

Geography, Transportation, and Endogenous Trade Costs

Giulia Brancaccio*, Myrto Kalouptsi†, Theodore Papageorgiou‡§

Cornell University, Harvard University, CEPR and NBER, Boston College

July 2019

In this paper we study the role of the transportation sector in world trade. We build a spatial model that centers on the interaction of the market for (oceanic) transportation services and the market for world trade in goods. The model delivers equilibrium trade flows, as well as equilibrium trade costs (shipping prices). Using detailed data on vessel movements and shipping prices, we document novel facts about shipping patterns; we then flexibly estimate our model. We use this setup to demonstrate that the transportation sector (i) attenuates differences in the comparative advantage across countries; (ii) generates network effects in trade costs; and (iii) dampens the impact of shocks on trade flows. These three mechanisms reveal a new role for geography in international trade that was previously concealed by the frequently-used assumption of exogenous trade costs. Finally, we illustrate how our setup can be used for policy analysis by evaluating the impact of future and existing infrastructure projects (e.g. Northwest Passage, Panama Canal).

Keywords: transportation, trade costs, geography, matching function estimation, shipping, network effects, import-export complementarity, trade imbalances, trade elasticity, maritime infrastructure

*Cornell University, Department of Economics, 464 Uris Hall, Ithaca NY 14850, USA, giulia.brancaccio@cornell.edu

†Harvard University, Department of Economics, Littauer Center 124, Cambridge MA 02138, USA, myrto@fas.harvard.edu

‡Boston College, Department of Economics, Maloney Hall 343, Chestnut Hill, MA, 02467, USA, theodore.papageorgiou@bc.edu

§We are thankful to Pol Antras, Costas Arkolakis, Aris Christou, Manolis Galenianos, Paul Grieco, Ken Hendricks, Fabian Lange, Robin Lee, Marc Melitz, Guido Menzio, Ariel Pakes, Nicola Rosaia and Jim Tybout for their many helpful comments. We also gratefully acknowledge financial assistance from the Griswold Center for Economic Policy Studies at Princeton University. This paper previously circulated under the title “Geography, Search Frictions and Endogenous Trade Costs” (NBER Working Paper 23581). The material pertaining to search frictions has been moved to the companion paper “Efficiency in Decentralized Markets”.

1 Introduction

Whether by sea, land or air, the entirety of trade in goods is carried out by the transportation sector. With world trade at full steam, the transportation sector has become central in everyday life. Yet, little is known about how the market for transportation services interacts with the market for world trade in goods.

In this paper, we study how this interaction shapes trade flows, trade costs, the propagation of shocks and the allocation of productive activities across countries. As we demonstrate, the transportation sector (i) attenuates differences in the comparative advantage across countries, reallocating production from net exporters to net importers; (ii) creates network effects in trade costs; and (iii) dampens the impact of shocks on trade flows. These three mechanisms reveal a new role for geography in international trade that was previously shrouded by the common assumption of exogenous trade costs.

The transportation sector includes several different segments which can be split into two categories: those that operate on fixed itineraries, much like buses, and those that operate on flexible routes, much like taxis. Containerships, airplanes and trains primarily belong to the first group, while trucks, gas and oil tankers, and dry bulk ships to the second. Here we focus on oceanic shipping and, in particular, dry bulk shipping; 80% of world trade volume is carried by ships and dry bulkers carry about half of that.¹ Dry bulk ships are the main mode of transportation for commodities, such as grain, ore, and coal. They are often termed the “taxis of the oceans,” as an exporter has to search for an available vessel and hire it for a specific voyage, with prices set in the spot market in a decentralized fashion. Despite the operational differences of the various transport modes, we argue that the core economic mechanisms discussed in this paper hold for most, if not all of them.

We leverage detailed micro-data on vessel movements, as well as rich data on contracts between exporters and shipowners to uncover some novel facts. First, satellite data of ships’ movements reveal that most countries are either large net importers or large net exporters; related to this, at any point in time a staggering 42% of ships are traveling without cargo (termed “ballast”). This natural trade imbalance is a key driver of trade costs. Indeed, transportation prices are largely asymmetric and depend on the destination’s trade imbalance: all else equal, the prospect of having to ballast after offloading leads to higher prices. For instance, shipping from Australia to China is 30% more expensive than the reverse: as China mostly imports raw materials, ships arriving there have limited opportunities to reload. This

¹Source: International Chamber of Shipping and UNCTAD (2015). Seaborne trade accounts for about 70% of trade in terms of value.

phenomenon is pervasive in most, if not all modes of transportation: trucks, trains, container air and ocean shipping, all exhibit similar price asymmetries that correlate with trade imbalances (the direction of the imbalance, however, may be the opposite). In fact, the US-China trade deficit in manufacturing has incentivized US exports of low value cargo, such as scrap or hay, to fill up the empty backhauls.²

In addition, we compute the elasticity of trade with respect to shipping costs, by regressing trade by country pair on the corresponding shipping price. To do so, we employ a novel instrument inspired by the insight that the shipping price exporters face depends on how attractive their destination is to the ship in terms of future loading opportunities, as discussed above. The estimated trade elasticity indicates that the transport sector has a substantial impact on world trade, especially given the large fluctuations in shipping prices.

Inspired by these facts, we build a spatial model that centers on the interaction of the market for transport and the market for world trade in goods, in the spirit of the search and matching literature. The globe is split into a number of regions that trade with each other. Geography enters the model both through regions' location in space, as well as their natural inheritance in commodities of different value. In each region, available ships and exporters participate in a random matching process. When matched with an exporter, ships transport the exporter's cargo to its destination for a negotiated price, and restart there. Ships that do not get matched decide whether to wait at their current location or ballast elsewhere to search there. Exporters that get matched have their cargo delivered and collect its revenue, while exporters that do not get matched wait at port. Finally, a large number of potential exporters decide whether and where to export, thus replenishing the exporter pool seeking transportation.

We derive the equilibrium trade costs (shipping prices), as well as an expression for the equilibrium bilateral trade flows, that is reminiscent of a gravity equation. As ships are forward-looking, trade costs depend on the attractiveness of both the origin and the destination, a rich object that captures a region's location, freight values, matching probabilities, as well as its neighbors' attractiveness. This insight applies beyond dry bulk. Although other transport modes require different modeling assumptions regarding their

²In 2005, about 60% of the containers sent via ships from Asia to North America came back empty (Drewry Consultants), and those "that did come back full were often transported at a steep discount for lack of demand (...) Shippers are so eager to fill their vessels for the return voyage to East Asia that they accept many types of unprofitable cargo, like bales of hay." Similarly, "airlines had become so eager to put something in their cargo holds on the inbound journey to China that rates go as low as 30 to 40 cents a kilogram, compared with \$3 to \$3.50 a kilogram leaving China [...] Very bluntly speaking, they're flying in empty and flying out full." (The International Herald Tribune, 01/30/2006)

In trucking, "[b]ackhaul practices are extremely important in explaining differences in prices; i.e. the truck companies are compensating their expenses on the empty backhaul in the first leg of the trip." (The World Bank, 2012). Even Uhaul moving trucks are priced similarly, "Rent a moving truck from Las Vegas to San Jose and you'll pay about \$100. In the opposite direction, the same truck will cost you 16 times that. (...) What accounts for the difference? (...) With so many people leaving the Bay Area, there are not enough rental trucks to go around." (SFgate, 02/15/2018).

operational practices, in equilibrium, prices are formed by the optimizing behavior of forward-looking transport agents and unavoidably depend, as above, on the attractiveness of origins and destinations, as well as that of their neighbors.³

Next, we estimate the model using the collected data. We first estimate the matching function capturing the trading process between ships and exporters, which gives the number of matches as a function of the number of agents searching on each side of the market. A sizable literature has estimated matching functions in different contexts (e.g. labor markets, taxicabs).⁴ Here, we adopt a novel approach to flexibly recover both the matching function, as well as searching exporters, which, unlike ships and matches, are not observed. Our approach draws from the literature on nonparametric identification (Matzkin, 2003) and, to our knowledge, we are the first to apply it to matching function estimation. Not imposing a functional form is important, since the shape of the underlying matching function is directly linked to welfare (see Brancaccio et al. (2019b)); furthermore, this approach is agnostic as to the nature of the meeting process thus allowing us to recover exporters flexibly.

We then estimate the remaining primitives including ship costs, the values of exporters’ cargo and exporter entry costs. In particular, we recover ship sailing and port costs from the optimal ballast choice probabilities, via Maximum Likelihood, following the dynamic discrete choice literature (Rust, 1987). Then, we obtain exporter valuations directly from observed prices. Finally, we use trade flows to recover exporter costs by destination.

Why is it important to account for the transport sector to study international trade in goods? We illustrate the role of endogenous trade costs through three experiments.

First, we compare our setup to one with “iceberg” trade costs that are exogenous and depend only on distance and the cargo’s value, as is the case in canonical trade models. We find that the transport sector mitigates differences in the comparative advantage across countries, reducing world trade imbalances. Indeed, under endogenous trade costs net exporters (importers) export less (more) than under exogenous trade costs, leading to a reallocation of productive activities from net exporters to net importers. This happens because of ships’ equilibrium behavior and in particular the strength of their bargaining position at different regions. Net exporters offer loading opportunities to ships, thus allowing them to command

³For instance, a model for container shipping (or other modes operating on fixed itineraries), would likely feature a small number of shipping firms that decide on their itineraries, as well as route pricing. The itinerary decision involves similar trade-offs as the ballast choice here; e.g. what network of countries can be serviced or how to manage backhaul trips. Prices similarly depend on the number and valuations of exporters, ship supply and crucially, the entire route serviced including the backhaul trips.

⁴For instance, in labor markets, data on unemployed workers, vacancies and matches delivers the underlying matching function. In the market for taxi rides, one observes taxi rides, but not hailing passengers; in recent work, Buchholz (2019) and Frechette et al. (2019) have used a parametric assumption on the matching function, to recover the passengers.

high prices, which in turn restrains their exports. The converse holds for net importers. This argument extends to a country's neighbors: a net exporter close to other net exporters offers even more options to ships and prices are even higher, which inhibits the neighborhood's exports.

Second, we illustrate that the transport sector dampens the impact of shocks on trade flows by considering a fuel cost shock. A decline in the fuel cost has a direct and an indirect effect. The direct effect is straightforward: as costs fall, shipping prices also fall and thus exports rise. The novel indirect effect is that a decline in fuel costs improves a ship's bargaining position, as it makes ballasting less costly. This dampens the original decline in prices and the increase in exporting. Indeed, the overall increase in world trade would have been 40% higher if ships were not allowed to optimally adjust their behavior, in response to a 10% decline in fuel costs.

Third, we explore the spatial propagation of a macro shock: a slow-down in China. Besides the direct effect to countries whose exports rely heavily on the Chinese economy, the optimal reallocation of ships over space differentially filters the shock in neighboring vs. distant regions. Because of the slow-down, fewer ships offload in China, reducing ship supply in the region. Although this impacts negatively China's own exports by raising prices, it benefits distant countries, such as Brazil, because ships reallocate there.

Finally, we consider the role of maritime infrastructure on world trade, as an illustration of how our setup can be used for policy evaluation. To do so, we examine the opening of the Northwest Passage: the melting of the arctic ice would reduce the travel distance between Northeast America/Northern Europe and the Far East. Although the shock is local, it has global effects: as Northeast America becomes a more attractive ballasting choice, ships have a stronger bargaining position and demand higher prices, pushing exports down everywhere else. Moreover, we consider the impact of three natural and man-made passages: the Panama Canal, the Suez Canal and Gibraltar and show that all passages substantially increase world trade and welfare.

Related Literature We relate to three broad strands of literature: (i) trade and geography; (ii) search and matching; (iii) industry dynamics.

First, our paper endogenizes trade costs and so it naturally relates to the large literature in international trade studying the importance of trade costs in explaining trade flows between countries (e.g. Anderson and Van Wincoop, 2003, Eaton and Kortum, 2002). In much of the literature, trade costs are treated as exogenous and follow the iceberg formulation of Samuelson (1954). Here, we consider what happens to trade flows when the equilibrium of the transport market is taken into account, so that trans-

port prices (an important component of trade costs, at least as large or larger than tariffs; Hummels, 2007) are determined in equilibrium, jointly with trade flows. Moreover, related to some of our empirical findings, Waugh (2010) has argued that asymmetric trade costs are necessary to explain some empirical regularities regarding trade flows across rich and poor countries.

We also contribute to a literature that has considered the role and features of the (container) shipping industry; e.g. Koopmans (1949) contains an early consideration of endogenous shipping prices; Hummels and Skiba (2004) explore the relationship between product prices at different destinations and shipping costs; Hummels et al. (2009) explore market power in container shipping; Ishikawa and Tarui (2015) theoretically investigate the impact of “backhaul” and its interaction with industrial policy; Cosar and Demir (2018) and Holmes and Singer (2018) study container usage; Asturias (2018) explores the impact of the number of shipping firms on transport prices and trade; Wong (2019) incorporates container shipping prices featuring a “round-trip” effect in a trade model, similar to the model of Behrens and Picard (2011). Finally, recent work has explored the matching of importers and exporters under frictions (Eaton et al., 2016, and Krolkowski and McCallum, 2018).

Our paper is also related to both older and more recent work on the role of geography in international trade (e.g. Krugman, 1991, Head and Mayer, 2004, Allen and Arkolakis, 2014), as well as the impact of transportation infrastructure and networks (e.g. Donaldson, 2018, Allen and Arkolakis, 2016, Donaldson and Hornbeck, 2016, Fajgelbaum and Schaal, 2017). We extend this literature by demonstrating that the transport sector reveals a new role for geography through three novel mechanisms (it attenuates differences in the comparative advantage across countries, creates network effects and dampens the impact of shocks).

Second, our paper relates to the search and matching literature (see Rogerson et al., 2005 for a survey). Our model is a search model in the spirit of the seminal work of Mortensen and Pissarides (1994), where firms and workers (randomly) meet subject to search frictions and Nash bargain over a wage. An important addition in our case is the spatial nature of our setup: there are several interconnected markets at which agents (ships) can search. Such a spatial search model was first proposed by Lagos (2000) (and analyzed empirically in Lagos, 2003) in the context of taxi cabs. We borrow heavily from his model; the key difference is that prices are set in equilibrium, while in the taxi market prices are exogenously set by regulation. This is crucial, as the role of endogenous trade costs is at the core of our paper.⁵ Finally,

⁵There are several other differences between our setup and that of Lagos (2000) and Lagos (2003): we model also the demand side (exporters/passengers); we allow for the potential of frictions in each region, while Lagos (2000) assumes that matching is frictionless locally; we allow for several sources of heterogeneity in different regions (travel/port costs, distances, matching rates); we allow trade to be imbalanced, while Lagos (2000) relies on taxi flows in and out of each location to be equal- this distinction is also crucial.

as discussed above, our paper also contributes to the literature on matching function estimation (see Petrongolo and Pissarides (2001) for a survey).

Third, we relate to the literature on industry dynamics (e.g. Hopenhayn, 1992, Ericson and Pakes, 1995). Consistent with this research agenda, we study the long-run industry equilibrium properties, in our case the spatial distribution of ships and exporters. Moreover, our empirical methodology borrows from the literature on the estimation of dynamic setups (e.g. Rust, 1987, Bajari et al., 2007, Pakes et al., 2007; applications include Ryan, 2012 and Collard-Wexler, 2013). Buchholz (2019) and Frechette et al. (2019) also explore dynamic decisions in the context of taxi cabs' search and shift choices respectively. Finally, Kalouptsi (2014) has also looked at the shipping industry, albeit at the entry decisions of shipowners and the resulting investment cycles in new ships, while Kalouptsi (2018) focuses on industrial policy in the Chinese shipbuilding industry.

The rest of the paper is structured as follows: Section 2 provides a description of the industry and the data used. Section 3 presents the facts. Section 4 describes the model. Section 5 lays out our empirical strategy, while Section 6 presents the estimation results. Section 7 demonstrates the importance of endogenous trade costs, while Section 8 assesses the role of maritime infrastructure projects. Section 9 concludes. The Online Appendix contains additional tables and figures, proofs to our propositions, as well as further data and estimation details.

2 Industry and Data Description

2.1 Dry Bulk Shipping

Dry bulk shipping involves vessels designed to carry a homogeneous unpacked dry cargo, for individual shippers on non-scheduled routes. Bulk carriers operate much like taxi cabs: a specific cargo is transported individually by a specific ship, for a trip between a single origin and a single destination. Dry bulk shipping involves mostly commodities, such as iron ore, steel, coal, bauxite, phosphates, but also grain, sugar, chemicals, lumber and wood chips; it accounts for about half of total seaborne trade in tons (UNCTAD, 2015) and 45% of the total world fleet, which includes also containerships and oil tankers.^{6,7}

There are four categories of dry bulk carriers based on size: Handysize (10,000–40,000 DWT), Handy-

⁶As already mentioned, bulk ships are different from containerships, which carry cargo (mostly manufactured goods) from many different cargo owners in container boxes, along fixed itineraries according to a timetable. It is not technologically possible to substitute bulk with container shipping.

⁷It is not straightforward to obtain information on the share of world trade value carried by bulkers. However, mining, agricultural products, chemicals and iron/steel jointly account for about 30% of total trade value (WTO, 2015).

max (40,000–60,000 DWT), Panamax (60,000–100,000 DWT) and Capesize (larger than 100,000 DWT). The industry is unconcentrated, consisting of a large number of small shipowning firms (Kalouptsi, 2014): the maximum fleet share is around 4%, while the firm size distribution features a large tail of small shipowners. Moreover, shipping services are largely perceived as homogeneous. In his lifetime, a shipowner will contract with hundreds of different exporters, carry a multitude of different products and visit numerous countries.

Trips are realized through individual contracts: shipowners have vessels for hire, exporters have cargo to transport and brokers put the deal together. Ships carry at most one freight at a time: the exporter fills up the hired ship with his cargo. In this paper, we focus on spot contracts and in particular the so-called “trip-charters”, in which the shipowner is paid in a per day rate.⁸ The exporter who hires the ship is responsible for the trip costs (e.g. fueling), while the shipowner incurs the remaining ship costs (e.g. crew, maintenance, repairs).

2.2 Data

We combine a number of different datasets. First, we employ a dataset of dry bulk shipping contracts, from 2010 to 2016, collected by Clarksons Research. An observation is a transaction between a shipowner and a charterer for a specific trip. We observe the vessel, the charterer, the contract signing date, the loading and unloading dates, the price in dollars per day, as well as some information on the origin and destination.

Second, we use satellite AIS (Automatic Identification System) data from exactEarth Ltd (henceforth EE) for the ships in the Clarksons dataset between July 2010 and March 2016. AIS transceivers on the ships automatically broadcast information, such as their position (longitude and latitude), speed, and level of draft (the vertical distance between the waterline and the bottom of the ship’s hull), at regular intervals of at most six minutes. The draft is a crucial variable, as it allows us to determine whether a ship is loaded or not at any point in time.

We also use the ERA-Interim archive, from the European Centre for Medium-Range Weather Forecasts (CMWF), to collect global data on daily sea weather. This allows us to construct weekly data on the wind speed (in each direction) on a 0.75° grid across all oceans. Finally, we employ several time series from Clarksons on e.g. the total fleet and fuel prices, as well as country-level imports/exports, production and commodity prices from other sources (e.g. Comtrade, IEA).

⁸Trip-charters are the most common type of contract. Long-term contracts (“time-charters”), however, do exist: on average, about 10% of contracts signed are long-term.

	Mean	St. Dev	Median	Min	Max
Contract price per day (thousand US dollars)	13.9	8.6	12	3.2	43
Contract trip price (thousand US dollars)	291	304	178	30.5	1,367
Contracts per ship	2.9	2.2	2	1	10
Trip duration (weeks)	2.89	1.36	2.95	0.5	5.44
Days between contract signing and loading date	6.39	7.12	5	0	30
Prob of ship staying at port conditional on not signing a contract	0.77	0.12	0.76	0.59	0.95

Table 1: Summary statistics. The contract price per day is reported by Clarksons. To create the price per trip, we multiply price per day with the average number of days required to perform the trip. Contracts per ship counts the number of contracts observed for each ship in the Clarksons dataset. To proxy for trip duration, we compute the nautical distance in miles and divide it by the average speed observed in the EE data. The probability of staying at port is calculated from the EE data by computing the frequency at which waiting ships that did not find a contract in a given week remain at port instead of ballasting elsewhere. We have 12,007 observations of shipping contracts and 393,058 ship-week observations at which the ship decides to either ballast someplace or stay at its current location.

Summary statistics Our final dataset involves 5,398 ships, which corresponds to about half the world fleet, between 2012 and 2016.⁹ We end up with 12,007 shipping contracts, for which we know the price, as well as the exact origin and destination (see the Online Appendix for our data matching procedure).¹⁰ As shown in Table 1, the average price is 14,000 dollars per day (or 290,000 dollars for the entire trip), with substantial variation. Trips last on average 2.9 weeks. Contracts are signed close to the loading date, on average six days before. The most popular loaded trips are from Australia, Brazil and Northwest America to China, while the most popular ballast trip is from China to Australia (5.7% of ballast trips). We have 393,058 ship-week observations at which the ship decides to either ballast someplace or stay at its current location. Ships that do not sign a contract, remain in their current location with probability 77%, while the remaining ships ballast elsewhere. Finally, Clarksons reports the product carried in about 20% of the sample. The main commodity categories are grain (29%), ores (21%), coal (25%), steel (8%) and chemicals/fertilizers (6%).

3 Facts

In this section, we present some novel facts about the transport sector and trade: we first discuss the implications of trade imbalances (section 3.1) and then, we quantify the impact of transport costs on world

⁹We drop the first two years (until May 2012) of vessel movement data, as satellites are still launched at that time and the geographic coverage is more limited.

¹⁰The Clarksons contracts somewhat oversample the intermediate size categories (Handymax and Panamax) and younger ships.

trade (section 3.2). Throughout the paper, unless otherwise noted, we split ports into 15 geographical regions, depicted in Figure 1 in the Online Appendix.¹¹

3.1 Trade Imbalances

World trade in commodities is greatly imbalanced. Indeed, most countries are either large net importers or large net exporters. This is shown in Figure 1, which plots the difference between the number of ships departing loaded and the number of ships arriving loaded, over the sum of the two. Australia, Brazil and Northwest America are big net exporters, whereas China and India are big net importers.

The fact that global trade features such large imbalances is not as surprising; to a large extent, this is due to the different natural inheritance of countries. For instance, Australia, Brazil and Northwest America are rich in minerals, grain, coal, etc. At the same time, growing developing countries require imports of raw materials to achieve industrial expansion and infrastructure building. For instance, in recent years, Chinese growth has relied on massive imports of raw materials. As a result, commodities flow out of producers such as Australia and Brazil, towards China and India.

As a consequence of the imbalanced nature of international trade, ships spend much of their time traveling ballast, i.e. without cargo. Indeed, we find that 42% of a ship’s traveled miles are ballast, so that a ship is traveling empty close to half the time.¹² Finally, the trade imbalance is a key driver of the trade costs that exporters face. First, a quick inspection of the data reveals that there are large asymmetries in trade costs across space: for instance, a trip from China to Australia costs on average 7,500 dollars per day, while a trip from Australia to China costs on average 10,000 dollars per day.¹³ In fact, most trips exhibit substantial asymmetry: the average ratio of the price from i to j over the price from j to i (highest over lowest), is 1.6 and can be as high as 4.1.

We further investigate the determinants of trade costs by considering how shipping prices are associated with the attractiveness of the destination, such as its demand for shipping. Indeed, ships may demand a premium to travel towards a destination with low exports (e.g. China), to compensate for the difficulty of finding a new cargo originating from that destination. As shown in Column III of Table 2, shipping to a destination where the probability of a ballast trip afterwards is ten percentage points higher, costs 2.3%

¹¹The trade-off is that we need a large number of observations per region, while allowing for sufficient geographical detail. To determine the regions, we employ a clustering algorithm that minimizes the within-region distance between ports. The regions are: West Coast of North America, East Coast of North America, Central America, West Coast of South America, East Coast of South America, West Africa, Mediterranean, North Europe, South Africa, Middle East, India, Southeast Asia, China, Australia, Japan-Korea. We ignore intra-regional trips and entirely drop these observations.

¹²This percentage is lower for smaller ships— it is 32% for Handysize, 41% for Handymax, 45% for Panamax and 49% for Capesize.

¹³This price asymmetry has been documented also in container shipping; see e.g. Wong (2019) and references therein.




Figure 1: Trade Imbalances. Difference between exports (ships leaving loaded) and imports (ships arriving loaded) over total trade (all ships). A positive (negative) ratio indicates that a country is a net exporter (importer); a ratio close to zero implies balanced trade.

more on average. Similarly, a 10% increase in the average distance traveled ballast after the destination, is associated with a 1.7% increase in prices.

Finally, it is worth noting that the type of product carried affects the price paid and overall more valuable goods lead to higher contracted prices, as shown in Column III Table 2.

3.2 Trade Elasticity

Do shipping prices have an impact on world trade? In this section we address this question in the context of bulk shipping. Ideally, we would like to regress bilateral trade flows on shipping prices, i.e.,

$$\log Q_t^{i \rightarrow j} = \beta_0 + \beta_1 \log \tau_t^{i \rightarrow j} + \epsilon_{ijt}$$

where $Q_t^{i \rightarrow j}$ is the total trade value from country i to country j (in bulk commodities) at time period (month) t and $\tau_t^{i \rightarrow j}$ is the shipping price from i to j at t . Naturally, this regression is going to lead to biased estimates, as prices are likely correlated with the error, ϵ_{ijt} . Thus, an instrument is required.

The instrument we leverage is inspired by the insight that, as discussed above, the attractiveness of an exporter's destination impacts the shipping price it faces. Consider the trade flow from i to j , $Q_t^{i \rightarrow j}$; the instrument we use for the shipping price $\tau_t^{i \rightarrow j}$ consists of the tariffs levied on commodity exports *from* the destination j . For example, the price to ship goods from Indonesia to China is instrumented using the tariffs on raw materials on routes starting from China. These tariffs do not directly affect the flows

	I	II	III
	log(price per day)		
Probability of ballast		0.234** (0.030)	0.556** (0.081)
Avg duration of ballast trip (log)		0.166** (0.014)	0.065** (0.032)
Coal			0.088** (0.045)
Fertilizer			0.245** (0.051)
Grain			0.131** (0.048)
Ore			0.124** (0.045)
Steel			0.135** (0.049)
Constant	10.284** (0.103)	9.127** (0.099)	8.915** (0.408)
Destination FE	Yes	No	No
Origin FE	Yes	Yes	Yes
Ship type FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
Obs	11,014	11,011	1,662
R ²	0.694	0.674	0.664

** $p < 0.05$, * $p < 0.1$

Table 2: Shipping price regressions. The dependent variable is the logged price per day in USD. The independent variables include combinations of: the average frequency of ballast traveling after the contract's destination (Probability of ballast), the average logged duration (in days) of the ballast trip after the contract's destination, as well as ship type, origin, destination and quarter fixed effects (FE). The product is reported in only 20% of the sample, so the regression in column III has substantially fewer observations. The omitted product category is cement.

from Indonesia to China. However, they affect the value of a ship unloading in China. Indeed, tariffs on j 's exports lead to a reduction in shipments from j , thus dampening the demand for shipping services in j and making j a less attractive destination for ships. Therefore, when negotiating a price to ship goods from i to j , a ship demands a higher price in order to compensate for its reduced opportunities upon arrival at j .¹⁴ Similarly, we also use the tariffs levied on commodities imported at the exporter's origin, i , as an instrument for the price $\tau_t^{i \rightarrow j}$. These tariffs reduce i 's imports, leading to lower ship supply in origin i , and, thus, higher shipping prices to export from i to j .

Table 3 presents the results. Column I showcases this mechanism by regressing the per-day shipping prices on the tariffs levied on exports from j to its first and second biggest trading partners (tariff $j \rightarrow (1)$ and tariff $j \rightarrow (2)$), as well as on tariffs on i 's imports from its first and second biggest trading partners (tariff $(1) \rightarrow i$ and tariff $(2) \rightarrow i$). We run the regression in differences to control for any fixed, route-specific characteristics; we also control for GDP, tariffs on the route considered, as well as tariffs on all goods other than commodities (all in differences).¹⁵ The signs of the instruments are as expected: higher tariffs tend to increase shipping prices. These results are interesting per se, as they showcase that shipping prices between any two countries are affected by shipping conditions on other routes, creating inter-dependencies and network effects in trade costs; this mechanism, which is formalized in our model and is central to the paper, is quantitatively important. Finally, the F-stat of this regression suggests that we may have a weak instrument problem. Hence to obtain the trade elasticity below we use the LIML estimator, which has better finite sample properties than TSLS.

The estimated trade elasticity with respect to shipping prices, shown in Column II of Table 3, is equal to 1.03 and is statistically significant. In other words, a 1% increase in shipping prices leads to a 1.03% decline in trade flows. The corresponding Anderson and Rubin (1949) confidence interval is [-1.397, -0.769]. This elasticity indicates that the transport sector has a substantial impact on world trade, especially given the large observed fluctuations in shipping prices (for instance, shipping prices experienced an 8-fold increase in the late 2000s, see Kalouptsi, 2014).

A few recent papers have estimated the same elasticity for the case of container shipping. Asturias (2018), who uses population as an instrument, finds the elasticity to be about 5. Wong (2019), who uses

¹⁴This instrument is valid as it should not impact directly $Q_t^{i \rightarrow j}$. Recall that here we focus only on raw materials, hence the supply chain should not be a concern (e.g. the instrument would be problematic if j imports steel and exports cars and we considered tariffs on cars). Moreover, we control directly for the tariffs from i to j and the overall level of tariffs on all goods other than commodities.

¹⁵We obtain yearly country-level trade flows from Comtrade and tariffs from the World Bank (WITS) and we focus only on bulk commodities; yearly average shipping prices come from our Clarksons dataset. The results are robust if we add country fixed effects, or if we use the weighted average of tariffs instead.

	$\Delta \log \left(\tau_t^{i \rightarrow j} \right)$	$\Delta \log \left(Q_t^{i \rightarrow j} \right)$
	OLS	LIML
$\Delta \log \left(\tau_t^{i \rightarrow j} \right)$		-1.03^{**} (0.425)
$\Delta \log \left(\text{tariff}_t^{j \rightarrow (1)} \right)$	0.070* (0.040)	
$\Delta \log \left(\text{tariff}_t^{j \rightarrow (2)} \right)$	0.135** (0.027)	
$\Delta \log \left(\text{tariff}_t^{(1) \rightarrow i} \right)$	0.136 (0.096)	
$\Delta \log \left(\text{tariff}_t^{(2) \rightarrow i} \right)$	-0.012 (0.082)	
GDP of i and j		
Controls (changes of)	tariff on i 's import (non-commodities)	
	tariff on j 's export (non-commodities)	
	tariff on j 's import from i (commodities)	
Obs	452	452
R ²	0.129	-
F-stat	7.30	

$^{**} p < 0.05, ^{*} p < 0.1$

Table 3: Elasticity of trade with respect to shipping prices. Data on yearly bilateral country-level trade value and tariffs are obtained from the World Bank (WITS) for the period 2010-2016. We focus on trade and tariffs for bulk commodities. To construct tariffs, we consider the minimum between the most favored nation tariff and preferential rates, if applicable, and consider a weighted average across commodities. Shipping prices are calculated from the per-day prices in Clarksons contracts, averaged at the year and country-pair level. We group countries in EU-27 and exclude countries without no access to sea.

the round-trip effect as an instrument (in particular, for route i, j she uses a Bartik-style instrument to proxy for the predicted trade volume on route j, i) finds an elasticity of about 3. It is also worth comparing the elasticity of trade with respect to shipping prices to that with respect to tariffs, which is estimated to be between 1.5 and 5 on average (e.g. Simonovska and Waugh, 2014, Caliendo and Parro, 2015, Brandt et al., 2017 and Arkolakis et al., 2018). The two elasticities are comparable, with our estimate overall somewhat lower. Recall, however, that we use total rather than waterborne trade value; therefore our estimate should be considered a lower bound of the trade elasticity with respect to shipping prices, as it ignores substitution towards other modes of transportation.¹⁶

4 Model

In this section, motivated by the above findings, we introduce a spatial model that centers on the interaction between the market for transport and the market for world trade in goods. In each period, the timing is as follows: In each region, available ships and exporters participate in a decentralized matching process. Ships that get matched transport their exporter’s cargo to its destination for a negotiated price, and restart there. Ships that do not get matched decide whether to wait at their current location or ballast elsewhere and search there. Exporters that get matched have their cargo delivered and collect their revenue. Exporters that do not get matched wait at port. Finally, a large number of potential exporters decide whether and where to export, thus replenishing the exporter pool seeking transportation the following period.

We first lay out the model’s setup; we then present the agents’ value functions and derive the equilibrium objects of interest: trade costs (shipping prices) and trade flows (gravity equation). We close the section with a detailed discussion of our main assumptions.

4.1 Environment

Time is discrete. There are I locations/regions, $i \in \{1, 2, \dots, I\}$. There are two types of agents, exporters and ships. Both are risk neutral and have discount factor β .

At each location i and period t , there are e_{it} exporters/freights that need to be delivered to another location. An exporter obtains revenue (or valuation), r , from shipping the good. Every period, at each location i , $\mathcal{E}_i$ potential exporters decide whether and where to export. If they decide to export, they pay

¹⁶For instance, if we exclude EU countries, which can easily substitute from oceanic shipping to land shipping, the elasticity increases to 3.

production and export costs, κ_{ij} and draw their revenue r , from a distribution F_{ij}^r with mean $\bar{r}_{ij}$.

There are S homogeneous ships in the world.¹⁷ In every period, a ship is either at port in some region i , or it is traveling loaded or ballast, from some location i to some location j . A ship at port in location i incurs a per period waiting cost c_i^w , while a ship sailing from i to j incurs a per period sailing cost c_{ij}^s . The duration of a trip between region i and region j is stochastic: a traveling ship arrives at j in the current period with probability d_{ij} , so that the average duration of the trip is $1/d_{ij}$.¹⁸

Freights can only be delivered to their destination by ships and each ship can carry (at most) one freight. Following the search and matching literature, we model new matches every period, m_{it} , using a matching function, whereby the number of matches at time t in region i is

$$m_{it} = m_i(s_{it}, e_{it})$$

where s_{it} is the number of unmatched ships in region i . $m_i(s_{it}, e_{it})$ is increasing in both arguments. Let λ_{it} denote the probability that an unmatched ship in location i meets an exporter; $\lambda_{it} = m_{it}/s_{it}$. Similarly, let λ_{it}^e denote the probability with which an unmatched exporter meets a ship; $\lambda_{it}^e = m_{it}/e_{it}$. The matching function captures the frictional trading process in a parsimonious fashion. In other words, we do not explicitly model the meeting technology between exporters and ships, which “would introduce intractable complexities” (Petrongolo and Pissarides, 2001); instead, the matching function captures the several realities of the market, including information frictions, port infrastructure and heterogeneities. In a companion paper (Brancaccio et al., 2019b), we argue that the meeting process is indeed not frictionless and provide evidence of search frictions in bulk shipping that limit trade.

When a ship and an exporter meet, they either agree on a price to be paid by the exporter to the ship or they both revert to their outside options. The outside option of the exporter is to remain unmatched and wait for another ship, while the outside option of the ship is to either remain unmatched in the current region or to ballast elsewhere. The surplus of the match over the parties’ outside options is split via the price-setting mechanism. The price, τ_{ijr} , paid to the ship delivering a freight of valuation r from region i to destination j , is determined by generalized Nash bargaining, with $\gamma \in (0, 1)$ denoting the exporter’s bargaining power. The price is paid upfront and the ship commits to begin its voyage immediately to j .

Ships that remain unmatched decide whether to remain in their current region or ballast elsewhere

¹⁷We follow Kalouptsi (2014) and assume constant returns to scale so that a shipowner is a ship.

¹⁸It is straightforward to have deterministic trip durations instead. Our specification, however, preserves tractability and allows for some variability e.g. due to weather shocks, without affecting the steady state properties of the model.

subject to i.i.d. logit shocks. Exporters that remain unmatched survive with probability $\delta > 0$ and wait in their current region.

4.2 Equilibrium

The state variable of a ship in region i includes its current location i , as well as the state (s_t, e_t) , where $e_t = [e_{1t}, \dots, e_{It}]$ denotes the distribution of exporters over all regions, and s_t is an $I \times I$ matrix including the ships that are traveling from i to j , s_{ijt} , as well as the ships at port s_{it} , $i, j = 1, \dots, I$. The state variable of an exporter in i includes his location i , valuation r and destination j , as well as (s_t, e_t) . In this paper, we consider the steady state of our industry model, following the tradition of Hopenhayn (1992). More specifically, agents view the spatial distribution of ships and exporters, (s_t, e_t) , as fixed and make decisions based on its steady-state value.

Ships Let V_{ij} denote the value of a ship that starts the period traveling from i to j (empty or loaded), V_i the value of a ship that starts the period at port in location i , and U_i the value of a ship that remained unmatched at i at the end of the period (we suppress the dependence on the steady state values (s, e)). Then:

$$V_{ij} = -c_{ij}^s + d_{ij}\beta V_j + (1 - d_{ij})\beta V_{ij} \quad (1)$$

In words, the ship that is traveling from i to j , pays the per period cost of sailing c_{ij}^s ; with probability d_{ij} it arrives at its destination j , where it begins unmatched with value V_j ; with the remaining probability $1 - d_{ij}$ the ship does not arrive and keeps traveling.

A ship that starts the period in region i obtains:

$$V_i = -c_i^w + \lambda_i E_{j,r} (\tau_{ijr} + V_{ij}) + (1 - \lambda_i)U_i \quad (2)$$

In words, the ship pays the per period port wait cost c_i^w ; it gets matched with probability λ_i , in which case it receives the agreed upon price, τ_{ijr} , and begins traveling. The ship takes expectation over the type of exporter it meets, i.e. its revenue and destination. With the remaining probability, $1 - \lambda_i$, the ship does not find an exporter and gets the value of being unmatched U_i .

If the ship remains unmatched, it faces the choice of either staying at i or ballasting to another region; in the latter case, the ship can choose among all possible destinations. The unmatched ship's value

function is:

$$U_i(\epsilon) = \max \left\{ \beta V_i + \sigma \epsilon_i, \max_{j \neq i} V_{ij} + \sigma \epsilon_j \right\} \quad (3)$$

where the shocks $\epsilon \in \mathbb{R}^I$ are drawn from a type I Extreme Value (Gumbel) distribution, with standard deviation σ . In words, if the ship stays in its current region i , it obtains value V_i ; otherwise the ship chooses another region and begins its trip there.

Let P_{ii} denote the probability that a ship in location i chooses to remain there, and P_{ij} the probability it chooses to ballast to j . We have:

$$P_{ii} = \frac{\exp(\beta V_i / \sigma)}{\exp(\beta V_i / \sigma) + \sum_{l \neq i} \exp(V_{il} / \sigma)} \quad (4)$$

and

$$P_{ij} = \frac{\exp(V_{ij} / \sigma)}{\exp(\beta V_i / \sigma) + \sum_{l \neq i} \exp(V_{il} / \sigma)}. \quad (5)$$

Exporters We now turn to the value functions of exporters; we begin with existing exporters and then consider exporter entry. An exporter that is matched in location i receives his revenue, r and pays the agreed price, τ_{ijr} for a total payoff of $r - \tau_{ijr}$. The value of an exporter that remains unmatched, U_{ijr}^e , is therefore given by

$$U_{ijr}^e = \beta \delta \left[\lambda_i^e (r - \tau_{ijr}) + (1 - \lambda_i^e) U_{ijr}^e \right] \quad (6)$$

In words, the exporter receives no payoff in the period and survives with probability δ ; if so, the following period with probability λ_i^e he gets matched and receives $r - \tau_{ijr}$, while with the remaining probability $1 - \lambda_i^e$ he remains unmatched again.

Each potential entrant, makes a discrete choice between destinations, as well as not exporting, also subject to i.i.d. shocks $\epsilon^e \in \mathbb{R}^I$, distributed according to a type I Extreme Value (Gumbel) distribution. Therefore, a potential entrant solves:

$$\max \left\{ \epsilon_0^e, \max_{j \neq i} \left\{ E_r U_{ijr}^e - \kappa_{ij} + \epsilon_j^e \right\} \right\}$$

where we denote by 0 the option of not exporting and normalize the payoff in that case to zero.

Potential exporters' behavior is given by the choice probabilities:

$$P_{ij}^e \equiv \frac{\exp(U_{ij}^e - \kappa_{ij})}{1 + \sum_{l \neq i} \exp(U_{il}^e - \kappa_{il})} \quad (7)$$

and

$$P_{i0}^e \equiv \frac{1}{1 + \sum_{l \neq i} \exp(U_{il}^e - \kappa_{il})} \quad (8)$$

where $U_{ij}^e \equiv E_r U_{ijr}^e$. Therefore, the number of entrant exporters in i equals $\mathcal{E}_i (1 - P_{i0}^e)$.

Trade Costs (Shipping Prices) As discussed above, the rents generated by a match between an exporter and a ship, are split via Nash bargaining. This implies the surplus sharing condition:

$$\gamma [(\tau_{ijr} + V_{ij}) - U_i] = (1 - \gamma) [(r - \tau_{ijr}) - U_{ijr}^e] \quad (9)$$

where $U_i \equiv E_\epsilon U_i(\epsilon)$. We use this condition to solve out for the equilibrium price τ_{ijr} , in the following lemma:

Lemma 1. *The agreed upon price between a ship and an exporter with valuation r and destination j in location i is given by:*

$$\tau_{ijr} = (1 - \mu_i) (U_i - V_{ij}) + \mu_i r \quad (10)$$

where $\mu_i = (1 - \gamma) (1 - \beta\delta) / (1 - \beta\delta (1 - \gamma\lambda_i^e))$.

Proof. Substitute U_{ijr}^e in (9). □

In other words, the price is a convex combination of the exporter's revenue, r , and the difference between the ship's value of transporting the freight, V_{ij} , and its outside option, U_i . Consistent with the evidence in Table 2 exporters that have a higher value, r , pay higher prices.

Crucially, the price depends on ships' equilibrium behavior through the value of traveling from i to j , V_{ij} (which in turn depends on V_j), as well as the value of the "outside option," U_i . These objects are very rich, as they capture the attractiveness of both the origin i , as well as the destination j , which consists of numerous features. For instance, destinations that are unappealing to ships because there are few exporters relative to ships and the probability of ballasting afterwards is high, command higher prices (consistent with the evidence presented in Table 2). The same holds for destinations that are further away (low d_{ij}), have low value exporters or low matching probabilities. Moreover, V_j controls for conditions at all possible ballast destinations from j , as well as for conditions at all possible export destinations from j , revealing the importance of network effects. Similarly, U_i controls for the attractiveness of the origin (e.g. exporter revenues, nearby ballast opportunities, matching probability). As a result, the price between i

and j depends on all countries, rather than just i and j .¹⁹

Trade Flows From equation (7) total flows from i to j equal

$$\mathcal{E}_i P_{ij}^e = \mathcal{E}_i \frac{\exp(U_{ij}^e - \kappa_{ij})}{1 + \sum_{l \neq i} \exp(U_{il}^e - \kappa_{il})} = \mathcal{E}_i \frac{\exp(\alpha_i (\bar{r}_{ij} - \tau_{ij}) - \kappa_{ij})}{1 + \sum_{l \neq i} \exp(\alpha_i (\bar{r}_{il} - \tau_{il}) - \kappa_{il})}$$

where $\bar{r}_{ij}$ is the average revenue from exporting from i to j , $\tau_{ij} \equiv E_r \tau_{ijr}$ and $\alpha_i = \beta \delta \lambda_i^e / (1 - \beta \delta (1 - \lambda_i^e))$. To obtain this expression, we solve for U_{ijr}^e from (6) to obtain $U_{ijr}^e = \alpha_i (r - \tau_{ijr})$.

This equation is a “gravity equation”; it delivers the trade flow (in quantity rather than value) from i to j as a function of two components. First, the primitives $\{\lambda_i^e, \bar{r}_{ij}, \kappa_{ij}, \mathcal{E}_i\}$ not just for i and j but for all regions; this is reminiscent of the analysis in Anderson and Van Wincoop (2003) who show that the gravity equation in a trade model needs to include a country’s overall trade disposition.

Second, it is a function of the endogenous trade costs, τ_{ij} , for all j , which are the key addition here. In this model, trade costs introduce network effects between countries: indeed, τ_{ij} depends on all locations both through the outside option of the ship at the origin i , U_i , as well as the ballast and export opportunities from the destination j , captured by V_j . Overall, any change in the primitives affects trade flows both directly, but also indirectly through its impact on trade costs. We illustrate the importance of this mechanism in Section 7.

Steady State Equilibrium A steady state equilibrium consists of a distribution of ships and exporters over locations (s^*, e^*) , ship choices P_{ij} , exporter choices P_{ij}^e , and prices τ_{ijr} , that satisfy the following conditions:

- (i) Ship optimal behavior, P_{ij} follows (4) and (5)
- (ii) Potential exporter behavior, P_{ij}^e , follows (7) and (8)
- (iii) Prices τ_{ijr} are determined by Nash bargaining, according to (10)
- (iv) Ships and exporters satisfy the steady state equations (established in the proof of Proposition 1 in the Online Appendix):

$$s_i^* = \sum_j P_{ji} (s_j^* - m_j(s_j^*, e_j^*)) + \sum_{j \neq i} \frac{P_{ji}^e}{1 - P_{j0}^e} m_j(s_j^*, e_j^*) \quad (11)$$

¹⁹It is worth noting how the model’s main outputs would change if the matching function were assumed to be $m_i = \min\{s_i, e_i\}$, so that the market were frictionless. With more ships than exporters, the shipping price is given by $\tau_{ij} = U_i - V_{ij}$, as the price has to be such that ships are indifferent between loading and going to destination j and remaining unmatched. Therefore the properties of endogenous trade costs (dependence on distance, origin, destination and entire network of countries) are independent of the presence of search frictions and still hold.

$$e_i^* = \delta (e_i^* - m_i(s_i^*, e_i^*)) + \mathcal{E}_i(1 - P_{i0}^e) \quad (12)$$

$$s_{ij}^* = \frac{1}{d_{ij}} \left(P_{ij}(s_i^* - m_i(s_i^*, e_i^*)) + \frac{P_{ij}^e}{1 - P_{i0}^e} m_i(s_i^*, e_i^*) \right) \quad (13)$$

As shown in Proposition 1 in the Online Appendix, there exist (s^*, e^*) that satisfy equations (11) through (13).

4.3 Discussion

We close this section with a discussion on several of our assumptions and some caveats.

We begin our discussion with the *matching process*. In our model the matching function is local, so that an exporter meets a ship only if they are in the same region, much like taxis and passengers. This is a modeling assumption, as there is no technological or other constraint that prevents an exporter from meeting and matching with a ship in another region. Nonetheless, there are economic disincentives that make distant matching unlikely, suggesting this is a reasonable approximation.

More specifically, practitioners explain that contracts tend to be signed with ships that are nearby, by arguing that “a ship is not a train” and it cannot promise exact arrival times far in advance due to weather conditions and port congestion. These delays are costly for exporters. Moreover, ships that are already in the region of the exporter have a distinct cost advantage over ships in other regions, since they do not need to incur the additional cost of sailing empty to the exporter’s region. Given that ships are in oversupply during our time period, exporters are not willing to pay (and wait) in order to contract with ships that are far away.²⁰ Reassuringly, the data support the local matching function assumption. For instance, about 20% of the contracts specify different signing and loading regions. Furthermore, as shown in Table 1, contracts are signed just 6 days on average prior to the loading date. In addition, the satellite data reveals that ships enter the region within 12 days of loading, which is well before the signing date.

Also related to the matching process, we assume that exporter valuations are sufficiently high so that in equilibrium, when a ship and an exporter meet, they always agree to form a match. It is easy to see that for every origin-destination pair, there exists a threshold of exporter value, below which the match surplus becomes negative and meetings do not result in matches. In this case, the price that a ship demands to stop searching is too high for the low value exporters to pay and the match becomes

²⁰We formalize this using the model estimates produced in Section 6. In particular, we use our model to examine whether an exporter would benefit from searching in multiple regions. We find that when we allow an exporter to search in the “best” other region, in addition to the loading region, exporters always prefer to match with a ship in the loading region. The large majority of matches (>80%) still take place in the loading region. In addition, the price paid and the exporter value function are virtually unchanged.

unprofitable. Thus, the support of the distribution of revenues, F_{ij}^r , needs to be bounded below by this threshold. This assumption is reasonable in our context, since exporter revenue is an order of magnitude higher than transport prices. Using the model estimates produced in Section 6, we find that the surplus from matching is fairly high; on average it equals 371,270\$, and it remains high even if we focus on the lowest priced commodity, namely coal (the average surplus in this case is \$249,150, well above zero). This is not surprising: to get a sense of the magnitudes, the average value of a coal cargo is \$3,864,250, while its average shipping price is \$297,720.

It is also worth noting that we model ships as *ex-ante homogenous agents*. Indeed, the data suggests that ship heterogeneity is very limited. Ships do not specialize neither geographically, nor in terms of products: the majority of ships deliver cargo to 13 out of 15 regions and carry at least 2 of the 3 main products (coal, ore and grain). Moreover, neither shipowner characteristics, nor shipowner fixed effects have any explanatory power in price regressions, as shown in Table 1 in the Online Appendix, while ballast decisions of ships in the same region are concentrated around the same options. Finally, home-ports are not an important consideration for shipowners, as the crew flies to their home country every 6-8 months.

We now turn to the *steady state* assumption of our industry model, which follows the tradition of Hopenhayn (1992) and a large body of other work in macro, IO, trade, etc. In particular, we consider the steady state of our dynamic system, where agents' strategies and value functions depend only on their own state and the long-run average aggregate state, which is constant. This assumption renders the problem more tractable and allows us to derive simple(r) expressions for prices and trade flows.

It is worth connecting this assumption to the empirical exercise coming up. Naturally, the data exhibits variation over time. Strictly speaking, our empirical approach relies on the assumption that agents “play against the steady state”; i.e. their strategies rely on their own state and the long-run industry average.²¹ That said, the steady state assumption is not unreasonable for the data at hand, which covers a period that is uniformly characterized by ship oversupply and relatively low demand for shipping services without any major shocks. Moreover, given the short-lived nature of the ships' ballasting decisions it does not feel unreasonable that they would ignore aggregate long-run shocks when making these weekly choices; in addition, transition dynamics to a new steady state should be short given that a ship can travel to most ports in the world in a month.²²

In this paper, we do not model *ship entry and exit*; exit is overall very small, while due to long

²¹This is reminiscent of the interpretation of the Oblivious Equilibrium (OE) concept of Weintraub et al. (2008).

²²To test whether ships respond to transitory shocks, we check if weather shocks affect ships' ballast decisions and find that they do not.

construction lags in shipbuilding (two to six years), the fleet is fixed in the short run; see Kalouptsi (2014, 2018). Our counterfactual results therefore hold in the short/medium run.²³ We also do not model speed choice (see Adland and Jia, 2018).

Finally, in the trade literature, trade costs include transport costs, tariffs and other barriers. Here we focus on microfounding and endogenizing transport costs. In our setup, other trade barriers are included in κ_{ij} , the cost of entry into exporting. In addition, we do not consider the determination of commodity prices; in other words, we take exporter valuations r to be exogenous. Determining this object in equilibrium within our setup is an interesting avenue for future research.

5 Estimation

In this section we lay out the empirical strategy followed to estimate the model of Section 4. Our empirical exercise consists of two distinct components: (i) estimation of the matching function and the searching exporters; (ii) estimation of ship travel and wait costs, exporter valuations and costs. We describe the empirical strategy for each component here, and present all results in Section 6.

5.1 Matching Function Estimation

A sizable literature has estimated matching functions in several different contexts (e.g. labor markets, marriage markets, taxicabs). For instance, in labor markets, one can use data on unemployed workers, job vacancies and matches to recover the underlying matching function. In the market for taxi rides, one observes taxi rides, but not hailing passengers; in recent work, Buchholz (2019), and Frechette et al. (2019) have used such data, coupled with a “parametric” assumption on the matching function to recover the hailing passengers.²⁴ In our data we observe ships and matches, but not searching exporters. Here, we adopt a novel approach to simultaneously recover both exporters, as well as a nonparametric matching function. Our approach contrasts with the literature that has imposed functional forms. Since parametric restrictions are directly linked to welfare, as shown for instance in Hosios (1990), they can be overly restrictive; for a detailed discussion of these issues see Brancaccio et al. (2019b). In addition, parametric assumptions on the matching function indirectly impose restrictions on the distribution of (unobserved)

²³It is straightforward to include a ship free entry condition in the model in order to consider longer-run counterfactuals. However in this case, we would need to take a stand on shipowners’ expectations of future demand.

²⁴Buchholz (2019) assumes an “urn-ball” matching function. Frechette et al. (2019) construct a numerical simulation of taxi drivers that randomly meet passengers over a grid that resembles Manhattan; this spatial simulation essentially corresponds to the matching function, and can be inverted to recover hailing passengers.

exporters that we recover, whereas our approach allows for more flexibility.²⁵

Suppose we have a sample $\{s_{it}, m_{it}\}_{t=0}^T$ for each market i . The unknowns of interest are the I matching functions $m_i(\cdot)$ and the exporters e_{it} , for all i, t ; henceforth, we suppress the i subscript to ease notation. Our approach relies on the literature on nonparametric identification (Matzkin, 2003) and nonseparable instrumental variable techniques (e.g. Imbens and Newey, 2009).

We assume that (i) the matching function $m(s, e)$ is continuous and strictly increasing in e ; (ii) the matching function exhibits constant returns to scale (CRS), so that $m(as, ae) = am(s, e)$ for all $a > 0$ and there is a known point $\{\bar{s}, \bar{e}, \bar{m}\}$, such that $\bar{m} = m(\bar{s}, \bar{e})$; (iii) the random variables s and e are independent. Assumption (i) is natural, as more exporters lead to more matches, all else equal. Assumption (ii) is a restriction that guarantees identification of both sets of unknowns and is discussed further below. Assumption (iii) is made for expositional purposes and is relaxed later on.

Let $F_{m|s}$ denote the distribution of matches conditional on ships, and F_e the distribution of exporters, e . Then at a given point $\{s_t, e_t, m_t\}$ we have:

$$\begin{aligned} F_{m|s=s_t}(m_t|s=s_t) &= \Pr(m(s, e) \leq m_t | s=s_t) \\ \text{monotonicity} &= \Pr(e \leq m^{-1}(s, m_t) | s=s_t) \\ \text{independence} &= \Pr(e \leq m^{-1}(s_t, m_t)) \\ &= F_e(e_t) \end{aligned} \tag{14}$$

The equation $F_e(e_t) = F_{m|s=s_t}(m_t|s=s_t)$ forms the basis for identification and estimation. Indeed note that this equation, along with the CRS assumption, allows us to recover the distribution $F_e(e)$, for all e : using the known point $\{\bar{s}, \bar{e}, \bar{m}\}$ and letting $a = e/\bar{e}$, for all e ,

$$\begin{aligned} F_e(a\bar{e}) &= F_{m|s=a\bar{s}}(m(a\bar{s}, a\bar{e}) | s=a\bar{s}) \\ &= F_{m|s=a\bar{s}}(a\bar{m} | s=a\bar{s}) \end{aligned} \tag{15}$$

We use (15) and vary a to trace out $\hat{F}_e(e)$, relying on a kernel density estimator for the conditional distri-

²⁵In Brancaccio et al. (2019a) we provide a detailed practitioner's guide to this approach in this and other contexts. For an application of this methodology in the context of labor markets see Lange and Papageorgiou (2019).

bution $\hat{F}_{m|s=a\bar{s}}(a\bar{m}|s=a\bar{s})$.²⁶ Once the distribution $\hat{F}_e$ is recovered, we obtain the number of exporters e_t from

$$e_t = \hat{F}_e^{-1} \left(\hat{F}_{m|s=s_t}(m_t|s=s_t) \right),$$

and the matching function at *any* point (s, e) from

$$m(s, e) = \hat{F}_{m|s}^{-1}(F_e(e)).$$

We choose the known point, $\{\bar{s}, \bar{e}, \bar{m}\}$, to be of the form $1 = m(\bar{s}, 1)$, so that one exporter is always matched when there are $\bar{s}$ ships. We set $\bar{s}$ iteratively, to be the lowest value such that $m_t \leq e_t$, for all t , thus obtaining a conservative bound on search frictions.

The intuition behind the identification argument is as follows: the observed correlation between s and m informs us on $\partial m(s, e)/\partial s$, since the sensitivity of matches to changes in ships is observed and s is independent of e by assumption; then, due to homogeneity, this derivative also delivers the derivative $\partial m(s, e)/\partial e$; and once these derivatives are known, integration leads to the matching function, which can be inverted to provide the number of exporters.

The CRS assumption is a reasonable starting point. In the labor literature, the majority of matching function estimates find support for constant returns to scale; as Petrongolo and Pissarides (2001) point out “divergences from constant returns are only mild and rare”. Nonetheless, to explore the robustness of our findings, in Section 6.1 we consider an alternative approach that does not make an assumption on the returns to scale, but instead relies on a restriction on the distribution F_e (Poisson).

Finally, as mentioned above, independence of ships and exporters is not a natural assumption in our setting. To relax it, we employ the literature on nonlinear IV techniques (e.g. Imbens and Newey, 2009, while for an application similar to ours see Bajari and Benkard, 2005). In particular, assume that an instrument z exists such that $s = h(z, \eta)$, with z independent of e, η . Under this formulation the endogeneity is driven by the correlation between η and e and, therefore, s is independent of e , conditional on η .

The approach now has two steps. In the first step, we recover η using the relationship $s = h(z, \eta)$; in

²⁶For instance, one can use a simple frequency estimator:

$$F_{m|s=s_t}(m_t|s=s_t) = \Pr(m \leq m_t|s=s_t) = \frac{\Pr(m \leq m_t, s=s_t)}{\Pr(s=s_t)} = \frac{\#1\{(m \leq m_t, s=s_t)\}}{\#1\{(s=s_t)\}}$$

where $1\{\cdot\}$ denotes the indicator function and $\#$ denotes the number of times. In practice, we use a Gaussian kernel density estimator.

practice we regress flexibly the number of ships s on the instrument, z , and set η equal to the residual. In the second step, we employ (14) conditioning on both s (as before) and η :

$$F_{m|s=s_t, \eta}(m_t|s=s_t, \eta) = F_{e|\eta}(e_t|\eta)$$

Similarly to above we recover the unknowns of interest e and $m(\cdot)$, by integrating both sides over η .

In this case, z consists of ocean weather conditions (unpredictable wind at sea) that shift the arrival of ships at a port without affecting the number of exporters.²⁷

5.2 Ship Costs and Exporter Revenues

We now turn to the ship cost parameters, $\{c_{ij}^s, c_i^w, \sigma\}$, for all i, j , as well as the exporter revenues $r \sim F_{ij}^r$ and production and export costs κ_{ij} , for all i, j . The estimation amounts to essentially matching the observed ballast decisions of ships to the model's predicted choice probabilities P_{ij} ; the observed prices τ_{ijr} , to the equilibrium Nash-bargained prices; and the observed trade flows (loaded trips) to the model's predicted equilibrium flows P_{ij}^e . Intuitively, ships' observed choices, conditional on observed prices, deliver the ship costs. Then, in equilibrium, prices inform us on the exporter's revenues. Finally, potential exporters make their entry decisions taking into account the (expected) prices they will face. From these decisions, we are able to back out the production and export cost.

Ship Costs Consider first the ship sailing costs, c_{ij}^s , wait costs, c_i^w , and the standard deviation of the logit shocks, σ . These parameters determine ships' optimal ballast choice probabilities, (4) and (5), given prices, through the value functions V_{ij} . Thus, we can estimate them via Maximum Likelihood using the observed ship choices. In particular, we use a nested fixed point algorithm to solve for the ship value functions at every guess of the parameter values, compute the predicted choice probabilities and then calculate the likelihood, as in Rust (1987). We provide the details of the approach in the Online Appendix, where we also prove that our value functions are well-defined using a contraction argument.

As is always the case in dynamic discrete choice, not all parameters are identified and some restriction needs to be imposed. Here, we have $I^2 + 1$ parameters and $I^2 - I$ choice probabilities, so we require $I + 1$

²⁷To proxy for the unpredictable component of weather, we divide the sea surrounding each region into 8 different zones (Northeast, Southeast, Southwest and Northwest both within 1,500 miles of the coast and between 1,500 and 2,500 miles from the coast), and we use the speed of the horizontal (E/W) and vertical (N/S) component of wind in each zone to proxy for weather conditions. We run a VAR regression of these weather variables on their lag component and season fixed effects and use the residuals, together with their squared term, as independent variables in the regression. The results are robust to the lag structure, as well as estimating jointly for neighboring zones.

restrictions; we show this formally, borrowing from the analysis of Kalouptsi et al. (2018) in the Online Appendix. The additional restrictions amount to using observed fuel prices to determine c_{ij}^s for some i, j ; see Section 6.2. Note also that the observed prices pin down the scale of payoffs (in dollars) and allow the identification of σ .²⁸

Exporter revenues and entry costs We next use data on shipping prices to estimate exporter revenues. Consider the equilibrium price (10) solved with respect to the exporter’s revenue r :

$$r = \frac{1 - \beta\delta(1 - \gamma\lambda_i^e)}{(1 - \gamma)(1 - \beta\delta)} \tau_{ijr} - \frac{\gamma(1 - \beta\delta(1 - \lambda_i^e))}{(1 - \gamma)(1 - \beta\delta)} (U_i - V_{ij}) \quad (16)$$

Note that the only unknowns in this expression are the revenue r and the bargaining coefficient γ , and that we have as many equations as the observed shipping prices. Indeed, λ_i^e is known from the previous step that estimates the matching function and searching exporters (λ_i^e is simply the average ratio of matches to exporters), β and δ are calibrated, while U_i and V_{ij} are known once $\{c_{ij}^s, c_i^w, \sigma\}$ have been estimated. Note that the valuation r is different for each matched exporter, while the bargaining coefficient parameter is a scalar, that is identical for all markets.

We first estimate the bargaining weight, γ by averaging equation (16) across regions and products; this average relationship links the average shipping price, to the average exporter revenue ($\bar{r} \equiv \sum_{ij} \bar{r}_{ij}/I^2$). We calibrate the average exporter revenue across regions, $\bar{r}$, to be equal to the average value of total trade in commodities (obtained from external trade data from Index Mundi) and solve (16) for γ .²⁹

Given this estimate for the bargaining weight, γ , we recover each exporter valuation r point-wise from (16) and obtain their distribution, F_{ij}^r , nonparametrically. Note that valuations are drawn from an origin-destination specific distribution, which allows for arbitrary correlation between a cargo’s valuation and destination.

We finally estimate the exporter entry costs, κ_{ij} , which capture both the cost of production, as well as

²⁸Unlike commonly used discrete choice models where only choices are observed and both the scale and level of utility need to be normalized, here we also observe a component of payoffs (prices in dollars) which allows us to relax the restriction on the scale of utility and identify σ .

²⁹In particular, we solve (16) with respect to γ and average over i, j, r so that:

$$\gamma = \frac{(1 - \beta\delta)(\bar{r} - \bar{\tau})}{\beta\delta E_{i,j,r} \lambda_i^e \tau_{ijr} + (1 - \beta\delta)\bar{r} - E_{i,j} (1 - \beta\delta(1 - \lambda_i^e)) (U_i - V_{ij})}$$

where $\bar{\tau}$ is the average observed price. To obtain $\bar{r}$, we first collect the average price of the five most common commodities (ore, coal, grain, steel, and fertilizers) from Index Mundi, and multiply it by the average tonnage carried by a bulk carrier (this is equal to the average vessel size times its utilization rate which equals about 65%). We then set $\bar{r}$ as their weighted average based on each commodity’s frequency in shipping contracts; we find $\bar{r}$ to equal 7 million US dollars. This approach of estimating the Nash bargaining weight is inspired by the empirical literature on oligopoly bargaining (e.g. Crawford and Yurukoglu, 2012, Ho and Lee, 2017).

any export costs beyond shipping prices. These are estimated from the exporters' entry decisions, given by the choice probabilities, P_{ij}^e , defined in (7) and (8). Indeed, we recover κ_{ij} from the following equation (Berry, 1994):

$$\ln P_{ij}^e - \ln P_{i0}^e = U_{ij}^e - \kappa_{ij}. \quad (17)$$

The only unknown in this equation is κ_{ij} . Recall that $U_{ij}^e = \alpha_i(\bar{r}_{ij} - \tau_{ij})$ and that both $\alpha_i = \beta\delta\lambda_i^e/(1 - \beta\delta(1 - \lambda_i^e))$ and the mean revenue $\bar{r}_{ij}$ have been estimated, while τ_{ij} is observed. Turning to the left hand side of (17), note that from the satellite data we have information on the proportions of loaded trips from i to j , i.e. $P_{ij}^e/(1 - P_{i0}^e)$. However, to obtain κ_{ij} we need both P_{ij}^e , as well as the fraction of potential exporters who choose *not* to export, P_{i0}^e , or equivalently, the number of potential exporters $\mathcal{E}_i$. We set $\mathcal{E}_i$ equal to the total production of the relevant commodities for each region i , thus assuming that a region's total production serves as an upper bound to the region's exports.³⁰

6 Results

In this section we present the results from our empirical analysis. We calibrate the discount factor to $\beta = 0.995$ and the exporter survival rate to $\delta = 0.99$. In our baseline estimation we ignore the different ship sizes, but our estimation results are similar when we consider Panamax alone or Handymax alone (these are the categories with sufficient data). Moreover, results are robust when we estimate the model separately by season.

6.1 Matching Function Results

We now present the estimates for the exporters and the matching function, obtained as described in Section 5.1. The matching function is estimated separately for each region i .

³⁰More precisely, let n_{it} denote the number of entrant exporters in period t and region i . Then, the number of exporters transitions as follows (see equation (1) in Online Appendix):

$$e_{it+1} = \delta(e_{it} - m_{it}) + n_{it}$$

so that in steady state

$$n_i = (1 - \delta)e_i + \delta m_i$$

or

$$n_i = \mathcal{E}_i(1 - P_{i0}^e) = (1 - \delta)e_i + \delta m_i$$

If $\mathcal{E}_i$ is known, we can solve this equation for P_{i0}^e since the right-hand-side is known. To determine $\mathcal{E}_i$, we collect annual country-level production data for grain (FAO), coal (EIA), iron ore (US Geological Survey), fertilizer (FAO) and steel (World Steel Association). To transform the production tons into a number of potential freights (i.e. shipments that fit in our bulk vessels), we first scale the production to adjust for the coverage of our data (we observe about half of the total fleet) and then divide by the average "active" ship size, taking into account a ship's utilization rate and the fact that a ship operates on average 340 days per year (due to maintenance, repairs, etc)

First, we present the results from the first stage regression of the number of ships on the unpredictable component of weather in surrounding seas for all regions. The results, shown in Table 4, indicate that ocean wind has a significant impact. This suggests that weather indeed affects trip duration and therefore weather shocks exogenously shift the supply of ships at port.^{31,32}

	Joint Significance	F-stat
North America West Coast	0	5.82
North America East Coast	0.03	2.22
Central America	0	4.78
South America West Coast	0.01	3.62
South America East Coast	0	5.78
West Africa	0	4.77
Mediterranean	0	4.85
North Europe	0.03	2.4
South Africa	0.03	2.54
Middle East	0	3.09
India	0	3.48
South East Asia	0	7.87
China	0	5.66
Australia	0	4.54
Japan-Korea	0	4.85

Table 4: First Stage, Matching Function Estimation. Regressions of the number of ships in each region on the unpredictable component of weather conditions in the surrounding seas. The first column reports the joint significance of the instruments and the second column the F-statistic. To proxy for the unpredictable component of weather, we divide the sea surrounding each region into 8 different zones (Northeast, Southeast, Southwest and Northwest both within 1,500 miles of the coast and between 1,500 and 2,500 miles from the coast), and we use the speed of the horizontal (E/W) and vertical (N/S) component of wind in each zone to proxy for weather conditions. We run a VAR regression of these weather variables on their lag component and season fixed effects and use the residuals, together with their squared term, as independent variables in the regression.

Figure 2 presents the weekly average number of exporters in each region. Not surprisingly perhaps, exporters are concentrated in Australia, the East Coast of North and South America and Southeast Asia, which are all rich in raw materials. India, Africa and Central America have the fewest exporters.

The matching function has reasonable properties. Exporters have substantially higher chances of

³¹This is consistent with contracting practices, where the ship promises a range of fuel consumption based on different weather conditions. Moreover, regressing the log of trip duration on wind speed, both upon arrival at the destination and upon departure at the origin confirms that the weather strongly affects trip duration; indeed, trips are 11% longer on average under bad weather conditions.

³²The first stage F-stat values shown in the last column of Table 4 suggest that we may have a weak instrument problem. As the second stage here is non-parametric, many of the available solutions do not necessarily apply. As a robustness, for each region i , we use unpredictable shocks to trip duration towards i as an alternative instrument. In particular, while shipowners' traveling decisions do not react to transitory weather shocks, the duration of a trip does. Moreover, in this case we can sign the first stage effect: longer trip duration *towards* region i should decrease the number of ships in i , s_i . The instrument has the correct sign and the F-stat is mostly higher than 10. In addition, the recovered exporters and matching functions are robust to our baseline results.




Figure 2: Average weekly number of estimated exporters.

finding a match than ships, consistent with our sample period of high ship supply and low demand. The matching rate for ships (exporters) declines as the market gets crowded with ships (exporters). Moreover, there is significant heterogeneity in the matching function across regions. To illustrate the importance of the nonparametric approach, we compare our estimated matching function to a typical parametric specification, the Cobb-Douglas. When estimating a Cobb-Douglas specification for the matching function, we find that the exporters recovered under this, more restrictive, assumption are different both in magnitude and in the relative ranking of regions, compared to the nonparametric case.³³

We close this section by discussing the robustness of our results to the CRS assumption. To do so, instead of assuming CRS, we impose a parametric assumption on the distribution of exporters, F_e . In particular, we assume that exporters are distributed Poisson, as we can then interpret the number of freights e_t as the number of arrivals of exporters at port every week. We estimate the number of exporters and the matching function following an analogous strategy, based on Matzkin (2003).³⁴ The

³³The Cobb-Douglas specification, $m_{it} = A_i s_{it}^{\alpha_i} e_{it}^{1-\alpha_i}$, is not straightforward to estimate, as e_{it} is not observed. We rewrite it as,

$$\log(m_{it}) = \log(A_i) + (1 - \alpha_i) \log e_{it} + \alpha_i \log s_{it} = \alpha_{i0} + \epsilon_{it} + \alpha_i \log s_{it}$$

and we recover α_i using an instrument (the unpredictable component of weather) for s_{it} .

³⁴Recall our basic equation, $F_{m|s}(m|s) = F_e(e)$. If $F_e(e)$ is Poisson with parameter ρ , we have:

$$F_{m|s}(m|s) = F_e(e) = \exp(-\rho) \sum_{k=0}^e \frac{\rho^k}{k!}$$

If ρ were known, we could solve this equation for e_{it} , all i, t , since the right-hand-side is known. We determine ρ iteratively by requiring again that the inequalities $m_{it} \leq e_{it}$ always hold.

results, shown in Figure 2 in the Online Appendix, are robust to this alternative restriction. Moreover, we find that the implied degree of homogeneity of the matching function under the Poisson distributed exporters is roughly equal to one, even though we do not assume CRS.

6.2 Ship Costs and Exporter Revenues

In our baseline specification, we construct seven groups for the sailing cost c_{ij}^s , roughly based on the continent and coast of the origin; and we estimate port wait costs c_i^w , for all i .³⁵ Note that c_{ij}^s is the *per week* sailing cost from i to j and its major component is the cost of fuel. Following the discussion on identification in Section 5.2, we set this cost for one of the groups (for trips originating from the East Coast of North and South America) equal to the average weekly fuel price (40,000 US dollars). Moreover, since the fuel cost is paid by the exporter when the ship is loaded, we add it to the observed prices.

The first two columns of Table 5 report the results.³⁶ Sailing costs are fairly homogeneous. Port wait costs are more heterogeneous and large, ranging between 90,000 and 290,000 US dollars per week. Consistent with industry narratives, waiting at port is costly, both due to direct port and security fees, as well as the rapid depreciation of the ship’s machinery and electronics and antifouling costs caused by the accumulation of microorganisms during immobility. Ports in the Americas are the most expensive, while ports in China, India, and Southeast Asia are the cheapest. The standard deviation of the logit shocks, σ , is estimated at about 16,000 US dollars, roughly 5% of the average trip price including the fuel cost payment. This suggests that the logit shocks do not account for a large part of utility or ballast decisions. As shown in Figure 3 in the Online Appendix, the model’s fit is very good, as our predicted choice probabilities are very close to the observed ones.³⁷

In the left panel of Figure 3 we plot the average exporter revenues across origins, while the third column of Table 5 reports the estimates. We estimate the bargaining coefficient to equal $\gamma = 0.3$. There is substantial heterogeneity in exporter revenues across space. South and North America have the highest revenues, while China, Japan, and Southeast Asia have the lowest. This ranking is reasonable, as for instance, Brazil exports grain which is expensive, whereas Southeast Asia exports mostly coal, which is

³⁵The seven groups are: (i) Central America, West Coast Americas; (ii) East Coast Americas; (iii) West and South Africa; (iv) Mediterranean, Middle East and North Europe; (v) India; (vi) Australia and Southeast Asia; (vii) China, Japan and Korea.

³⁶The standard errors are computed from 200 bootstrap samples with the resampling done at the ship level. We combine these bootstrap samples with those of the matching function to incorporate the error from the matching function estimation.

³⁷As our data comes from a period of historically low shipping prices, our estimated value functions are negative. This is partly due to the fact that we are not modeling ships’ expectations, so the value function does not take into that under mean-reverting demand for seaborne trade prices will eventually go up (see Kalouptsi, 2014). If we compute the equilibrium under higher exporter revenues that lead to prices closer to the ones observed before 2010, the ship value function indeed becomes positive.




Figure 3: Exporter valuations. The left panel of the figure plots the estimates for the average exporters' valuation across regions. The right panel correlates the average exporters' valuation with the share of exports in grain (source, Comtrade). The size of the circle proxies the number of observations.

one of the cheapest commodities. We generalize this example by focusing on grain, the most expensive frequently shipped commodity. In particular, using data from Comtrade, we explore whether regions that have a high share of grain exports tend to have higher estimated revenues. The results, shown in the right panel of Figure 3, reveal that indeed there is a positive correlation between the two. Of course, there may be other factors determining the valuation of an exporter such as inventory costs, just in time production, etc. On average, the price τ_{ij} is equal to about 5% of the mean valuation $\bar{r}_{ij}$, consistent with other estimates in the literature (e.g. UNCTAD, 2015, Hummels et al., 2009). Finally, the estimated exporter costs, κ_{ij} , exhibit substantial heterogeneity across destinations from a given origin, as well as across origins. On average κ_{ij} is the same order of magnitude as the average valuation $\bar{r}_{ij}$, reminiscent of a free entry condition into exporting. Moreover, we find that exporter costs are lower between an origin i and a destination j if the same language is spoken at i and j , which is reasonable since κ_{ij} includes both production costs, as well as other exporting costs, as discussed in Section 4.3.

	Port Costs c^w	Sailing Costs c^s	Exporters Valuations $\bar{r}$	Preference Shock σ
North America West Coast	227.65 (8.77)	46.75 (0.36)	9,047.58 (497.1)	
North America East Coast	272.3 (4.31)	- -	14,639.72 (261.11)	
Central America	175.41 (5.06)	46.75 (0.36)	8,129.37 (400.22)	
South America West Coast	265.55 (7.77)	46.75 (0.36)	12,561.15 (425.65)	
South America East Coast	292.5 (5.23)	- -	16,813.61 (455.06)	
West Africa	145.3 (4.84)	47.65 (0.33)	9,452.09 (659.51)	
Mediterranean	121.89 (3)	46.16 (0.28)	4,364.56 (269.17)	
North Europe	122.48 (1.71)	46.16 (0.28)	5,761.4 (202.04)	
South Africa	220.11 (7.28)	47.65 (0.33)	8,621.94 (420.55)	
Middle East	118.45 (2.14)	46.16 (0.28)	5,409.67 (252.23)	
India	97.23 (1.8)	45.93 (0.28)	4,800.29 (366.23)	
South East Asia	93.14 (1.02)	40.99 (0.28)	1,734.75 (81.99)	
China	91.07 (0.98)	40.89 (0.25)	2,708.65 (160.27)	
Australia	193.29 (2.85)	40.99 (0.28)	5,929.6 (160.19)	
Japan-Korea	100.41 (1.9)	40.89 (0.25)	2,863.02 (214.54)	
				16.53 (0.1070)

Table 5: Ship costs and exporter valuation estimates. All the estimates are in 1,000 USD. Standard errors computed from 200 bootstrap samples. The sailing cost for the East Coast of North and South America is set equal to the weekly fuel cost at 40,000 US dollars.

7 The Role of Endogenous Trade Costs

In this section, we illustrate the importance of endogenous trade costs. In particular, we demonstrate that the transportation sector (1) mitigates differences in the comparative advantage across countries reallocating productive activities from net exporters to net importers; (2) creates network effects in trade

costs; and (3) dampens the impact of shocks on trade flows. We illustrate these mechanisms, by studying how our setup compares to one with exogenous “iceberg” trade costs, the impact of a fuel cost shock and the spatial propagation of a macro shock (a Chinese slow-down). These mechanisms also shape the policy analysis considered in the next section.

Exogenous Trade Costs

We begin by showcasing how endogenous trade costs attenuate differences in the comparative advantage across countries (mechanism 1) and generate network effects (mechanism 2) by comparing our setup to one with exogenous trade costs. Typical trade models assume that trade costs are “iceberg”, so that a percentage of the traded good’s value is lost in transportation; this percentage is often a function of the distance between the origin and destination (Samuelson, 1954). In our setup, this amounts to assuming that the price paid for transportation from origin i to destination j is a function of only the distance, $1/d_{ij}$, and the average exporter valuation $\bar{r}_{ij}$. We thus construct a measure of exogenous trade costs by flexibly regressing the observed shipping prices between i and j on distance and $\bar{r}_{ij}$. This ensures that the overall level of shipping prices is the same. We then consider the trade patterns under these (counterfactual) trade costs, using the exporters’ choices, (7) and (8).

Strikingly, we find that in this world of exogenous iceberg trade costs, trade is substantially less balanced: trade imbalances are 24% higher on average. Indeed, net exporters experience a 9.4% increase in their exports, (and up to 24% for Australia), while net importers experience a 4.7% decline in their exports (and up to 22% for India). In other words, the transportation sector acts as a smoothing factor for world trade.

The core mechanism is related to the ship’s equilibrium behavior and in particular the strength of its bargaining position. Consider a ship that is located near a large net exporter, such as Brazil. As loading chances are high, the ship’s bargaining position upon meeting an exporter (i.e. its outside option) is strong and the ship can extract a high price, which in turn, tends to restrain Brazilian exports compared to a world with exogenous trade costs.³⁸ In contrast, consider a ship located near a net importer, such as India. The ship is unlikely to reload there. This lack of options puts the ship in a weak bargaining position and forces it to accept lower prices, which in turn increases Indian exports compared to the case of exogenous trade costs. The left panel of Figure 4 presents the change in each region’s exporting under exogenous trade costs and reveals that indeed net exporters experience a disproportionately large increase

³⁸In addition, the likely destinations are net importers where ships’ value are low. This further increases prices faced by next exporters like Brazil and dampens exports and trade imbalances.




Figure 4: Exogenous Trade Costs: the left panel presents the change in exports. The right panel correlates the change in exports under exogenous trade costs and a region’s centrality, defined as the distance-weighted sum of the imbalance of all other regions, i.e. $\sum_j (1/d_{ij}) (\text{exports}_j - \text{imports}_j)$.

in their exports.

This implies that in a world of endogenous trade costs, differences in the comparative advantage across countries are attenuated, as productive activities are reallocated from (efficient) net exporters to (inefficient) net importers. In our setup this corresponds to production decreasing in high $\bar{r}$ /low κ countries which face higher trade costs all else equal and increasing in low $\bar{r}$ /high κ ones, which face lower trade costs. Indeed, total net value of trade increases by 10% under exogenous trade costs. In addition, we find that this argument extends to a region’s neighborhood; this is because trade costs depend on the entire network of neighboring countries. Indeed, a net exporter close to other net exporters offers even more options to ships and prices are even higher, which inhibits the neighborhood’s exports. Hence, a country’s own trade imbalance, as well as the imbalance of its neighborhood, are crucial factors determining its trade disposition. To demonstrate these neighborhood effects, we consider a centrality measure for each region that consists of the weighted sum of trade imbalance in all regions, where the weights are given by the distance (i.e. $\sum_j (1/d_{ij}) (\text{exports}_j - \text{imports}_j)$). The right panel of Figure 4 correlates the change in exports to this centrality measure and shows that the overall imbalance of a neighborhood matters: regions whose neighborhood is overall a net-exporting one, offer high outside options to ships, which pushes prices up and thus exports down (and vice versa).

A Shock to Fuel Costs

Here we explore how world trade reacts to a fuel cost shock. This shock directly affects the main (variable) cost of transportation, c^s , and as a result also changes ship optimal behavior. This exercise illustrates how the transport sector dampens the impact of the shock (mechanism 3), while at the same time reallocating production across countries (mechanism 1).^{39,40}

Consider a 10% decrease in c^s . The left panel of Figure 5 presents the resulting change in world exports, while the right panel presents the impact of the shock on several outcomes of interest. A decline in c^s has a direct and an indirect effect. The direct effect is straightforward: as costs fall, shipping prices also fall and thus exports rise.⁴¹ Indeed, we see that exports increase everywhere, on average by 4.4%. The indirect effect is that the decline in sailing costs lowers the cost of ballasting. This implies that the ships' outside option, U , is now higher, which leads, all else equal, to an increase in prices that dampens the direct effect (mechanism 3). This is intuitive: reduced sailing costs imply that ships are less "tied" to their current region (and indeed ballast miles increase by 17%), and, as their bargaining position is stronger, they negotiate higher prices. The dampening is substantial: as shown in the second column in the right panel of Figure 5, if ships were not allowed to optimally adjust their behavior, the increase in trade would have been 41% higher.

In addition, as fuel costs decline and the importance of the transport sector is reduced, trade is driven to a greater extent by comparative advantages. Indeed, as shown in the right panel of Figure 5, net exporters experience an increase in exports of about 5%, while net importers experience an increase of about 3% (mechanism 1). This is largely driven by the indirect effect: ships ending a trip at net importing regions, are now less likely to wait there given the lower sailing cost; their outside option is higher and they can command higher prices that further reduce exports from these regions. Indeed, if ships were not allowed to optimally adjust their behavior, the increase in the exports of net importing regions would have been more than twice as high.

³⁹For a more detailed analysis of the impact of fuel costs and ship fuel efficiency see Brancaccio et al. (2018).

⁴⁰In this and all remaining counterfactuals, we compute the steady state spatial equilibrium distribution of ships and exporters. In the Online Appendix, we provide the computational algorithm employed. The use of nonparametric techniques in the estimation of the matching function may require substantial extrapolation in the counterfactuals; reassuringly, we find that the counterfactual matches and ships are always strictly within the range of our data.

⁴¹Formally, there is a direct increase in the surplus of all matches, since now a match between a ship and a freight is more valuable. Using the ship and freight value functions, the match surplus is given by

$$S_{ijr} = r - \frac{c_{ij}^s}{1 - \beta(1 - d_{ij})} + \frac{d_{ij}\beta}{1 - \beta(1 - d_{ij})}V_j - U_i - U_{ijr}^e.$$

A decline in c_{ij}^s , holding everything else constant, directly increases S_{ijr} . All else equal, this reduces export prices, τ , which in turn increases the value of an unmatched exporter, U_{ijr}^e , and thus induces more entry into the export market.




Figure 5: Fuel Cost Shock: impact of a 10% decline in c^s . The left panel presents the change in exports. In the right panel, the first column presents the total effect of the shock, while the second column presents the direct effect that does not allow ships to optimally adjust their behavior.

Chinese Slow-down

Finally, we explore the spatial propagation of a macro shock: a slow-down in China. This experiment showcases how the transport sector creates network effects (mechanism 2), while again reallocating production (mechanism 1).

We consider a reduction in the revenue of freights going to China, $\bar{r}_{i,china}$, by 10%. The left panel of Figure 6 plots the change in exports, while the right panel collects some statistics. We begin by looking at China itself. Chinese exports decline by 11%, even though they are not directly affected by the change in $\bar{r}_{i,china}$: the entirety of this decline is driven by the transport sector. Indeed, endogenous trade costs create a complementarity between imports and exports: the high Chinese imports, led to a large number of ships ending their trip in China and looking for a freight there, which in turn reduced trade costs for Chinese exporters (mechanism 1). Therefore, when imports decline, fewer ships end up in China and Chinese exporters are hurt.

Next, note that, as China is a large importer that trades with multiple countries, world exports naturally decline. Indeed, China's large trading partners, such as Australia, Indonesia and Brazil, experience a substantial decline in their exports; total exports decline by 11%. However, in addition to this direct effect, the optimal reallocation of ships over space differentially filters the shock in neighboring vs. distant regions (mechanism 2). Even though distant countries, such as Brazil, also experience a decline in exports, they benefit from the reallocation of ships across space. In particular, prior to the shock, a

large fraction of the fleet was located in the Southeast Pacific region, with ships traveling between China and its neighbors, mainly Australia and Indonesia. Following the shock, these ships reallocate to other parts of the world, pushing up exports there all else equal and dampening the overall decline from the direct effect. To see this, the right panel of Figure 6 shows that if ships were not allowed to optimally adjust their behavior, the decline in exports for distant countries would have been more than double (e.g. Brazilian exports would have fallen by 21% rather than 9%). This underscores the importance of being close to a large net importer like China: exporting countries in that “pocket” of the world, gained not just by directly exporting to China, but also indirectly from the increased supply of ships in that region.⁴²




Figure 6: Chinese Slow-down: impact of a 10% decline of the revenue from exporting to China. The left panel presents the change in exports. In the right panel, the first column presents the total effect of the slow-down, while the second column presents the direct effect that does not allow ships to optimally adjust their behavior. China’s neighbors include Japan/S.Korea, Australia, Southeast Asia and India. China’s non-neighbors include all other regions.

8 The Role of Maritime Infrastructure

How much do large maritime infrastructure projects contribute to world trade? We use our estimated setup to address this question by first evaluating a future project, the Northwest Passage. We then examine the impact of three natural and man-made passages: the Panama Canal, the Suez Canal, and Gibraltar.

The Northwest Passage is a sea route connecting the northern Atlantic and Pacific Oceans through the Arctic Ocean, along the northern coast of North America. This route is not easily navigable due to

⁴²China’s neighbors do not see their exports decline even further, because when Chinese imports decline, ships’ outside options fall everywhere, reducing shipping prices all else equal and dampening the decline in exports (mechanism 3).

Arctic sea ice; with global warming and ice thinning, it is likely that the passage will soon be open for shipping. The opening of the Northwest passage would reduce the distance between East Coast of North America and the Far East, as well as Northern Europe and the Far East.




Figure 7: Northwest Passage: The left panel presents the change in exports under the opening of the Northwest passage. In the right panel, the first column presents the total effect of the opening, while the second column the direct effect that does not allow ships to optimally adjust their behavior.

To simulate the impact of this new route, we reduce the nautical distance between Northeast America and Northern Europe to and from China/Japan/S.Korea by 30%.⁴³ The left panel of Figure 7 presents the resulting change in exports by region, while the right panel collects some statistics.

Not surprisingly, Northeast America sees its exports increase by 8%, while Northern Europe by 1.2%. Interestingly, exports from China and Japan/S.Korea are only marginally affected (0.3% for Japan, and -0.1% for China). On one hand, the import-export complementarity pushes exports up; but on the other hand ballasting is now less costly for ships: when in China or Japan ships can now ballast to the East Coast of North America more cheaply. The ships' higher outside option tends to increase prices and decrease exporting (mechanism 3).

Figure 7 reveals that other countries, not directly affected by the opening of the Northwest Passage, experience changes in their trade. This illustrates how network effects lead to the propagation of local shocks (mechanism 2). For instance, Brazil, Northwest America and Australia see their exports fall by 1%, even though none of these countries are directly affected by this passage. Indeed, ships that used to ballast to these regions now choose to ballast to Northeast America, that experiences a 17% increase in ships ballasting there. Overall, global welfare increases by 1.84%.

⁴³We calculate this change in travel distance via the Northwest Passage from Ostreng et al. (2013).

	Change in Exports	Max	Min	Most Affected	Change in Welfare
Suez	−3.51%	4.14%	−25.95%	Middle East	−5.25%
Panama	−3.23%	1.31%	−28.16%	Northeast America	−3.28%
Gibraltar	−6.37%	2.57%	−44.73%	Mediterranean	−5.03%

Table 6: Impact of Suez, Panama and Gibraltar on world expots. The real and counterfactual maritime routes among regions are calculated using the Dijkstra’s Algorithm. These are combined with ships’ average speed to compute trip duration with and without the corresponding passage.

We also quantify the impact of three existing passages (Suez, Panama, Gibraltar) by considering the change in world trade in their absence. These passages reduce nautical distance and thus the duration of specific trips. Table 6 presents the results. All passages substantially increase world trade and welfare. Removing the Suez Canal reduces trade by 3.5% and up to 26% in the Middle East, while welfare falls by 5.25%. Removing the Panama Canal leads to a decline in world trade of 3%, but up to 28% in the Northeast America. Welfare falls by 3.3%. Finally, Gibraltar seems to be the most critical one, as removing it would reduce world trade by close to 7% and up to 44% in the Mediterranean with an overall welfare decline of 5%.

9 Conclusion

In this paper we focus on the importance of the transportation sector in world trade. We build a spatial model where both trade flows and trade costs are equilibrium objects. Different experiments showcase that the transportation sector unveils a new role for geography through three mechanisms: it mitigates differences in the comparative advantage across countries, creates network effects and dampens the impact of shocks. While our empirical implementation focuses on bulk shipping, similar mechanisms are present in most, if not all, modes of transportation. We also demonstrate our setup’s potential to be used for policy evaluation by considering the quantitative impact of maritime infrastructure projects, such as the opening of the Northwest Passage. It is straightforward to use our setup in other counterfactuals, such as tariffs, trade wars, environmental regulations and port infrastructure. Finally, embedding our setup within a general equilibrium framework that endogenizes product prices is an exciting avenue for future research.

References

- ADLAND, R. AND H. JIA (2018): “Dynamic Speed Choice in Bulk Shipping,” *Maritime Economics & Logistics*, 20, 253–266.
- ALLEN, T. AND C. ARKOLAKIS (2014): “Trade and the Topography of the Spatial Economy,” *Quarterly Journal of Economics*, 129, 1085–1140.
- (2016): “The Welfare Effects of Transportation Infrastructure Improvements,” *mimeo*, Yale University.
- ANDERSON, J. E. AND E. VAN WINCOOP (2003): “Gravity with Gravitas: A Solution to the Border Puzzle,” *American Economic Review*, 93, 170–192.
- ANDERSON, T. AND H. RUBIN (1949): “Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations,” *Annals of Mathematical Statistics*, 20, 46–63.
- ARKOLAKIS, C., N. RAMONDO, A. RODRIGUEZ-CLARE, AND S. YEAPLE (2018): “Innovation and Production in the Global Economy,” *American Economic Review*, 108, 2128–2173.
- ASTURIAS, J. (2018): “Endogenous Transportation Costs,” *mimeo*, School of Foreign Service in Qatar, Georgetown University.
- BAJARI, P. AND C. L. BENKARD (2005): “Demand Estimation with Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach,” *Journal of Political Economy*, 113, 1239–1276.
- BAJARI, P., C. L. BENKARD, AND J. LEVIN (2007): “Estimating Dynamic Models of Imperfect Competition,” *Econometrica*, 75, 1331–1370.
- BEHRENS, K. AND P. M. PICARD (2011): “Transportation, Freight Rates, and Economic Geography,” *Journal of International Economics*, 85, 280–291.
- BERRY, S. (1994): “Estimating Discrete-Choice Models of Product Differentiation,” *RAND Journal of Economics*, 25, 242–262.
- BRANCACCIO, G., M. KALOUPTSIDI, AND T. PAPAGEORGIOU (2018): “The Impact of Oil Prices on World Trade,” *mimeo*, Harvard University.
- (2019a): “A Guide to Estimating Matching Functions in Spatial Models,” *mimeo*, Harvard University.
- BRANCACCIO, G., M. KALOUPTSIDI, T. PAPAGEORGIOU, AND N. ROSAIA (2019b): “Efficiency in Decentralized Transport Markets,” *mimeo*, Harvard University.
- BRANDT, L., J. V. BIESEBROECK, L. WANG, AND Y. ZHANG (2017): “WTO Accession and Performance of Chinese Manufacturing Firms,” *American Economic Review*, 107, 2784–2820.

- BUCHHOLZ, N. (2019): “Spatial equilibrium, search frictions and dynamic efficiency in the taxi industry,” *mimeo*, Princeton University.
- CALIENDO, L. AND F. PARRO (2015): “Estimates of the Trade and Welfare Effects of NAFTA,” *The Review of Economic Studies*, 82, 1–44.
- COLLARD-WEXLER, A. (2013): “Demand Fluctuations in the Ready-Mix Concrete Industry,” *Econometrica*, 81, 1003–1037.
- COSAR, A. K. AND B. DEMIR (2018): “Shipping inside the Box: Containerization and Trade,” *Journal of International Economics*, 114, 331–345.
- CRAWFORD, G. S. AND A. YURUKOGLU (2012): “The Welfare Effects of Bundling in Multichannel Television Markets,” *American Economic Review*, 102, 643–685.
- DONALDSON, D. (2018): “Railroads of the Raj: Estimating the Impact of Transportation Infrastructure,” *American Economic Review*, 108, 899–934.
- DONALDSON, D. AND R. HORNBECK (2016): “Railroads and American Economic Growth: A “Market Access” Approach,” *Quarterly Journal of Economics*, 131, 799–858.
- EATON, J., D. JINKINS, J. TYBOUT, AND D. XU (2016): “Two-sided Search in International Markets,” *mimeo*, Penn State University.
- EATON, J. AND S. KORTUM (2002): “Technology, Geography and Trade,” *Econometrica*, 70, 1741–1779.
- ERICSON, R. AND A. PAKES (1995): “Markov-Perfect Industry Dynamics: A Framework for Empirical Work,” *The Review of Economic Studies*, 62, 53–82.
- FAJGELBAUM, P. D. AND E. SCHAAL (2017): “Optimal Transport Networks in Spatial Equilibrium,” *NBER working paper*, 23200.
- FRECHETTE, G. R., A. LIZZERI, AND T. SALZ (2019): “Frictions in a Competitive, Regulated Market Evidence from Taxis,” *American Economic Review*, forthcoming.
- HEAD, K. AND T. MAYER (2004): “The Empirics of Agglomeration and Trade,” in *Handbook of Regional and Urban Economics*, ed. by J. V. Henderson and J. F. Thisse, Elsevier, vol. 4.
- HO, K. AND R. S. LEE (2017): “Insurer Competition in Health Care Markets,” *Econometrica*, 85, 379–417.
- HOLMES, T. J. AND E. SINGER (2018): “Indivisibilities in Distribution,” *NBER working paper*, 24525.
- HOPENHAYN, H. A. (1992): “Entry, Exit, and Firm Dynamics in Long Run Equilibrium,” *Econometrica*, 60, 1127–1150.

- HOSIOS, A. J. (1990): “On the Efficiency of Matching and Related Models of Search and Unemployment,” *Review of Economic Studies*, 57, 279–298.
- HUMMELS, D. (2007): “Transportation Costs and International Trade in the Second Era of Globalization,” *The Journal of Economic Perspectives*, 21, 131–154.
- HUMMELS, D., V. LUGOVSKYY, AND S. SKIBA (2009): “The Trade Reducing Effects of Market Power in International Shipping,” *Journal of Development Economics*, 89, 84–97.
- HUMMELS, D. AND S. SKIBA (2004): “Shipping the Good Apples Out? An Empirical Confirmation of the Alchian-Allen Conjecture,” *Journal of Political Economy*, 112, 1384–1402.
- IMBENS, G. W. AND W. K. NEWHEY (2009): “Identification and Estimation of Triangular Simultaneous Equations Models Without Additivity,” *Econometrica*, 77, 1481–1512.
- ISHIKAWA, J. AND N. TARUI (2015): “Backfiring with Backhaul Problems: Trade and Industrial Policies with Endogenous Transport Costs,” *mimeo, University of Hawaii*.
- KALOUPSIDI, M. (2014): “Time to Build and Fluctuations in Bulk Shipping,” *American Economic Review*, 104, 564–608.
- (2018): “Detection and Impact of Industrial Subsidies, the Case of Chinese Shipbuilding,” *Review of Economic Studies*, 85, 1111–1158.
- KALOUPSIDI, M., P. T. SCOTT, AND E. SOUZA-RODRIGUES (2018): “Identification of Counterfactuals in Dynamic Discrete Choice Models,” *NBER working paper*, 21527.
- KOOPMANS, T. C. (1949): “Optimum Utilization of the Transportation System,” *Econometrica*, 17, 136–146.
- KROLIKOWSKI, P. M. AND A. H. MCCALLUM (2018): “Goods-Market Frictions and International Trade,” *Federal Reserve Bank of Cleveland, Working Paper no. 16-35R*.
- KRUGMAN, P. (1991): “Increasing Returns and Economic Geography,” *The Journal of Political Economy*, 99, 483–499.
- LAGOS, R. (2000): “An Alternative Approach to Search Frictions,” *The Journal of Political Economy*, 108, 851–873.
- (2003): “An Analysis of the Market for Taxicab Rides in New York City,” *International Economic Review*, 44, 423–434.
- LANGE, F. AND T. PAPAGEORGIOU (2019): “Beyond Cobb-Douglas: Flexibly Estimating Matching Functions with Unobserved Matching Efficiency,” *mimeo, McGill University*.

- MATZKIN, R. L. (2003): “Nonparametric Estimation of Nonadditive Random Functions,” *Econometrica*, 71, 1339–1375.
- MORTENSEN, D. T. AND C. A. PISSARIDES (1994): “Job Creation and Job Destruction in the Theory of Unemployment,” *The Review of Economic Studies*, 61, 397–415.
- OSTRENG, W., K. M. EGER, B. FLOISTAD, A. JORGENSEN-DAHL, L. LOTHE, M. MEJLAENDER-LARSEN, AND T. WERGELAND (2013): *Shipping in Arctic Waters: A Comparison of the Northeast, Northwest and Trans Polar Passages*, Berlin, Heidelberg: Springer.
- PAKES, A., M. OSTROVSKY, AND S. T. BERRY (2007): “Simple Estimators for the Parameters of Discrete Dynamics Games,” *The RAND Journal of Economics*, 38, 373–399.
- PETRONGOLO, B. AND C. A. PISSARIDES (2001): “Looking Into the Black Box: A Survey of the Matching Function,” *Journal of Economic Literature*, 39, 390–431.
- ROGERSON, R., R. SHIMER, AND R. WRIGHT (2005): “Search-Theoretic Models of the Labor Market: A Survey,” *Journal of Economic Literature*, 43, 959–988.
- RUST, J. (1987): “Optimal replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55, 999–1033.
- RYAN, S. P. (2012): “The Costs of Environmental Regulation in a Concentrated Industry,” *Econometrica*, 80, 1019–1062.
- SAMUELSON, P. A. (1954): “The Transfer Problem and Transport Costs, II: Analysis of Effects of Trade Impediments,” *The Economic Journal*, 64, 264–289.
- SIMONOVSKA, I. AND M. E. WAUGH (2014): “The Elasticity of Trade: Estimates and Evidence,” *Journal of International Economics*, 92, 34–50.
- THE WORLD BANK (2012): “Road Freight in Central America: Five Explanations for High Costs of Service Provision,” Tech. rep., The World Bank, Washington, DC.
- UNCTAD (2015): “Review of Maritime Transport,” Tech. rep., United Nations Conference on Trade and Development, Geneva, Switzerland.
- WAUGH, M. (2010): “International Trade and Income Differences,” *American Economic Review*, 100, 2093–2124.
- WEINTRAUB, G. Y., C. L. BENKARD, AND B. V. ROY (2008): “Markov Perfect Industry Dynamics with Many Firms,” *Econometrica*, 76, 1375–1411.

WONG, W. F. (2019): “The Round Trip Effect: Endogenous Transport Costs and International Trade,”
mimeo, University of Oregon.

WTO (2015): “International Trade Statistics,” Tech. rep., World Trade Organization, Geneva, Switzerland.