Defying Gravity: The Substitutability of Transportation in International Trade

Lux, Matthias

New York University

May 2011

Defying Gravity: The Substitutability of Transportation in International Trade*

Matthias Lux†

Cornerstone Research

February 2012

Abstract

What role do individual modes of transportation play in international trade? To study this question, I develop a model of international trade that incorporates a role for transportation and thus allows me to study mode-specific trade flows. I use a novel data set to estimate the complete model for a sample of 79 countries distinguishing air, sea, and surface transportation. The estimated model implies that surface transportation is mostly used for trade over short distances, whereas air and sea transportation dominate long-distance trade. Furthermore, the different modes of transportation display a high degree of substitutability. Using counterfactual analysis I show the implications for the roles played by the different modes of transportation. Long-distance modes are more important for poor countries because in order for them to realize gains from trade they need access to technologically advanced but far-away markets. Rich countries, on the other hand, can substitute long-distance trade more easily for trade with neighboring countries without changing the gains from trade much. As a consequence, reducing the estimated asymmetries in mode-specific trade costs for only one long-distance mode, either air or sea, can reduce income differences in the sample by about 35%.

*I am grateful to Jonathan Eaton, Virgiliu Midrigan, and Michael Waugh for their continued advice and encouragement. I would also like to thank Gaurab Aryal, Andrei Levchenko, Cecilia Parlatore Siritto, Andres Rodriguez–Clare, Felix Tintelnot, and Mu–Jeung Yang, as well as audiences at the PSU Macro/International workshop, the NYU Macro Student Lunch, and the NYU International/Development Seminar for helpful discussions and comments, and James Feyrer for generously providing me with his data on bilateral nautical distances. Large parts of the project were completed while I was visiting the economics department at PSU. I gratefully acknowledge the department's hospitality.

†Cornerstone Research, Two Embarcadero Center, 20th Floor, San Francisco, CA 94111, mlux@cornerstone.com. The views expressed are solely those of the author and do not necessarily reflect the views of Cornerstone Research.

1. Introduction

For many countries, transportation related charges have become a larger obstacle to realizing further gains from trade than tariffs. Figure 1 shows that almost all countries importing into the US face higher transportation charges than tariffs. The average ratio of transportation charges to tariffs in the sample is 3.2. As a consequence of this, World Bank (2009) argues for investments in transportation infrastructure to lower these transportation mode-specific trade barriers. In 2010, the World Bank invested $ 9 billion dollars or 15% of its total lending into such infrastructure projects.

But quantitative models of international trade are largely silent on the interaction of trade and transportation. These models generally assume that there is just an aggregate trade cost to be paid if two countries engage in trade.[1] Therefore, these models cannot be used to study the effects of changes in transportation mode-specific trade costs on trade and, ultimately, welfare. In particular, they cannot be used to understand the returns to the infrastructure investments meant to decrease transportation mode-specific trade barriers.

In this paper, I develop a model of international trade that incorporates a role for the mode of transportation and thus allows to address the implications of change in mode-specific trade costs. I then use this model to study the role different modes of transportation play in international trade. The framework rests on the static multi-country model of Ricardian comparative advantage developed in Eaton and Kortum (2002). There are two sectors: a final good sector that produces a non-traded consumption good and an intermediate good sector producing a continuum of tradeable varieties used in the production of the final good. Each country can produce each intermediate variety choosing from a menu of available production technologies. Each productivity on the menu corresponds to a mode of transportation with which the good can be exported. Trading goods between countries is subject to iceberg trade costs where these trade costs can differ across the different modes of transportation. As in standard Ricardian models an importer chooses the minimal price when deciding from where to source a good. However, in my model an importer can choose both the source country and the mode of transportation when choosing the minimal price, whereas standard models only allow a choice of the exporting country. Allowing for the choice of the mode of transportation allows me to study mode-specific trade flows and assess the implications of changing mode-specific trade costs.

The modelling approach taken allows me to incorporate an arbitrary number of transportation modes, whereas the previous literature has only considered the case of

[1] See, for example, Anderson and van Wincoop (2003), Eaton and Kortum (2002), and Chaney (2008).

two modes (cf. Harrigan (2009)). This generalization turns out to be important because the data set I use to estimate the complete model suggests to distinguish at least three different modes. This data set reports mode-specific bilateral manufacturing trade flows for a sample of 79 countries; the sample year is 2005. The three different modes of transportation I distinguish are air, vessel, and surface transportation. Relative to models of aggregate trade flows the estimation of the model is complicated due to the fact that the gravity equation in the model is non-linear in the mode-specific trade costs. As a consequence, I jointly estimate the model's system of equations using non-linear least squares. In the estimation, the mode-specific trade cost functions are allowed to differ along two dimensions. On the one hand the coefficients on geographic controls like distance can differ: for example, the distance elasticities of air and surface transportation are not restricted to be identical. On the other hand, the geographic controls entering the trade cost functions can differ: the distance travelled by a ship transporting goods from, say, Germany to Italy is much larger than the distance flown by an airplane between the same countries. In addition, I allow for a mode- and country-specific exporter fixed effect similar to Waugh (2010). The differences in the trade cost functions across the different modes of transportation are crucial in matching mode-specific trade flows. It turns out that the distance elasticity of surface transportation is by far the largest, and that the cost of air transportation barely rises in distance. Furthermore, the large contiguity effect usually found in the literature is mostly caused by surface transportation; being contiguous to a country has a much smaller effect on vessel transportation and barely any on air transportation.

To further highlight the role played by different modes of transportation I calculate the share of gains from trade attributable to each individual mode. That is, I calculate the loss in gains from trade if a given mode was not available. The average welfare loss for shutting off trade by sea is with 2.9% largest, followed by surface transportation with an average loss of 2.4%. Air transportation is the least important one with an average loss of 0.8%. All these losses, however, are small compared to the average loss of 11.1% entailed in a move to autarky. The reason is that modes are strongly substitutable for each other. Another interesting feature is that this counterfactual highlights the implications of the geographic distribution of technology levels: on the globe, high technology countries are mostly clustered together and so are low technology countries. As a consequence, access to long-distance modes of transportation like trading by sea is more important for poor countries than for rich countries. On the other hand, access to surface transportation is more important for rich countries than for poor countries.

I then use the model to conduct another counterfactual, aimed at investigating the role of transportation in reducing income differences. Recently, Waugh (2010) has shown that the systematic asymmetries in aggregate trade costs explain up to 30% of income differences in the country sample he uses. I use the model to investigate the extent to which the asymmetries in mode-specific trade costs can reduce income differences. The results show that reducing the asymmetries in air or sea transportation alone can reduce income differences by about 35%. To put this number into perspective, the income differences are reduced by 60% when moving to free trade. The reason for this strong role played by a single mode of transportation in realizing gains from trade is, again, the substitutability among the modes estimated in the model.

The arguments in this paper contribute to the large literature that tries to determine the many different sources through which gains from trade arise. Most closely related are Fogel (1964) and Donaldson (2008) in that they also evaluate the gains arising from different means of transportation. Fogel (1964) investigates the effect of the railroads connecting the US east and west coast that were built in the 19th century and concludes that they did not lead to a significant increase in trade flows. Instead, they mostly led to substitutions away from the system of inland waterways used before to the newly built railroads. Donaldson (2008) investigates the effect of the railroads built by the British in 19th century India. He concludes that the railroads led to a considerable welfare gain and that about 90% of these gains occurred as gains from trade. As in these papers, I concentrate on the role played by different modes of transportation in trade. The difference is that I do not concentrate on a particular infrastructure project but provide a framework in which one can discuss the effects of different infrastructure projects more generally.

By estimating mode-specific trade cost functions the paper also contributes to the large literature that studies the determinants of international trade costs. Anderson and van Wincoop (2004) provide an excellent recent survey of this literature. Most of the literature studies aggregate trade costs between countries and evaluates their determinants. One of the exceptions that study mode-specific trade costs is Hummels (2001) who develops an empirical discrete choice model of mode-specific trade flows to determine the effect of delivery time on trade costs. He estimates his model using US import data and finds that each day saved in shipping time is worth about 0.8% of the value of the shipment. However, his approach does not allow one to estimate complete bilateral trade cost functions. Furthermore, since he does not specify a full general equilibrium model it is impossible to judge the contributions of individual modes to the overall gains from trade. Harrigan (2009) develops a complete model of mode-specific trade flows. His main concern is with the degree to which faster transportation can act as a source of comparative advantage. He derives a set of implications from the model and tests them using US import data. However, as mentioned above his modelling approach does not generalize to more than two modes. But the quantitative importance of multiple modes is an important feature of the data set I use to estimate my model.

The aim in modelling mode-specific trade flows in an explicit general equilibrium setting is to establish a link between mode-specific trade costs and welfare. Such a model closes a gap in two other strands of the literature. First, there are many studies that investigate the link between trade costs and transportation infrastructure. They generally find that improving transportation infrastructure lowers mode-specific trade costs. For example, Clark, Dollar, and Micco (2004) use micro data from the U.S. Import Waterborne Databank to investigate the determinants of maritime transport costs. They conclude that improving port efficiency from the 25th to the 75th percentile of their efficiency index decreases shipping costs by 12%. Another example is Limao and Venables (2001). They investigate the impact of infrastructure on trade costs and infer that dropping from the median to the 75th percentile on the distribution of infrastructure quality raises transport costs by 12%. But since these studies do not specify a general equilibrium model of mode-specific trade flows, they are unable to link the reduction in mode-specific trade costs to welfare, which is the natural measure for judging investment projects, for example. The model presented in this paper fills this gap by providing just such a link. The only other model I am aware of that links infrastructure improvements to welfare gains is Donaldson (2008), already mentioned above. However, one property of his model is that all trade between two regions uses the exact same mode of transportation. That is to say, there is exactly one cost-minimizing choice between to regions. But looking at modern trade flows between two countries, it is evident that most country pairs employ a mixture of different modes of transportation. The model developed in this paper explicitly incorporates such mode-specific mixtures.

The other part of the literature where a link between mode-specific trade costs and welfare is of interest is a small literature that tries to determine the effects of stronger competition among international carriers on international trade costs. Hummels, Lugovskyy, and Skiba (2009) investigate the role of shipping cartels in inhibiting international trade. They find that the market power exerted by shippers explains a large part of the variation in trade costs. A back-of-the-envelope calculation shows that reducing the market power would boost trade volumes by 6% to 15%. Micco and Serebrisky (2006) investigate the role of increased competition in international air transportation and the relation to air shipping costs. They investigate open-skies agreements (OSAs) that liberalize air transportation markets and conclude that these agreements reduce air shipping costs by about 9% for developed countries. They also find that for less developed countries these OSAs do not have a discernible cost effect. From a normative perspective, the ultimate object of interest in waging whether to liberalize shipping cartels or the aviation industry should be the change in a country's welfare. With the help of the model developed in this paper, it is possible to estimate these welfare gains and determine the desirability of deregulation. Furthermore, the model offers a potential explanation for the different effects of OSAs on developed and developing countries.

2. The Empirics of Mode-Specific Trade Flows

In this section I first introduce the data used in the paper. To the best of my knowledge, the enlarged data set on mode-specific trade flows has not been used before in the literature. I then highlight four properties of the data that will guide the development of the structural model and provide insights into the identification of some of the key parameters.

2.1. The Data

To study mode-specific trade flows in a multi-country setting I have to go beyond the data set usually used for studying mode-specific questions in international trade. Almost all papers investigating these solely rely on the US Imports and Exports of Merchandise.[2] But as a consequence, the US is always on one side of the observed trade flows, either as an importer or as an exporter. This does not allow the identification of country specific components for any other country but the US, which makes the complete specification of a general equilibrium model of international trade impossible. Therefore, I combine this data set with a novel data set from Eurostat, the European Union's statistical agency. The EU data set contains all external trade flows of the 27 EU members disaggregated by the mode of transportation. External trade flows are trade flows between EU countries and non-EU countries. The fact that only external flows are contained in this data set is due to the difference in customs requirements for recording intra- and extra-EU trade flows. The EU data set distinguishes nine different modes of transportation. However, to make the data set compatible with the US data, I only use three modes: air, sea, and surface. What I call surface is thus an aggregate of the remaining seven modes. These are trade by road, rail, inland waterways, fixed mechanism, postal, and unknown mode of transportation. Fixed mechanism transportation refers to goods that do not need external transportation, for example air planes, boats, and trucks. In 2005, the average bilateral share of rail and road based transportation in

[2] See, for example, Hummels (2001), Harrigan (2009), and Hummels and Schaur (2010). Clark, Dollar, and Micco (2004) rely on the US Waterborne Database which also only records trade with the US as one partner.

the category of surface based transportation was 80%.[3] Thus, this category can really be thought of as mostly reflecting trade by surface transportation.

The data on bilateral trade flows are collected from the statistical agencies of the 27 EU member countries. This gives the data set a particular structure, depicted in figure 2. EU member countries and the US take on both the role of reporters and partners, whereas all other countries in the sample are only partners. The distinction is the valuation of the reported trade flows: imports to reporting countries are reported including freight and insurance – what is commonly referred to as "cost, insurance, freight" or *c.i.f.* – whereas exports from reporting countries are registered excluding these additional charges. This is commonly referred to as "free alongside ship" or *f.a.s.* This difference is due to the particular nature of customs forms used in international trade: exporters are only required to report the value of the goods transported, whereas importers also record the cost for shipping and insurance.[4] The summary statistics presented in this section ignore this distinction, but the estimation procedure developed below will take this difference into account.

As explained above, the EU data set only reports trade with external partners. Therefore only trade flows between any two countries in different blocks in figure 2 are contained in the sample but not within any one block. For example, the sample contains the trade flows between Germany and Canada but not between Germany and Belgium. Because of the use of both nautical and great circle distances in the estimation procedure I also exclude all landlocked countries. This leaves 23 reporting countries, consisting of the US and 22 European Union members. The EU countries not represented in the sample are Austria, the Czech Republic, Hungary, and Slovakia. In addition, Belgium and Luxembourg have been combined into one country. There are 56 partner countries from all parts of the world in the rest of the sample. The sample year is 2005.

Since the model developed later is based on the Ricardian idea of comparative advantage, it is best thought of as describing trade in manufactured goods. Therefore I restrict attention to manufacturing trade flows in what follows.[5]

2.2. Four Properties of the Data

In this subsection I describe four properties of the data. The first two show that the choice of the transportation mode does not solely depend on country- or good-specific factors. The third fact demonstrates that it is in fact geography that interacts very differ-

[3] See Lux (2010) for a more detailed discussion of this data set.

[4] See Hummels and Lugovskyy (2006) for a related discussion of the difference between *c.i.f.* and *f.a.s.* flows in international trade.

[5] See the data appendix for a discussion of the concordance used.

ently with the different transportation modes. The fourth property documents changes in mode-specific transportation charges over time that will inform the inference on the substitutability of transportation modes.

Fact I: Countries Alone Do Not Determine the Transportation Choice

Table 1 shows the summary statistics for the bilateral mode-specific shares $g_{ni}^m = \frac{X_{ni}^m}{X_{ni}}$ in the data set. X_{ni}^m is the spending of country n on imports from country i that are transported via mode m and $X_{ni} = \sum_m X_{ni}^m$, where the summation is over the three modes air, vessel, and surface. Trade by vessel is the most important mode of transportation with an average share of just over 60%. Air transportation is the second most important mode with a share of just over 20%. Surface transportation has an average share of almost 19%. The variation of these bilateral mode-specific shares reported in the table, measured as the coefficient of variation, is large for all three modes. For air and surface transportation the coefficient of variation is 115.6% and 137.7%, respectively. Even for maritime trade the variation is about 50%. The first and third quartile of each share distribution are also reported and further corroborate the significant amount of variation in mode-specific shares.

Table 1 calculates the statistics across all bilateral pairs. To understand the role of different modes of transportation at the country-level, figures 3 to 5 plot the median of the export share g_{ni}^m per exporter and mode against the (log) GDP per capita. The figures show that most countries use all modes of transportation for their exports and confirm the impression from table 1. They also demonstrate the variation in the use of different modes across different countries. Air transportation plays a bigger role in exports of rich countries but maritime trade is more important for poor countries. Surface transportation is insignificantly correlated with an exporter's GDP per capita in the data set.

The summary statistics show that all three modes are actively used in international trade and that there is a lot of variation in bilateral shares. The exporter level analysis further details this variation. It shows that the variation of the summary statistics is not exclusively caused by variation across different countries but that there is an active use of multiple modes of transportation even at the exporter's level. This shows that the composition of mode-specific trade cannot be explained solely by country specific characteristics.

Fact II: Goods Alone Do Not Determine the Transportation Choice

The second stylized fact concerns the question whether the goods traded determine the choice of the mode of transportation or whether there is substitutability of modes for a given good. Since only the US data set contains information on good-specific trade flows, I restrict the analysis of this point to the US Exports of Merchandise. To investigate this point, I first calculate the mode-specific trade shares of US exports per commodity. I then compute the Herfindahl index of these mode-specific shares for every commodity across different modes, i.e.

$$HI(j) = \sum_m s(j)_m^2,$$

where

$$s(j)_m = \frac{x(j)_m}{x(j)}.$$

$x(j)_m$ is the value of US exports of good j that is transported by mode m and $x(j) = \sum_m x(j)_m$ is the total value of US exports of good j. A good is a HS10 category. Figure 6 plots the histogram of the Herfindahl indices over commodities. The lower bound of $1/3$ represents an equal distribution of mode-specific shares for the commodity, implying that the good is exported using all three modes equally. The upper bound of one signals that all trade is concentrated in one single mode. If goods were transported with only one mode of transportation, the histogram would show that the Herfindahl indices for all commodities would be concentrated at one. But the histogram shows that only just under 4% of all goods have an index of nearly one. The bulk of goods have an index below 0.6. A trade-weighted average of these Herfindahl indices gives a value of 0.57.

The histogram is only a count of commodities. To investigate the importance of the goods being exported with one predominant mode and their contribution to total export values more closely, table 3 reports the share in overall US Exports that falls upon goods that are exported with any mode-specific share – air, sea, or surface – above a certain threshold. For example, about 0.6% of all US exports in 2005 were exported with one mode-specific share above 99.5% and just under 5% of all exports had one mode-specific share above 95%. Thus the vast majority of goods are not automatically linked to a mode of transportation but rather exported with different modes of transportation to different destinations, thus giving a central role to the bilateral geography. I investigate the role of geography in shaping mode-specific trade flows in the next fact.

Fact III: Geography Affects Transportation Modes Differently

To get a better understanding of the interaction of different modes of transportation with geography I estimate a naive gravity regression as follows:

$$\ln X_{ni}^{m} = \alpha + \alpha_1 \log\left(\frac{Y_n}{N_n}\right) + \alpha_2 n_n + \alpha_3 \left(\frac{Y_i}{N_i}\right) + \alpha_4 n_i + d_{ni}^{k} + l_{ni} + b_{ni} + c_{ni} + \varepsilon_{ni}^{m}. \quad (1)$$

X_{ni}^{m} is the value of trade between n and i that is transported by mode m, α is a constant, Y_k is the GDP of country k, N_k country k's population, d_{ni}^{k} a distance dummy, l_{ni} a dummy for n and i speaking a common language, b_{ni} one for sharing a common border, and c_{ni} one for being on the same continent. It is a naive gravity regression because the value of trade between n and i by mode m potentially depends on the costs of other alternative modes of transportation. The degree to which this matters depends on the substitutability of different transportation modes. This naive gravity regression does not take this substitution possibility into account, but the structural model developed below will provide a way to do this.

Table 2 shows the results of the estimation of equation (1). The results show that richer countries trade more but that air is the mode most strongly influenced by this effect. What is more, this effect is stronger for the exporter than for the importer for all three modes. With respect to distance, the results show that the different modes have very different profiles. Air is least affected by distance and surface is most strongly affected by it. The other geographic controls show that being contiguous increases trade by surface a lot but has an insignificant influence on the other two modes. Similarly, speaking a common language does not foster trade by surface but rather trade by air and sea. All of these effects are qualitatively in line with results from aggregate gravity regressions. The new feature here is the quantitative variation across different modes.

The geographic controls used here are the ones commonly used as proxies for trade costs in gravity regressions (e.g. Anderson and van Wincoop (2004)). The differences across the different modes of transportation then seem to suggest that the transportation modes have different cost profiles. Taking this together with the second fact reported above leads me to model the choice of different transportation modes as being caused by the different bilateral characteristics of the trading countries as oppsed to being solely caused by good-specific characteristics.

Together with the substitutability of transportation modes even at the HS10 level documented above then suggests that agents minimize the costs across transportation means associated with exporting a given good to a certain destination.

Fact IV: Transportation Costs Vary Across Time

To understand the substitutability of different modes in response to the different cost profiles documented above, it is necessary to observe some exogenous variation in mode-specific trade costs. One source that reports such mode-specific charges are the US Imports of Merchandise. This data set reports the charges paid for transportation and insurance for each import separately from the total import value. Figure 7 plots the average (across goods and exporters) ad-valorem equivalents of these charges for air and sea transportation separately from 1995 to 2005. The time series plot shows that there has been a considerable amount of variation in these charges over time. Most noticeably, there was a large spike in air transportation charges in 2002. This spike is a consequence of the tightened security measures after the terrorist attacks of September 11th, 2001. This exogenous shock to transportation charges is what will inform the estimation of the substitutability between different modes in the fourth section.

3. Modelling the Mode of Transportation

In this section I develop a model of international trade that incorporates a choice of the mode of transportation and thus allows one to study mode-specific trade flows. The model is a multi-country Ricardian framework based on Eaton and Kortum (2002) (EK, henceforth) with a production structure similar to Alvarez and Lucas (2007).

3.1. The Economic Environment

Consider a world of $i = 1, \dots, N$ countries, each with a measure L_i of consumers. Each consumer supplies one unit of labor inelastically and only has preferences over the non-traded final good.

In each country, there is a representative firm producing this non-traded final good. The firm has access to the following Cobb-Douglas production technology:

$$Q_i = L_{i,f}^{\alpha} q_{i,f}^{1-\alpha}.$$

Here, $L_{i,f}$ is the amount of labor used in the production of the final good in country i, $\alpha \in [0, 1]$ is the labor share, common across countries, and $q_{i,f}$ is an aggregate intermediate good. The firm's objective is to minimize the production cost of producing Q_i by choosing labor and the aggregate intermediate input, taking prices as given. The price of the final good is denoted by $P_{i,f}$. The aggregate intermediate good is assembled from a continuum of tradeable intermediate goods, $q_i(j)$, according to the CES aggregator

$$q_i = \left(\int_0^1 q_i(j)^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}},$$

where $\sigma > 0$ is the elasticity of substitution among the different varieties.

Each intermediate good $j \in [0,1]$ can be produced in each country under perfect competition. To produce quantity $q_i(j)$ in country i, labor l_i and the aggregate intermediate good q_i are combined according to

$$q_i(j) = X_i(j) l_i^{\beta} q_i^{1-\beta},$$

where $\beta \in [0,1]$ is the labor share in intermediate good production. Across goods j production technologies only differ by the productivity term $X_i(j)$. Firms then minimize the cost of supplying good j given the wage w_i and the price of the aggregate intermediate P_i. Assuming free factor mobility within each country, the unit cost of the input bundle for an intermediate variety j is identical for all varieties and given by

$$c_i = B w_i^{\beta} P_i^{1-\beta},$$

where $B = \beta^{-\beta}(1-\beta)^{1-\beta}$.

Because the intermediate goods are tradeable, firms will source their supply of intermediate j from the lowest cost supplier. I assume that trade across countries is subject to iceberg trade costs, so that an amount $\tau > 1$ of a good needs to be shipped for one unit to arrive.[6] When choosing to source an intermediate variety from another country, an importer can decide which of M different modes of transportation to use. The price of intermediate j produced in country i and delivered to country n via mode m is therefore

$$P_{ni}^m(j) = \frac{c_i}{X_i(j)} \tau_{ni}^m(j).$$

$\tau_{ni}^m(j)$ represents the iceberg trade cost. Not only does it depend on the origin country i and destination n, but also on the mode m used for transportation and the variety j being ordered. I assume that the trade cost can be decomposed into two components:

$$\tau_{ni}^m(j) = \tau_{ni}^m \tau_i^m(j).$$

τ_{ni}^m is a component that is independent of the good being shipped and only depends on the mode of transportation and (n,i)-specific characteristics, such as geography.

[6] Commonly, this formulation is attributed to Samuelson (1954). However, von Thuenen (1826) already proposes such a treatment of transportation costs albeit in an economic geography framework.

$\tau_i^m(j)$ represents the particular costs of shipping good j via mode m, independent of the destination. This cost can vary across exporters. Examples for this cost would be special packaging requirements or a good's bulkiness that makes it more or less costly to ship with mode m. Furthermore, different countries can have differing abilities at adjusting good j to mode m, thus allowing $\tau_i^m(j)$ to differ by exporter.

Since the focus here is on aggregate determinants of mode-specific transportation costs, it is convenient to define an *effective productivity* $Z_i^m(j) = \frac{X_i(j)}{\tau_i^m(j)}$. This effective productivity is an adjustment of the basic ability with which country i can produce good j by its ability of preparing good j for transportation with mode m. The price of delivering good j from country i to country n via mode m can then be rewritten as

$$P_{ni}^m(j) = \frac{c_i}{Z_i^m(j)} \tau_{ni}^m. \tag{2}$$

When choosing from where and how to source good j, importers in country n will choose the lowest cost supplier, so that the actual price of good j in country n is given by

$$P_n(j) = \min_{i,m} P_{ni}^m(j). \tag{3}$$

Note that in addition to choosing the lowest cost producer importers can now also choose the mode of transportation. In this aggregate approach to modelling mode-specific trade flows the benefits of using one mode over another exclusively stem from a lower price, as evident in (3). I thus abstract from intertemporal motives of the transportation choice as modeled in Hummels (2001) or Hummels and Schaur (2010). Instead, I treat them as components $\tau_i^m(j)$ of which I only describe the aggregate behavior through the particular distributional assumption on $Z_i^m(j)$, to which I will turn next.

To facilitate the aggregation of these good-specific demands, I make an assumption on the distribution of effective productivities $Z_i^m(j)$. Because of the common component $X_i(j)$ these effective productivities are correlated over modes of transportation m for a given variety j. Therefore, I assume that the vector $\boldsymbol{Z_i}(j) = (Z_i^m(j))_{m=1}^M$ is distributed according to

$$F_i(\mathbf{z}) = \exp\left[-T_i\left(M^{-1}\sum_{m=1}^{M}(z_m^{-\theta})^{\frac{1}{1-\rho}}\right)^{1-\rho}\right], \tag{4}$$

where $T_i > 0$, $\theta > \max\{1, \sigma - 1\}$, and $\rho \in [0, 1)$. M is the total number of transportation modes available. It is helpful for the interpretation of this distribution to view it as univariate marginals combined by a copula. In particular, (4) is the combination of

Fréchet marginals

$$u(z) = \exp\left[-T\frac{z^{-\theta}}{M^{1-\rho}}\right]$$

coupled by a Gumbel-Hougaard copula

$$\varphi(\boldsymbol{u}(z)) = \exp\left[-\left(\sum_l (-\ln u_l(z))^{\frac{1}{1-\rho}}\right)^{1-\rho}\right],$$

where $\boldsymbol{u}(z) = (u_l(z))_l$. To verify this decomposition, plug the marginals back into the copula to obtain (4). This decomposition shows that θ governs the dispersion of the productivity draws. The larger θ, the lower the dispersion of productivities. T_i influences the mean productivity level: a higher T_i leads to larger productivity draws on average. The association between the different mode-specific draws is entirely determined by the copula. Nelsen (2006) shows that for this copula, ρ corresponds to Kendall's τ, a rank correlation statistic. For $\rho = 0$, the productivities are independent, where as for $\rho \to 1$ the draws are perfectly dependent. I assume that the productivities $\boldsymbol{Z}_i(j)$ are independently distributed across countries.

A distribution similar to (4) is also mentioned by EK and used by Ramondo and Rodriguez-Clare (2009) to describe productivities in a model of multinational production, albeit without making the connection to copulas. The main difference is the presence of the norming factor M^{-1}, which is crucial to include in the current framework.

Proposition 1 summarizes some important implications of the economic structure as laid out above, the proof of which can be found in the appendix.

Proposition 1. *With the above economic structure,*

i) the share of goods that n buys from i is given by

$$\pi_{ni} = \frac{\Phi_{ni}}{\Phi_n}, \tag{5}$$

where $\Phi_n = \sum_j \Phi_{nj}$ and

$$\Phi_{nj} = T_j(c_j\tau_{nj})^{-\theta},$$

and

$$\tau_{nj} = \left[\frac{1}{M}\sum_m (\tau_{nj}^m)^{-\frac{\theta}{1-\rho}}\right]^{\frac{1-\rho}{\theta}}; \tag{6}$$

ii) within the goods that n buys from i, the share that is being transported by mode m

is given by

$$\gamma_{ni}^m = \frac{(\tau_{ni}^m)^{-\frac{\theta}{1-\rho}}}{\sum_r (\tau_{ni}^r)^{-\frac{\theta}{1-\rho}}}; \tag{7}$$

iii) the distribution of prices of goods actually sold by country i in country n and shipped via mode m is independent of the source and the mode of transportation;

iv) the price index of the intermediate good aggregate in country n is given by

$$P_n = \chi \Phi_n^{-\frac{1}{\theta}} \tag{8}$$

with $\chi = \left[\Gamma\left(\frac{\theta+1-\sigma}{\theta}\right)\right]^{\frac{1}{1-\sigma}}$ and $\Gamma(\cdot)$ being the gamma function.

To close the model I assume balanced trade.[7] Following the same logic as in Alvarez and Lucas (2007) and Waugh (2010), the wages can be solved for using

$$w_i L_i = \sum_j w_j L_j \pi_{ji}. \tag{9}$$

Thus, for any given set of parameters, the endogenous prices w_i and P_i can be solved for using (8) and (9).

Property $iii)$ of proposition 1 is instrumental in connecting the model's parameters to observed trade flows. Just as in EK, because the average spending in country n on goods bought from country i and transported by mode m is equal over all sources and modes, the fraction of goods country n buys from country i via mode m is also the fraction of its expenditure on these goods:

$$\frac{X_{ni}^m}{X_n} = \gamma_{ni}^m \pi_{ni}. \tag{10}$$

Summing over m implies

$$\frac{X_{ni}}{X_n} = \pi_{ni}. \tag{11}$$

At the same time,

$$\frac{X_{ni}^m}{X_{ni}} = \gamma_{ni}^m \tag{12}$$

where X_{ni}^m is the *c.i.f.* value of goods that country n imports from country i via mode m, $X_{ni} = \sum_m X_{ni}^m$ and $X_n = \sum_i X_{ni}$. The share country n spends on goods from country i

[7] This is mostly an assumption of convenience. It would be easy to introduce trade imbalances between countries as in Dekle, Eaton, and Kortum (2008).

is larger the smaller country i's relative input costs are, the better its relative technology, and the smaller its aggregated mode-specific trade costs are. The mode-specific share, however, is solely determined by the bilateral mode-specific trade cost relative to the aggregated bilateral mode-specific trade costs.

3.2. A Closer Look at the Model

In this subsection I discuss several properties of the model. First, I explore the determinants of the transportation choice. Then I show how the model nests Ricardian models that study aggregate trade flows. Lastly, I discuss a simple welfare statistic and under what circumstances this statistic is sufficient to deduce the gains from trade.

3.2.1. The Transportation Choice

To understand the choice of which mode country n uses when ordering its goods from country i, it is useful to remember the trade-off for each individual good j in the set of goods that n buys from i, Ω_{ni}. The possible prices of each good are given by

$$P^m_{ni}(j) = \frac{c_i \tau^m_{ni}}{Z^m_i(j)}, \quad j \in \Omega_{ni}.$$

The importers are trading off a higher effective productivity against the mode-specific trade costs. Equation (7) shows for each mode m the fraction of goods transported by that mode, i.e. the fraction of goods for which the trade-off was resolved in favor of mode m. Because of property $iii)$ of proposition 1, this is also equal to the fraction of n's expenditures on goods from i that are transported with mode m. Taking the ratio of these expenditure shares for two modes, say air and vessel, gives

$$\frac{X^a_{ni}}{X^v_{ni}} = \left(\frac{\tau^a_{ni}}{\tau^v_{ni}}\right)^{-\frac{\theta}{1-\rho}}. \tag{13}$$

This shows that the elasticity of substitution between air and vessel shipments is governed both by θ and ρ. Using the copula interpretation of the multivariate Fréchet distribution makes the interpretation of the elasticity parameter $\frac{\theta}{1-\rho}$ clear. If θ is large, effective productivities are less dispersed. In that case, small differences in mode-specific trade costs lead to larger changes in mode-specific trade flows. Similarly, if the correlation between effective productivities is large – if the good specific components $\tau^m_i(j)$ do not play such a large role compared to the core productivity $X_i(j)$ – mode-specific productivities within a certain variety j are very similar, even if the dispersion of productivities across varieties j may be large. Producers can then exploit the smallest differences between mode-specific transportation costs because it is cheap to fit the good for another mode.

The mode-specific trade shares γ_{ni}^m derived in (7) show that in the model relative bilateral mode-specific trade costs between countries determine the choice of the mode of transportation. In particular, one exporter can have very different mode-specific shares with each of his trade partners according to the bilateral mode-specific trade costs between them. In this way, the model is able to capture differences that might arise from different geographic characteristics of the modes of transportation: whereas trucks and railroads might be very cheap for a destination close by, it seems reasonable to expect ships to be the preferred means of transportation for longer distances as already suggested by the naive gravity regressions in section two.

Looking at changes in mode-specific trade costs and the response of trade flows, it follows from (10) that

$$\frac{\partial \ln \left(\frac{X_{ni}^m}{X_n} \right)}{\partial \ln \tau_{ni}^m} = -\frac{\theta}{1-\rho}(1-\gamma_{ni}^m) - \theta \gamma_{ni}^m (1-\pi_{ni}) < 0. \tag{14}$$

An increase in trade costs for mode m will decrease country n's spending on goods delivered by mode m from country i. This decrease happens at two margins. The first margin is an *internal* substitution. At this margin the price change triggers a substitution of modes away from m but keeping the source country i fixed. This is the first summand on the right hand side. It shows that the elasticity is larger the smaller the mode-specific share γ_{ni}^m is and the higher ρ. A high ρ means that the core productivity term $X_i(j)$ is the dominant determinant of $Z_i^m(j)$. Thus, the modes are very close substitutes in terms of costs, and a small change in one mode's trade costs triggers a large change away from that mode. This margin is smaller if the mode-specific share γ_{ni}^m is already high. A high mode-specific share means that m is already the lowest cost mode for delivering goods to n, which results in mode m dominating the composite trade cost τ_{ni} defined in (6). Changing the mode-specific trade cost, then, does not change the relative price very much, so that the response of the mode-specific trade share $\frac{X_{ni}^m}{X_{ni}}$ is small.

The second term represents the *external* margin. It shows the substitution away from i as a supplier. As discussed above, a high γ_{ni}^m means that mode m dominates the bilateral composite trade cost τ_{ni}. Therefore, a change in the mode-specific trade cost leads to a strong response of the bilateral composite trade cost, which is what determines the share country n spends on goods from country i. Correspondingly, the external margin is stronger the larger γ_{ni}^m is. Following a similar argument, if i is already the main supplier for goods to n – represented through a high π_{ni} – a change in the trade cost does not change the relative prices very much, so that the elasticity of the external margin is lower.

Lastly, both margins are influenced in the same way by θ. Remember that a larger θ implies a lower variance of the distribution. Thus, there are smaller differences in productivities. As a result, small changes result in larger substitutions of modes and sources.

As for the overall change in flows between n and i we have

$$\frac{\partial \ln \left(\frac{X_{ni}}{X_n}\right)}{\partial \ln \tau_{ni}^m} = -\theta \gamma_{ni}^m (1 - \pi_{ni}) \tag{15}$$

which is exactly the external margin discussed above. Changing a mode-specific trade cost should only affect goods that are being transported with that mode, so that the external margin discussed in (14) is also the overall change in trade between i and n as shown in (15).

Note that both elasticities are non-constant: they are non-linear and increasing in the mode-specific trade share. Thus, phases of rapid changes in mode-specific trade costs do not have to translate into rapid changes in trade flows if that mode has a very low share. On the other hand, as the mode becomes more important and takes up a larger share in the bilateral trade relationship smaller price changes can have larger effects on trade flows. The introduction of Open Skies Agreements (OSAs) seems to offer one example of this. Micco and Serebrisky (2006) are puzzled that OSAs result in larger changes in trade flows for middle- and high-income countries than for low-income countries. Since low-income countries have lower air shares (cf. Lux (2010)), the present model would predict exactly such a difference.

3.2.2. Deconstructing Gravity

To derive the gravity equation in this model, define a country's total sales as

$$Q_i := \sum_r X_{ri} = \sum_r \Phi_{ri} \frac{X_r}{\Phi_r} = c_i^{-\theta} T_i \underbrace{\sum_r \left[M^{-1} \sum_m \left((\tau_{ri}^m)^{-\theta} \right)^{\frac{1}{1-\rho}} \right]^{1-\rho} \frac{X_r}{\Phi_r}}_{=:\Lambda_i}.$$

Λ_i can be interpreted as country i's market access. Using this and the definition of Φ_n in (11) leads to the gravity equation

$$X_{ni} = \frac{X_n Q_i}{\Phi_n \Lambda_i} \left[M^{-1} \sum_m \left((\tau_{ni}^m)^{-\theta} \right)^{\frac{1}{1-\rho}} \right]^{1-\rho}. \tag{16}$$

Trade between two countries is determined both by the importer's total spending X_n conditional on $\Phi_n \Lambda_i$ and the exporter's total sales Q_i. The strength of the competition in the import market, summarized through Φ_n, and the exporter's market access, Λ_i, influence trade negatively, as do geographic barriers τ_{ni}^m. The composite bilateral trade cost τ_{ni} defined in (6) can then be interpreted as a theoretically consistent aggregator of mode-specific trade costs; it is the aggregate bilateral trade cost index. An interpretation of approaches like EK or Waugh (2010) is that in modelling aggregate trade flows they concentrate on modelling τ_{ni} and ignore the aggregation implicitly involved. To get a better understanding of the trade cost index τ_{ni}, log-linearize (6) around $\tau_{ni}^m = 1$:

$$\ln \tau_{ni} = \sum_m \gamma_{ni}^m \ln \tau_{ni}^m. \tag{17}$$

According to the model, aggregate trade costs are approximately a weighted average of the mode-specific trade costs where the weights are equal to the bilateral mode-specific trade shares γ_{ni}^m. Note that (16) collapses to EK's expression for $\tau_{ni}^m \equiv \tau_{ni}$ for all m: if trade costs do not differ across the modes of transportation, then the two models are identical.

Note that the gravity equation (16) is no longer log-linear in the trade costs, and in particular no longer log-linear in distance. To the extent that distance influences the modes differently, such a log-linear gravity expression is misspecified, as can also be seen from (17).

3.2.3. A Sufficient Statistic Approach?

Arkolakis, Costinot, and Rodriguez-Clare (2010) show that in most standard trade models it is possible to investigate the welfare gains from trade and also the gains from a given change in trade costs through two simple statistics: the trade elasticity and the share of expenditures on domestic goods. Their characterization of standard trade models includes models in the tradition of EK and also Melitz (2003). However, the present framework violates this characterization as long as changes in mode-specific trade costs are concerned because the assumption of a common and constant trade elasticity is violated.[8] To see this, note that the trade elasticity is

$$\frac{\partial \ln \left(\frac{X_{ni}^m}{X_{nn}^m} \right)}{\partial \ln \tau_{ni}^m} = -\frac{\theta}{1-\rho}(1-\gamma_{ni}^m) - \theta \gamma_{ni}^m.$$

[8]The restriction to mode-specific trade costs is important. Because the model nests EK, the arguments of Arkolakis, Costinot, and Rodriguez-Clare (2010) go through even in the current framework as long as only changes in the aggregate trade cost index τ_{ni} are concerned.

Whereas in standard trade models this elasticity is constant, and equal to $-\theta$ in EK, in the disaggregated model this elasticity is dependent on the bilateral pair through γ_{ni}^m. As a consequence the argument in Arkolakis, Costinot, and Rodriguez-Clare (2010) that allows them to reduce welfare gains to a function of the home expenditures and the trade elasticity does not hold for $\rho > 0$. For $\rho = 0$, however, the expression collapses to $-\theta$ and the argument is again applicable. The intuition is that in this case modes are not more interdependent than individual countries; for the sake of the welfare gains, mode-specific trade flows can be viewed as separate countries. What causes the breakdown of their result is the fact that the substitution elasticities between modes within a country and across countries can and do differ, as I will show in the next section.

4. Estimation

In this section I discuss the estimation of the model and present the results. Given the non-linearity of the model already discussed in connection with the gravity equation, the model's estimation has to be based on a system of non-linear equations. I first derive this system of equations and discuss the estimation strategy. Then I determine θ and ρ, two parameters that have to be determined outside of the main estimation procedure. Lastly, I present the results and discuss their robustness.

4.1. Estimating the Model

To determine the set of parameters to be estimated, I assume a log-linear form for the mode-specific trade cost function. In particular, I assume

$$\begin{aligned} \ln \tau_{ni}^m &= f_i^m + \alpha_1^m b_{ni} + \alpha_2^m l_{ni} + \alpha_3^m c_{ni} + \alpha_k^m d_{ni}^{k,m} \\ &= f_i^m + \boldsymbol{\alpha}_m' \boldsymbol{s}_{ni}, \end{aligned} \tag{18}$$

where the second line implicitly defines $\boldsymbol{s}_{ni}$ and $\boldsymbol{\alpha}_m$. b_{ni} is a dummy variable that is one if n and i share a common border, l_{ni} is a dummy that is one if n and i share a common language, c_{ni} is a dummy that is one if n and i are on the same continent, and $d_{ni}^{k,m}$ is one if the distance between countries n and i lies in the k-th interval. Note that distance is a mode-specific regressor in the estimation: I use nautical distances for trade by vessel and great circle distances for trade by air and surface. Nautical distances measure the shortest path across water between the largest ports of any two countries and are thus generally larger than great circle distances. Figure 8 plots the histogram of the ratio of bilateral great circle distance to nautical distance. The histogram shows that the relative distances are concentrated around one, but that there is a substantial left tail for which the great circle distance is much smaller than the nautical distance. The nautical distances are from Feyrer (2009); see the data appendix for a more detailed description of their construction. The effects on the other regressors are allowed to vary by mode of transportation. f_i^m is a mode-exporter specific fixed effect in the trade cost function. Its inclusion is motivated by the arguments in Waugh (2010).

The set of parameters to be estimated consists of the trade cost function parameters $\boldsymbol{\alpha}_m$ and the fixed effects f_i^m for the three modes, the price index parameters Φ_i, and θ and ρ. Conditional on θ and ρ, estimating the other parameters starts from (10). Some simple algebra leads to

$$\frac{X_{ni}^m}{X_n}\frac{X_i}{X_{ii}} = M^{-(1-\rho)}\frac{\Phi_i}{\Phi_n}\left(\tau_{ni}^m\right)^{-\frac{\theta}{1-\rho}}\left[\sum_r \left(\tau_{ni}^r\right)^{-\frac{\theta}{1-\rho}}\right]^{-\rho} \qquad \forall\, m \tag{19}$$

for a given country pair (n, i). Theoretically, it is possible to estimate the trade cost and price index parameters based on this set of equations. But this requires that all mode-specific trade flows are *c.i.f.* trade flows. Remember, though, that for every pair of countries I observe the flows from the reporter to the partner as *f.a.s.* and the flows from the partner to the reporter as *c.i.f.*[9] Exploiting the assumption from above that the observed component of trade costs does not depend on any good specific characteristics, it holds that $X_{ni}^m = \tau_{ni}^m Y_{ni}^m$, where Y_{ni}^m is *f.a.s.* spending on imports from country i to country n via mode m. It is then possible to rewrite (19) as

$$\frac{Y_{ni}^m}{X_n}\frac{X_i}{X_{ii}} = M^{-(1-\rho)}\frac{\Phi_i}{\Phi_n}\left(\tau_{ni}^m\right)^{-\frac{\theta}{1-\rho}-1}\left[\sum_r \left(\tau_{ni}^r\right)^{-\frac{\theta}{1-\rho}}\right]^{-\rho}. \tag{20}$$

Thus for every country pair in the sample I now have six equations: three for the flows from the partner to the reporter, i.e. (19), and three for the other direction, i.e. (20).

To derive the actual equations used to estimate the model define

$$\tilde{X}_{ni}^m = \ln\left(\frac{X_{ni}^m}{X_n}\frac{X_i}{X_{ii}}M^{1-\rho}\right)$$

and

$$\tilde{Y}_{in}^m = \ln\left(\frac{Y_{in}^m}{X_i}\frac{X_n}{X_{nn}}M^{1-\rho}\right).$$

Then for any tuple (n, i) there are three equations – for air, sea, and surface – with *c.i.f.* trade flows and three equations with *f.a.s.* flows:

[9] Every country pair I observe must, of course, always consist of one reporter and one partner.

$$\tilde{X}_{ni}^{m} = \ln \Phi_i + \frac{-\theta}{1-\rho}\left(f_i^m + \boldsymbol{\alpha}_m' \boldsymbol{s}_{ni}\right) - \rho \ln \sum_r \exp\left[\frac{-\theta}{1-\rho}\left(f_i^r + \boldsymbol{\alpha}_r' \boldsymbol{s}_{ni}\right)\right] - \ln \Phi_n + \varepsilon_{ni}^m \tag{21}$$

and

$$\tilde{Y}_{in}^{m} = \ln \Phi_n + \left(\frac{-\theta}{1-\rho} - 1\right)\left(f_n^m + \boldsymbol{\alpha}_m' \boldsymbol{s_{ni}}\right) - \rho \ln \sum_r \exp\left[\frac{-\theta}{1-\rho}\left(f_n^r + \boldsymbol{\alpha}_r' \boldsymbol{s}_{ni}]\right)\right] - \ln \Phi_i + \varepsilon_{in}^m, \tag{22}$$

I jointly estimate this system via non-linear least squares. The Φ_i are captured through fixed effects and so are the f_i^m. The $\boldsymbol{s}_{ni}$ only consist of observable components. The parameters to be estimated are Φ_i, f_i^m, and $\boldsymbol{\alpha}_m$. Because of the use of a spline for the effect of distance, I normalize Φ_{US} and f_{US}^m for all m to one and zero, respectively. θ and ρ cannot be reliably estimated in this system of non-linear equations. Although they are theoretically identified, I have found that the use of fixed effects to capture Φ_i and f_i^m makes it impossible to estimate them reliably. I discuss the identification of them in the next section.

Given the system of equations, define the error term as

$$\boldsymbol{u}_{ni}' := \begin{pmatrix} \varepsilon_{ni,c}^a & \varepsilon_{ni,c}^v & \varepsilon_{ni,c}^l & \varepsilon_{ni,f}^a & \varepsilon_{ni,f}^v & \varepsilon_{ni,f}^l \end{pmatrix} \tag{23}$$

where the subscript c signals *c.i.f.* flows and f signals *f.a.s.* flows. I assume that $\boldsymbol{u}_{ni} \sim (\mathbf{0}, \boldsymbol{\Omega})$ $i.i.d.$ over the tuples (n, i). This assumption does not allow for any correlation of trade flows of the same exporter or importer beyond the explicitly modelled correlation through Φ_i or f_i^m. However, it does allow for an arbitrary correlation structure of the flows between any two countries. Because θ and ρ are estimated parameters but are used in the estimation of the system of equations, the asymptotic approximation to the standard errors is incorrect. Instead, I determine the standard errors of the estimation via bootstrapping. The bootstrap is based on $B = 500$ replications.

After having identified the parameter Φ_i, it is possible to identify T_i according to

$$T_i = \frac{X_{ii}}{X_i}\frac{\Phi_i}{c_i^{-\theta}}, \tag{24}$$

where the wages used to calculate c_i are determined through the balanced trade condition (9) using observed trade shares X_{ni}/X_n as in Waugh (2010). The balanced trade condition uses workforce data based on Heston, Summers, and Aten (2009); the exact procedure is described in the data appendix. The price indices needed for c_i are calculated using (8) and the estimated Φ_i.

4.2. Determining θ and ρ

Estimating θ and ρ starts from (13). Taking logs of that expression gives

$$\ln\left(\frac{X_{ni}^{a}}{X_{ni}^{v}}\right) = -\frac{\theta}{1-\rho}\ln\left(\frac{\tau_{ni}^{a}}{\tau_{ni}^{v}}\right).$$

The problem with estimating $-\frac{\theta}{1-\rho}$ is that τ_{ni}^{m} is generally unobserved. Using proxies as in (18), however, does not allow a separate identification of the elements of $\boldsymbol{\alpha}_m$ from (18) and $-\frac{\theta}{1-\rho}$. As argued in section 2, the observable transportation charges contained in the US Imports of Merchandise data set constitute a component of τ_{ni}^{m} and vary over time.

In line with the trade cost function assumed in (18), one estimation approach would be

$$\ln\left(\frac{X_{US,i}^{a}}{X_{US,i}^{v}}\right) = \frac{f_i^a}{f_i^v} + \alpha_\delta \ln\left(\frac{\delta_{US,i}^{a}}{\delta_{US,i}^{v}}\right) + \alpha_1 b_{US,i} + \alpha_2 l_{US,i} + \alpha_3 c_{US,i} + \alpha_k^a d_{US,i}^{k,a} + \alpha_l^a d_{US,i}^{l,v} + \xi_{US,i}, \quad (25)$$

where $\xi_{US,i}$ is an error term assumed to satisfy the standard assumptions. The same controls as in (18) are added since the transportation charges $\delta_{US,i}^{m}$ only represent part of the overall trade costs. The charges $\delta_{US,i}^{m}$ are the ad-valorem equivalent of the import charges reported in the data set. Comparing (18) and (25) it becomes clear that there is no possibility of reliably estimating the exporter specific trade cost component f_i^m. The problem is that there is no variation over the importer; the transportation charges are only reported for the US as an importer. On the other hand, an estimation utilizing the time dimension of the data would lead to inaccurate estimates given the relatively short time span of 11 years. Assuming that f_i^m does not change over short horizons of time, it is possible to improve on (25) by estimating it in a differenced form (I drop the importer subscript to avoid cluttered notation):

$$\Delta \ln\left(\frac{X_{i,t}^{a}}{X_{i,t}^{v}}\right) = \alpha_\delta \Delta \ln\left(\frac{\delta_{i,t}^{a}}{\delta_{i,t}^{v}}\right) + \Delta\xi_{i,t}, \quad (26)$$

with $\Delta y_t = y_t - y_{t-1}$. This formulation assumes that the coefficients on time-invariant proxies like contiguity and distance do not change over short intervals of time. Some evidence for this assumption with respect to distance can be found in Disdier and Head (2008). In this differenced form, the assumption on the error term is now (cf. Cameron and Trivedi (2005))

$$E_t\left[\left(\Delta \ln \frac{\delta_{i,t}^a}{\delta_{i,t}^v} \Delta\xi_{i,t}\right)\right] = 0, \tag{27}$$

where $E_t(\cdot)$ denotes the conditional expectation on time t information. This estimation approach is similar to the demand equation specification used by Broda and Weinstein (2006) to estimate demand elasticities. But instead of specifying a supply equation, I exploit the exogenous shock of the 9/11 attacks to identify the substitution elasticity as discussed below.

Figure 9 shows the scatter plot of $\Delta \ln\left(\frac{X_{i,t}^a}{X_{i,t}^v}\right)$ against $\Delta \ln\left(\frac{\delta_{i,t}^a}{\delta_{i,t}^v}\right)$, using annual data from 1995 to 2005. The sample of exporters is the same country sample as the one used in the main estimation. As predicted by the model, there is a negative relationship between the two variables. Table 4 contains the regression results. The first column contains the basic regression (26). The estimate for α_δ is significant with a value of -6.7. The result is quite robust to other specifications. The parameter estimate barely changes when time fixed effects are included. This confirms that there are not different growth rates of the two mode-specific trade values over this time period that could influence the result. One worry from the inspection of figure 9 is that the outliers might dominate the result. To control for these outliers, column three reports the results of an iterative least squares procedure that reweighs observations in each iteration depending on their influence on the estimates. The results show that the outliers are not driving the estimate of α_δ. The estimate slightly increases for the robust procedure.

One further worry might be that the variation in transportation charges is not exogenous, i.e. that the identifying assumption (27) is violated. To investigate this furter I estimate the coefficient by splitting the sample in pre-1999 and post-2000. The variation in the latter half is strongly driven by the increase in air transportation charges in the wake of the 9/11 attacks documented in figure 7 and discussed above. The estimation shows that the coefficient for these two sub-samples are -6.3 and -7, respectively, which frame the estimate of -6.7 obtained using the whole sample.

Disentangling θ and ρ is more difficult. The difficulty of estimating the aggregate trade elasticity is common to this class of Ricardian models (cf. EK and Waugh (2010) but also Fieler (2009)). The original idea of EK to estimate θ is via an arbitrage condition on goods' prices in different countries. This condition states that

$$\frac{p_n(l)}{p_i(l)} \leq \tau_{ni} \quad \forall\, l\, \forall\, n, i.$$

With this condition it is possible to estimate τ_{ni} and thus obtain an estimate of θ independent of the effects of geographic proxies. Simonovska and Waugh (2009) argue that this approach actually upward biases the estimate of θ because of the inherent bias in the first-order statistic used to estimate τ_{ni}. They correct for this bias in an EK model and arrive at an estimate of $\theta = 4.22$. This approach is not correct, however, if goods can be transported with different modes of transportation. The reason is that the no-arbitrage condition changes to

$$\frac{p_n(l)}{p_i(l)} \leq \min_m \tau_{ni}^m \quad \forall\, l\, \forall\, n, i.$$

Because $\min_m \tau_{ni}^m \leq \tau_{ni}(\tau_{ni}^m)$ where $\tau_{ni}(\tau_{ni}^m)$ is the choice-theoretically consistent trade cost index derived above, they overestimate the true θ as it would be estimated using the disaggregated model presented here. Thus, their estimate is an upper bound on the value of the trade elasticity. EK suggest an alternative estimation procedure relying on wage data instead of price data. This approach is valid in the current context, as well. With this method they reach a value of $\theta = 3.6$. This is generally considered to be a lower bound on the parameter.

It turns out that the difference between the mode-specific trade costs is not very large, so that $\min_m \tau_{ni}^m - \tau_{ni}(\tau_{ni}^m)$ is likely to be small and the resulting remaining bias in the estimate of Simonovska and Waugh (2009) as well. This suggests using a value of $\theta = 4$, which implies a $\rho = 0.4$ for an estimate of $\frac{\theta}{1-\rho}$ of 6.7.

4.3. Estimation Results

I drop all zero trade flows from the sample.[10] This leads to a drop of about 8.5% of observations. The remaining sample contains 7308 observations, which are roughly equally distributed over the six equations.

4.3.1. Model Fit

To understand the fit of the model I use the correlation between the model's predicted bilateral trade shares and the data. The correlations are reported in table 5. For the aggregate bilateral trade share $\frac{X_{ni}}{X_n}$ the correlation is 0.59. At this stage of aggregation the model is essentially identical to EK in the formulation of Waugh (2010). As he discusses in the paper, the fit of the model in explaining trade flows is very accurate measured through the model's R^2. Stating the correlation measure here is meant to serve as an anchor for evaluating the model's ability to fit mode-specific trade flows.

Table 5 also reports the correlation between data and model predictions for mode-

[10] Helpman, Melitz, and Rubinstein (2008) investigate the potential sample selection bias introduced by this method and find a negligible bias.

specific trade shares $\frac{X_{ni}^m}{X_n}$, i.e. the predicted γ_{ni}^m. Sea based trade shares have the highest correlation with 0.65. The correlation for air shares is 0.6. The model's fit is worst for surface based trade flows; the correlation between predicted and actual trade shares drops to 0.44. Given that the last is an agglomeration of several flows, this does not seem too unexpected.

Comparing the fit for the bilateral aggregate trade share and the mode-specific trade shares reveals that the model describes the mode-specific trade flows about as well as the aggregate ones. Thus, an accurate description of mode-specific trade flows is possible even without concentrating on goods' characteristics to explain the transportation choice.

4.3.2. Parameter Estimates

Table 6 shows the estimated coefficients of the mode-specific trade cost functions for contiguity, sharing a common language, and being on the same continent. The coefficient estimates are translated into ad-valorem cost equivalents using the relation $100 \times (\exp(\hat{\alpha}_j) - 1)$. The reported significance levels are based on bootstrapped standard errors with 500 replications. In line with usual estimates of these coefficients (cf. Anderson and van Wincoop (2004)) all three characteristics lead to a reduction in trade costs. What is new here is the separation across modes of transportation. Not surprisingly, sharing a common border has the strongest effect on surface transportation: it lowers trade costs by 46.2% compared to just about 12.3% for trade by vessel. Air transportation profits from a common border through a reduction of 21.8%. A common language is most helpful in lowering trade costs for air transportation, surface based transportation profits the least from it. However, the variation across modes is considerably lower than compared to the estimates for contiguity. Being based on a common continent is, not surprisingly, most helpful for surface based trade with a cost reduction of ca. 24% and least cost-reducing for shipping. The estimates are very precise based on the reported significance levels. Overall, this deconstruction of the effects of trade cost proxies into their differential effects through the modes of transportation shows the large and intuitive heterogeneity that is lost in the usual aggregate approach to estimating trade costs.

Figure 10 plots the estimated distance splines $\hat{\alpha}_k^m$ over the distance categories. I choose the mid-point of each interval on the x-axis. The intervals are reported in table 7. They have been chosen such that roughly an equal amount of observations falls into each category to maximize the precision of the estimates. The bands around the splines are 95% confidence intervals. The confidence intervals are based on bootstrapped standard errors, as above. As for the other trade cost proxies, disaggregating the effect of distance by the mode of transportation reveals a rich variation that seems intuitive. Distance has the largest effect on surface based transportation, which seems quite intuitive given that most of this category entails goods transported by trucks and railroad. The effect on air transportation is relatively modest. Going from trade below 2500 km to over 11000 km raises air transportation costs only by about 69% compared to 177% for surface based transportation. Again, this seems very much in line with common intuition: once an airplane is in the air the marginal kilometer should be rather cheap compared to a truck. The distance effect on shipping is somewhere in between these two extremes. It rises in the beginning, but flattens off for larger distances.

Figure 11 plots the estimated exporter fixed effects in the trade cost function for air transportation, f_i^a, against the log of the GDP per capita. In line with what Waugh (2010) finds there is a strong negative correlation between the two variables. Figure 12 and 13 plot the fixed effects for vessel and surface transportation, respectively, and the same pattern emerges: rich countries face lower barriers to exporting and do so across different modes of transportation. To further investigate the variation of the exporter fixed effects, I calculate the coefficient of variation both within countries across modes and across countries for each mode. The mean of the variation within countries and across modes is 0.31 with a median of 0.25. The variation across countries is around three times as large: 0.78 for air, 0.93 for sea, and 1.08 for surface transportation.

Figure 14 plots the estimated technology parameters T_i against GDP per capita. The technologies are strongly positively associated with GDP per capita. This strong covariation is also found in Waugh (2010). Japan has the highest estimated technology closely followed by Korea. The countries with the lowest technologies are predominantly African countries; Ghana has the lowest estimate.

4.3.3. Difference to an Aggregate Model

How different are the estimates from an aggregate model if the variation across modes is as large as shown in table 6? To investigate the answer to this question, I simulate aggregate trade flows using the estimates from the model. I then estimate an aggregate specification just as in Waugh (2010) using these simulated data. The last row of table 6 reports the results from this experiment. Not surprisingly, the aggregate estimates lie in between the disaggregated estimates. Figure 15 combines the distance estimates from the disaggregated model with the estimates from the aggregated model. As in the case of the trade cost function parameters, the aggregate distance effect is a combination of the effects of the three modes. There are two things to note about the results.

First, the large aggregate contiguity effect is what Overman, Redding, and Venables (2004) and Hummels (2007) call the puzzling fact that most countries tend to trade with their neighbors. Disaggregating the model and estimating mode-specific trade costs reveals that the contiguity effect is largely driven by surface trade. From this observation, one plausible explanation of this effect might be based on different transportation economies of scale associated with the different modes. Assume that the different economies of scale are captured in large fixed transportation costs associated with maritime trade and much smaller fixed transportation costs for surface trade. Then many more small shipments will be send via surface transportation whereas only sufficiently large shipments are traded between countries where maritime trade is the cheaper option. The estimates show that surface trade is only feasible for short distances; the distance profile is much steeper for surface than for the other modes. If the amount of small shipments is a sufficiently large share of overall trade flows, this mechanism could generate the large estimated contiguity effect.

The second interesting thing to note concerns the disaggregated distance profiles in figure 15. There is a sense in the gravity literature that distance only matters for shorter distances and not so much for larger distances. The disaggregated estimates show that this is the combination of surface trade having both a large distance elasticity and being predominantly used for regional trade with neighbors or on the same continent. This leads to a sharp increase in the aggregate estimate of the distance effect. For larger distances, however, the other two modes become more dominant and the distance effect is dominated by the distance profile of air and vessel. The estimates show that trade in these modes reacts much less to distance which leads to the observed flattening of the aggregate distance effect.

4.4. Robustness

In this subsection I discuss some additional results that are meant to highlight the robustness of the benchmark estimation.

4.4.1. Weighted NLS

The benchmark estimation is not efficient since it does not explicitly exploit the assumed correlation structure of $\mathbf{\Omega}$. To determine whether a weighted NLS procedure would yield more efficient – and generally different – estimates, I re-estimate the model. At this second step I use a weighted NLS procedure, where the estimate of $\mathbf{\Omega}$ is based on the residuals $\hat{\boldsymbol{u}}_{ni}$. The standard error of the regression when comparing the weighted and the unweighted estimate are virtually indistinguishable. The sample is large enough for there to be no efficiency loss to the unweighted procedure.

4.4.2. The Importance of ρ

To determine the sensitivity of the estimates with respect to ρ, I re-estimate the model with two different values of ρ: $\rho = 0.2$ and $\rho = 0.6$, keeping $\theta = 4$. Table 8 reports the resulting correlations between predicted mode-specific shares $\frac{X^m_{ni}}{X_n}$ and the data. The correlations barely change with the different values of ρ, which hints at it being poorly identified in the model itself. Note that Ramondo and Rodriguez-Clare (2009) experience a very similar result in their model of multinational production, using a very similar correlated Fréchet distribution.

4.4.3. Great Circle Distances

The benchmark model uses actual nautical distances for the maritime trade flows. It seems clear that in many cases the great circle distance is only a poor proxy for the distance actually travelled by a ship. Indeed, the discussion of figure 8 above has already established the extent of this difference in the data set. However, it is much less clear what implications this difference in distances has for the study of trade flows. In order to determine the quantitative importance of this difference, I re-estimate the model using only great circle distances. The fit of the model is basically unchanged; the standard error of the regression is only 0.7% larger. Table 9 shows the results for the coefficients of the trade cost function. Although there are some differences between the two models, they all are extremely minor. Figure 16 plots the resulting distance effects of the different modes. Here, there is basically no change with respect to surface or air transportation. On this metric, maritime trade has the lowest distance sensitivity of all modes. For distances beyond 6000 km, the distance effects of air and vessel transportation are basically indistinguishable. The fact that the relative distance effects become more similar with increasing distance means that the air share grows relative to their vessel share: more distant partners have higher air shares relative to the vessel shares than closer partners. This is reminiscent of the effects Harrigan (2009) finds. Here, however, this effect occurs at a more aggregate level and for a wider sample of countries.

5. Two Counterfactuals

Before discussing the counterfactuals, some as yet unspecified parameters have to be determined. Since these parameters are not readily estimateable from the data I choose to calibrate them in line with previous work. Table 10 summarizes these parameters and the values I choose. β is the value added share in intermediate good production. Using data on value added and gross manufacturing production from the UNIDO data I calculate a value of 0.31 for the countries in my sample. Waugh (2010) uses a minimally larger share of 0.33. σ is the elasticity of substitution between intermediate varieties. The value does not play any role other than in the constant χ of the price index and only needs to satisfy the restriction $\theta > \max\{1, \sigma - 1\}$. I choose $\sigma = 4.5$. α controls value added in the final good production. Alvarez and Lucas (2007) discuss plausible values in the range of 0.7 to 0.8. I follow Waugh (2010) and pick $\alpha = 0.75$.

5.1. The Role of Transportation

In the first series of counterfactuals I take the benchmark model as estimated above and compare this to worlds in which one of the three modes is not available. The aim of the counterfactual is to try to understand the contribution of each individual mode to the gains from trade. To gauge this contribution, I calculate for each counterfactual the statistic

$$\Delta_{i,m} = \frac{\frac{\omega_{i,m}}{\omega_i} - 1}{\frac{\omega_{i,\text{aut}}}{\omega_i} - 1}. \tag{28}$$

$\omega_i = \frac{w_i}{P_{i,f}}$ denotes the real wage of country i in the benchmark case, $\omega_{i,m}$ denotes the real wage when mode m is made prohibitively expensive, and $\omega_{i,\text{aut}}$ is the real wage under autarky. The statistic $\Delta_{i,m}$ denotes the fraction of the gains from trade that are foregone if mode m is eliminated. Figures 17 to 19 show a scatter plot of these statistics against (log) GDP per capita for the three different modes along with the best fit lines. Figure 17 shows that richer countries rely slightly more heavily on air transportation for realizing their gains from trade. But the slope of this relationship is rather small. Furthermore, it shows that most countries do not rely very much on trade by air. The gains from trade would be on average only 8% lower if air transportation was impossible. Figure 18 shows that the picture is very different for trade by sea. First, the importance of sea transportation is large for most countries: gains from trade would be on average 34% lower if sea transportation was not possible. Furthermore, the relation with a country's development level is reversed. Poor countries rely much more heavily on access to sea transportation than rich countries. The reason for this becomes clear when looking at the role played by surface transportation in realizing gains from trade as depicted in figure 19. This figure makes clear that rich countries rely more heavily on surface transportation than poor countries.

The intuition for this pattern is best understood by thinking about European and

African countries. If European countries' access to far-away markets is made more expensive by eliminating air transportation, some trade will shift towards neighboring countries. But for European countries, these are still technologically advanced markets that offer many opportunities to exploit comparative advantages. Thus, the second-best price for most goods will not be much higher than it was and the welfare gains from trade are not strongly influenced. The story is different for African countries. If African countries are forced to trade more with their neighbors because long-distance trade has become more expensive, this means going from trading with technologically advanced but far-away markets to trading with close but technologically much less advanced markets. This results in sharply increasing prices for goods and thus a larger negative change in the real wage.

This counterfactual highlights the role of different modes of transportation in overcoming the adverse distribution of technological achievements across the globe. Technologically advanced countries are bunched together and so are technologically disadvantaged countries. As a result, the poor countries rely on long-distance modes of transportation to profit from trade much more heavily than rich countries.

Whereas the discussion so far has focused on relative gains from the modes, table 11 summarizes the level effects. The first column shows that the average welfare change across modes is very different. Shutting off air transportation decreases the average real wage by only 0.8% whereas shutting off sea or surface transportation has much larger effects: -2.9% and -2.4%, respectively. Note that these changes do not add up to the average loss of 11% when going to autarky. The reason is the strong substitutability. To highlight this, the second column shows the drop in world trade. Comparing this to the share of world trade transported by each affected mode prior to the counterfactual in column three shows that the drop is always much smaller. That is, a lot of trade is actually retained and shifts to other modes of transportation.

5.2. The Role of Transportation in Decreasing Income Differences

Recently, Waugh (2010) has argued that exporter-specific fixed effects in the trade cost functions can explain up to 30% of income differences. Given the strong substitutability of transportation modes highlighted in the first set of counterfactuals, I now examine the ability of mode-specific exporter fixed effects to account for income differences. To do this, I compare income differences in the benchmark model to a model where I reduce the mode-specific fixed effects in the trade cost functions to the US level for a particular mode. That is, I set $f_i^m = \min\{f_i^m, f_{US}^m\}$. Table 12 shows the resulting changes in income differences, measured as the variance of the log of the real wage. The results show that reducing the exporter specific barriers for either air or sea transportation reduces income differences by about 35% in the sample. Reducing the barrier for surface transportation only reduces income differences by a little over 20%. The difference is easily understood with the intution already developed above. As discussed in Waugh (2010), the reduction of income differences through a reduction of asymmetries in trade costs is caused by poor countries starting to import a lot more goods to compensate for their technological disadvantage. But if the asymmetry is reduced in surface transportation, poor countries can only trade more cheaply with other countries nearby, which are mostly technologically disadvantaged, as well. If, on the other hand, the asymmetries are reduced in long-distance modes like air or sea transportation, poor countries can access technologically advanced far-away markets and thus exploit much larger gains from trade. The last column shows that reducing all mode-specific exporter fixed effects simultaneously reduces income differences by about 40%, only slightly more than in the case of air or sea transportation. Again, because the substitutability between different transportation modes is so strong, reducing the asymmetries in only one long-distance mode already allows countries to realize large additional gains from trade.

To put these numbers in perspective, I determine the reduction of income differences achieved in a world of free trade, i.e. where $\tau_{ni}^{m} \equiv 1$ for all modes m. In such a world, income differences shrink to about 40% of what they are in the benchmark case. Thus, reducing the exporter specific barriers in sea transportation alone achieves half the reduction in income differences compared to a world of free trade.

As a further check I calculate the change in income differences resulting from reducing the fixed effect to the minimum within each country. The results are reported in the third column. This exercise is meant to show that the reduction in income differences does not come from the reduction of within country variation of these fixed effects but really from the much larger across country variation in export barriers.

6. Conclusion

In this paper I have developed and estimated a model of mode-specific international trade flows to study the role played by individual modes of transportation in international trade and the effects of changes to mode-specific as opposed to aggregate trade costs. Although the model is able to accomodate any number of transportation modes, I concentrate on the three most important ones: air, sea, and surface transportation. I estimate the model using a novel data set. The resulting estimates of the trade cost functions are quite intuitive, showing that air and vessel transportation are important for long-distance trade and surface transportation mostly used for short-distance trade.

I then use counterfactual analysis to show two things. First, I decompose the gains from trade into their mode-specific components. To do this I compare the benchmark model to a world where one mode is prohibitively expensive. The most important lesson is that modes are strongly substitutable: even removing the possibility to trade by sea, the most widely used mode, reduces welfare on average by only about 3% compared to the average loss of 11% when moving to autarky. The second implication is the different importance of transportation modes for different countries. Poor countries rely much more heavily on access to long-distance transportation in realizing gains from trade than rich countries. The reason lies in the geographic distribution of technologies. Since rich and poor countries form clusters on the globe, eliminating long-distance trade for rich countries is not as damaging as removing it for poor countries. Rich countries can substitute much of their trade towards their neighbors who are still technologically advanced enough to exploit large gains from trade. Poor countries, on the other hand, face a much steeper price increase when forced to trade more with their technologically disadvantaged neighbors.

In a second counterfactual I investigate the role of individual modes of transportation in reducing income differences. The high substitutability estimated implies that removing asymmetries in air or sea transportation alone can reduce income differences in the sample by about 35%. In comparison, moving to a world of free trade reduces income differences by about 60%.

One limitation of the model developed here is that it is static. However, there is evidence that the delivery time matters for international trade, cf. Evans and Harrigan (2005) or Alessandria, Kaboski, and Midrigan (2008). An interesting avenue for future research would thus be to study the role of different transportation modes in international trade in a dynamic context. The considerable challenge here is to generalize these dynamic trade models to a multi-country framework.

A Data Appendix

A1. Gross Manufacturing Production

The data for gross manufacturing production come from three different sources. If available, I use the UNIDO Industrial Statistics Database 2008.[11] However, these data are not available for all countries in the sample. Where necessary, I supplement the data with gross manufacturing production from the OECD STAN database[12] or impute it with UN National Accounts Statistics on value added in manufacturing.[13]

Exchange rate adjustments for the OECD data are made using the exchange rates from the Penn World Tables 6.3.

A2. Trade Data

To limit the trade flows to manufacturing I employed the concordance suggested in Maskus (1991).

Mode-Specific Trade Data: The data for bilateral trade flows disaggregated by the mode of transportation come from two sources. The data on flows involving the U.S. are from the US Imports/Exports of Merchandise. The second data source is the `XTNET` data base from Eurostat.[14] To convert the trade data quoted in Euro into US dollars, I use the exchange rates as reported in the Penn World Tables, edition 6.3.

Aggregate Trade Data: Trade data on aggregate trade flows are from the update data set based on Feenstra, Lipsey, and Bowen (1997).

A3. Geographic Data

The geographic data used for the trade cost function all come from the CEPII data base.[15] The great circle distance is measured in kilometers.

The data for bilateral nautical distances have been generously provided by James Feyrer. For details of the calculation of these distances, see Feyrer (2009). The data set reports bilateral distances as days for a round trip. To convert them into kilometers, I assume a vessel operating 24 hours per day at 20 knots.

A4. Labor Force Data

Using data on GDP per capita, population, and GDP per worker, the labor force for each country in the sample is recovered as in Caselli (2005). The data are from Heston, Summers, and Aten (2009).

B Proof of Proposition 1

The proof is based on arguments similar to the ones in Ramondo and Rodriguez-Clare (2009). First, note that the distribution of prices for goods n buys from i is given by

[11] http://www.unido.org/index.php?id=o3472
[12] http://www.oecd.org/document/54/0,3343,en_2649_34445_21573686_1_1_1_1,00.html
[13] http://unstats.un.org/unsd/snaama/Introduction.asp
[14] http://epp.eurostat.ec.europa.eu/newxtweb/
[15] http://www.cepii.fr/anglaisgraph/bdd/distances.htm

$$
\begin{aligned}
G_{ni}(p) := \mathbb{P}[P_{ni}(j) \leq p] &= 1 - \mathbb{P}\left[\forall\, m:\ P^m_{ni}(j) \geq p\right] \\
&= 1 - \mathbb{P}\left[\forall\, m:\ Z^m_i(j) \leq \frac{c_i \tau^m_{ni}}{p}\right] \\
&= 1 - \exp\left[-\Phi_{ni} p^\theta\right],
\end{aligned}
\tag{B.1}
$$

where

$$
\Phi_{ni} = T_i c_i^{-\theta} \left[\frac{1}{M} \sum_m \left((\tau^m_{ni})^{-\theta}\right)^{\frac{1}{1-\rho}}\right]^{1-\rho}. \tag{B.2}
$$

With this, the price distribution of goods in n is

$$
G_n(p) = 1 - \prod_i (1 - G_{ni}(p)) = 1 - \exp[-\Phi_n p^\theta]. \tag{B.3}
$$

where $\Phi_n = \sum_i \Phi_{ni}$. Now, since $G_{ni}(p)$ is independent across i, the same reasoning as in EK leads to the conclusion that the measure of goods n buys from i is equal to

$$
\pi_{ni} = \frac{\Phi_{ni}}{\Phi_n}. \tag{B.4}
$$

This proves (5). To see what share n buys from i via mode m, let us focus on one particular mode, say $m = a$. Then the share is the mass of prices for which

$$
\begin{aligned}
& P^a_{ni}(j) \leq P^m_{nl}(j) \quad \forall\, m\, \forall\, l \\
\Leftrightarrow \quad & Z^m_l(j) \leq \frac{c_l \tau^m_{nl}}{c_i \tau^a_{ni}} Z^a_i(j) \quad \forall\, m\, \forall\, l
\end{aligned}
\tag{B.5}
$$

Define the set A as the above event. Also, define $a_{nl,m} = \frac{c_l \tau^m_{nl}}{c_1 \tau^a_{n1}}$ and set, without loss of generality, $i = 1$. Then

$$
\begin{aligned}
\mathbb{P}(A) &= \int_0^\infty F(dz, a_{n1,v} z, \ldots, a_{nl,m} z, \ldots) \\
&= \int_0^\infty \frac{\theta}{M} z^{-\theta-1} T_1 \left((c_1 \tau^a_{n1})^\theta\right)^{1-\frac{1}{1-\rho_1}} \left(\frac{1}{M} \sum_m (c_1 \tau^m_{n1})^{-\frac{\theta}{1-\rho}}\right)^{-\rho} \exp\left[-\Phi_n (c_1 \tau^a_{n1})^\theta z^{-\theta}\right] dz \\
&= M^{-1} \left(T_1 c_1^{-\theta} (\tau^a_{n1})^{-\theta}\right)^{\frac{1}{1-\rho}} \left(T_1 c_1^{-\theta} \left(M^{-1} \sum_m (\tau^m_{n1})^{-\frac{\theta}{1-\rho}}\right)^{1-\rho}\right)^{-\frac{\rho}{1-\rho}} \frac{1}{\Phi_n} \\
&\int_0^\infty \theta z^{-\theta-1} (c_1 \tau^a_{n1})^\theta \Phi_n \exp\left[-\Phi (c_1 \tau^a_{n1})^\theta z^{-\theta}\right] dz \\
&= \left[\frac{M^{-(1-\rho)} T_1 c_1^{-\theta} (\tau^a_{n1})^{-\theta}}{\Phi_{n1}}\right]^{\frac{1}{1-\rho}} \frac{\Phi_{n1}}{\Phi_n} \\
&= \frac{(\tau^a_{n1})^{-\frac{\theta}{1-\rho}}}{\sum_m (\tau^m_{n1})^{-\frac{\theta}{1-\rho}}} \frac{\Phi_{n1}}{\Phi_n},
\end{aligned}
\tag{B.6}
$$

because $-\frac{x}{1-x} = 1 - \frac{1}{1-x}$ and

$$F(dx_{11}, x_{21}, \ldots, x_{MN}) = \frac{\theta}{M} T_1 \left(\frac{1}{M} \sum_m (x_{m1})^{-\frac{\theta}{1-\rho}} \right)^{-\rho} (x_{11})^{-\frac{\theta}{1-\rho}-1} \exp\left[-\sum_i T_i \left(M^{-1} \sum_m (x_{mi})^{-\frac{\theta}{1-\rho}} \right)^{1-\rho} \right] dx_{11}. \tag{B.7}$$

Now $\frac{\Phi_{n1}}{\Phi_n}$ is the share of goods n buys from $i = 1$, so

$$\gamma_{n1}^a := \frac{(\tau_{n1}^a)^{-\frac{\theta}{1-\rho}}}{\sum_m (\tau_{n1}^m)^{-\frac{\theta}{1-\rho}}} \tag{B.8}$$

is the share of goods n buys from $i = 1$ that are transported via mode $m = a$. This proves (7). Then the distribution of prices in n for goods that are actually imported from country $i = 1$ with mode $m = a$ is

$$\begin{aligned} \mathbb{P}\left[\{P_{n1}^a(j) \leq p\} \mid \{P_{n1}^a(j) \leq P_{nl}^m(j) \, \forall m \, \forall l\}\right] &= \frac{\mathbb{P}\left[\{P_{n1}^a(j) \leq p\} \cap \{P_{n1}^a(j) \leq P_{nl}^m(j) \, \forall m \, \forall l\}\right]}{\mathbb{P}\left[\{P_{n1}^a(j) \leq P_{nl}^m(j) \, \forall m \, \forall l\}\right]} \\ &= \frac{1}{\mathbb{P}(A)} \mathbb{P}\left[\{Z_1^a(j) \geq \frac{c_1 \tau_{n1}^a}{p}\} \cap \{Z_l^m(j) \leq \frac{c_l \tau_{nl}^m}{c_1 \tau_{n1}^a} Z_l^a(j) \, \forall m \, \forall l\} \right] \\ &= \frac{1}{\mathbb{P}(A)} \int_{\frac{c_1 \tau_{n1}^a}{p}}^{\infty} F(dz, \ldots, a_{nl,m} z, \ldots) \\ &= \frac{1}{\mathbb{P}(A)} \exp\left[-\Phi_n (c_1 \tau_{n1}^a)^\theta z^{-\theta} \right] \Big|_{\frac{c_1 \tau_{n1}^a}{p}}^{\infty} \\ &= 1 - \exp[-\Phi_n p^\theta] \\ &= G_n(p). \end{aligned} \tag{B.9}$$

This proves the third claim. Applying once again the same logic as in EK, it can then be shown that

$$P_n = \underbrace{\left[\Gamma\left(\frac{\theta + 1 - \sigma}{\theta} \right) \right]^{\frac{1}{1-\sigma}}}_{\chi} \Phi_n^{-\frac{1}{\theta}}. \tag{B.10}$$

This proves (8).

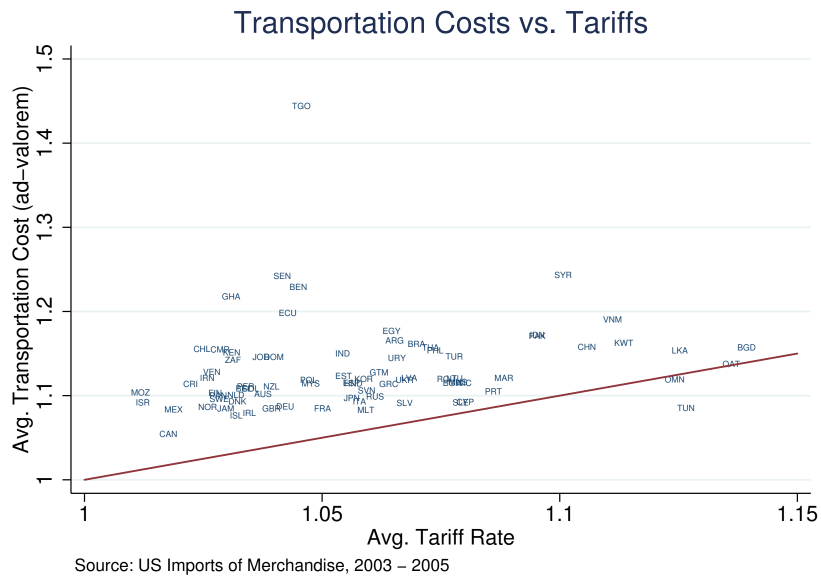

Figure 1: Average ad-valorem transportation charges against average tariffs for imports into the US. The red line is the 45° line.

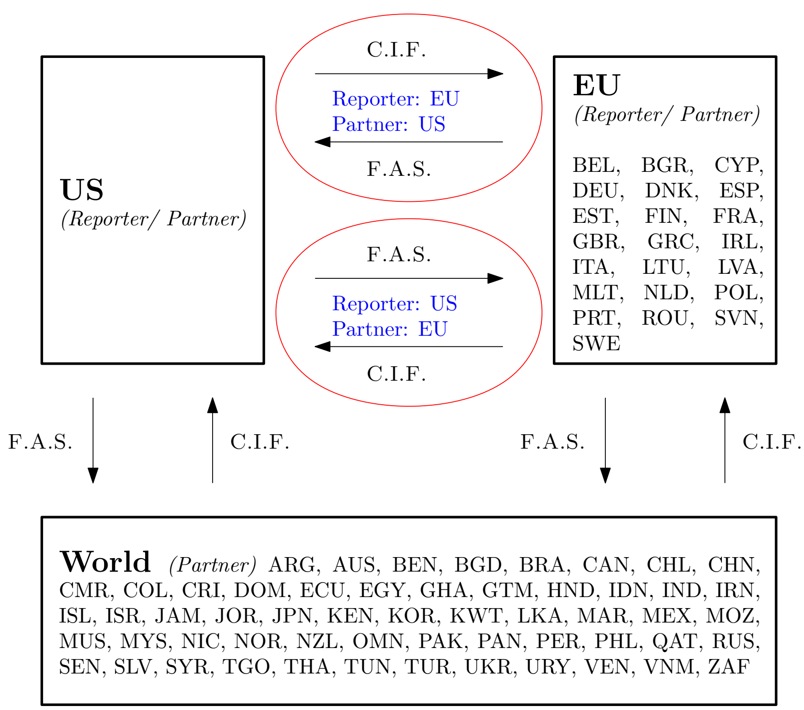

Figure 2: Structure of the combined data set of mode-specific trade flows. Trade is only observed between countries in different blocks but not within.

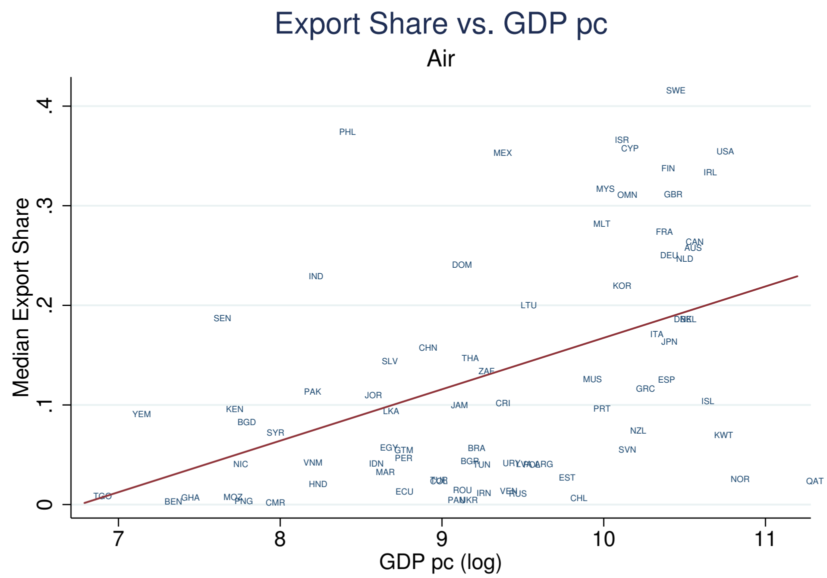

Figure 3: The graph shows the median share of exports by air against the (log) GDP per capita. The solid line shows the best fit.

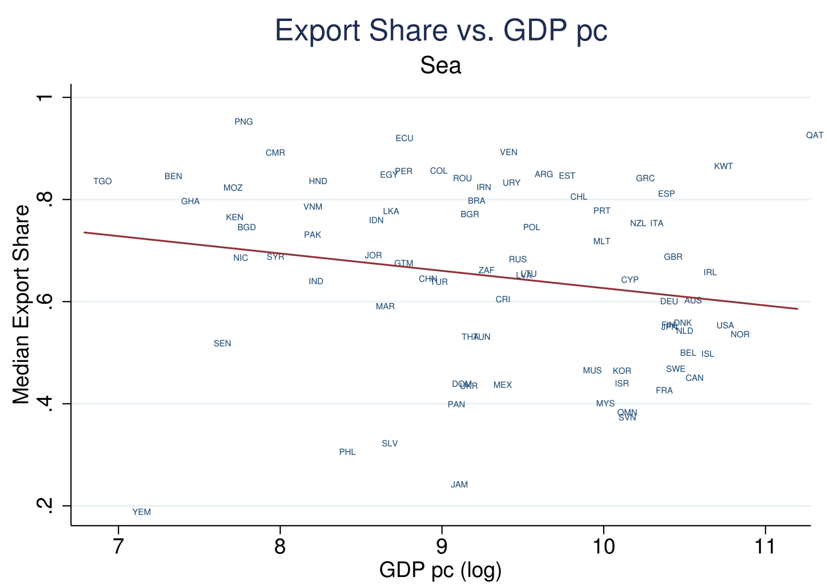

Figure 4: The graph shows the median share of exports by sea against the (log) GDP per capita. The solid line shows the best fit.

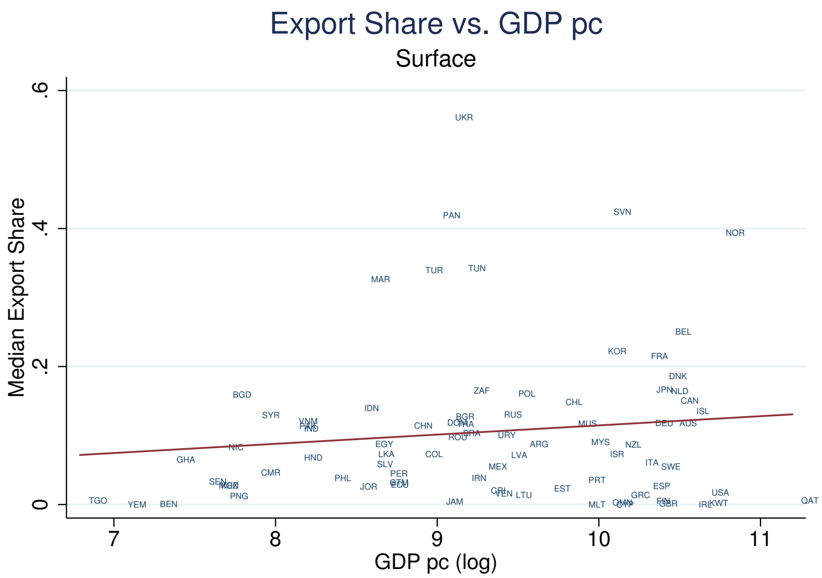

Figure 5: The graph shows the median share of exports by surface transportation against the (log) GDP per capita. The solid line shows the best fit.

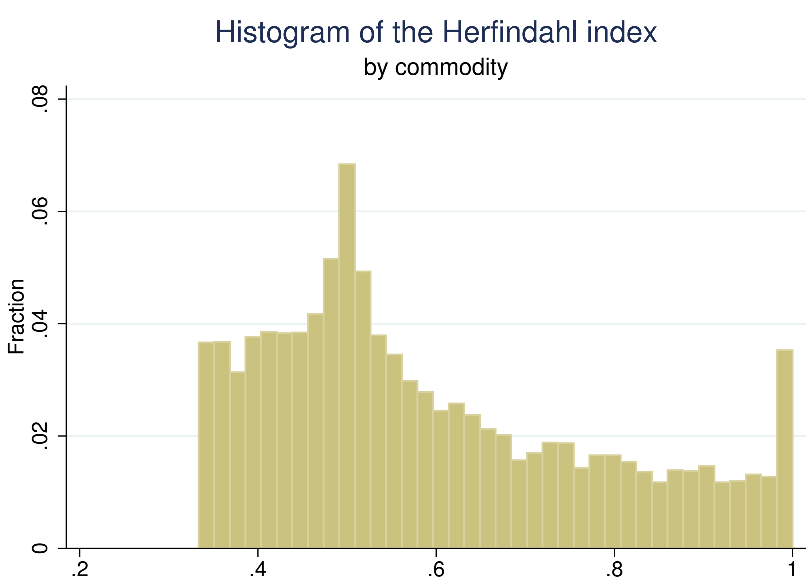

Figure 6: The Herfindahl index of mode-specific US export shares per HS 10 commodity. Data are from the US Exports of Merchandise.

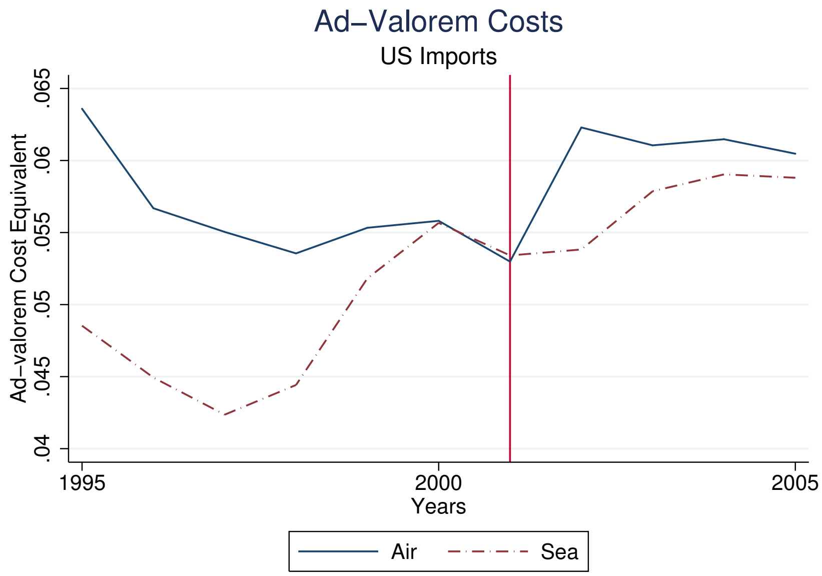

Figure 7: Average ad-valorem transportation charges for imports into the US over time. The data are from the US Imports of Merchandise.

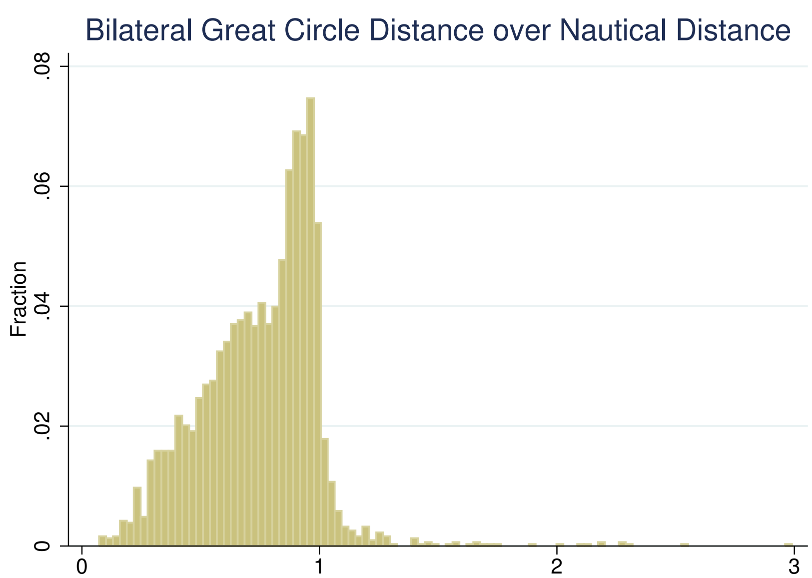

Figure 8: The histogram of the ratio of bilateral great circle distance to bilateral nautical distance.

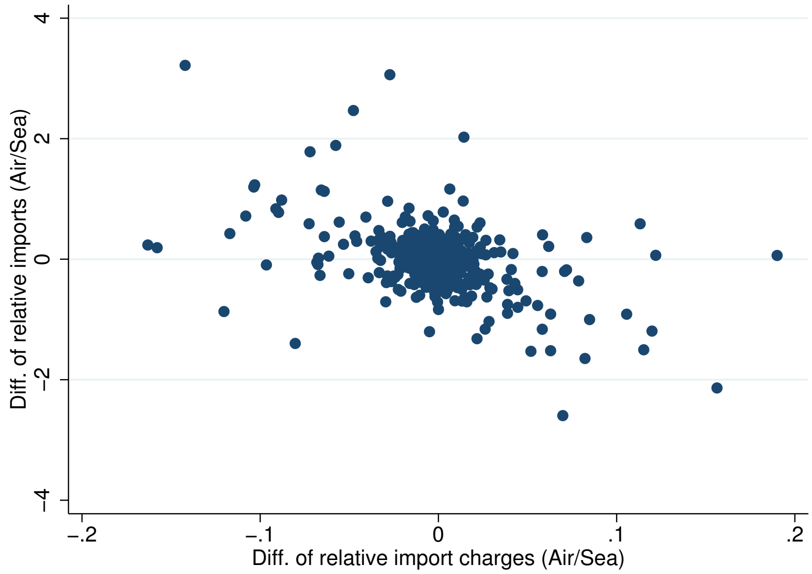

Figure 9: $\Delta \ln \left(\frac{X^a_{i,t}}{X^v_{i,t}} \right)$ versus $\Delta \ln \left(\frac{\delta^a_{i,t}}{\delta^v_{i,t}} \right)$ for 1996 to 2005 (annually) over exporters i. Data are from the US Imports of Merchandise.

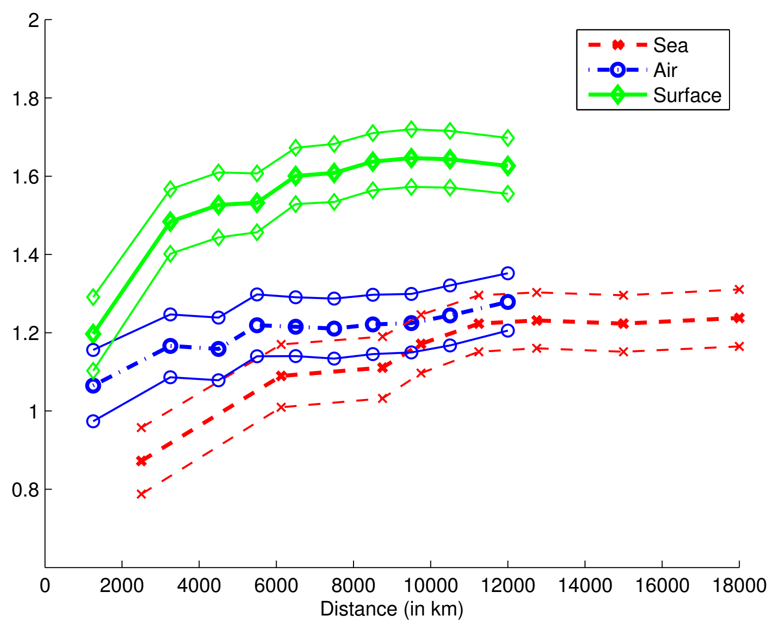

Figure 10: $\hat{\alpha}^m_k$ against the mid point of the k-th distance interval. The bounds around the estimates are 95% confidence intervals based on bootstrapped standard errors.

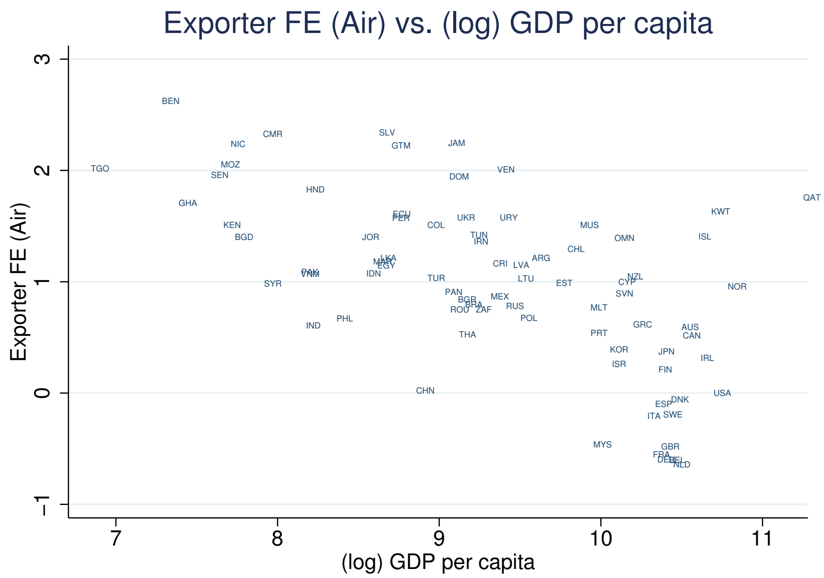

Figure 11: Estimated $\hat{f}_i^a$ against the log of GDP per capita.

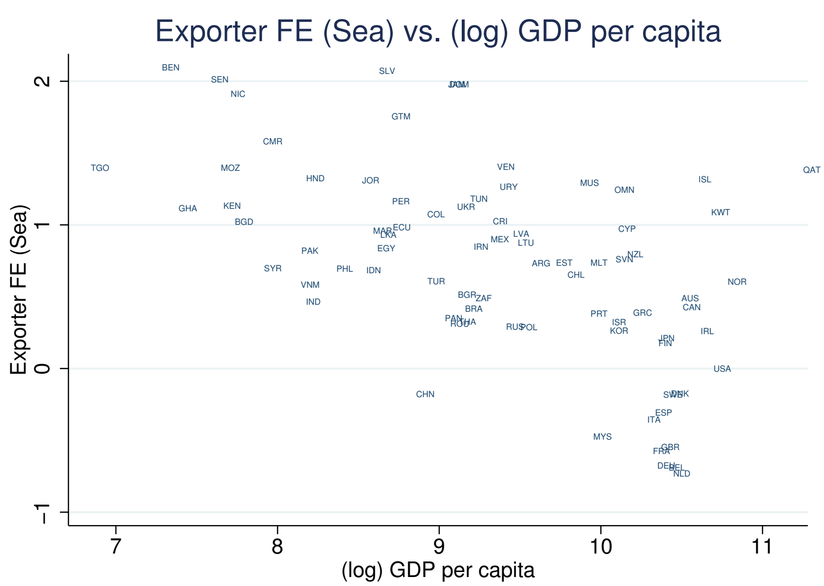

Figure 12: Estimated $\hat{f}_i^v$ against the log of GDP per capita.

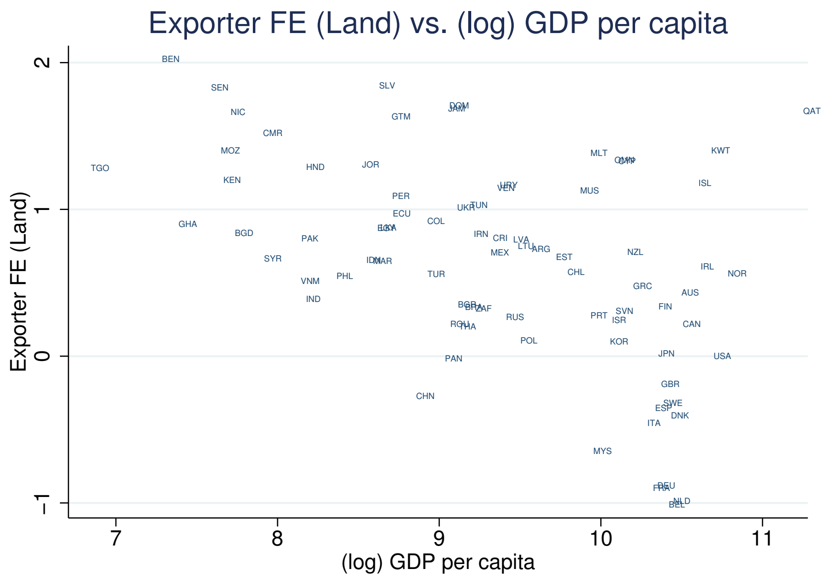

Figure 13: Estimated $\hat{f}_i^s$ against the log of GDP per capita.

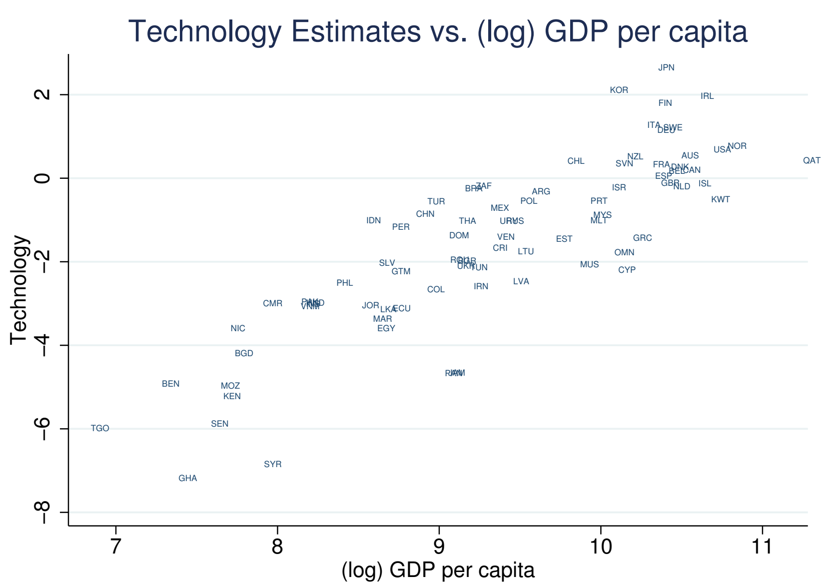

Figure 14: Estimated $\hat{T}_i$ against the log of GDP per capita.

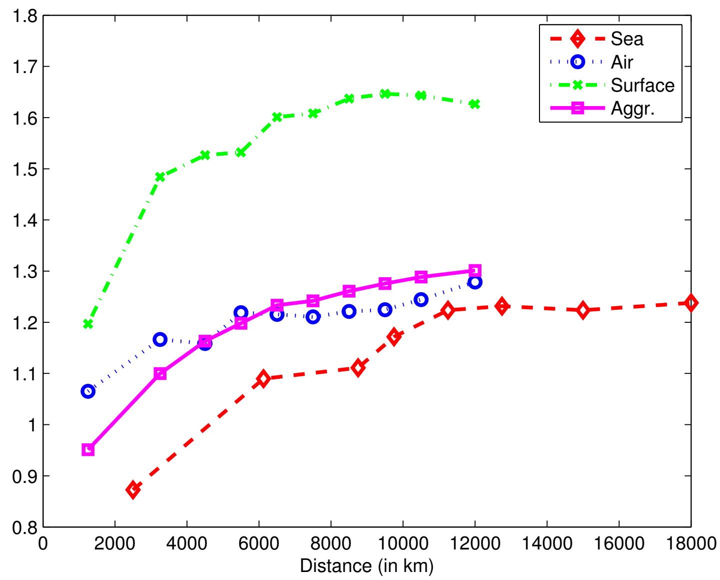

Figure 15: $\hat{\alpha}_k^m$ against the mid point of the k-th distance interval together with the distance estimates from an aggregate model. The distance for air, surface, and the aggregate estimate is the great circle distance, whereas the vessel estimate is plotted against bilateral nautical distance.

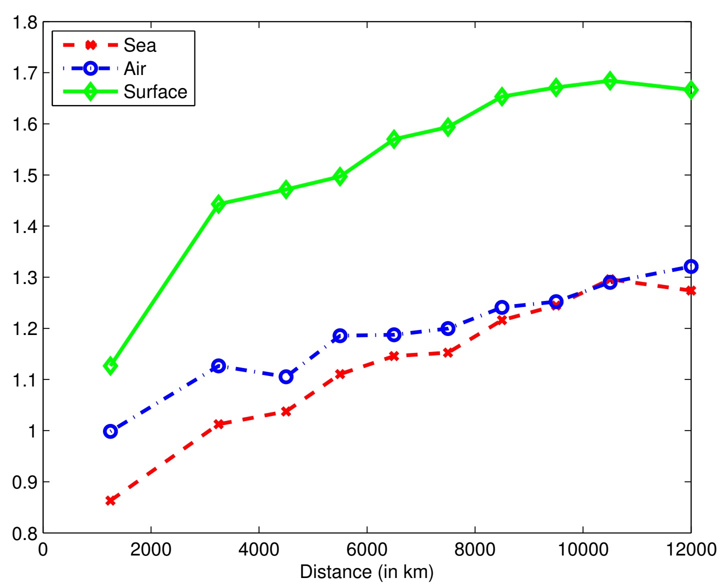

Figure 16: $\hat{\alpha}_k^m$ against the mid point of the k-th distance interval. The estimates are obtained using only great circle distance.

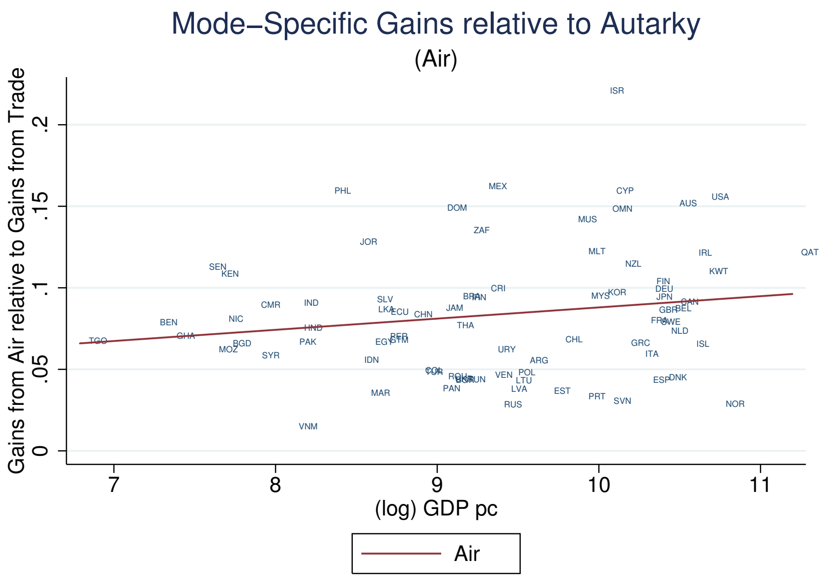

Figure 17: Plotting $\Delta_{i,m}$ for air transportation by country against (log) GDP pc. See (28) for a definition. The straight line is the best fit.

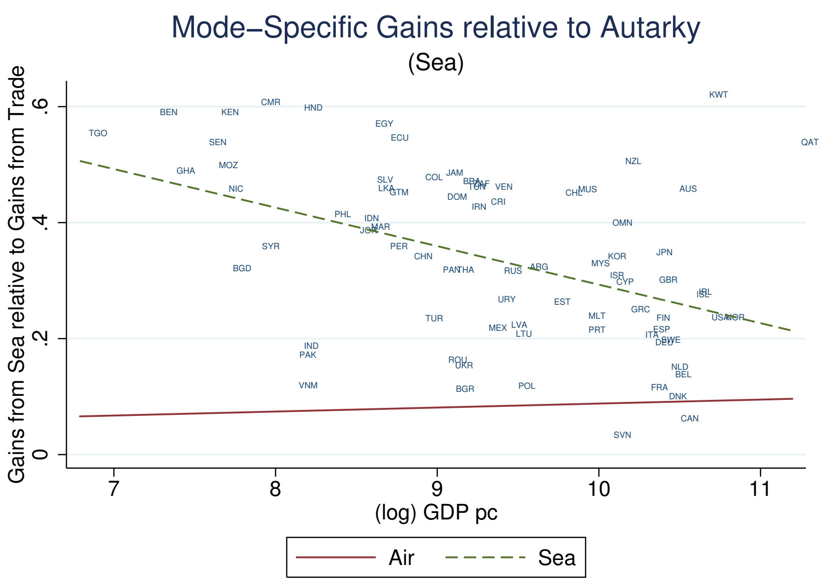

Figure 18: Plotting $\Delta_{i,m}$ for sea transportation by country against (log) GDP pc. See (28) for a definition. The lines are best fit lines.

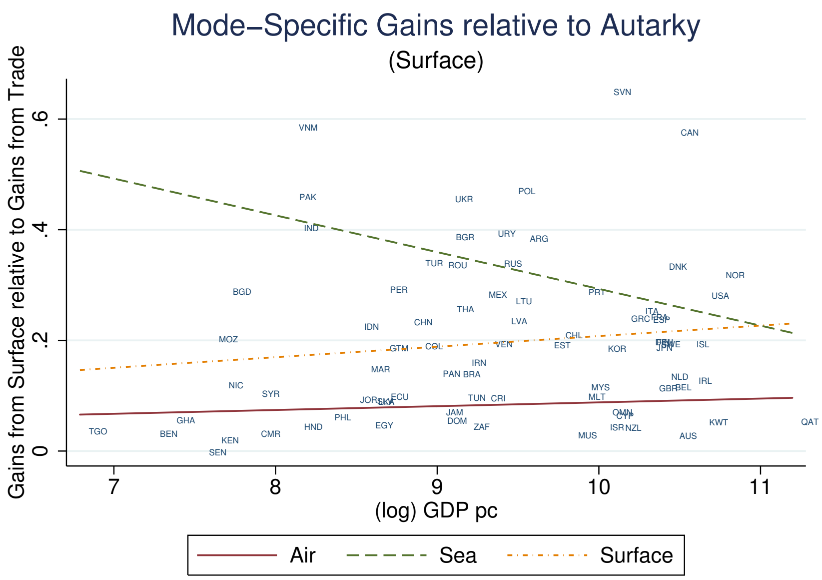

Figure 19: Plotting $\Delta_{i,m}$ for surface transportation by country against (log) GDP pc. See (28) for a definition. The lines are best fit lines.

Table 1: Mode-Specific Bilateral Shares

	Mean	CV	25%	75%
Air	20.8%	115.7%	2.6%	29.7%
Sea	60.3%	49.5%	38.7%	85.7%
Land	18.9%	137.7%	0.5%	26.6%

Notes: Summary statistics of the bilateral mode-specific trade shares for the countries contained in the sample. The sample year is 2005.

Table 2: Naive Gravity Regressions

	The Modes		
Regressors	Air	Sea	Surface
(log) GDP pc Imp.	1.57^{**}	1.28^{**}	1.21^{**}
(log) Pop. Imp.	1.04^{**}	0.89^{**}	0.83^{**}
(log) GDP pc Exp.	2.45^{**}	1.56^{**}	1.35^{**}
(log) Pop. Exp.	1.23^{**}	1.17^{**}	1.13^{**}
$(0, 2500]$	–	–	–
$(2500, 4000]$	-0.45^{*}	-0.84^{**}	-2.35^{**}
$(4000, 5000]$	-0.48^{*}	-1.05^{**}	-2.46^{**}
$(5000, 6000]$	-0.61^{**}	-1.11^{**}	-2.1^{**}
$(6000, 7000]$	-0.53^{**}	-1.54^{**}	-2.65^{**}
$(7000, 8000]$	-0.63^{**}	-1.34^{**}	-2.48^{**}
$(8000, 9000]$	-0.59^{**}	-1.34^{**}	-2.4^{**}
$(9000, 10000]$	-0.47^{*}	-1.4^{**}	-2.34^{**}
$(10000, 11000]$	-0.96^{**}	-1.95^{**}	-2.84^{**}
$(11000, \infty)$	-1.52^{**}	-1.77^{**}	-2.75^{**}
Com. Language	1.31^{**}	0.91^{**}	-0.28
Contiguous	-0.51	-0.72	2.6^{**}
Com. Continent	0.13	0.05	0.92^{**}
Constant	-60.78^{**}	-43.59^{**}	-40.2^{**}
R^2	0.71	0.69	0.52
No. Obs.	2477	2487	2212

Notes: Results of the estimation of (1) using robust standard errors. ** denotes significance at the 1% level, * denotes significance at the 5% level.

Table 3: Value Shares in Exports

	Mode-Specific Share Threshold			
	> 0.995	> 0.99	> 0.98	> 0.95
Value Share in Exports	0.6%	1%	1.8%	4.9%

Notes: The aggregate value share of exported goods with one good-specific mode-specific share above the given threshold. Data are from the US Exports of Merchandise. A good is a HS10 classification.

Table 4: Estimation of $-\theta/(1-\rho)$

	OLS	OLS	Robust	≤ 1999	≥ 2000
$-\frac{\theta}{1-\rho}$	-6.7	-6.7	-6.9	-6.3	-7
	(0.00)	(0.00)	(0.00)	(0.02)	(0.00)
Time Effect	no	yes	no	no	no
No. Obs	775	775	775	308	390
R^2	0.19	0.25	–	0.19	0.22

Notes: Estimates of equation (26). p-values are given in parentheses.

Table 5: Correlations of Predicted Shares and Data

	Modes			
	Air	Sea	Land	Aggregate
Correlation	0.6	0.65	0.46	0.58

Notes: Correlations of predicted bilateral mode-specific trade shares, $\widehat{X^m_{ni}/X_n}$, and data, X^m_{ni}/X_n and for aggregate bilateral trade shares.

Table 6: Trade Cost Estimates

	Contiguity	Common Language	Common Continent
Air	$-21.8\%^{**}$	$-20.6\%^{**}$	$-18.7\%^{**}$
Sea	$-12.3\%^{*}$	$-15.8\%^{**}$	$-10.6\%^{**}$
Surface	$-46.2\%^{**}$	$-11.6\%^{**}$	$-23.9\%^{**}$
Aggr.	$-40.3\%^{**}$	$-16.9\%^{**}$	$-16.5\%^{**}$

Notes: Estimates of the mode-specific trade cost functions expressed as ad-valorem equivalents. The last two rows report the values from estimating an aggregate model with data simulated according to the model. ** denotes significance at the 1% level, * denotes significance at the 5% level. The standard errors are based on bootstrapped standard errors.

Table 7: The Distance Intervals

Air & Surface	Vessel
$[0, 2500)$	$[0, 5000)$
$[2500, 4000)$	$[5000, 7500)$
$[4000, 5000)$	$[7500, 9000)$
$[5000, 6000)$	$[9000, 105000)$
$[6000, 7000)$	$[105000, 12000)$
$[7000, 8000)$	$[12000, 13500)$
$[8000, 9000)$	$[13500, 16500)$
$[9000, 10000)$	$[16500, \infty)$
$[10000, 11000)$	
$[11000, \infty)$	

Notes: The distance intervals used in the paper. Air and surface uses great circle distance, vessel transportation is in nautical distances.

Table 8: Robustness with respect to ρ

	Air	Sea	Land
Mode-Specific Shares ($\rho = 0.4$)	0.6	0.65	0.46
Mode-Specific Shares ($\rho = 0.2$)	0.53	0.61	0.52
Mode-Specific Shares ($\rho = 0.6$)	0.59	0.64	0.48

Notes: Correlations of predicted bilateral mode-specific trade shares, $\widehat{X^m_{ni}/X_n}$, and data, X^m_{ni}/X_n for different values of ρ. $\rho = 0.4$ is the benchmark case.

Table 9: Trade Cost Estimates for Great Circle Distance

	Contiguity	Common Language	Common Continent
Air	$-22\%^{**}$	$-19.9\%^{**}$	$-12\%^{*}$
Sea	$-14.6\%^{**}$	$-15.6\%^{**}$	-5.6%
Surface	$-46\%^{**}$	$-10.9\%^{**}$	$-17.7\%^{**}$

Notes: Estimates of the mode-specific trade cost functions using great circle distance for all modes. The last two columns report the values from estimating an aggregate model with data simulated according to the model. ** denotes significance at the 1% level, * denotes significance at the 5% level, and ^{a} denotes significance at the 10% level. Standard errors are bootstrapped.

Table 10: Calibrated Parameters

	Value	Source
β	0.31	UNIDO
σ	4.5	$\sigma < \theta + 1$
α	0.75	Alvarez and Lucas (2007)

Notes: Values for the parameters of the model that are calibrated.

Table 11: Summary Statistics of the Role of Transportation

		Aggr. Trade	
Counterfactual	Avg. Welfare Change	$\tilde{X}/X$	X^r/X
Air	-0.8%	94.9%	87.8%
Sea	-2.9%	64.7%	44%
Surface	-2.4%	80.2%	68.3%
Autarky	-11.1%	0%	–

Notes: X denotes total world trade and $\tilde{X}$ denotes total world trade in the counterfactual. X^r is the total trade of the unaffected modes prior to the counterfactual.

Table 12: Real Wage Differences relative to Benchmark

Counterfactual	Air	Sea	Surface	All
$f_i^m = f_{US}^m$	66.4%	64.3%	78%	60.7%
$\tau_{ni}^m \equiv 1$	–	–	–	39.3%
$f_i^m = \min_l f_i^l$	99.2%	98.3%	100.5%	99%

Notes: Ratios of the variances of log real wages. The variance of the log of the real wage in the benchmark case is 0.21.

References

ALESSANDRIA, G., J. KABOSKI, AND V. MIDRIGAN (2008): "Inventories, Lumpy Trade, and Large Devaluations," mimeo.

ALVAREZ, F., AND R. E. LUCAS (2007): "General Equilibrium Analysis of the Eaton-Kortum Model of International Trade," *Journal of Monetary Economics*, 54(6), 1726 – 1768.

ANDERSON, J. E., AND E. VAN WINCOOP (2003): "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review*, 93(1), 170 – 192.

——— (2004): "Trade Costs," *Journal of Economic Literature*, XLII, 691 – 751.

ARKOLAKIS, C., A. COSTINOT, AND A. RODRIGUEZ-CLARE (2010): "New Trade Models, Same Old Gains?," mimeo.

BRODA, C., AND D. E. WEINSTEIN (2006): "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, 121(2), 541–585.

CAMERON, A. C., AND P. K. TRIVEDI (2005): *Microeconometrics – Methods and Applications*. New York: Cambridge University Press.

CASELLI, F. (2005): "Accounting for Cross-Country Income Differences," in *Handbook of Economic Growth*, ed. by P. Aghion, and S. Durlauf, chap. 9, pp. 679 – 741. Amsterdam: Elsevier.

CHANEY, T. (2008): "Distorted Gravity: The Intensive and Extensive Margin of International Trade," *American Economic Review*, 98(4), 1707 – 1721.

CLARK, X., D. DOLLAR, AND A. MICCO (2004): "Port Efficiency, Maritime Transport Costs, and Bilateral Trade," *Journal of Development Economics*, 75, 417 – 450.

DEKLE, R., J. EATON, AND S. KORTUM (2008): "Global Rebalancing with Gravity: Measuring the Burden of Adjustment," *IMF Staff Papers*, 55, 511–540.

DISDIER, A.-C., AND K. HEAD (2008): "The Puzzling Persistence of the Distance Effect on Bilateral Trade," *The Review of Economics and Statistics*, 90(1), 37 – 48.

DONALDSON, D. (2008): "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure," mimeo.

EATON, J., AND S. KORTUM (2002): "Technology, Geography, and Trade," *Econometrica*, 70(5), 1741 – 1779.

EVANS, C. L., AND J. HARRIGAN (2005): "Distance, Time, and Specialization: Lean Retailing in General Equilibrium," *American Economic Review*, 95(1), 292 – 313.

FEENSTRA, R., R. LIPSEY, AND H. BOWEN (1997): "World Trade Flows, 1970 - 1992, with Production and Tariff Data," NBER Working Paper 5910.

FEYRER, J. (2009): "Trade and Income - Exploiting Time Series in Geography," NBER Working Paper 14910.

FIELER, A. C. (2009): "Non-Homotheticity and Bilateral Trade: Evidence and a Quantitative Explanation," mimeo.

FOGEL, R. W. (1964): *Railways and American Economic Growth*. Baltimore: Johns Hopkins Press.

HARRIGAN, J. (2009): "Airplanes and Comparative Advantage," mimeo.

HELPMAN, E., M. MELITZ, AND Y. RUBINSTEIN (2008): "Estimating Trade Flows: Trading Partners and Trading Volumes," *Quarterly Journal of Economics*, 123(2), 441 – 487.

HESTON, A., R. SUMMERS, AND B. ATEN (2009): "Penn World Table 6.3," Discussion paper, Center for International Comparisons of Production, Income, and Prices at the University of Pennsylvania.

HUMMELS, D. (2001): "Time as a Trade Barrier," mimeo.

——— (2007): "Transportation Costs and International Trade in the Second Era of Globalization," *Journal of Economic Perspectives*, 21(3), 131 – 154.

HUMMELS, D., AND V. LUGOVSKYY (2006): "Are Matched Partner Trade Statistics a Usable Measure of Transportation Costs?," *Review of International Economics*, 14(1), 69 – 86.

HUMMELS, D., V. LUGOVSKYY, AND A. SKIBA (2009): "The Trade Reducing Effects of Market Power in International Shipping," *Journal of Development Economics*, 89(1), 84 – 97.

HUMMELS, D., AND G. SCHAUR (2010): "Hedging Price Volatility Using Fast Transport," *Journal of International Economics*, 82(1), 15–25.

LIMAO, N., AND A. J. VENABLES (2001): "Infrastructure, Geographical Disadvantage, Transport Costs, and Trade," *World Bank Economic Review*, 15(3), 451 – 479.

LUX, M. (2010): "The Mode of Transportation in International Trade," mimeo.

MASKUS, K. (1991): "Comparing International Trade Data and Product and National Characteristics Data for the Analysis of Trade Models," in *International Economic Transactions: Issues in Measurement and Empirical Research*, ed. by P. Hooper, and J. Richardson, vol. 55 of *National Bureau of Economic Research Studies in Income and Wealth*. Chicago: University of Chicago Press.

MELITZ, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71(6), 1695 – 1725.

MICCO, A., AND T. SEREBRISKY (2006): "Competition Regimes and Air Transport Costs: The Effects of Open Skies Agreements," *Journal of International Economics*, 70(1), 25 – 51.

NELSEN, R. B. (2006): *An Introduction to Copulas*. New York: Springer, 2nd edn.

OVERMAN, H. G., S. J. REDDING, AND A. J. VENABLES (2004): "The Economic Geography of Trade, Production, and Income: A Survey of Empirics," in *Handbook of International Trade*, ed. by E. K. Choi, and J. Harrigan, chap. 12, pp. 353 – 387. Hoboken: Wiley-Blackwell.

RAMONDO, N., AND A. RODRIGUEZ-CLARE (2009): "Trade, Multinational Production, and the Gains from Openness," mimeo.

SAMUELSON, P. A. (1954): "The Transfer Problem and Transport Costs, II: Analysis of Effects of Trade Impediments," *Economic Journal*, 64(254), 264–289.

SIMONOVSKA, I., AND M. E. WAUGH (2009): "The Elasticity of Trade: Estimates and Evidence," mimeo.

VON THUENEN, J. H. (1826): *Der isoierte Staat in Beziehung auf Landwirtschaft und Nationaloekonomie*. Hamburg: Perthes.

WAUGH, M. E. (2010): "International Trade and Income Differences," *American Economic Review*, 100(5), 2093–2124.

WORLD BANK (2009): *World Development Report: Reshaping Economic Geography*. Washington D.C.: The World Bank.