The Impact of Oil Prices on World Trade

Giulia Brancaccio[*], Myrto Kalouptsidi[†], Theodore Papageorgiou[‡]

Princeton University, Harvard University, CEPR and NBER, McGill University

March 2018

In this paper we investigate the importance of fuel oil costs in determining world trade. We use detailed data on ship movements across the globe and transaction-level shipping prices, along with a dynamic model describing the world shipping industry, to measure the elasticity of trade with respect to ship fuel costs. We find that the average estimated elasticity is 0.35, but ranges from 0.1 to about 1.2 depending on the level of the fuel cost. The pass-through of fuel costs to exporters is low, at 0.17. Strikingly, the trade elasticity features a pronounced asymmetry in low vs. high oil prices. As fuel costs decline, the elasticity plateaus. This "flattening out" of the elasticity is attributed to the equilibrium of the transportation sector and the changes in the relative bargaining positions of ships and exporters in particular. Finally, we use the estimated elasticity to assess the importance of ship design on trade flows: if the large fuel efficiency gains achieved in the 1980s had not been realized, trade would be 12% lower today.

Keywords: fuel costs, shipping, world trade, trade elasticity, oil prices, fuel efficiency, fuel cost pass-through

[*]Princeton University, Department of Economics, Fisher Hall, Princeton, NJ 08544, USA, giulia@princeton.edu

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

[‡]McGill University, Department of Economics, 855 Sherbrooke St West, Room 519, Montreal, QC, H3A 2T7, Canada, theodore.papageorgiou@mcgill.ca

1 Introduction

How significant an effect do oil prices have on trade flows? Many have claimed that oil price spikes have the potential to put a break on world trade by increasing transport costs.[1] Indeed, oil prices determine ship fuel costs, which constitute the core (variable) cost component of the transportation sector. In this paper, we compute the world trade elasticity with respect to oil prices. We are able to isolate the impact of oil shocks on transportation costs (vs. other channels such as production input costs or income effects) through the use of a simple structural trade model that explicitly incorporates the transportation sector. We show that modeling the transport sector is crucial in understanding the mechanism behind the oil price pass-through to exporters and the patterns of the estimated elasticity.

We focus on oceanic shipping, which accounts for the large majority of international trade, and in particular on trade in bulk commodities, such as minerals, grain, ores and chemicals, which in turn accounts for about half of all seaborne trade in tonnage (UNCTAD, 2015); however, our findings may hold more generally. Bulk ships are often thought of as the "taxis of the ocean", as the industry's structure and operation resembles that of taxicabs. For this sector of the transportation industry, we are able to collect data on shipping contracts between shipowners and exporters that correspond to specific trips; each transaction includes the shipping price as well as the origin and destination of the trip. We also obtain AIS ship movement data that inform us on ships' sequences of loaded and empty trips. This dataset was first used in our prior work, Brancaccio et al. (2017), henceforth BKP.

We employ the dynamic spatial search model built in BKP. This model is in the spirit of Mortensen and Pissarides (1994) and Lagos (2000) and centers on the behavior of ships and exporters. Ships are homogeneous and contract with exporters for individual trips. Fuel costs are the main variable cost of ships and they are captured in our setup through ships' cost of sailing. When a ship and an exporter meet in a particular region of the world, they bargain over the price to transport the exporter's cargo from this location to the exporter's desired destination. The ship then travels the distance and restarts in the destination by searching for a new cargo. In every period, if a ship does not find a match with an exporter, it makes a choice: it can either wait at its current location, or it can choose a destination to ballast to (i.e. travel empty). Exporters have one cargo to ship. Potential exporters decide whether to export or

[1] For instance, a report by The Economist (Wood, 2008) states "For the countries of Asia, where the price of transporting goods to the West has a significant impact on their attractiveness as a manufacturing location, these [the rising oil prices] are serious issues. In an era of high and volatile oil prices, is the world not as flat as has been suggested?". Similarly, in 2008 Paul Krugman writes in his blog "higher fuel prices are putting the brakes on globalization: if it costs more to ship stuff, there will be less shipping". In his book "Why your world is about to get a whole lot smaller: oil and the end of globalization", Rubin (2009) argues that spiking oil prices will curtail long-distance shipping and traveling.

not, as well as which destination to export to. Once this entry decision is made, they wait at port until they can match with a ship that will transport their cargo. The model parameters (the matching function describing the meeting process between exporters and ships, the valuations and entry costs of exporters and the ship port costs) are estimated using our datasets as in BKP.

A key feature of the data that our model is designed to capture, is the impact of trade imbalances on shipping prices and trade flows. Indeed, as discussed in BKP, most countries are either net importers or net exporters of bulk commodities. This asymmetry is also reflected in shipping prices: it is more expensive to hire a ship when the destination is a net importing country, as the ship's profitability is low there (it has to either wait for a cargo or incur the fuel costs of ballasting elsewhere). These features are important in determining the pass-through of oil prices to shipping prices, as well as the world trade elasticity with respect to oil prices.

We compute the trade elasticity by considering how changes in the fuel costs pass-through to shipping prices and in turn affect trade flows. Our estimated trade elasticity is 0.35 at the average observed fuel cost level. However, it depends crucially on the level of the oil price and in fact ranges between 0.1 and 1.2. We also estimate that the pass-through of fuel costs to exporters is low, as the elasticity of prices with respect to fuel costs equals 0.17 at the average observed fuel cost. This suggests that ship variable costs are not a good approximation of transport costs and that modeling the transport sector is important in understanding trade flows and trade costs.

A striking feature of the estimated trade elasticity is its pronounced asymmetry in low and high levels of fuel costs. Indeed, the elasticity gets steeper as the fuel cost increases, while it plateaus as the fuel cost decreases. This asymmetry is generated by the equilibrium of the transportation sector and in particular, the changes in the relative bargaining positions of ships and exporters. Naturally, as the fuel cost declines, trade increases, since fuel costs are a key input in transportation costs. In addition, however, fuel costs are a key determinant of ships' relative bargaining position. As the fuel cost declines, the world becomes more "flat", since distance matters less, and ships can reallocate cheaply across different regions. Therefore, they are less "tied" to their current location and are able to extract higher prices than otherwise. This dampens the price decline and mutes the increase in trade disproportionally at low fuel costs. In contrast, when fuel costs are high, it is costlier for ships to change locations and exporters at their current location have a stronger bargaining position, leading to large increases in trade. This effect is particularly pronounced in net exporting regions, where the high likelihood of finding a load makes ships almost certain to stay put when oil prices are high. As a result, as the world becomes flat under low fuel costs, the increase in trade shrinks because of the ships' strong bargaining position.

To further document this mechanism, we explore a relevant testable prediction of our model. When fuel costs decline and distance is of less importance, ship values tend to equalize over space. A ship in an unloading region is now not much worse off than a ship in a loading region, as ballasting from the former to the latter is cheaper. As the dispersion in ship value functions declines, so does the dispersion in shipping prices which depend on the ship's value at the destination. We test this using the observed shipping prices and oil prices and find that indeed as oil prices increase, shipping prices equalize across space.

Finally, we use our estimates to assess how much recent trends in fuel efficiency of ship design have affected trade flows. Ship design is affected by a number of factors over time, such as long-term trends in the shipbuilding industry, technological improvements and environmental policies. For instance, the 1980s witnessed large improvements in fuel efficiency (following the oil price shocks of the 1970s), while the last 10 to 20 years witnessed a reversion, with fuel efficiency deteriorating by about 13%. We compute that the efficiency gains achieved in the 1980s by world shipyards led to a decline in shipping prices of 5.5% and an increase in trade by 12%. On the other hand, the recent deterioration in fuel efficiency design resulted in a 5.6% reduction in world trade. These calculations showcase the policy relevance of our estimated trade elasticity with respect to fuel costs, as a number of environmental regulations imposing fuel efficiency targets are currently discussed by international organizations.

The literature quantifying the impact of oil prices on trade, as well as transport costs, is relatively thin. von Below and Vezina (2016) (see also work cited therein) employ a gravity equation that incorporates oil prices to measure their impact on trade through transportation costs. The estimated elasticity is fairly large, between -1.2 and -1.8. A challenge with this approach is that it is potentially difficult to separate the impact of oil shocks operating through transport costs, as opposed to other facets of the economy (the authors do adopt an IV approach to handle this). Hummels (2007) measures the elasticity of freight costs with respect to fuel costs and estimates it to about 0.2 to 0.3, while a few recent papers have explored the elasticity of trade with respect to freight rates, as well as the role of the transportation sector overall (e.g. Hummels (2007); Asturias (2016); Wong (2017), BKP). For instance, Wong (2017) estimates the trade elasticity with respect to shipping prices to about -3 and Limao and Venables (2001) report a similar estimate.[2] Bridgman (2008) considers a trade model with an energy-using transportation sector to investigate the 1970s and 1980s oil shocks and finds that they had a significant impact on world trade.

[2]North (1958) and Estevadeordal et al. (2003) also find that changes in transportation prices have been historically important determinants of world trade.

Rubin (2009) argues that a shortage in oil and increasing oil prices will limit globalization.

The approach adopted here to measure the trade elasticity is different from most of the existing literature; we employ our structural model to isolate the impact of oil shocks only through the transportation sector rather than other channels (e.g. higher input costs, income effects, or correlated changes in demand). In addition, our model guides us in understanding the shape of the elasticity, as well as the mechanism of the oil price pass-through to exporters.

Finally, there is a strand of literature that investigates how oil prices affect trade more broadly, besides its impact through transport prices (e.g. Backus and Crucini (2000), Kilian et al. (2009), Chen and Hsu (2012) and others). In addition, a number of papers in the field of maritime economics have explored ships' optimal speed choices, as well as fuel efficiency, emissions and environmental policies on energy efficiency in different ship segments (e.g. Adland and Jia (2016a), Adland and Jia (2016b), Adland et al. (2017) and Adland et al. (2017)). Lastly, more broadly, our paper contributes to the large literature on trade costs (e.g. Anderson and Van Wincoop, 2003) and geography (e.g. Krugman, 1991), while our methodology borrows from the theoretical and empirical literature on dynamic models (Rust (1987); Hopenhayn (1992); Ericson and Pakes (1995)).

The paper proceeds as follows. In Section 2 we discuss the key features of the oceanic bulk shipping industry and the data. In Section 3 we present the model, as well as the procedure employed to estimate it. In Section 4 we present our main results on the trade elasticity with respect to fuel costs and the oil price pass-through to exporters, while in Section 5 we investigate our setup's prediction regarding the relationship between price dispersion and oil prices. In Section 6 we ask how much the recent improvements in ship fuel efficiency have contributed to trade growth. Finally, in Section 7 we conclude.

2 Industry and Data

Bulk shipping involves large oceanic carrier vessels (larger than 10,000 DWT capacity) that carry mostly commodities and raw materials, such as grain, iron ore, steel, coal, chemicals, etc. The industry is very unconcentrated with a large number of small firms. These ships operate much like taxi cabs: a shipowner contracts with a cargo owner for a specific trip; the ship is filled up with this exporter's load and it delivers the cargo at the agreed upon destination. The ship then restarts in that destination by looking for a new contract. Similar to taxi cabs, bulk shipping services are considered fairly homogeneous. Unlike taxis however, prices are not regulated; they are negotiated between the shipowner and the exporter, and mediated by one or multiple shipbrokers.

We exploit two main databases. The first consists of a sample of contracts between shipowners and exporters obtained from Clarksons Research. Each observation is a contract for a trip and it specifies the origin and destination of the trip, the loading and signing dates, the ship, the shipowner and cargo owner and finally the price. We use this dataset to obtain information on trip prices.

The second dataset, obtained from ExactEarth Ltd, contains ship movements collected from satellites (we observe the location of about half of the world fleet every six minutes). The AIS data also report the ship's draft (i.e. the distance between the bottom of the ship's hull and the waterline) which allows us to distinguish loaded from empty movements. We employ this dataset to construct ship movement histories of loaded and empty trips. For a more detailed description of the industry we refer the reader to BKP and Kalouptsidi (2014); BKP also provides more details on the data, as well as summary statistics and notable data patterns.

A prevailing feature of the data is the large trade imbalances and their impact on shipping prices. Indeed, as documented in BKP, most countries are either large net importers or large net exporters of the commodities carried in bulk vessels. China and India are the biggest importers, while Brazil, Australia and North America are the biggest exporters. These trade imbalances translate into asymmetric ship hiring rates at different regions of the world: although a ship in Brazil is very likely to find a cargo, a ship in China is much less so. As a result, a ship would much rather unload a cargo in Brazil than China, as her options are much better in Brazil and this is reflected in the prices the ship agrees upon. Indeed, the price to unload in China is substantially higher than the price to unload in Brazil. Shipping prices exhibit pronounced asymmetries, reflecting the world's natural geography (distances and natural inheritance) and its impact on ship profitability at different regions of the world.

3 Model

We next provide a description of the dynamic spatial search model of the global shipping industry introduced in BKP. We refer the interested reader to BKP for further details.

3.1 Environment

Time is discrete. There are I regions in the world and two types of agents: exporters (or freights) and ships. Both exporters and ships are risk neutral and share a discount factor β. Let f_i denote the freights awaiting transportation in region i. Freights are heterogeneous in (i) their destination, j (endogenized below); (ii) their valuation, v, denoting the value of the cargo upon delivery to the destination.

Ships are homogeneous and carry at most one freight at a time. Every period a ship is either sailing towards destination j, full or empty, at a per-period sailing cost c_{ij}^s, which is primarily the fuel cost; or it is waiting in port i at cost c_i^u. A ship sailing from i to j arrives at its destination with probability ξ_{ij}, so that the average trip duration equals $1/\xi_{ij}$.

Ships at port i meet exporters originating from port i randomly. Every period the number of matches in location i is determined by a matching function $m_i(f_i, s_i)$, where s_i is the number of unmatched ships in location i, and f_i is the number of unmatched freights. Let $\lambda_i = m_i/s_i$ denote the probability that a ship meets a freight and $\lambda_i^f = m_i/f_i$ the probability that a freight meets a ship. We follow the bulk of the search literature, and assume that the surplus of a meeting is split via generalized Nash bargaining. This determines the price τ_{ijv} that a freight with valuation v pays the ship in order to be transported from i to j. Let $\gamma \in (0,1)$ denote the freight's bargaining power.[3]

Timing is as follows:

1. Ships and freights match.

2. Unmatched ships draw additive iid preference shocks $\epsilon = [\epsilon_1, ..., \epsilon_I] \in \mathbb{R}^I$ from a Type I extreme value distribution and decide whether to (i) stay in their current region and wait for freight; or (ii) ballast toward some destination j. Each preference shock is associated with one of these choices.

3. Ships already traveling from i to j, arrive with probability ξ_{ij}. Unmatched ships that decided to ballast begin traveling to their chosen destination. Existing unmatched exporters disappear with probability $1 - \delta$.

4. In each region i, $\mathcal{E}_i$ potential exporters decide whether and to which destination to export to. The exporters that do enter the market, pay an entry cost κ_{ij}, draw their valuations v from a distribution F_{ij}^v, and join the pool of unmatched exporters the following period.

3.2 Equilibrium

In this section, we derive firm behavior. We focus on the steady state equilibrium of this framework.

[3] We implicitly assume that freight valuations, v, are large enough so that all meetings are converted into matches, i.e. the match surplus is always positive. Given that in this context, valuations are an order of magnitude greater than the shipping price, this assumption seems quite reasonable.

Ships The value of a ship traveling from i to j, W_{ij}, is given by,

$$W_{ij} = -c_{ij}^{s} + \xi_{ij}\beta U_j + (1 - \xi_{ij})\,\beta W_{ij} \tag{1}$$

i.e. the ship pays the per period travel cost, c_{ij}^{s}, then with probability ξ_{ij} it arrives at its destination and begins the following period unmatched at j, and with the complement probability the following period it is still traveling. U_j is the value of a ship that is unmatched in port j at the start of the period and is given by,

$$U_i = -c_i^{u} + \lambda_i E_{j,v} V_{ijv} + (1 - \lambda_i)\,J_i \tag{2}$$

i.e. the ship pays the port cost, c_i^{u}, and with probability λ_i it meets some freight with destination j and value v, while with the complement probability it remains unmatched and needs to make a decision. We examine each case in turn.

The value of a ship that is matched with a freight with value v and destination j is

$$V_{ijv} = \tau_{ijv} + W_{ij}$$

i.e. the ship receives the shipping price, τ_{ijv}, which is the outcome of Nash bargaining, and immediately begins its journey towards j, obtaining value W_{ij}, defined in (1) above.

The value of a ship that remains unmatched in i at the end of the period is

$$J_i(\epsilon) = \max\left\{\beta U_i + \sigma\epsilon_{ii}, \max_{j\neq i} W_{ij} + \sigma\epsilon_{ij}\right\}$$

i.e. the ship can either remain in i and obtain value U_i the following period, defined in (2) or it can choose a destination j and ballast there, obtaining value W_{ij}. The terms ϵ_{ij} capture unobserved idiosyncratic ship costs or measurement error and have standard deviation σ. Finally let $J_i \equiv E_\epsilon J_i(\epsilon)$ denote the "ex ante" value of an unmatched ship, i.e. *before* drawing ϵ_{ij}.

Freights The value of an unmatched freight in market i with destination j and valuation v at the end of the period is

$$J_{ijv}^{f} = \delta\beta\left(\left(1 - \lambda_i^{f}\right) J_{ijv}^{f} + \lambda_i^{f} V_{ijv}^{f}\right) \tag{3}$$

i.e. conditional on surviving (which occurs with probability δ), with probability $1-\lambda_i^f$ it remains unmatched the following period, and with the complement probability it is matched with a ship and receives

$$V_{ijv}^f = v - \tau_{ijv} \tag{4}$$

i.e. its value v minus the shipping price, τ_{ijv}.

Shipping Price Nash bargaining implies the surplus sharing condition:

$$\gamma \left(V_{ijv} - J_i\right) = \left(1-\gamma\right)\left(V_{ijv}^f - J_{ijv}^f\right)$$

Using this condition and substituting in the value functions, allows us to solve out for the shipping price:

$$\tau_{ijv} = \frac{\gamma\left(1-\beta\delta\left(1-\lambda_i^f\right)\right)}{1-\beta\delta\left(1-\gamma\lambda_i^f\right)}\left(J_i - W_{ij}\right) + \frac{\left(1-\gamma\right)\left(1-\beta\delta\right)}{1-\beta\delta\left(1-\gamma\lambda_i^f\right)} v \tag{5}$$

The equilibrium price depends positively on v, so that more valuable freights pay higher prices. Moreover, note that the price rises with travel time, $1/\xi_{ij}$, (through the traveler's value W_{ij} defined in equation (1)), as well as the the per-period travel costs, c_{ij}^s. It depends negatively on the value of an unmatched ship at the destination, U_j: freights that are headed towards destinations that offer a high value to ships pay less. From the value of U_i, such destinations may have low port costs, c_i^u, attractive ballast opportunities, high matching probability, λ_i, freights with high values, v and/or freights that are headed to attractive destinations (see equation (2)). Finally, the equilibrium price depends on the ship's outside option, i.e. its value if unmatched in i, J_i. Ships matched in regions with high matching probabilities, high value freights or attractive ballast opportunities have higher outside options and as a result command higher prices.

Freight Entry Finally, we characterize the freight entry decision. Every period there are $\mathcal{E}_i$ potential exporters and the value of each is

$$J_i^{ef} = \max\left\{\epsilon_0^f, \max_{j\neq i}\left\{E_v J_{ijv}^f - \kappa_{ij} + \epsilon_j^k\right\}\right\} \tag{6}$$

i.e. each one chooses between its outside option of not exporting, in which case the payoff has been normalized to zero, and exporting to one of the $j \neq i$ destinations. If it chooses to export, it pays a destination-specific production and exporting cost, κ_{ij}, it then draws value, v from F_{ij}^v, and obtains the value of an unmatched freight, J^f_{ijv}, given by (3). In addition, each potential entrant draws a vector of additive iid choice-specific preference shocks, ϵ^f, from a type I extreme value distribution.

3.3 Model Estimation and Results

We now provide a brief description of the estimation strategy followed in BKP. We estimate (i) the matching function and the freights searching for ships, (ii) the port costs, c^u_i, (iii) the distributions of freight values, F^v_{ij} and (iv) the production and exporting costs, κ_{ij}.

Matching Function Estimation In contrast to the majority of empirical work on estimating matching functions, we do not observe all necessary ingredients, $\{s, f, m\}$; indeed, we lack data on one side of the market, namely searching freights.[4] Our goal is thus to use data on the number of ships searching and the number of matches to recover the matching function and the number searching freights.

In order to overcome this lack of data, as well as to avoid imposing functional form assumptions on the matching function, we use results from the literature on non-parametric identification, namely Matzkin (2003). We provide an intuitive description of the approach and refer the reader to BKP for the formal treatment. As is common in the literature, assume that the matching function is strictly increasing in the number of freights. Then, from

$$m_{it} = m_i\left(s_{it}, f_{it}\right)$$

if we knew $m_i\left(.\right)$, as well as m_{it} and s_{it}, we could invert the matching function and recover f_{it}. Conversely, if we observed f_{it}, but not $m_i\left(.\right)$, we could use a nonparametric estimator to recover the matching function. In our case, however, we observe neither $m_i\left(.\right)$, nor f_{it}.

First assume that s_{it} and f_{it} are independent. Then, by leveraging the correlation between m_{it} and s_{it} in the data, we can identify the elasticity of the matching function with respect to s_{it}: since s_{it} and f_{it} are independent, for a given shift in s_{it}, the resulting change in m_{it} is informative about the underlying elasticity. In practice, we relax the independence assumption by using an instrument: we leverage unexpected shocks to sea weather that exogenously shocks the arrivals of ships in a market.

Second assume the matching function exhibits constant returns to scale. This assumption is consistent with the literature on matching function estimation which has largely failed to reject constant returns to scale (see e.g. Petrongolo and Pissarides, 2001). This assumption implies that knowing the elasticity of

[4]See for instance Petrongolo and Pissarides (2001); two exceptions to the literature are Buchholz (2016) and Frechette et al. (2016).

the matching function with respect to s_{it} is sufficient to back out its elasticity with respect to f_{it}. Now that we have essentially recovered the shape of the matching function, we can invert the function to back out the unobserved number of freights, f_{it}, at any point in our sample.

Ship Costs Consider next the ship conditional choice probabilities; the probability that a ship chooses to remain at market i rather than ballast elsewhere is

$$p_{ii} = \frac{\exp(\beta U_i/\sigma)}{\exp(\beta U_i/\sigma) + \sum_{j\neq i} \exp(W_{ij}/\sigma)},$$

while the probability that it chooses to ballast to market j is

$$p_{ij} = \frac{\exp(W_{ij}/\sigma)}{\exp(\beta U_i/\sigma) + \sum_{j\neq i} \exp(W_{ij}/\sigma)}.$$

These model implied probabilities depend on the ship values, U_i and W_{ij}, which in turn depend on travel costs, c^s_{ij} and port costs, c^u_i, as well as prices, τ_{ijv} (see (1) and (2)).

The travel costs, c^s_{ij} consist primarily of the cost of fuel (Stopford, 2009). We thus calibrate their level using the average weekly price of fuel in our sample, and calculate that weekly travel costs are equal to about \$69,100.[5] We allow travel costs to differ for each pair (i,j) as follows: c^s_{ij} takes one of seven values based on the continent and coast of the origin; we set one of the c^s_{ij} (East Coast of North and South America) to equal to the average fuel cost of 69,100\$ and estimate the rest as described below.[6] Note also that since the fuel cost is paid by the exporter when the ship is loaded, we add it to the observed prices; this generates an asymmetry in the cost of traveling full and traveling empty.

We back out the parameters of interest $\left\{c^u_i, c^s_{ij}, \sigma\right\}$ via Maximum Likelihood using ships' observed choice probabilities, as well as the observed prices, following Rust (1987). In particular, for a given set of parameters, we solve for the value functions, compute the implied conditional choice probabilities, calculate the likelihood and update the parameters until the latter is maximized (we calibrate the discount factor to $\beta = 0.995$).

[5]Over our sample period ships face an average fuel price of 470\$ per meter ton. Assuming ships travel at an average speed of 13mph, the average daily consumption of fuel is around 20 tons per day (Stopford (2009)). This adds up to a total expenditure for fuel of \$69,100 per week.

[6]We cluster ports into 15 regions (see Table 1). Moreover, the seven groups for c^s_{ij} 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. We have also experimented with other groupings, such as (i) $c^s_{ij} = c^s$ for all i and j; (ii) coarser and finer groups; (iii) c^s_{ij} clustered by the distance between i and j, as well as clustered by the weather between i and j (both capture nonlinear effects of distance on the sailing cost).

Distribution of Freight Values and Freight Entry Costs For each observed price, we can back out the implied valuation v, using the price function, (5). In particular, note that we now have estimates for the freight matching probability, λ^f (it is given by m/f and f is estimated along with the matching function as described above) and the ship values, W_{ij} and J_i since ship costs are now known. Using data on the average value of trade in commodities and the average shipping price, we are able to estimate the bargaining power coefficient, γ, as well (the freight survival probability is calibrated to $\delta = 0.99$). We can now solve for v point-wise from the price equation and obtain the distributions $F_{ij}^v(.)$.

Finally, we recover the production and exporting costs, κ_{ij}, from the observed trade flows. Indeed, the probability that a potential exporter in i chooses to export in j is given by

$$G_{ij} = \frac{\exp\left(J_{ij}^f - \kappa_{ij}\right)}{1 + \sum_{l \neq i} \exp\left(J_{ij}^f - \kappa_{il}\right)} \tag{7}$$

where the values J_{ij}^f are now known from equations (3) and (4). Using our data to compute the number of loaded trips from i to j as well as external data on the total production of relevant commodities for each region i, we can estimate G_{ij} for all j, as well as G_{i0}. We recover κ_{ij} from (7) as in Berry (1994).

Results The main parameter estimates are given in Table 1. Briefly, we find that the fuel cost c_{ij}^s exhibits relatively low variation over space, consistent with the fact that fuel price dispersion is low. Port costs are large and vary substantially across regions. Freight valuations also differ significantly across regions: Brazil and then Australia are the highest value exporters. We also find that freight valuations correlate with the commodities produced in each region; for instance, countries that produce a lot of grain, which is the highest value commodity, tend to have higher valuations. Finally, the estimated matching function and freight distribution are reasonable: we find a high correlation between the retrieved number of the searching exporters and the realized exports (measured via Comtrade data). For further details on the parameter estimates, see BKP.

	Port Costs c_u	Sailing Costs c_s	Exporters Valuations μ_v	Preference Shock σ
North America West Coast	2.458	0.693	79.605	
	(0.07)	(0.002)	(2.038)	
North America East Coast	2.271	-	103.145	
	(0.021)	-	(2.229)	
Central America	1.846	0.693	73.161	
	(0.022)	(0.002)	(3.007)	
South America West Coast	1.996	0.693	59.063	
	(0.026)	(0.002)	(1.679)	
South America East Coast	2.563	-	125.877	
	(0.027)	-	(3.001)	
West Africa	1.421	0.64	57.838	
	(0.015)	(0.002)	(2.658)	
Mediterranean	1.637	0.568	59.87	
	(0.018)	(0.003)	(2.475)	
North Europe	1.399	0.568	49.199	
	(0.009)	(0.003)	(1.959)	
South Africa	2.478	0.64	99.074	
	(0.035)	(0.002)	(2.907)	
Middle East	1.273	0.568	44.203	
	(0.007)	(0.003)	(2.355)	
India	1.48	0.624	84.722	
	(0.014)	(0.003)	(4.2)	
South East Asia	1.67	0.56	72.282	
	(0.008)	(0.002)	(3.324)	
China	1.438	0.558	66.382	
	(0.01)	(0.002)	(3.61)	
Australia	2.635	0.56	70.507	
	(0.025)	(0.002)	(2.543)	
Japan-Korea	1.53	0.558	55.589	
	(0.022)	(0.002)	(2.514)	
				0.117
				(0.0008)

Note: all the estimates are in 100,000 USD. Standard errors computed from 500 bootstrap samples.

Table 1: Ship costs and exporter valuation estimates. The sailing cost for the East Coast of North and South America is set equal to 0.69 (the fuel cost).

4 Fuel Cost and Trade Elasticity

We now consider how a change in the cost of fuel affects maritime trade. We can think of this as the impact of an oil price shock on trade through transport costs. The overall effect is shown in Figure 1. The left panel plots the percent change in world exports, defined as the sum of realized matches between freights and ships over all ports, against the percent change in fuel costs captured by c_{ij}^s in our setup. The right panel plots the corresponding elasticity. The change in trade is substantial: the elasticity is estimated at 0.35 at the average fuel cost in our data and ranges from 0.1 to 1.2. This elasticity is lower than the results of the literature summarized in the Introduction.

A striking feature of the estimated trade elasticity, is its pronounced asymmetry in low and high levels of fuel costs. Indeed, the elasticity gets steeper as the fuel cost increases, while it plateaus at low fuel cost levels. For instance, consider a symmetric decline and increase of 30% from the weekly fuel expenditure in our sample of about 69,100\$. When the fuel cost increases by 30%, total trade falls by about 15%; in contrast, when the fuel cost declines by 30%, total trade grows by about 8%, which is substantially less. As we argue below, this asymmetry is generated by the equilibrium of the transportation sector and in particular, the relative bargaining position of ships and exporters.

Since fuel costs are a core input in the transportation "production", higher fuel costs naturally lead to higher shipping prices, all else equal. Fuel costs, however, do not act just as inputs for ships; they also determine their relative bargaining position. When fuel costs increase, ships become "captive" to their current location: as it is costly to ballast elsewhere, ships are constrained in their ability to exploit opportunities in other regions. This limits the competition for ships between exporters across different regions and allows exporters to pay lower shipping prices than otherwise. In contrast, when fuel costs decrease, it is cheaper for ships to reallocate in space. Ships are less "tied" to their current region and it is now the exporters that are in a weaker bargaining position and thus forced to agree to higher prices than otherwise. This indirect effect of fuel price shocks, coupled with the direct effect of fuel cost on the cost of transporting a cargo, determines the level and shape of the elasticity plotted in Figure 1.

Formally, from the price equation (5), a decline in the fuel cost c_{ij}^s, directly reduces the value of a traveling ship, W_{ij}, by reducing its cost (see equation (1)). At the same time however it affects ships' outside options, captured by the J_i term (see equation (3)). This effect tends to "dampen" the overall reduction in prices, and freight entry, when fuel costs fall.

In order to understand the asymmetry in Figure 1 note that the effect of fuel cost shocks on the ships'

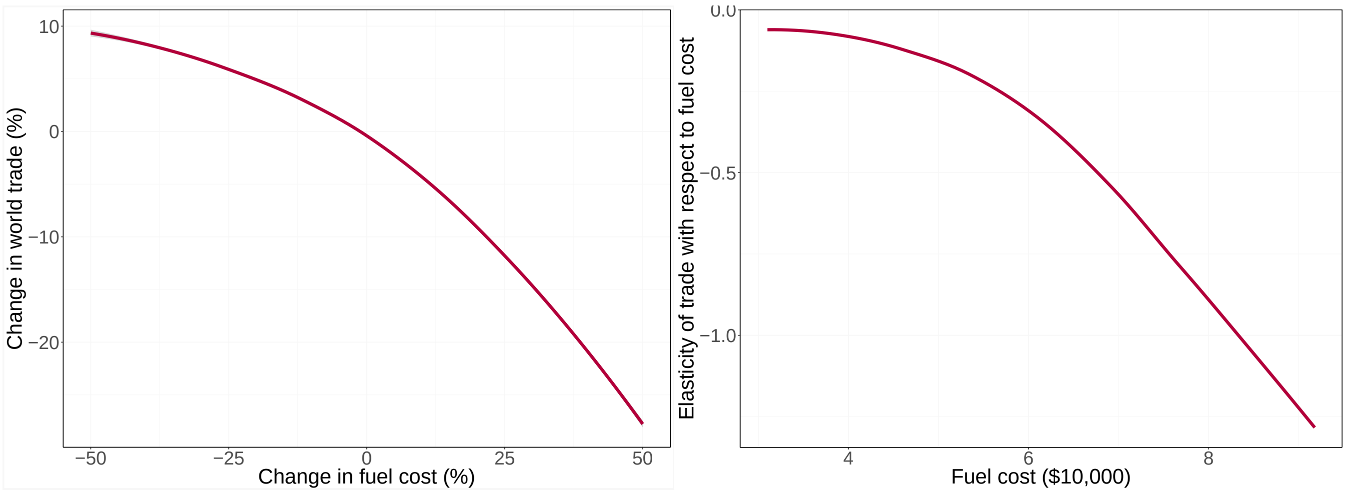

Figure 1: This figure uses the estimated model to compute changes in total trade, defined as the sum of realized matches of freights and ships over all ports, in reaction to changes in the fuel cost. The left panel shows the total change in world trade for different shocks to fuel costs, while the right panel plots the corresponding elasticity. As described in Section 3.3, in the baseline estimation we calibrate the c^s_{ij} to be equal to the average weekly fuel expenditure within our sample period, which is around $69,100. To produce the figure above, we use the model to simulate the equilibrium level of exports associated with different counterfactual values of fuel expenditure c^s_{ij}. The range of counterfactual costs that we consider is from $26,000 to $80,000, which corresponds to plus and minus 50% changes to our baseline cost.

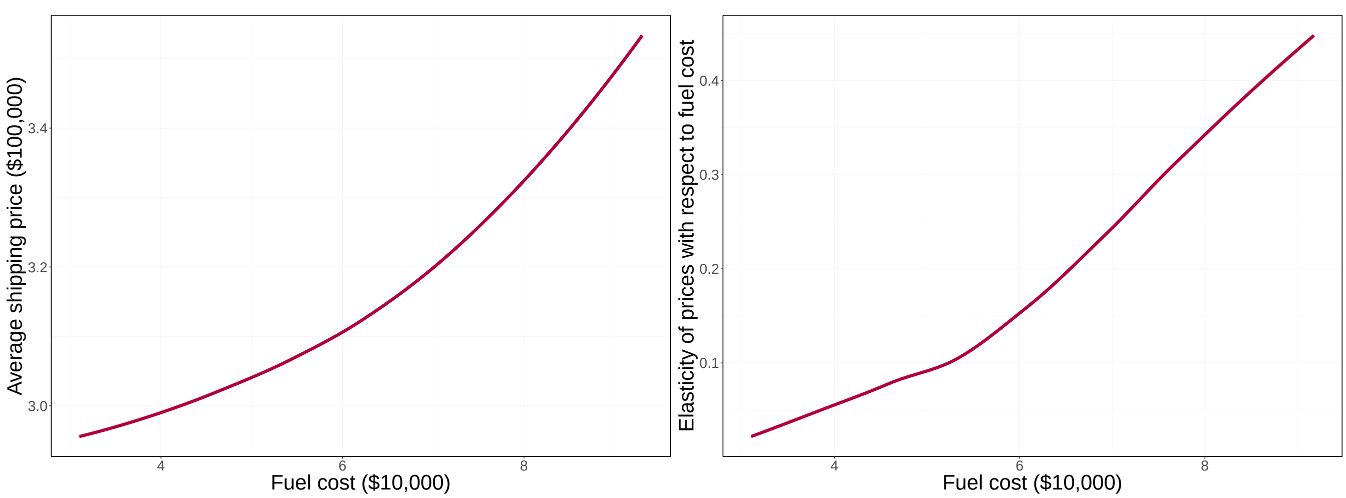

Figure 2: This figure uses the estimated model to compute shipping prices as a function of the fuel price. The left panel shows the average shipping prices for different fuel costs, while the right panel plots the corresponding elasticity. As described in Section 3.3, in our baseline estimation we calibrate c^s_{ij} to be equal to the average weekly fuel expenditure within our sample period, which is around $69,100. To produce the figure above, we use our model to simulate the equilibrium value of shipping prices (τ_{ij}) associated with different counterfactual values of fuel expenditure c^s_{ij}. The range of counterfactual costs that we consider is from $26,000 to $80,000, which correspond to plus and minus 50% changes to our baseline cost.

bargaining position is different for different levels of the fuel cost, c_{ij}^s. In a world with high fuel costs, ships' bargaining position reacts less to changes in the fuel cost. Indeed, at high fuel costs, the ship is likely to stay put rather than ballast, especially when in loading regions (net exporters). As a result, the price decline is not dampened as above and the increase in trade is sharper. Formally, when fuel costs are high, the outside option of ships J_i, is roughly equal to value of remaining unmatched in the current region, U_i, especially in exporting countries. A change in c_{ij}^s has little direct effect on U_i and therefore J_i, and the dampening impact of c_{ij}^s on prices is very small: most of the reduction in c_{ij}^s shows up directly in prices which leads to a correspondingly large increase in freight entry in net exporting countries.

Consider now a world with low fuel costs. In contrast to the situation above, it is now relatively cheap for ships to reallocate in space. In this "flat" world, a ship is likely to choose the ballasting rather than the waiting option, and as a result, declines in the fuel cost have a large impact on its outside option. This mutes substantially the increase in trade and the elasticity of trade shrinks dramatically. Formally, the outside option of ships, J_i, is more likely to equal to one of the traveler values W_{ij} and a reduction in c_{ij}^s, can now have a sizable increase in , J_i , leading to a dampening in the price decline and a smaller increase in freight entry. It is worth noting that the same mechanisms are also present in an economy without search frictions, since ships' outside options are affected in a similar manner.

We also examine the pass-through of fuel cost shocks to shipping prices. Figure 2 displays shipping prices as a function of fuel costs, as well as the elasticity of shipping prices with respect to fuel costs. The elasticity is equal to 0.17 at the average fuel cost level and ranges from 0.03 to 0.43. Naturally, it features the same asymmetry. The level of the elasticity suggests that the pass-through of fuel costs to exporters is relatively low. Our estimates are similar to those found in Hummels (2007). The low pass-through, as well as the varying range of the elasticity suggests that transport costs are not well approximated by fuel costs and that modeling the equilibrium of the transportation sector is important in understanding the nature of trade costs.

We explore further our proposed mechanism by investigating separately the behavior of net exporters vs. net importers. Figures 3 and 4 plot exports, shipping prices and trade elasticities for different fuel costs levels for net exporting and net importing regions separately. The graphs are telling: at high levels of the fuel cost, for net exporters decreases in fuel cost are associated with substantially steeper increases in trade than for net importers. Indeed, ships in net exporting countries almost always prefer to stay put rather than ballast away when the fuel costs are high. Therefore, they benefit less from decreases in fuel prices – their bargaining position remains weak and any decline in fuel costs leads to large reductions in

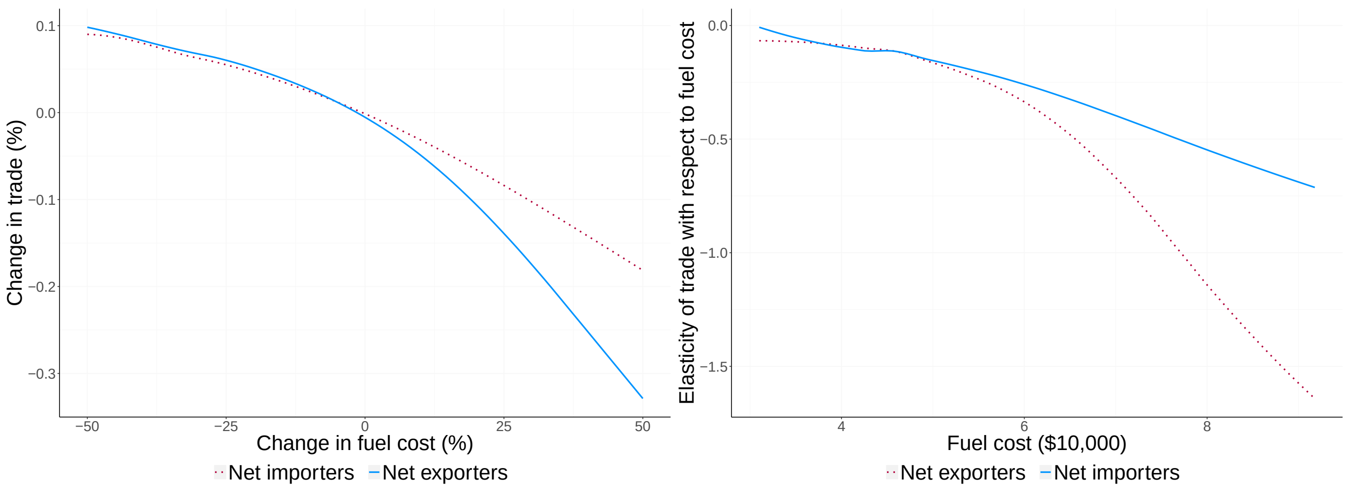

Figure 3: This figure uses the estimated model to compute changes in trade, defined as the sum of realized matches of freights and ships over all ports, in reaction to changes in fuel price for net importers and net exporters. The left panel shows the total change in trade for different shocks to fuel costs, while the right panel plots the corresponding elasticity. As described in Section 3.3, in our baseline estimation we calibrate c_{ij}^s to be equal to the average weekly fuel expenditure within our sample period, which is around \$69,100. To produce the figure above, we use our model to simulate the equilibrium values of trade associated with different counterfactual values of fuel expenditure c_{ij}^s. The range of counterfactual costs that we consider is from \$26,000 to \$80,000, which correspond to plus and minus 50% changes to our baseline cost. We classify a region as a net importer if the number of incoming cargoes is higher than the number of outgoing cargoes and vice versa.

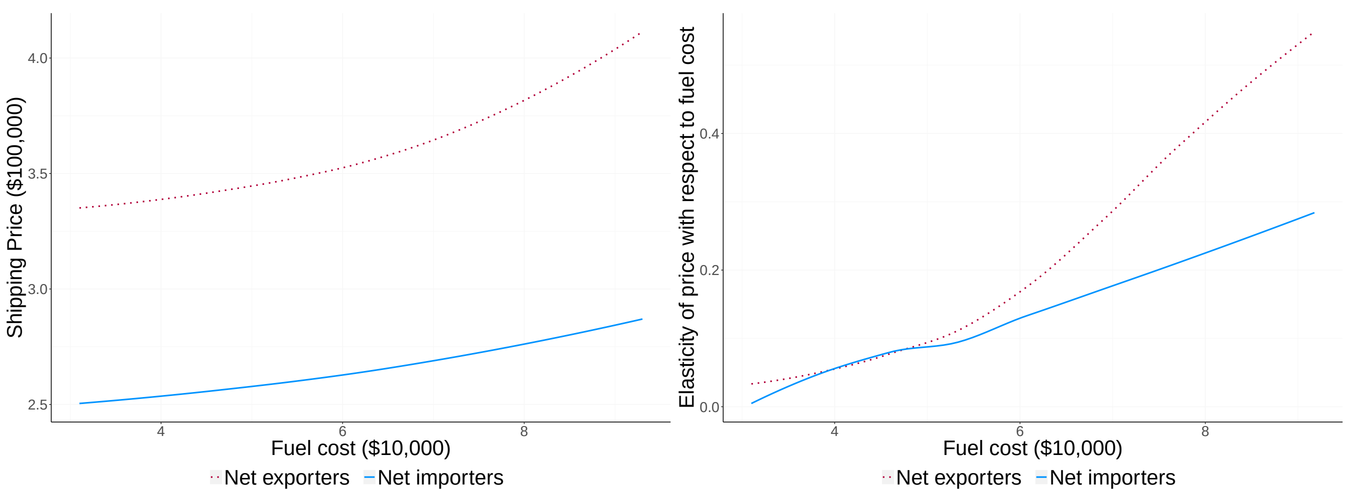

Figure 4: This figure uses the estimated model to compute shipping prices as a function of the fuel price for net importers and net exporters. The left panel shows the average shipping prices for different shocks to fuel costs, while the right panel plots the corresponding elasticity. As described in Section 3.3, in our baseline estimation we calibrate c_{ij}^s to be equal to the average weekly fuel expenditure within our sample period, which is around \$69,100. To produce the figure above, we use our model to simulate the equilibrium value of shipping prices (τ_{ij}) associated with different counterfactual values of fuel expenditure c_{ij}^s. The range of counterfactual costs that we consider is from \$26,000 to \$80,000, which correspond to plus and minus 50% changes to our baseline cost. We classify a region as a net importer if the number of incoming cargoes is higher than the number of outgoing cargoes and vice versa.

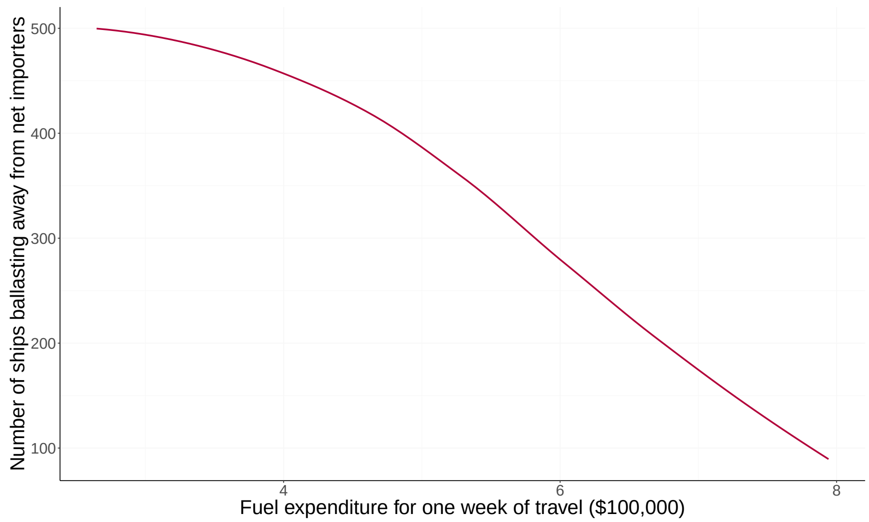

Figure 5: This figure uses the estimated model to compute the number of ships ballasting away from net importers as a function of the fuel price. As described in Section 3.3, in our baseline estimation we calibrate c_{ij}^s to be equal to the average weekly fuel expenditure within our sample period, which is around $69,100. To produce the figure above, we use our model to simulate the equilibrium number of ballasting ships associated with different counterfactual values of fuel expenditure c_{ij}^s. The range of counterfactual costs that we consider is from $26,000 to $80,000, which correspond to plus and minus 50% changes to our baseline cost We classify a region as a net importer if the number of incoming cargoes is higher than the number of outgoing cargoes.

shipping prices and thus large increases in trade. Therefore, a decline in fuel cost when the fuel cost is high, disproportionally benefits net exporters, widening the gap between countries. As fuel costs further decline, however, the world becomes flat and importers and exporters become more similar to each other, as ships can costlessly reallocate.

Finally, it is worth noting that, as fuel costs decline ships do ballast more and there is a reallocation across space from net importers to net exporters. However at low levels of fuel costs, further declines lead to small changes in reallocation: at these levels most opportunities are taken up and there are considerably fewer ships waiting at net importers that are able to take advantage of the reduction in fuel cost to ballast to net exporters. This is shown in Figure 5 where the number of ships ballasting away from net importers flattens out at sufficiently low levels of the fuel cost. This also implies a much lower trade elasticity to oil shocks, as shown in Figure 1. The world fleet utilization is close to full capacity, given the geographical constraints.

5 Oil Shocks and Price Dispersion

In this section we search for descriptive evidence that the equilibrium in the transportation sector and the relative bargaining position of ships and exporters are important determinants of world trade elasticities with respect to oil shocks. To do so, we consider a prediction of our setup regarding the correlation between fuel oil prices and the cross-sectional dispersion of shipping prices.

Indeed, consistent with the narrative of the previous section, as fuel costs decline and the world becomes flat, ship values tend to equalize over space. For instance, a ship in China is now not much worse off than a ship in Australia, as ballasting from China to Australia is cheaper. As the dispersion in ship value functions declines, so does the dispersion in shipping prices

To make this argument more precise, consider again the price equation, (5), for a given origin i: differences in prices across different destinations j are driven by differences in the value function of a traveling ship, W_{ij} (holding constant the distribution of freight valuations). As c^s_{ij} falls, W_{ij} tend to equalize across *destinations*, j, both because differences in distance are less important, but also because differences in the ship's value, U_j, across different destinations matter less. As discussed above, ending a trip in a destination that is attractive to ships is less important than it used to be due to the cheaper ballasting.

In addition, again from the price equation, (5), holding constant differences in freight valuations and matching probabilities, any difference in the shipping price across *origins*, i, is driven by differences in the value of an unmatched ship, J_i.[7] From the discussion above, as fuel costs fall, so do the differences in J_i across i. Nash bargaining implies that all else equal, differences in prices should also shrink.

It is worth emphasizing that the decline in the dispersion of prices *per-day* both across destinations, but also across origins, is driven by the convergence in the ship valuations across space, resulting from a lower cost of ballasting. In a model where prices depended exclusively on distances and ships' outside options did not affect them, the dispersion of prices *per-day* should not depend on the level of fuel costs.[8]

Therefore, a testable implication of our framework is that shipping price dispersion should be increasing the fuel cost. Tables 2 and 3 report the results from a regression of the dispersion of shipping prices on fuel cost. Table 2 shows that as fuel prices increase, the dispersion of per-day shipping prices for trips with different origins increases. Similarly, Table 3 shows that as fuel prices increase, the dispersion of shipping prices per-day for trips with different destinations increases. In both cases, the coefficient of the fuel cost

[7]Note that, as shown in equation (1), conditions at the origin do not affect the value of W_{ij}.

[8]It should be noted that search frictions are not necessary for this result. In a world without search frictions, but where prices are still endogenous, ships' outside options continue to matter and are affected in a similar manner.

is positive and significant, and robust to the inclusion of different time fixed effects.

	I	II	III
	Dispersion of per-day-shipping price across different origins, log		
Fuel Price (log)	0.288***	0.741***	0.597**
	(0.090)	(0.194)	(0.236)
FE	None	Year	Quarter × year
Observations	175	175	175
R^2	0.056	0.735	0.917

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 2: This table reports the estimates from a regression of the dispersion of shipping prices across trip origins on fuel costs. For all the regressions, we compute the average shipping price per-day that different exporting countries face within a month. The dependent variable in all regressions is the standard deviation of this average shipping price across origins (in logs). The main independent variable is the monthly fuel price (in logs). In column I we report the raw correlation, while in columns II and III we add year and quarter-year fixed effects respectively.

6 Fleet Fuel Efficiency and World Trade

In this section, we compute how much recent trends in the fleet's fuel efficiency have affected trade flows. Figure 6, which is taken from Faber and Hoen (2015), plots an index of energy efficiency over the last 50 years, for different types of ships.[9] Energy efficiency for bulk carriers improved dramatically (by about 25%) in the 1980s and after a short stable period began deteriorating. This observed reversion in energy efficiency is striking and not noted in other transport sectors, such as aircraft or trucks.

Although it may seem unintuitive at first, the reversion makes sense if one considers the determinants of design fuel efficiency. Indeed, ship design is determined endogenously by the "long-run" market equilibrium

[9] In order to assess design fuel efficiency, Faber and Hoen (2015) use data from the IHS Maritime World Register of Ships and from Clarksons World Fleet Register to compute each ship's Estimated Index Value (EIV) and compare it with the EIVs of ships that entered the fleet between 1999 and 2008 (reference line). A value above the reference line means that a ship emits more CO_2 per ton-mile under standard conditions, and it is therefore less efficient than the average comparable ship between 1999 and 2008. The EIV computation takes into account the ship's capacity (deadweight tonnage), main engine power, auxiliary power and design speed. According to the authors, the EIV is a simplified version of the EEDI shown in Figure 6. The interested reader should refer to Faber and Hoen (2015) for more details.

	I	II	III
	Dispersion of per-day-shipping price across different destinations, log		
Fuel Price (log)	0.228**	0.812***	0.766***
	(0.091)	(0.200)	(0.229)
FE	None	Year	Quarter × year
Observations	175	175	175
R^2	0.035	0.718	0.923

$^{**}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 3: This table reports the estimates from a regression of the dispersion of shipping prices across trip destinations on fuel costs. For all the regressions, we compute the average shipping price per-day that different exporting countries face within a month. The dependent variable in all regressions is the standard deviation of this average shipping price across destinations (in logs). The main independent variable is the monthly fuel price (in logs). In column I we report the raw correlation, while in columns II and III we add year and quarter-year fixed effects respectively.

in the shipping and shipbuilding markets; as such it is affected by market conditions (predominantly world trade), fuel costs, as well as environmental policies. For instance, the sharp improvements in efficiency in the 1980s follow the oil crisis of the 1970s. The lag in the shipyards' reaction is also consistent with the time required to produce novel ship designs (Faber and Hoen (2015)). The current deterioration in ship efficiency may originally have been due to the massive increase in trade starting in the 1990s-2000s: at the time, shipyards faced severe capacity constraints and opted for simple and quick to build designs (Kalouptsidi (2014), Kalouptsidi (2017), Faber and Hoen (2015)), while high freight rates made high fuel costs less painful for shipowners. Nowadays, the persistently low fuel costs reduce the shipowners' willingness-to-pay for high energy efficiency.

Overall, the combination of (i) long-term trends in the shipbuilding industry; (ii) technological improvements; and (iii) environmental policies lead to time-varying ship designs in terms of fuel efficiency. Here, we use our estimates for the trade elasticity to calculate (i) the gains in trade because of the improvement in design fuel efficiency since the early 1980s; and (ii) the reduction in trade brought about by the more recent deterioration in fuel efficiency. Our estimated elasticities suggest that the 25% improvement in fuel efficiency since the 1980s have led to a decline in shipping prices by 5.5% and a corresponding increase

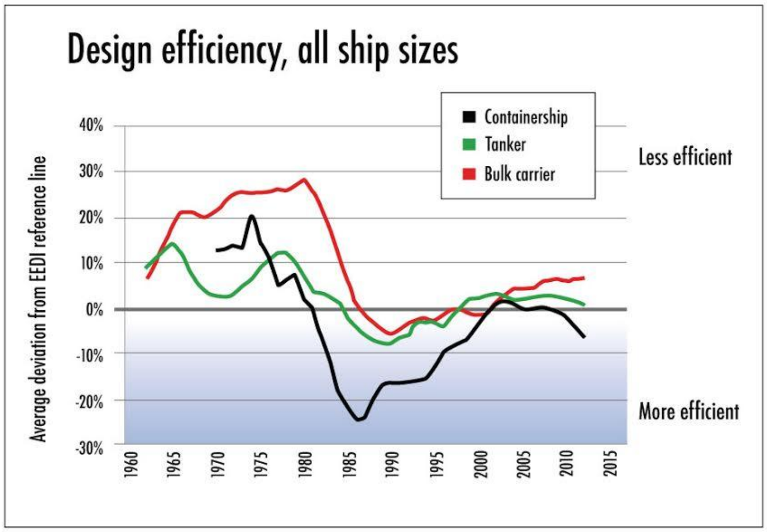

Figure 6: This figure displays the fleet's Energy Efficiency Design Index (EEDI) compared to the EEDIs of ships that entered the fleet between 1999 and 2008 (reference line). The graph is taken from Faber and Hoen (2015). See also footnote 9.

in trade by 11.8%. On the other hand, since the 1990s efficiency has deteriorated by 13%, resulting in a 5.6% reduction in world trade, and a 2.5% increase in shipping prices.

In summary, consistent with the trade elasticity estimated above, we find that ship design is an important determinant of trade. As global institutions are currently considering the phasing in of substantial environmental policies (e.g. the International Maritime Organization (IMO) is currently phasing in limits on sulphur in ship fuel) it is our hope that our methodology and estimates can be of use in the cost-benefit analysis of determining the optimal levels of environmental standards.

7 Conclusions

In this paper, we measure the world trade elasticity with respect to oil prices. We show that the elasticity is substantially asymmetric with respect to high and low fuel costs. In particular, as fuel costs decline, the elasticity becomes flatter. This "flattening-out" of the elasticity is attributed to the stronger bargaining position of ships vs. exporters as distances becomes less important. Indeed, as fuel costs decline and the world becomes flat, ships can leverage their ability to reallocate cheaply and command higher prices. Their stronger bargaining position puts a break on trade growth. The trade elasticity with respect to fuel costs is of great relevance in a number of policy debates regarding environmental, trade and other regulations.

Our approach is not free of caveats. We do not model all the margins along which ships can react to fuel costs. For instance, ships can adjust their speed in response to oil shocks and in particular they can "slow-steam" during oil price spikes. Moreover, we only consider our model in steady state; therefore, firms do not form expectations about oil price fluctuations, nor do they act dynamically with respect to such beliefs. Our elasticity should be considered a short or medium-term one, as we do not model the investment decisions of shipowners in purchasing vessels, nor the reaction of the shipbuilding sector that may change the design of ships. All these issues form interesting avenues for future research. Finally, we consider the impact of oil shocks solely through transportation costs. This is the stated goal of the paper; yet, oil is an key input in production of several industries and understanding its overall impact in the economy is an important question of interest.

References

Adland, R., G. Fonnes, H. Jia, O. D. Lampe, and S. P. Strandenes (2017): "The Impact of Regional Environmental Regulations on Empirical Vessel Speeds," *Transportation Research Part D*, 53, 37–49.

Adland, R. and H. Jia (2016a): "Dynamic Speed Choice in Bulk Shipping," *Maritime Economics & Logistics*, 1–14.

——— (2016b): "Vessel Speed Analytics using Satellite-Based Ship Position Data," *Proceedings of the 2016 IEEE IEEM*, 1299–1303.

Anderson, J. E. and E. Van Wincoop (2003): "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review*, 93, 170–192.

Asturias, J. (2016): "Endogenous Transportation Costs," *mimeo, School of Foreign Service in Qatar, Georgetown University.*

Backus, D. K. and M. J. Crucini (2000): "Oil Prices and the Terms of Trade," *Journal of International Economics*, 50, 185–213.

Berry, S. (1994): "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics*, 25, 242–262.

Brancaccio, G., M. Kalouptsidi, and T. Papageorgiou (2017): "Geography, Search Frictions and Endogenous Trade Costs," *NBER working paper*, 23581.

BRIDGMAN, B. (2008): "Energy Prices and the Expansion of World Trade," *Review of Economic Dynamics*, 11, 904–916.

BUCHHOLZ, N. (2016): "Spatial equilibrium, search frictions and efficient regulation in the taxi industry," *mimeo, Princeton University.*

CHEN, S.-S. AND K.-W. HSU (2012): "Reverse Globalization: Does High Oil Price Volatility Discourage International Trade," *Energy Economics*, 34, 1634–1643.

ERICSON, R. AND A. PAKES (1995): "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *The Review of Economic Studies*, 62, 53–82.

ESTEVADEORDAL, A., B. FRANTZ, AND A. M. TAYLOR (2003): "The Rise and Fall of World Trade, 1870-1939," *Quarterly Journal of Economics*, 118, 359–407.

FABER, J. AND M. HOEN (2015): "Historical Trends in Ship Design Efficiency," Tech. rep., CE Delft and Transport & Environment, Delft, Netherlands.

FRECHETTE, G. R., A. LIZZERI, AND T. SALZ (2016): "Frictions in a Competitive, Regulated Market Evidence from Taxis," *mimeo, New York University.*

HOPENHAYN, H. A. (1992): "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica*, 60, 1127–1150.

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

KALOUPTSIDI, M. (2014): "Time to Build and Fluctuations in Bulk Shipping," *American Economic Review*, 104, 564–608.

——— (2017): "Detection and Impact of Industrial Subsidies, the Case of Chinese Shipbuilding," *forthcoming, Review of Economic Studies.*

KILIAN, L., A. REBUCCI, AND N. SPATAFORA (2009): "Oil Shocks and External Balances," *Journal of International Economics*, 77, 181–194.

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.

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

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.

North, D. (1958): "Ocean Freight Rates and Economic Development 1750-1913," *The Journal of Economic History*, 18, 537–555.

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.

Rubin, J. (2009): *Why Your World Is About to Get a Whole Lot Smaller: Oil and the End of Globalization*, New York: Random House.

Rust, J. (1987): "Optimal replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55, 999–1033.

Stopford, M. (2009): *Maritime Economics*, London and New York: Routledge, 3rd ed.

UNCTAD (2015): "Review of Maritime Transport," Tech. rep., United Nations Conference on Trade and Development, Geneva, Switzerland.

von Below, D. and P.-L. Vezina (2016): "The Trade Consequences of Pricey Oil," *IMF Economic Review*, 64, 303–318.

Wong, W. F. (2017): "The Round Trip Effect: Endogenous Transport Costs and International Trade," *mimeo, University of Wisconsin-Madison*.

Wood, J. (2008): "Fuelling Global Trade: How GDP Growth and Oil Prices affect International Trade Flows," Tech. rep., The Economist Intelligence Unit, London, New York, Hong Kong.