the highest number of unique routes are big trading partners like China, the United Kingdom, Germany, and France. Countries with the lowest unique routes are smaller trading partners like Chad, Uzbekistan, Tajikistan, and Bhutan.

Figure 8 shows a clear positive relationship between trade with the US and unique routes. This is counter to the intuition of Stylized Fact 2. While the average shipment is most likely to go direct for countries with stronger trade ties to the US, the overall variation in routes is also larger.

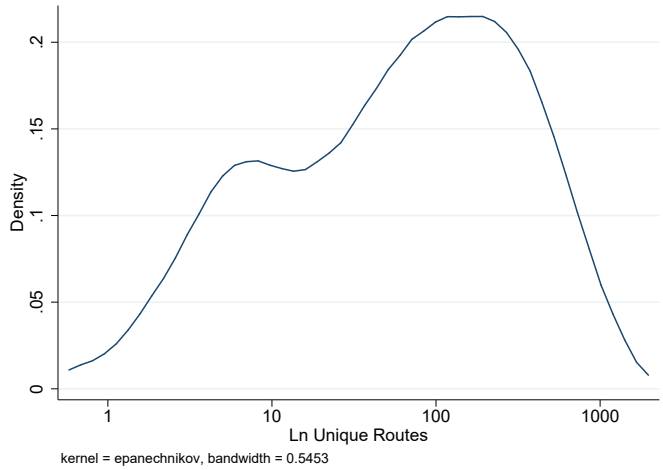


Figure 7: Number of Unique Routes by Origin-Destination Pair

These relationships can be summarized in our third stylized fact:

Stylized Fact 3 *There is significant variation in routes across origin-destination pairs. In particular, origin countries with more trade have a higher number of unique routes and a smaller number of stops (more direct trade).*

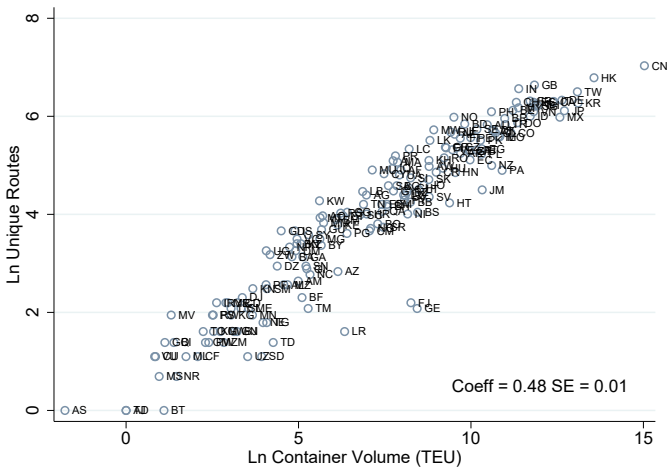


Figure 8: Positive Correlation between Unique Routes and Trade

4. Theoretical Framework

In this section, we develop a model of global trade where shipments are sent indirectly, through endogenously formed transport networks. Entrepôts emerge as ports through which goods flow but which are neither goods' origin or destination. Scale economies at these ports or along routes can further reduce costs at, and concentrate shipments through, an entrepôt.

To find closed-form solutions for trade costs and volumes, we rely heavily on the work of Allen and Arkolakis (2019). The model here extends their results to a Ricardian trade setting. In our model, consumers in each country choose the lowest-cost supplier for each product. Simultaneously, perfectly competitive producers from each potential origin country use labor to produce their good and choose an export route from the set of all possible routes. Crucially, route costs borne by producers are idiosyncratic for each product; bad shocks both make shipment through a particular route less likely as well as make final sales to destinations less likely.

Throughout, we maintain a production and consumption setting that is as general as possible, allowing for any number of goods, industries, and input-output linkages. However we impose certain restrictions on cost heterogeneity in order to generate moment conditions that can be matched to the data to yield estimates of leg-specific shipping costs.

4.1. Setup

We embed the Allen and Arkolakis (2019) model in a general Eaton and Kortum (2002) framework, allowing for a complex production function and trade and consumption in intermediates, as well as final goods.

Unlike Eaton and Kortum (2002), in our setting, the technologies available to firms in each industry and country is non-stochastic. Instead, firms additionally solve an optimal route choice problem, where idiosyncratic variation in a products' optimal route generates random variation in the price of each intermediate product-origin pair.

Consumption

In each country j, consumers consume goods $\omega_n \in \Omega_n$ from each of N industries n.

$$U_j = U_j(C_j)$$

where U_j 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.

The utility function itself is allowed to vary across destinations, and the number of goods in each industry need not be a continuum, but can be. Within each industry and product category, goods are homogenous and normal.

Production

Goods can be 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. The production function is given by

$$q_{in}(\omega) = f_{in}(z_{in}, L_{in}, Q_{in})$$

where f_{in} is a continuous, twice differentiable country-industry-specific production function, z_{in} is a production technology common to industry n and country i, L_{in} is a vector of non-tradable 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 homogenous.

Crucially, cost minimization leads all firms in the same country and industry to have the same marginal cost of production. However, the production process itself can be arbitrarily different across industries and countries.

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

where P_i is the matrix of prices of all goods ω in industries n in i and 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

In each country i, a competitive fringe of firms produces each product ω in the continuum of each industry n.

In order to sell goods abroad at any destination $j \in J$, a firm producing product ω in industry n must pay tariffs κ_{ijn} and iceberg transport costs $\tau_{ijn}(\omega)$ after optimally choosing the route between i and j to minimize the shipping costs incurred. Because markets are competitive, firms selling from i to j price their goods at marginal cost. Observed prices for these products at j are therefore

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

Producers and consumers purchasing a good ω in industry n at location j source the lowest cost supplier globally.[13]

Shipping

In order to be competitive at their export destinations, producers seek to minimize shipping costs by choosing the lowest cost shipping route available for their product. A shipping route r is comprised of a series 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 $k = K_r$).

[13] Here, non-tradable goods can be assumed to have infinite transport costs.

Following Allen and Arkolakis (2019), we assume that moving from stop to stop involves iceberg transport costs as well as product- and route-specific idiosyncratic cost shocks $\epsilon_{ijnr}(\omega)$,[14] drawn from the Fréchet distribution such that $F_{ijn}(\epsilon)$, the cumulative distribution function of the idiosyncratic draws is:[15]

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

where shape parameter $\theta > 0$ captures the randomness or dispersion in the choice of routes from i to j. A higher $\epsilon_{ijnr}(\omega)$ draw means that industry n has a lower cost for route r when shipping from i to j.

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

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

where $\tilde{\tau}_{ijr}$ is the product of all leg-specific costs t_{k_r-1,k_r} and is common to all products taking the same route r. Importantly, while iceberg transport costs are usually a function of distance (including in Allen and Arkolakis (2019)), we place no structure on the leg-specific trade costs. Instead, we allow trade costs to be a flexible function of exogenous and endogenous variables: $t = f(X_{exog}, X_{end})$. In our estimation these will later include distance and scale or route volume.

Each product's ultimate shipping cost from i to j will be the minimum transport cost route over a set of all other routes for origin i, destination j, and product ω in industry n:[16]

$$\tau_{ijn}(\omega) = \min\{\tau_{ijnr_1}(\omega), ..., \tau_{ijnr_R}(\omega)\}$$

[14] 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 θ.

[15] This distribution is the same for each product across industries so we drop the product-industry subscripts n

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

4.2. Equilibrium

We use the properties of the Fréchet distribution to find expressions for two observables: (1) the equilibrium mass of products that will be shipped from any origin to any destination through a specific leg and (2) the total volume of trade between any country pair, including transshipments. We then close the model by imposing goods and labor market clearing conditions, the former generating a within-industry gravity equation, the latter generating an expression for wages.

Route volume

Firms from origin country i select the lowest-cost route before consumers in j select the lowest-cost intermediate good supplier across all the origin countries. We observe 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 then use the properties of the Frechét distribution to consider the probability that a given country and route r' will be selected as the lowest cost route-supplier combination for good ω conditional on price p:

$$G_{jn\omega}(p) \equiv \Pr\left\{\min_{i\in I, r\in R_{ij}\backslash r'} p_{ijnr}(\omega) > p)\right\} = 1 - \exp\left\{-p^{\theta}\cdot\sum_{i}\left[(c_{in}\kappa_{ijn})^{-\theta}\cdot\sum_{r\in R_{ij}}\tilde{\tau}_{ijr}^{-\theta}\right]\right\}$$

and further find two analytic expressions. First, 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:

$$\pi_{ijnr\omega} \equiv \Pr\left\{p_{ijnr\omega} \leq \min_{i'\in I\backslash i,\, r'\in R_{ij}\backslash r} p_{i'jnr'\omega})\right\} = \frac{\left[c_{in}\kappa_{ijn}\cdot\tilde{\tau}_{ijr}\right]^{-\theta}}{\sum_{i'\in I}\left[(c_{i'n}\kappa_{i'jn})^{-\theta}\cdot\sum_{r'\in R_{i'j}}\tilde{\tau}_{i'jr'}^{-\theta}\right]} \tag{2}$$

By the law of large numbers this is also the share of all goods sold in j in industry n

that come from i and take route r.[17] Next, we can find the unconditional (pre-selection) average transport cost from i to destination j:

$$\tau_{ijn} = \gamma^{-1/\theta} \left(\sum_{r \in R_{ij}} \tilde{\tau}_{ijr}^{-\theta} \right)^{-\frac{1}{\theta}} \tag{3}$$

where γ is the function $\Gamma(t) = \int_0^\infty x^{t-1} \exp^{-x} dx$ evaluated at $\left(\frac{1+\theta}{\theta}\right)^{-\theta}$.

We define two matrices following Allen and Arkolakis (2019). First $A = [a_{ijn} \equiv t_{ijn}^{-\theta}]$, where each element is a function of the leg-specific trade cost $a_{ijn} \equiv t_{ij}^{-\theta}$. Second, we define the matrix B as $B \equiv (I - A)^{-1}$. Using these and substituting for the definition of $\tilde{\tau}_{ij}$ (equation (1)) 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 which passes through leg k, l as

$$\pi_{ijn}^{kl} = \left[(c_{in}\kappa_{ijn})^{-\theta} \cdot b_{nik} a_{nkl} b_{nlj} \right] \cdot \Phi_{jn}^{-1} \tag{4}$$

where $\Phi_{jn} = \sum_{i'} (c_{i'n}\kappa_{i'jn})^{-\theta} \cdot b_{ni'j}$ is a multilateral resistence term that accounts for the average costs, openness, and connectivity of competitors from all other countries i'. This equation is the direct analogue to equation (7) in Allen and Arkolakis (2019). The key distinction here is Φ_{jn}, which accounts for the Ricardian selection of lowest-price sources for each good ω. Intuitively, the traffic flowing on a given leg responds both to that leg's effectiveness in reducing route costs as well as to comptitive forces which make trades increasingly less likely to be pursued for more expensive routes. However, multilateral resistence is j, i-level, and therefore enters proportionately into traffic flows for all k, l-pairs – a fact that will be crucial for our estimation.

Furthermore from summing across industries, origins, and destinations, we can recover the share of observed global shipping that passed through leg k, l.

[17] Recall the number of goods in each industry is set arbitrarily large so that the law of large numbers will hold.

$$\pi^{kl} = \sum_n a_{nkl} \cdot \sum_j b_{nlj} \frac{\Phi_{kn}}{\Phi_{jn}} \tag{5}$$

Equations (4) and (5) correspond the shares of goods passing through leg k to l, including shipments bound for l and those continuing onward to other destinations. Because they account both for optimal route selection and competition on price, they correspond to observable volumes after route selection and competition among producers.

The sum of products sold in j in industry n from country i is equal to the share of all products sold in j in industry n that come from i and take route r, summed across all r routes:

$$\pi_{ijn} \equiv \sum_r \frac{\left[c_{in}\kappa_{ijn} \cdot \tilde{\tau}_{ijr}\right]^{-\theta}}{\sum_{i' \in I}\left[\left(c_{i'n}\kappa_{i'jn}\right)^{-\theta} \cdot \sum_{r' \in R_{i'j}} \tilde{\tau}_{i'jr'}^{-\theta}\right]} = \frac{\left(c_{in}\kappa_{ijn}\tau_{ijn}\right)^{-\theta}}{\Phi_{jn}} \tag{6}$$

4.3. Predictions

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

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

The first term on the right is the effect of t_{kl} on trade with i through a change in the total volume consumed at j in industry n. The first term in parentheses is the effect through any changes to the production costs at i, which can happen if the price of inputs changes or through a change in wages. The second term in parentheses the effect through trade costs between i and j in industry n, and the final term is the effect through multilateral resistance. This final term is the only term that can be negative. Intuitively, a reduction in trade costs between k and l can increase consumption at j, reduce ex-

pected trade costs between i and j, and reduce production costs at i. All of these result in an increase in trade volumes between i and j. However, a reduction in trade costs between k and l also stiffens competitions at j. If this last effect is large enough, it can overturn the sign of the first three.

The relationship is positive if and only if the following condition is true:

$$\epsilon_{X_{jn},t_{kl}} + \epsilon_{c_{in},t_{kl}} + \epsilon_{\tau_{in},t_{kl}} > -\epsilon_{\Phi_j,t_{kl}} \tag{7}$$

That is, if the elasticities of consumption at j ($\epsilon_{X_{jn},t_{kl}}$), production costs at i ($\epsilon_{c_{in},t_{kl}}$), and trade costs between i and j ($\epsilon_{\tau_{in},t_{kl}}$) with respect to t_{kl} are larger than the elasticity of multilateral resistance at j with respect to t_{kl} ($\epsilon_{\Phi_j,t_{kl}}$). Rearranging terms, we have $\frac{\partial X_{ijn}}{\partial t_{kl}} > 0$ if and only if the following holds:

$$\epsilon_{X_{jn},t_{kl}} + \left[\epsilon_{c_{in},t_{kl}} + \epsilon_{\tau_{in},t_{kl}}\right](1-\pi_{ijn}) > \sum_{i'i'} (\epsilon_{c_{i'n},t_{kl}} + \epsilon_{\tau_{i'jn},\tau_{kl}})\pi_{ijn} \tag{8}$$

The sum of the effects on production and transport costs between all other countries i' (other than i) and j has to be less than a function of the effects on production and transport cost at i and the overall propensity of consumption at j to grow.

This last expression shows most clearly that the effect of a decline in trade costs between k and l has the potential to negatively affect trade flows between i and j. In particular, if the shift differentially favors trade and production costs from other countries to j, trade volumes from i to j will suffer.

Finally, because the elasticity $\epsilon_{\tau_{in},t_{kl}}$ is equivalent to the proportion of trade from i to j that goes through k,l, a decline in t_{kl} is more likely to positively affect X_{ij} the more k,l is used in proportion to other routes and the higher that proportion is relative to the same proportion for other countries i'. The higher the same proportion is in other countries, the more likely trade volumes between X_{ij} will fall with a fall in trade costs t_{kl}.

4.4. Estimation Equation

Using equations (4) and (6) we can calculate the probability of any good traveling through leg k,l conditional on being sold from i to j. If X_{ijn} is the total volume of trade between i

and j in industry n, we can express the total volume of traffic between k and l in a given industry as

$$\Xi_n \equiv \sum_i \sum_j X_{ijn} \cdot b_{ikn} a_{kln} b_{ljn} b_{ijn}^{-1} \tag{9}$$

Note this equation is identical to what is established in Allen and Arkolakis (2019). Conditional on the observed trade volumes, the contribution of trade between i and j to traffic between k and l is invariant to multilateral resistance, tariffs, or technology. This is despite the significant differences between our economic frameworks. In particular, expensive trade routes here suffer from Ricardian selection at destination markets, where the route's impact on prices make them uncompetitive. Yet, this does not impact the estimation of trade costs.

The intuition for this result is that Ricardian selection, tariffs, and multilateral resistance all operate through adjusting the total volume of trade, but do not differentially favor one route from an origin i to a destination j. Put differently, any change to a cost in one country will affect trade from that country and others proportionally on all routes.

Equation (9) gives us a relationship between trade volumes, trade costs, and traffic for a given industry. To map our model into the data we make one final assumption: there is a set of industries $\bar{N}$ for which trade costs are identical and all trade $X_{\bar{N}} \equiv \sum_{n \in \bar{N}} X_n$ and traffic $\Xi_{\bar{N}} \equiv \sum_{n \in \bar{N}} \Xi_n$ are observable. Summing equation (9) over industries $n \in \bar{N}$ yields

$$\Xi_{\bar{N}} = \sum_i \sum_j X_{i\bar{N}} \cdot b_{\bar{N}ik} a_{\bar{N}kl} b_{\bar{N}lj} b_{\bar{N}ij}^{-1} \tag{10}$$

Equation (9) tells us that to accurately measure transport costs, we need only observables for transportation and 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. Conversely, we need not observe all traffic nor account for all trade to estimate transport costs.

We will operationalize equation (10) in our estimation by using observed total con-

tainerized traffic volumes, isolating the set of industries that are containerized and assuming that all trade in those industries are in fact by container, no other trade is by container, and all transport costs are identical across containerized industries.

4.5. Closing the model

Factor and goods market clearing and balanced trade conditions close the model. Because they are unnecessary for this stage, we refrain from writing these until we impose further structure.

5. Trade Cost Estimation

5.1. Recovering Trade Costs

To estimate the welfare impact of counterfactuals that alter the geography of container trade, we need estimates for point-to-point container trade costs. These costs are fully determined by the matrix A. This direct cost matrix A is essential as it allow for the estimation of physical trade costs τ_{ij} between any two trading nations i and j for use in counterfactual analysis through the identity $A = (I - B)^{-1}$ and Equation 10.

In an ideal world, estimation would recover the trade costs that rationalize (a) observed bi-lateral containerized trade flows and (b) observed bi-lateral containerized traffic flows. Jointly, these two objects imply a set of port-to-port iceberg trade costs in the matrix A that perfectly fit the world. With these two observable objects, we can nonlinearly solve Equation 10 in terms of the matrix A, using the identity $A = (I - B)^{-1}$.

The beauty of this procedure is that it is agnostic to the exact specification of any particular trade model that generates trade flows X. By conditioning estimation on these flows X, all origin, destination, and origin-destination factors are controlled for. In particular, items such as all origin-destination tariffs and non-tariff barriers are all accounted for.[18]

[18]This does not mean that we can disentangle the two, rather we can directly account for these factors collectively.

We first use 2014 US Customs data on containerized and non-containerized shipments to construct the share of each HS 4-digit commodity code that travels by container. This gives a share between 0 and 1, all commodities with a containerized share above 80% are labeled as a containerized good. In practice we find a bimodal distribution, with some commodities being never containerized (oil and iron ore) and other always containerized (washing machines and children's toys). With this data, we then collapse the 2014 EORA International Input-Output database at the country level, segregating containerized and non-containerized commodities.[19] With this database, we can extract out a trade matrix for containerized commodities, X, for all trade partners. We do this for all 126 independent countries in the database.

The second piece of the puzzle, the traffic matrix Ξ, is a bit more involved. Direct implementation in estimation difficult for a few reasons. First, all traffic flows may not be observed. While our data gives us a complete view of seaborne containerized trade, it omits land based trade and internal within-country trade. This is a problem for about 10% of potential trade routes for countries that are near each other. For example, containerized trade from Canada to the US can occur via truck or train, and will not be observed in our traffic matrix Ξ, even though it appears in the trade matrix X.

In light of this limitation, we assume a functional form that accounts for this omission. We need a mapping between observables, such as distance and traffic flows to trade costs. We estimate the determinants of iceberg trade costs (the terms t_{kl}) over every observable trade leg from port k to port l. This methodology follows the spirit of Allen and Arkolakis (2019).

In particular, we seek to understand the role played by both scale and traditional distance in determining trade costs. We consider the exponential mapping:

$$a_{ij} = t_{ij}^{-\theta} = \frac{1}{1 + \exp\left(Z\beta\right)} \in [0, 1] \,.$$

The term a_{ij} is an element of the trade cost matrix A. This trade cost matrix is a trans-

[19] We are implicitly shutting down the substitution between containerized and non-containerized trade. This is supported by out bi-modal distribution of goods. Furthermore, while in the past, some ports did not handle containerized trade, by 2014, nearly all ports had some ability to handle containerized trade, relegating break-bulk trade to a small portion of all trade volumes.

form of the raw underlying direct shipping costs t_{ij} between country i and j and the dispersion of trade cost shocks θ. The matrix Z is a vector of covariates, including both traditional gravity variables, as well as variables that can effectively shift shipping costs. The vector β denotes the importance and direction of these covariates on shipping costs. This functional form is useful as it maps from the real numbers to the unit interval, as is required by our theoretical foundation. However, it is also parsimonious as it considers a linear projection.

We parsimoniously consider the following relationship for international trade:

$$Z\beta = \beta_0 + \beta_1 \log \text{sea distance}_{ij} + \beta_2 \log \text{traffic}_{ij} + \beta_3 \log \text{traffic}_{i\cdot} + \beta_4 \log \text{traffic}_{\cdot j} + \beta_5 \mathbb{I}_{backhaul} + \beta_6 \mathbb{I}\{i, j \in \text{Land}$$

Coefficient β_0 is an intercept, β_1 considers the role of sea distance between the nearest principal port.[20] Coefficient β_2 considers the role of port-to-port traffic on shipping costs; positive estimates indicate economies of scale, negative estimates indicate congestion effects. Work in highways, like Allen and Arkolakis (2019), find significant negative economies of scale, highlighting congestion effects.[21] Coefficient $\beta_3$and β_4 considers the total incoming and outgoing traffic at ports i and j respectively, directly considering the economies and diseconomies of scale in loading and unloading. Coefficient β_5 considers the role of the backhaul problem from Wong 2017, where ship capacity is fixed by the shipping direction with the highest demand. Shipping costs are lower on the so called "return route." The indicator variable simply consider if the traffic from i to j is more or less than the traffic from j to i. Finally β_6 consider an indicator variable for two countries that share a land border.[22]

Estimation seeks to find the vector β that match observed container traffic with ex-

[20]This variable is largely constant within country, with the exception of a small set of countries. For example, in the United States, we use Los Angeles for Asia, New York for Europe and Africa, and Houston for parts of Latin American and the Caribbean. Results are robust to considering other notions of distance.

[21]We can also consider congestion on a bilateral shipping route, however congestion may also occur at the port-level. In robustness exercises, we explore these issues.

[22]We do not estimate the diagonals a_{ii}, the cost of shipping from one's own port back to itself. We assume that these costs do not change in the counterfactual and we only estimate our data on international trade data, abstracting away from domestic trade. See Allen and Arkolakis (2019) for estimates of internal road transport trade costs.

Table 1: Trade Cost Estimates

Coefficient	Estimate	Standard Errors
β_0 (intercept)	5.36	.
β_1 (log distance)	0.10	.
β_2 (log route traffic)	-0.97	.
β_3 (log outgoing port traffic)	0.27	.
β_4 (log incoming port traffic)	0.38	.
β_5 (two-way trade)	[omitted for now]	.
β_6 (land borders)	-.4272	.

Notes: Results are highly preliminary and subject to change. Positive coefficients indicate higher trade costs.

pected container traffic. We denote this expected traffic,

$$\hat{\Xi}\left(A\left(\beta\right);X\right),$$

which is constructed from estimates of β as well as trade data X.

We then find estimates of parameters that minimize the differences of expected traffic and observed traffic for countries that do not share a land border:

$$\arg_{\beta}\min\sum_{ij\neq\text{land borders}}\left|\Xi_{ij}-\Xi_{ij}^{expected}\left(A\left(\beta\right),B\left(A\left(\beta\right)\right),X\right)\right|.$$

As noted before, sharing a land border indicates that we may not fully observe traffic of containerizable goods on that given leg. We thus cannot match traffic data from these countries with our expected traffic. In a forthcoming appendix, we explore an alternative estimate technique that relies on a set of alternative moments.

Table 1 shows the results of our estimation. Positive values for β indicates increases in trade costs and negative values indicate decreases in trade cost. We find that distance increases trade cost and increased shipping traffic decreases trade costs, after fully accounting for the role of total trade volumes in X.

Notably, we find container trade exhibits significant scale effects on the route-level. This result is very different from Allen and Arkolakis (2019), who find significant congestion costs in domestic road trade. This likely reflects the true underlying nature of the shipping industry versus the trucking industry. While shipping lanes can get crowded or "full", they do not experience anywhere near the same level of congestion as highway shipping. Furthermore, while it is difficult for one truck driver to control more than one or two container loads, the largest sea vessels can have capacities of 25,000 containers, without significant increases in operating costs. On the other hand, port level trade exhibit significant congestion costs, there are costs to having to many incoming and outgoing ships.[23]

Figure 9 plots our resulting matrix of pairwise trade costs. Routes drawn in thicker and ligher colors are lower-cost routes. Shorter and more heavily trafficed routes are the cheapest. The effect of scale is observable here: Syria to France is one of the highest cost legs, significantly higher than Singapore to Gibraltar, a much longer distance. Even among the subset of bilateral pairs for which we observe traffic, the triangle inequality is violated 280 times.

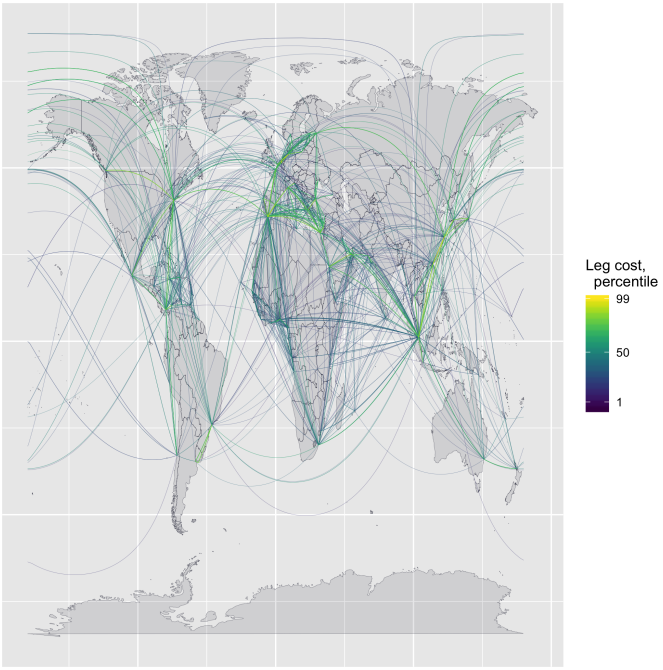
Figure 10 plots bilateral incoming and outgoing trade costs for Singapore and Lebanon separately. Singapore is not only well-connected but both as an origin and destination has some of the cheapest legs. Singapore ships to Lebanon which has both fewer and shorter connections.

Figure 11 plots country-level averages of the expected trade cost (from the B-matrix). Entrepôts such as Egpyt, Panama, and (not visible) Singapore and Gibraltar have generally cheaper trade costs, as does China, due to the scale of shipping as well as acecss to nearby low-cost entrepôt (Korea, Singapore, and Japan).

One way to check our estimation is to take our model and compare our expected traffic flows to actual traffic flows. We do this in Figure 12. In this figure, we compare actual global observed container traffic shares with the container traffic predicted with our β variables. We include both a best fit line and a 45 degree line. If our model perfectly

[23]Current analysis follows Allen and Arkolakis (2019) and abstracts away from endogeneity and model mis-specification concerns. Forthcoming analysis adopts an instrumental variables approach, instrumenting using expected trade conditional of 1960 economic fundamentals. Results are quantitatively similar.

Figure 9: Trade Cost Estimates, All Legs



Notes: Figures are preliminary and subject to change.

fit the world, all observations would lie on the 45 degree line. In general, we fit the data extremely well, with an correlation between the expected and expected shares (in logs) of .97.

With our estimates of trade costs of international containerized shipping in hand, we now turn to a full model of world trade, that incorporates a containerized shipping sector, with non containerized traded goods (both physical and services), as well as a non-traded local product. This model will allow for both estimating the role of scale and entrepots on local and global welfare. It will also allow the gauge the benefit from scale technology to rationalize investments in transportation infrastructure.

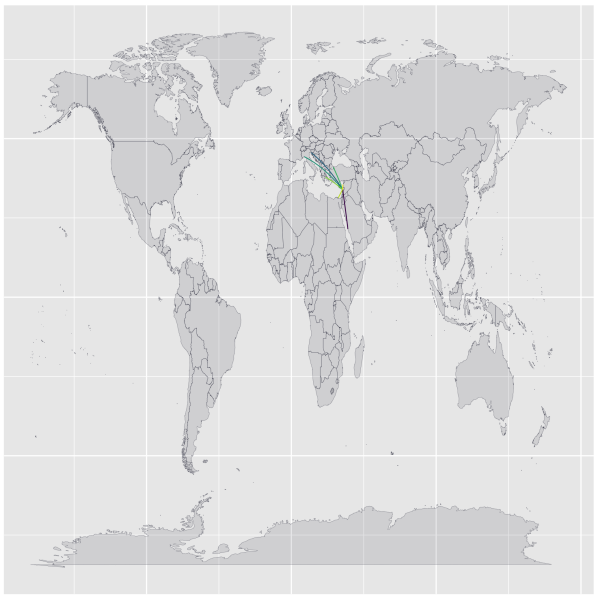
Figure 10: Trade Costs by Country



(a): Singapore, Origin



(b): Singapore, Destination



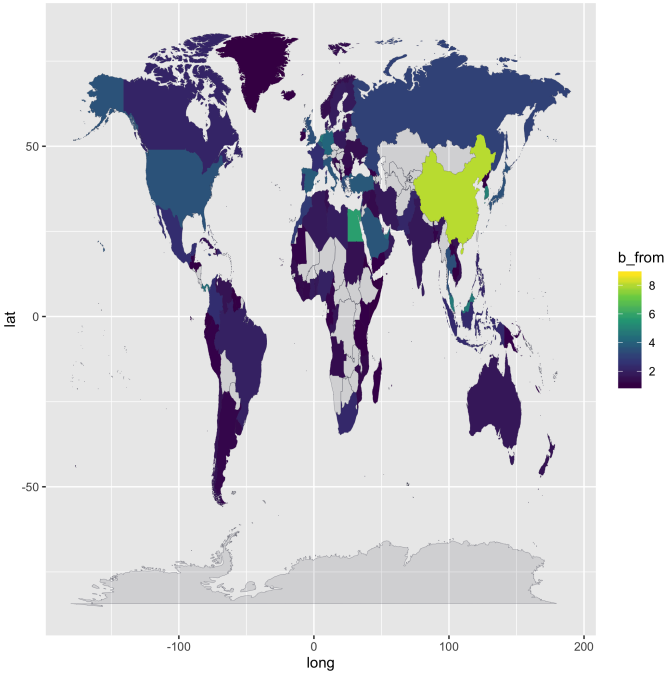
(a): Lebanon, Origin



(b): Lebanon, Destination

Notes: Lighter values indicate lower trade costs. Figures are preliminary.

Figure 11: Expected Trade Costs, Country Average



Notes: Lighter values indicate lower expected trade costs. Figures are preliminary.

6. Counterfactuals

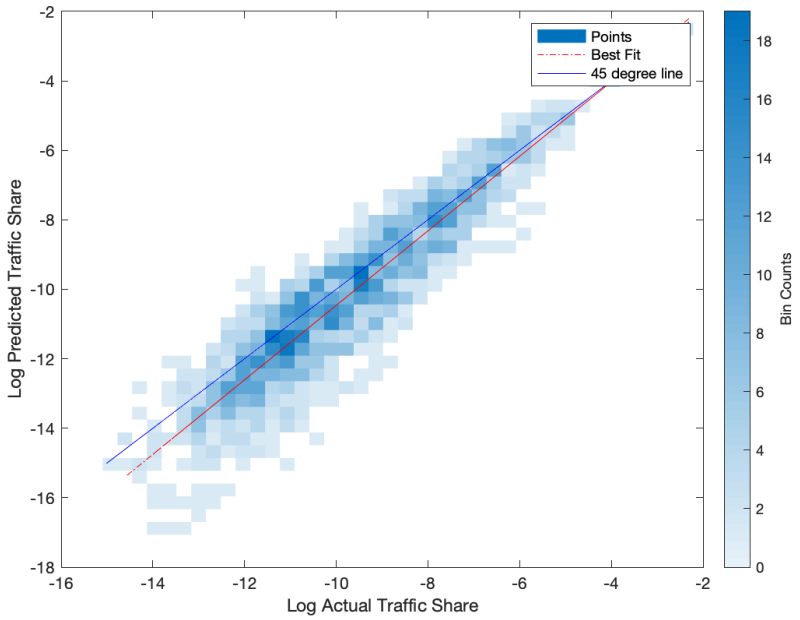
6.1. Quantitative General Equilibrium Model

We now introduce structural assumptions into our general framework to deliver a quantifiable general equilibrium in order to estimate counterfactuals. Our quantifiable model of entrepôt uses the Caliendo and Parro (2014) framework of trade. We assume there are three sectors: one nontradable, one tradable and containerized and one tradable and non-containerized, with a continuum of goods in each sector and trade in intermediates as well as final goods.

6.1.1. Setup

As in Caliendo and Parro (2014), intermediate goods ω_n are used in the production of composite goods that are in turn used domestically both as final goods and as materials

Figure 12: Model Fit Comparisons



Notes: Figures are preliminary and subject to change.

for intermediate production by firms in all industries. We assume there are three sectors $n \in [c, nc, nt]$. Intermediates in the nt sector are only sourced domestically while ω_{nc} and ω_c goods are sourced internationally. Trade routes are modeled for all three sectors but will only be estimated for intermediates ω_c.

Unlike Eaton and Kortum (2002) or Caliendo and Parro (2014), the technologies available to firms in each sector and country is non-stochastic. Instead, firms must additionally solve an optimal route choice problem, where idiosyncratic variation in a products' optimal route generates random variation in the price of each intermediate product-origin pair.

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

$$U_i = \prod_n m_{in}^{\eta_n}; \text{where } \sum \eta_n = 1$$

where ρ_n is the industry-specific elasticities of substitution and η_n is the Cobb-Douglas industry share, $\sum_n \eta_n = 1$.Because each product ω in sector n is homogeneous, consumers in country j will choose the lowest-cost provider of each such product.

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

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

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

The marginal cost of production for firms is

$$c_{in} \equiv \frac{\Upsilon_{in} w_i^{\gamma_{in}} \prod_{n'}^{N} P_{in'}^{\gamma_{in}^{n'}}}{z_{in}} \tag{11}$$

where $P_{in'}$ is the price of a composite good from sector n' and constant $\Upsilon_{in} = \prod_{n'}^{N} \left(\gamma_{in}^{n'}\right)^{\gamma_{in}^{n'}} \left(\gamma_{in}\right)^{\gamma_{in}}$.

Composite goods production In each country i, composite goods in industry n are produced using a CES aggregate of intermediates Ω_n, purchased and sold domestically at marginal cost. In traded industries, intermediates are sourced internationally from lowest-cost suppliers. For a given product, the price available to composite producers in j is

$$p_{jn}(\omega) = min_{i,r} \left\{c_i \kappa_{ijn} \tilde{\tau}_{ijn} / \epsilon_{ijnr}(\omega)\right\} \tag{12}$$

where $\tilde{\tau}_{ijn}$ is the expected trade cost from i to j in industry n and $\epsilon_{ijnr}(\omega)$ is the Frechét

draw for a given route.

When compared to the relevent equation in Caliendo and Parro (2014), it's clear that idiosyncratic route draws are generating the stochastic price dispersion usually assumed to be idiosyncratic TFP. Using the standard aggregation, the resulting price at j of the composite in industry n is expected to be

$$P_{jn} = A_n \left[\sum_{i=1}^{I} c_i^{-\theta_n} \kappa_{ijn}^{-\theta_n} \tilde{\tau}_{ijn}^{-\theta_n} \right] \tag{13}$$

where A_n is a constant.

6.1.2. Equilibrium in changes

The general equilibrium can be defined in changes using four conditions in hat algebra which dicatate how equilibrium-set parameters adjust to a change in trade costs $\hat{t}_{kl} = t'_{kl}/t_{kl}$ through a change in expected trade costs $\hat{\tau}_{ijn} = \tau'_{ijn}/\tau_{ijn}$

First, production costs at i in industry j respond to a shock to a given t_{kl} according to the equation:

$$\hat{c}_{in} = \hat{w}_i^{\gamma_{in}} \prod_{k=1}^{N} \hat{P}_{ik}^{\gamma_{ink}} \tag{14}$$

Next, a change in price at i in industry n relative to shock to t_{kl} is:

$$\hat{P}_{in} = \left[\sum_{i=1}^{J} \pi_{ijn} [\hat{\tau}_{ijn} \hat{c}_{in}]^{-\theta_n} \right]^{-1/\theta_n} \tag{15}$$

Accordingly, trade volumes between i and j in industry n accoding to standard changes through production and transport costs:

$$\hat{\pi}_{ijn} = \left[\frac{\hat{c}_{in} \hat{\tau}_{ijn}}{\hat{P}_{in}} \right]^{-\theta_n} \tag{16}$$

Trade volumes similarly adjust

$$X'_{in} = \sum_{k=1}^{N} \gamma_{ink} \sum_{j=1}^{I} \frac{\pi'_{ijn}}{1+\kappa_{ijn}} X'_{jk} + \alpha_{in} I'_i \tag{17}$$

Trade is balanced to a deficit shifter such that

$$\sum_{n=1}^{N} \sum_{i=1}^{I} \frac{\pi'_{ijn}}{1+\kappa_{ijn}} X_{in} - D_i = \sum_{n=1}^{N} \sum_{i=1}^{I} \frac{\pi'_{jin}}{1+\kappa_{jin}} X_{jn} \tag{18}$$

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

Counterfactuals will adjust t_{kl} for some set of k, l pairs. In our baseline counterfactuals, we adjust some leg costs while leaving others constant. This assumes no scale effects of trade. In subsequent work, we will shock leg pairs and jointly solve for new trade volumes, new traffic, new endogenous trade costs.

6.2. Counterfactual Estimation

As a preliminary exercise, we run four counterfactuals to understand the effects of changing trade cost on particular shipping legs, but leaving the rest of the network unchanged.[24] In particular we take the underlying trade cost matrix mapping, A, and alter it such that trade costs decrease by 10% for all incoming and outgoing containerized traffic from a particular country. This reflects investments in port infrastructure, helping alleviate port-level congestion effects. Two counterfactuals look at the role played by larger countries - China and the United Kingdom. Two more counterfactual look at the role played by a smaller entrepts, Singapore and Belgium. These paired counterfactuals show how similar changes in transport costs on different legs can have very different and far reaching consequences.

[24] In this preliminary exercise, we do not change scale economies, allowing for us to retain the use of our current equilibrium. Furthermore, we are currently only considering two sectors, a tradable containerized sector and a non-tradable sector. Future iterations will allow for both scale economies and a tradable non-containerized sector.

Table 2: UK vs Belgium: Trade Cost Reduction Effects - Containerized Trade

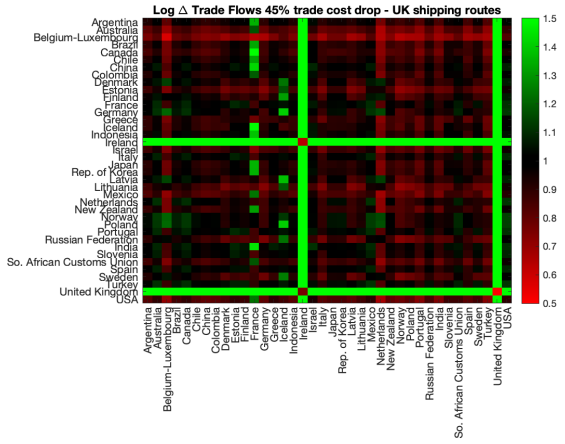
	Change in UK direct trade costs by 10%		Change in Belgian direct trade costs by 10%	
Rank	Trade Increase	Exporter	Trade Increase	Exporter
1	16.72%	United Kingdom	32.47%	Belgium-Luxembourg
2	16.44%	Ireland	27.86%	Netherlands
3	13.54%	Iceland	22.35%	Germany
4	11.08%	Netherlands	17.04%	United Kingdom
5	10.67%	Germany	16.25%	Norway
6	10.57%	Belgium-Luxembourg	15.21%	Denmark
7	7.69%	France	14.04%	France
	1.02%	World Average (Weighed)	1.51%	World Average (Weighed)

6.3. Europe: the UK vs Belgium

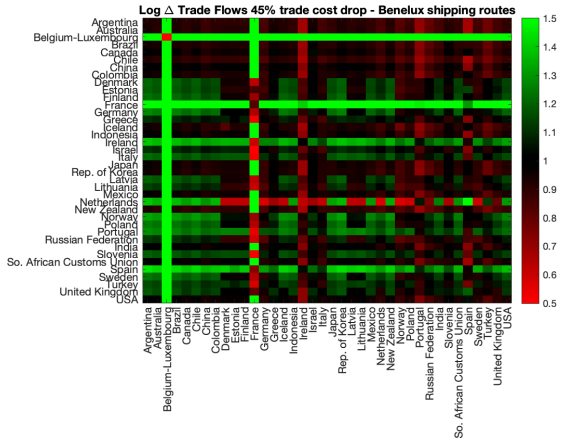
In our first counterfactual, we simulate what occurs if the shipping cost of all direct routes from Felixstowe and other UK ports decrease by 10%. This simulates what would happen in a counterfactual world, where Britain actually invests in infrastructure and becomes more open. Table 2 shows the countries whose exports increase the most when container trade costs into and out of the UK decrease by 10%. Figure 13, Panel (a) shows what happens to trade flows between nations if all shipping if all legs that either originate or end in the UK see a 10% cost decrease. Rows represent exports and columns represent imports. In general, we see that trade to and from the UK increases. However, we also see a massive spillover effect for one country, Ireland. Most Irish trade transits the UK at some point. Ireland not only get the general equilibrium effect of buying cheaper British goods, they also gain a market access effect. This is due to the Irish now having lowered trade costs with the UK, and through the UK, lowered trade costs with every other country in the world. However, there is a second effect here. In red are trade routes that shrink. These routes now see a trade diversion effect. While the UK in the baseline had substantial trade, it was not an entrepot hub. When the UK sees this

Figure 13: Shipping Volume Consequences

(a): Trade Cost Drop: UK



(a): Trade Cost Drop: Belgium



massive decrease in trade costs, more intra-European trade now is goes to and from the UK, cutting out other producers of goods.

Our second counterfactual looks at the role played by Belgium[25], which is substantially smaller than the United Kingdom, but highly connected with the rest of the world and home of Europe's second largest container port. This effect of reducing trade costs with Belgian ports is quite different from the United Kingdom. Belgium does not produce much, but it has massive traffic flowing though its transportation network. Its principal port, Antwerp, is the second largest in all of Europe and connects many European nations with the rest of the world. Decreases in trade cost. Figure 13, Panel (b) shows what occurs when trade costs with all direct routes to Belgian ports decreases by 10%. While the UK case caused mostly trade diversion to the UK, to the detriment of most of Europe, the effect is much more nuanced in Belgium. While Belgian trade does increase, most European bi-lateral pairs also see trade increases. Belgium now better connects the rest of Europe, enhancing trade. Furthermore, even though the Belgian economy is smaller than the British economy, and exports less, there are much larger aggregate trade effects from lowering Belgian versus UK trade costs.

[25]Belgium and Luxembourg are merged in our underlying data. We leave this as is.

Table 3: China vs Singapore: Trade Cost Reduction Effects - Containerized Trade

	Change in China direct trade costs by 10%		Change in Singapore direct trade costs by 10%	
Rank	Trade Increase	Exporter	Trade Increase	Exporter
1	14.95%	Dem. Peoples Rep. of Korea	8.92%	Singapore
2	14.24%	Singapore	7.97%	Malaysia
3	13.39%	Rep. of Korea	7.01%	Brunei Darussalam
4	13.00%	China	6.83%	Bangladesh
5	11.76%	Thailand	6.54%	Congo
6	11.20%	Congo	6.38%	Indonesia
7	10.99%	Cambodia	6.27%	Madagascar
	3.69%	World Average (Weighed)	1.98%	World Average (Weighed)

6.4. Asian Entrepots: China versus Singapore

We perform a similar exercise in Asia, but consider two polar opposites; China and Singapore. China exports 2.2 trillion in exports, while Singapore exports a mere 372 billion. In a conventional trade model, reducing China's export cost would have an order of magnitude larger effect on world trade flows than a smaller country. However Singapore is a highly connected nation, connecting South Asian nations with the rest of the world. See Table 3 for summary statistics. If we reduce the costs of direct containerized trade with China, world trade increases 3.70%. Similarly, if we reduce the costs of direct containerized trade with Singapore, world trade increases by 1.98%, even though Singapore is an order of magnitude smaller than China. Essentially, reducing trade costs with China simply helps China's export and import destinations. However the size of China means that these effects in aggregate are large. Reducing trade costs with Singapore has an additional effect of helping countries connect with each other, multiplying the effect of a small country, like Singapore. Even though its gross exports are an order of magnitude smaller than China, the effect of changing Singaporean trade costs are similar to that of changing Chinese trade costs.

References

Ahn, JaeBin, Amit K Khandelwal, and Shang-Jin Wei. 2011. "The role of intermediaries in facilitating trade." *Journal of International Economics*, 84(1): 73–85.

Allen, Treb, and Costas Arkolakis. 2019. "The welfare effects of transportation infrastructure improvements." National Bureau of Economic Research.

Anderson, James E, and Eric Van Wincoop. 2004. "Trade costs." *Journal of Economic literature*, 42(3): 691–751.

Anderson, James E, Mykyta Vesselovsky, and Yoto V Yotov. 2016. "Gravity with scale effects." *Journal of International Economics*, 100: 174–193.

Antras, Pol, and Arnaud Costinot. 2011. "Intermediated trade." *The Quarterly Journal of Economics*, 126(3): 1319–1374.

Asturias, Jose. 2013. "Endogenous Transportation Costs." Discussion paper, University of Minnesota.

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.

Blondel, Vincent D, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. "Fast unfolding of communities in large networks." *Journal of statistical mechanics: theory and experiment*, 2008(10): P10008.

Blonigen, Bruce A, and Wesley W Wilson. 2008. "Port efficiency and trade flows." *Review of International Economics*, 16(1): 21–36.

Blum, Bernardo S, Sebastian Claro, and Ignatius Horstmann. 2010. "Facts and figures on intermediated trade." *American Economic Review*, 100(2): 419–23.

Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papageorgiou. 2017. "Geography, Search Frictions and Endogenous Trade Costs." National Bureau of Economic Research.

Caliendo, Lorenzo, and Fernando Parro. 2014. "Estimates of the Trade and Welfare Effects of NAFTA." *The Review of Economic Studies*, rdu035.

Coşar, A Kerem, and Banu Demir. 2018. "Shipping inside the box: Containerization and trade." *Journal of International Economics*, 114: 331–345.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, geography, and trade." *Econometrica*, 70(5): 1741–1779.

Fajgelbaum, Pablo D, and Edouard Schaal. 2017. "Optimal transport networks in spatial equilibrium." National Bureau of Economic Research.

Francois, Joseph F, and Ian Wooton. 2001. "Trade in international transport services: the role of competition." *Review of International Economics*, 9(2): 249–261.

Ganapati, Sharat. 2018. "The Modern Wholesaler: Global Sourcing, Domestic Distribution, and Scale Economies."

Head, Keith, and Thierry Mayer. 2014. "Gravity Equations: Workhorse, Toolkit, and Cookbook." *Handbook of International Economics*, 4: 131.

Hendricks, Ken, Michele Piccione, and Guofu Tan. 1995. "The economics of hubs: The case of monopoly." *The Review of Economic Studies*, 62(1): 83–99.

Hendricks, Ken, Michele Piccione, and Guofu Tan. 1997. "Entry and exit in hub-spoke networks." *The RAND Journal of Economics*, 291–303.

Hendricks, Ken, Michele Piccione, and Guofu Tan. 1999. "Equilibria in networks." *Econometrica*, 67(6): 1407–1434.

Holmes, Thomas J, and Ethan Singer. 2018. "Indivisibilities in Distribution." National Bureau of Economic Research.

Hummels, David. 2007. "Transportation costs and international trade in the second era of globalization." *The Journal of Economic Perspectives*, 21(3): 131–154.

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.

Krugman, Paul. 1993. "The hub effect: or, threeness in interregional trade." *Theory, policy and dynamics in international trade*, 29–37.

Lazarou, Nicholas-Joseph. 2016. "Endogenous Hub Formations in International Trade." Forum for Research in Empirical International Trade Working Paper 1144.

Limao, Nuno, and Anthony J Venables. 2001. "Infrastructure, geographical disadvantage, transport costs, and trade." *The World Bank Economic Review*, 15(3): 451–479.

Rauch, James E. 1999. "Networks versus markets in international trade." *Journal of international Economics*, 48(1): 7–35.

Rauch, James E. 2001. "Business and social networks in international trade." *Journal of economic literature*, 39(4): 1177–1203.

Rauch, James E, and Vitor Trindade. 2002. "Ethnic Chinese networks in international trade." *Review of Economics and Statistics*, 84(1): 116–130.

Rua, Gisela. 2014. "Diffusion of containerization." *Finance and Economics Discussion Series Staff working paper 2014–88. Federal Reserve Board, Washington, DC.*

Wong, Woan Foong. 2019. "The Round Trip Effect: Endogenous Transport Costs and International Trade."

Appendix A Additional Figures and Tables

In figure A.1 below, we use the Louvain method created by Blondel et al. (2008) to identify the communities—groups of nodes that are more densely connected to one another than to other nodes—within the shipping network.

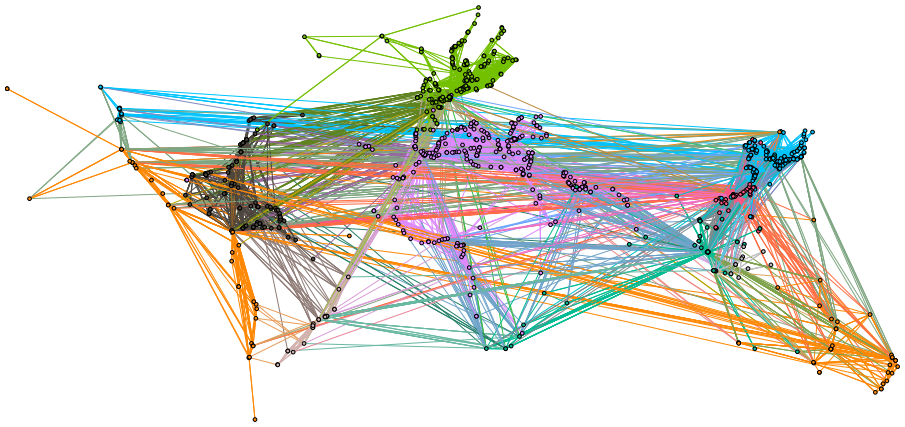


Figure A.1: Map of Global Port of Call Network

The dots are Source: Authors' calculations.

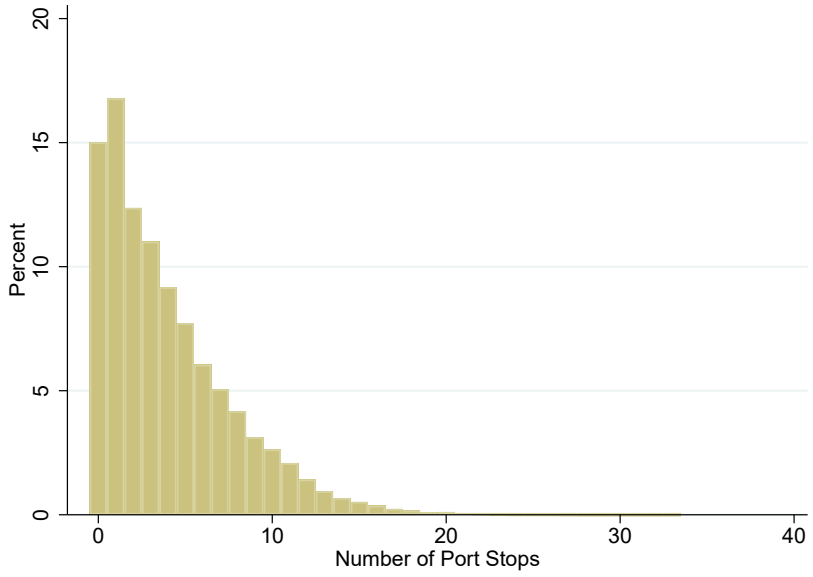
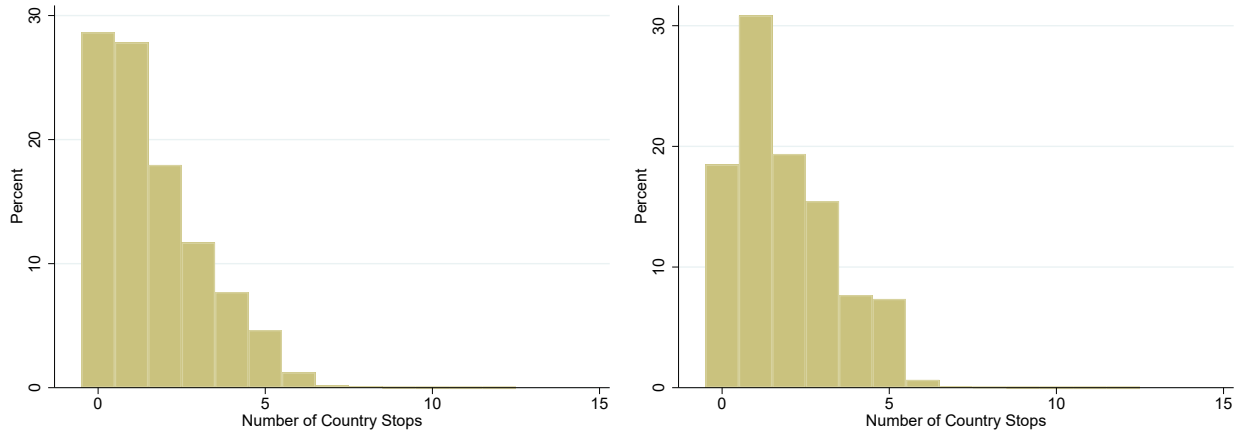


Figure A.2: Distribution of Port Stops per Container (TEU)

Panel (A) and (B) of Appendix Figure A.3 below show the histogram of number of stops per shipment using shipment weight and value as weights. Less than 30 percent by weight measured in tons and less than 20 percent by value are transported between an exporting origin and the US directly. The remaining trade makes at least one stop along the way, with the average shipment stopping at around 2 additional countries that are not involved in the bilateral trade (1.6 by weight and 1.9 by value).



Panel (A) Number of Countries per Ton

Panel (B) Number of Countries per USD Value

Figure A.3: Distribution of Third-Party Countries Involved in Bilateral Trade by Weight and Value