Economies of Scale and Alliances in Container Shipping

Lia Petrose *

(Job Market Paper)

November 1, 2025

Abstract

Containerized shipping is crucial for global trade: it is responsible for two-thirds of it by volume and four-fifths by value. Shipping carriers use large vessels that stop at a sequence of ports, aggregating demand across markets in order to take advantage of potential scale economies. Over the last two decades, ship size and market concentration has increased, driven both by mergers of top carriers and by alliances that offer services jointly. For example, the share of the world fleet operated by the top 10 carriers went from 45% in 2000 to over 80% in 2020. While this could allow carriers to reduce costs by exploiting economies of scale, it can also increase their ability to exercise pricing power. The effect of consolidation on consumers—and thus the merits of regulatory scrutiny—depends on the balance between these two forces. This paper studies the effect of consolidation on consumer welfare by estimating existing economies of scale and by testing alternative models of pricing behavior. I find that there are considerable reductions in capital costs associated with vessel capacity. I also reject models of joint pricing at the service or alliance level in favor of individual carrier pricing. The results point to alliances being an efficient method for realizing cost savings, with the associated increase in market power causing limited harm to consumers.

*MIT Department of Economics. Email: `lpetrose@mit.edu`
I thank my advisors Nikhil Agarwal, Dave Donaldson, and Nancy Rose for their guidance and careful engagement with all parts of the paper. I also thank Mert Demirer, Glenn Ellison, Sara Ellison, Tobias Salz, Ben Vatter, Mike Whinston, David Atkin, Arnaud Costinot, and numerous other participants of the MIT IO Lunch and IO Seminar for their feedback throughout the course of this project.

1 Introduction

Over the past 25 years, containerized ocean shipping has become the most prevalent way to conduct global trade.[1] Its format makes it especially amenable to the transport of final goods—as opposed to bulk commodities—meaning its share of global trade by value is greater even than its volume throughput would suggest. As a transportation service, it has two features that distinguish it from other modes. First, it features highly standardized units that can be arranged interchangeably across vessels. Second, these vessels follow fixed schedules—referred to as liner services—characterized by a sequence of ports and a journey duration. For example, an exporter would purchase container slots from an ocean carrier to get his goods from an origin to a destination port by a given week; the carrier can place those containers on any vessel calling at the destination port within the agreed period. This standardization favors large vessels arriving at high frequency to the served ports, requiring a significant upfront capital investment.

The fixed costs of committing such vessels to a service schedule introduce economies of scale. Historically, carriers have exploited these scale economies by explicitly coordinating on prices and quantities through conferences under exemption from antitrust law. However, as global trade has grown increasingly dependent on this industry, regulators became concerned that the block exemptions provide cover for anti-competitive practices, such as fully collusive pricing. For that reason, the European Commission announced the repeal of those exemptions in 2006, and the policy went into effect in 2008, coinciding with the financial crisis. The years that followed saw the largest decline in demand for container shipping since the energy crisis of 1973, interrupting a long period of sustained growth.

To deal with excess capacity on their books and increased uncertainty about where their competitors' capacity would be allocated, carriers subsequently consolidated the ownership and operations of their vessel fleets. The growth in vessel size magnified this effect, such that the cumulative capacity operated by the top ten carriers more than doubled between 2010 and 2018. For example, the largest vessel in 2000 could

[1] Total global throughput rose from just over two hundred million containers to almost nine hundred million containers over this period.

carry approximately eight thousand containers, measured in twenty foot equivalent units (TEU). By 2006, that number had increased to fifteen thousand, and the largest vessels today can carry over twenty-four thousand TEU.

Service alliances, which are global joint ventures among carriers, also gained prominence in this period. Two of the three largest carriers entered into a new alliance in 2012, and by 2015, the largest and remaining top carrier had also joined an alliance. While alliances, unlike conferences, are barred from discussing prices, member carriers can and do set capacity jointly on their alliance services. This entails selection of the port sequence and vessel fleet, as well as determining service characteristics like total journey duration and frequency of calls at port. Carriers then individually negotiate prices with their customers. Ideally, this form of consolidation is better for consumers because it allows firms to realize the same increasing returns while limiting their exercise of market power. In practice, the cost sharing enabled by alliances may also soften price competition.

This paper answers the question: to what extent do these forces offset one another, and what is the total effect on consumers? To do so, I quantify the extent of scale economies present in this industry, and then perform a test of conduct in order to distinguish the pricing assumption that best fits the observed market outcomes. The data are compiled from a variety of sources. First, I obtain bills-of-lading derived shipping transactions (quantities) at the shipment level from Panjiva. I combine these with market prices at the origin by destination by month level from Drewry Shipping Intelligence. Finally, I obtain rich data on shipping operations and costs from AXSMarine. Together, this database provides a near complete picture of carrier decisions on the margin. One drawback of the data from Drewry is that prices are averaged at the market level, limiting the source of variation used in estimation. However, the estimated own-price elasticity is robust to alternative specifications and exclusion restrictions (discussed in more detail in Section 4), alleviating some of this concern. From AXSMarine, I also obtain details on time charter contracts for vessels dating back decades. I use this database to obtain a realistic measure of the capital costs carriers face by vessel characteristics.

To measure the extent of economies of scale, I estimate the cost parameters allowing for increasing returns at the individual carrier or service level. Marginal costs are allowed to vary with fuel expenses, measured in terms of vessel consumption and monthly marine fuel prices, as well as capital costs, which are predicted vessel rental rates for the service's fleet. I find the presence of scale economies, most apparently through reduced capital costs associated with the opportunity cost of deploying vessels for the duration of the service. Holding those fixed, I conduct a pair-wise test of fit in the spirit of Rivers and Vuong (2002) between candidate pricing models. The test relies on sufficient variation in cost-side instruments to target the difference in markups predicted by the two assumptions on firm pricing behavior. The validity of the test also requires that the performance (in a goodness-of-fit sense) of the two models be sufficiently different: in the extreme case when they perform equally well, the test is degenerate. In this setting, the difference in market shares at the firm and alliance level offsets the large estimated price elasticity to generate enough dispersion in markups. I find strong evidence in favor of individual carrier pricing over joint service pricing.

Using the same testing framework, I also discriminate between specifications of firms' variable costs. I then conduct a series of counterfactual exercises progressively removing the cost benefits of alliance membership. Overall, the benefits from greater investment in capital costs, consequently reducing marginal costs, outweigh the harm even under fully collusive pricing.

The extent to which potential reductions in cost get passed through to consumers in the form of lower prices depends on the exercise of market power by carriers. Sound regulatory policy thus requires estimates of consumers' willingness to pay for shipping services offered by carriers operating in their market. Because data on shipping prices and costs have been difficult to access for researchers, empirical analysis of this type is scarce. This paper brings together rich data on carrier operations and shipping transactions to estimate demand for shipping and cost parameters capturing returns to scale. This provides the necessary foundation to test for anti-competitive conduct and thus assess the effect of consolidation on consumer welfare.

Related Literature: Briefly, this paper contributes to three areas of the empirical industrial organization literature. First, there is a growing body of methods for discriminating between candidate models of firm conduct. Going back to Bresnahan (1982), papers in this vein have used exclusion restrictions on the joint distribution of the idiosyncratic component of demand and supply to distinguish between conduct assumptions with different predictions of markups. These arguments were generalized by Berry and Haile (2014) and applied across a wide variety of contexts by following work. Concurrently, papers testing between models of conduct in a non-nested way have made use of the framework proposed by Rivers and Vuong (2002). Of these, this paper is closest to Backus, Conlon, and Sinkinson (2021). This paper adapts their method to account for returns to scale then applies those insights in order to answer a pressing question in an important setting using novel data.

Second, it relates to the large literature in empirical industrial organization on entry and pricing decisions in the presence of scale economies. For example, Berry, Carnall, and Spiller (2006) estimates economies of density in airlines using the linked nature of multi-stop tickets, Ryan (2012) studies the effects of regulation in a concentrated industry where capacity investments generate increasing returns, and Fan (2013) estimates scale economies in newspaper distribution and advertising.

Finally, it contributes to the growing body of work on the economics of the shipping industry—Kalouptsidi (2014), Brancaccio, Kalouptsidi, Papageorgiou, and Rosaia (2023), Ganapati, Wong, and Ziv (2024), Brancaccio, Kalouptsidi, and Papageorgiou (2024)—and separately, on inferring transportation costs directly from shipping transactions—Hummels et al. (2009), Brancaccio, Kalouptsidi, and Papageorgiou (2020), Shapiro (2016).

The rest of the paper proceeds as follows: Section 2 discusses the setting and institutional details, Section 3 describes the data, Sections 4 and 5 lay out the model and estimation results, Section 6 describes the testing framework and results. Finally, Section 7 assesses the welfare effects of counterfactual policies and Section 8 concludes.

2 Liner Shipping

Container companies are complex entities that provide commercial cargo shipping services along fixed routes with regular schedules. The largest service they offer, containerized shipping, is the primary means of transporting dry non-bulk commodities across borders. Since its introduction in the mid-twentieth century, the standardized container unit has become the most ubiquitous form of packaging goods to be transported over long distances. Similar to other major transportation industries, national or international regulators have played a role in the pricing and capacity allocation of the services provided. For liner shipping, the opening of the Suez canal, and the ensuing surplus of capacity, led to the establishment of associations which limited entry into routes to dues paying members. These associations eventually became the liner conferences prevalent by the end of the 20th century. Following trade liberalization, and especially China's entry into the World Trade Organization (WTO) in 2001, containerized throughput accelerated enough to attract the attention of competition authorities. Citing this growth, the EU Commission announced the repeal of the antitrust exemptions covering liner conferences in 2006, and the repeal went into effect in 2008, coinciding with the financial crisis. Carriers responded by consolidating container capacity in three ways. First, a series of mergers took place between 2012 and 2020 among the top ten carriers. Second, this top decile of firms all joined one of three global alliances. Finally, the continuing trend in increasing vessel size mechanically drove capacity consolidation especially for the top carriers.[2]

Figure 1 shows the membership growth of the major service alliances, and Figure 2 shows the trend in concentration of capacity over the same period. Concurrently, technological or navigational advances relaxed the limits placed on vessel size by choke points such as the Panama Canal, and median vessel size rose dramatically, further consolidating capacity. These changes motivate the question which is the focus of this paper: to what extent do reduced costs from potential returns to scale offset the welfare impact of the increased pricing power carriers now enjoy?

[2]Vessels typically have a single owner and single operator, though the two may be different entities. Because the largest vessels are the most capital intensive, they are disproportionately owned or operated by the largest carriers.

Figure 1: Timeline of alliance activity 2005-Present.

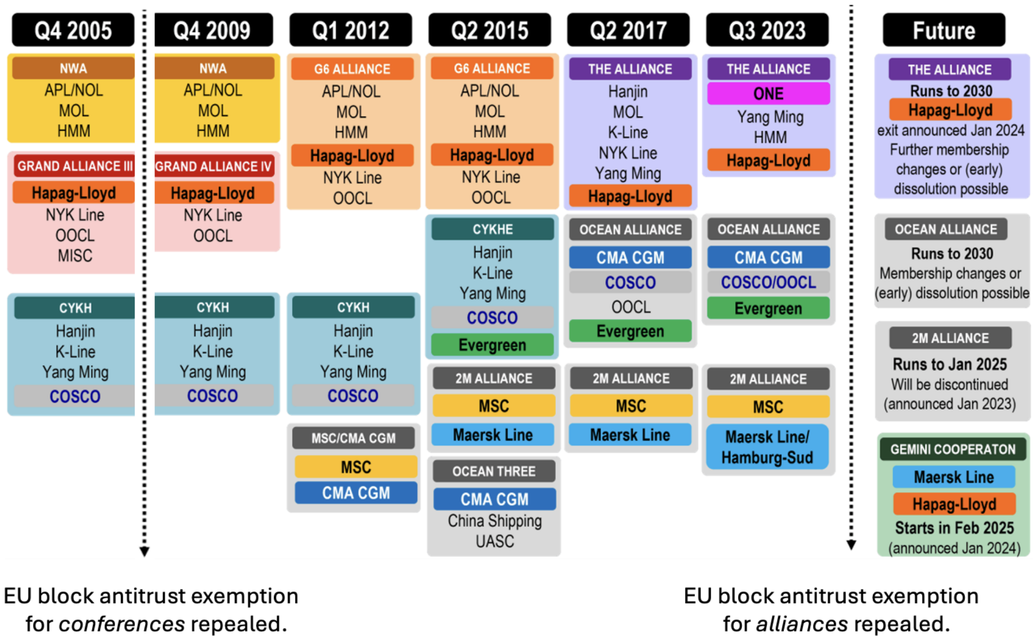

Note: Dashed arrows corresponding to 2006 and 2024 indicate changes in regulatory policy. *Source: Notteboom, Pallis, and Rodrigue 2022.*

Figure 2: Cumulative capacity (L) and share (R) of world fleet operated by the top 10 carriers.

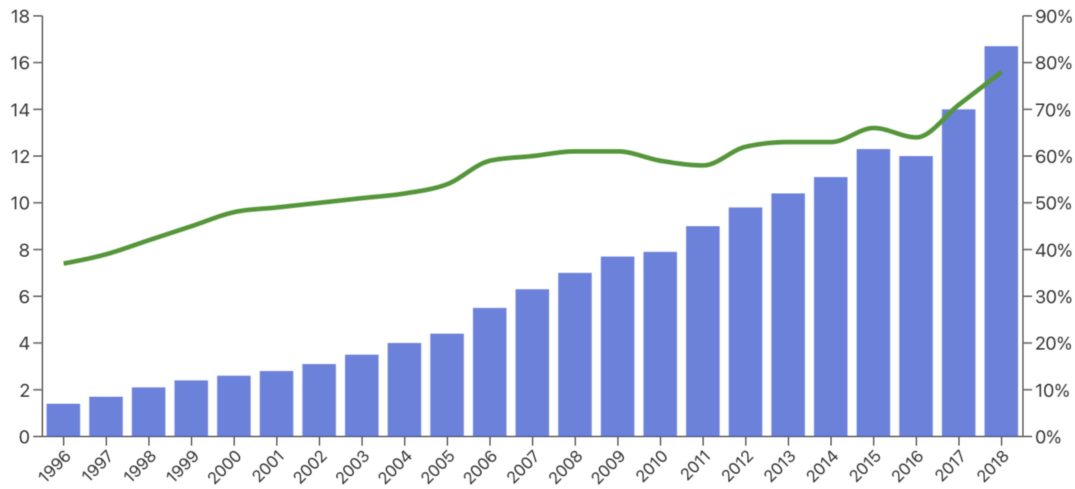

Note: The left axis shows the total capacity, in millions TEU, controlled by the top 10 carriers from 1996 to 2018. The right axis shows this capacity as a share of the world fleet. *Source: Alphaliner.*

3 Data

3.1 Data sources

To estimate the returns to scale, I construct a database of shipping transactions and costs between 2019 and 2024.

Prices and quantities: The core market level data are constructed from three sources. First, I obtain quantities of containers transported from an origin port to a destination port from Panjiva, a data vendor that aggregates bill-of-lading derived customs records. These data are at the shipment level and include the exporting and importing firm, the carrier operating the shipping service, and the number of containers corresponding to the shipment. In addition to the maritime shipments, I also obtain quantities of air cargo at the monthly level. I combine these with data on freight rates, i.e. the price of shipping a container in USD, at the origin-destination port by month level from Drewry Shipping Intelligence. The freight rates are composed of three parts—the base freight rate, origin and destination port handling fees, and a fuel adjustment surcharge—all averaged at the origin by destination by month level. I supplement these with monthly prices for a barrel of crude oil from Macrotrends, and monthly prices for three measures of marine fuel from the United States Department of Agriculture (USDA) transport database. Summary statistics of these variables are reported in Table 1.

Service costs: For carrier costs, I use data from Alphaliner, a vessel database provided and maintained by AXSMarine. Starting in 1999, AXSMarine has collected data on ship movements, ownership structure, and fixture rates for leasing or chartering contracts between ship owners and ship operators. Four of these databases contribute to the dataset used in this paper: the events database on movements and port calls, the services database on liner shipping service offerings by carriers, the vessels database on all vessel characteristics, and the fixtures database on time charter contracts. Importantly, I observe the alliance membership of all carriers and the roles they play in each service offering. Starting with events, I collect all port calls associated with a service offering. Then I average all characteristics for vessels listed under a service—such as age, capacity, speed, draft, and fuel consumption—into those of a service composite vessel. Separately, using the full period of the fixtures database, I fit a model of vessel lease prices by ship type using the age, speed, capacity of the vessel as well as the lease length and lease year. I then predict the lease cost for that composite vessel for a month long lease in the reference year 2018. This forms the main measure of capital costs for each service. To obtain service level fuel costs, I take the product of the average fuel consumption of a vessel on the service, the duration of the service, and the primary measure of fuel price (Fuel 1). Summary statistics of these variables are reported in Table 2.

Finally, I combine the market and services data into a single database of equilibrium container prices, volumes, service characteristics, and service costs, at the month by origin port by destination port level. Carriers who appear as part of a transaction in the market-level data are matched to services under which they are listed if the service is active in both the origin and destination port of that transaction. This allows for a many-to-many match across carriers and services. On average, a service is matched to three carriers and a carrier is matched to two services. To avoid heterogenous weights on observations from multiple matches, I normalize the total number of containers per carrier across all of its services. The four total quantity variables taken to supply estimation are created by summing this measure across carriers (by service) and markets (by route). The remaining cost variables—fuel and capital costs—are constructed out of raw service characteristics. This procedure is described in more detail in section 5.2.

3.2 Descriptive Evidence

Three patterns emerge from these data. First, the trend in concentration worsened for medium and long distance markets over the sample period, perhaps due to the bottlenecks during the 2020 global pandemic. Figure 3—which plots the average value of the Herfindahl-Hirschman Index (HHI) for markets falling within the same distance bins as above—shows it improving for those markets in the following years, but never recovering to its pre-pandemic levels.

Second, the initial difference between prices in markets served by alliances and those that are not—as seen in Table 1—is smaller after accounting for market characteristics

Table 1: Summary statistics of market-level data.

	Total	Service Type	
		Alliance	Non-alliance
Freight Price p_m thou usd	7.54 (4.44)	7.94 (4.46)	5.16 (3.42)
Quantities q_j thou teu	7.09 (8.28)	7.39 (8.13)	5.23 (8.9)
Oil Price usd/barrel	67.3 (24)		
Fuel 1 Price usd(*)/ton	439 (146)	452 (148)	365 (105)
Fuel 2 Price usd(*)/ton	701 (296)	726 (305)	552 (174)
Fuel 3 Price usd(*)/ton	573 (210)	589 (215)	477 (147)

(∗) Fuel prices are normalized by distance.

Table 2: Summary statistics of service-level data.

	Total	Service Type	
		Alliance	Non-alliance
Capacity K_j thou teu	10.2 (2.16)	10.9 (1.18)	5.64 (0.77)
Duration τ_j days	50.7 (6.66)	50.8 (6.09)	50 (9.35)
Consumption φ tons/day	122 (44)	109 (30.9)	200 (27.2)
Vessel Age years	12.1 (2.69)	11.4 (2.02)	16.6 (1.48)
Vessel Speed knots	23.6 (0.81)	23.4 (0.51)	25 (0.81)
Vessel Count	6.47 (1.06)	6.36 (0.97)	7.14 (1.32)
Capital cost C_j thou usd/vessel	20.1 (2.61)	21.1 (1.33)	14.6 (1.06)

and time trends in fuel costs. This can be seen in Figure 4 which overlays the conditional density of prices on the unconditional density by alliance presence. On the left, the distinction between prices in markets with and without an alliance presence becomes smaller when the raw price density is taken conditional on the origin-destination pair,

Figure 3: Average Herfindahl-Hirschman Index (HHI) of markets by distance.

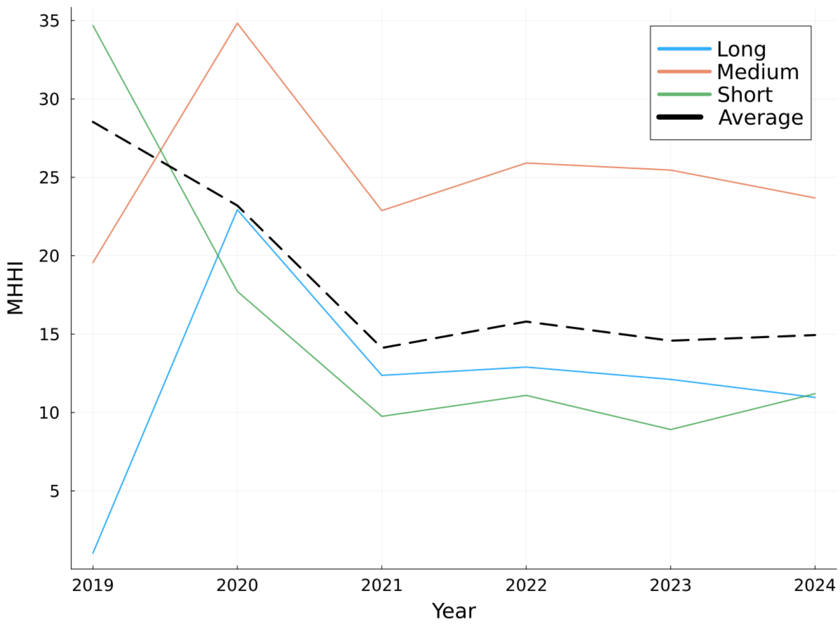

Note: Short distance markets fall within the first quartile of the distance distribution, medium between the first and third, and long distance markets fall within the last.

the origin-quarter, and the destination-quarter. This effect is stronger after accounting for the price of fuel, as shown in the right panel. These figures motivate the inclusion of fixed effects for those factors in sections 4 and 5.

Third, there are persistent differences in the characteristics of alliance and non-alliance services which translate to costs. Figure 5 shows the distribution of average vessel size for markets falling under short, medium, and long distance categories. The top left panel shows a gap between short distance markets and medium or long ones: vessels operating in short distance markets have lower average capacity. The following three panels consider each distance bin separately. The similarity between medium and long distance markets in aggregate conceals heterogeneity by alliance status. Relative to non-alliance services, these figures appear to suggest that alliance services dedicate larger vessels to short distance markets. However, the alliance vessel size distribution stays centered around the same point while that of non-alliance services moves from a relative shift left to a relative shift right across the three panels. This reflects the fact that alliance services are longer; thus, a single alliance service could serve short, medium, and long distance markets with the same set of vessels. Non-alliance services on the other hand tend to have fewer ports

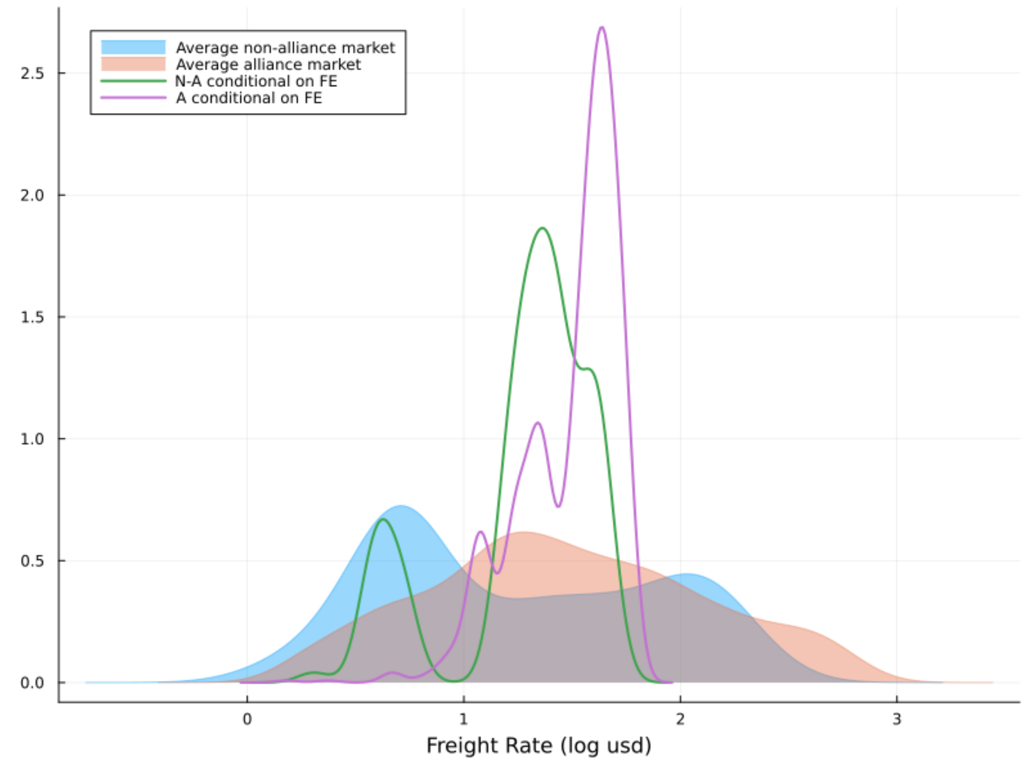

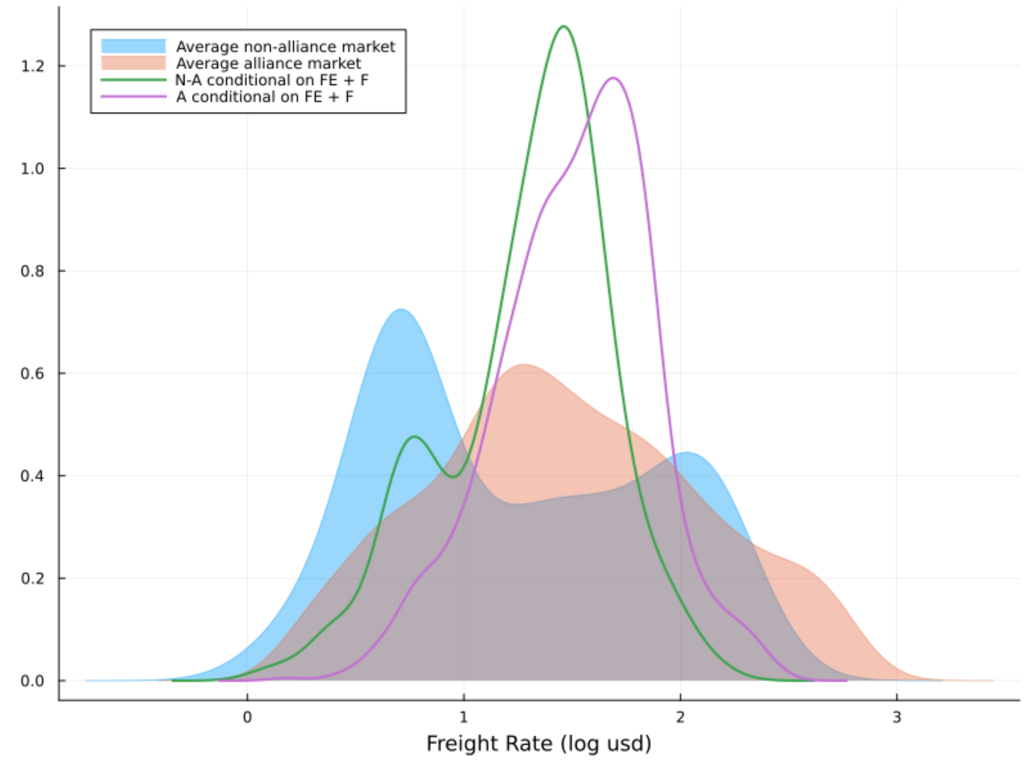

Figure 4: Distribution of p_{jm} and $\hat{p}_{jm}$ by alliance presence.

in their rotation, and display a larger standard deviation in their total journey duration, as confirmed in Table 2.

Figure 5: Mean vessel size by distance between origin and destination.

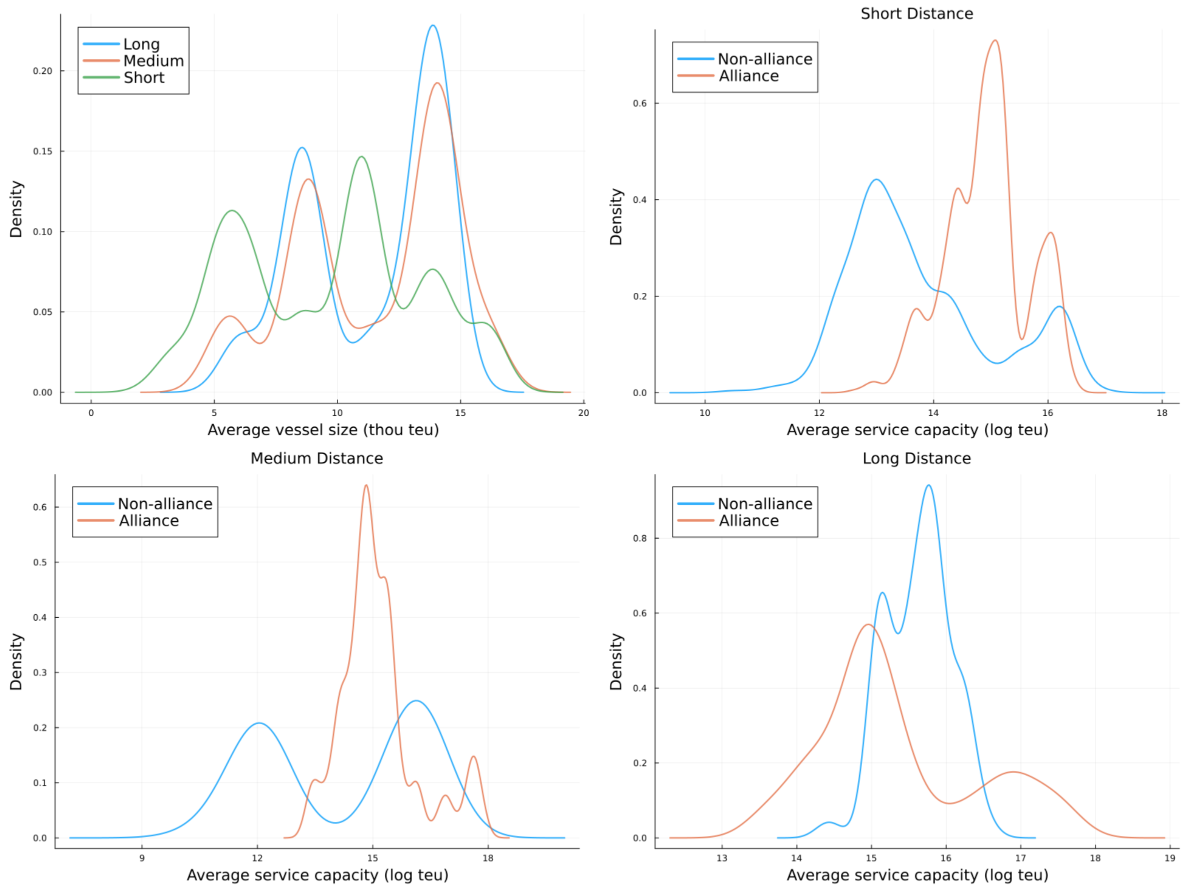

Differences in observable characteristics reflect differences in costs. For example, fuel costs are a direct function of the service duration and vessel consumption profile, and an indirect function of the year the vessel was built. Figure 6 shows the relationship between these vessel characteristics and implied fuel costs. In the left panel, the build years of alliance vessels are shifted right relative to their non-alliance counterparts.[3] The right panel shows the declining relationship between how much fuel a vessel consumes per day at sea and the year it was built, conditional on its carrying capacity. Thus, while both alliance and non-alliance services face fluctuations in the price of marine fuel, an alliance service with younger vessels operating in the same market is less exposed to those fluctuations.

While proxies for marginal costs fall with greater capacity, hinting at economies of

[3]The difference appears not as stark in this figure, but it translates to meaningful differences in average vessel age between the two types of services. The amplified peaks and dips in the early 2010's are evidence of the large capacity adjustments carriers made following the financial crisis.

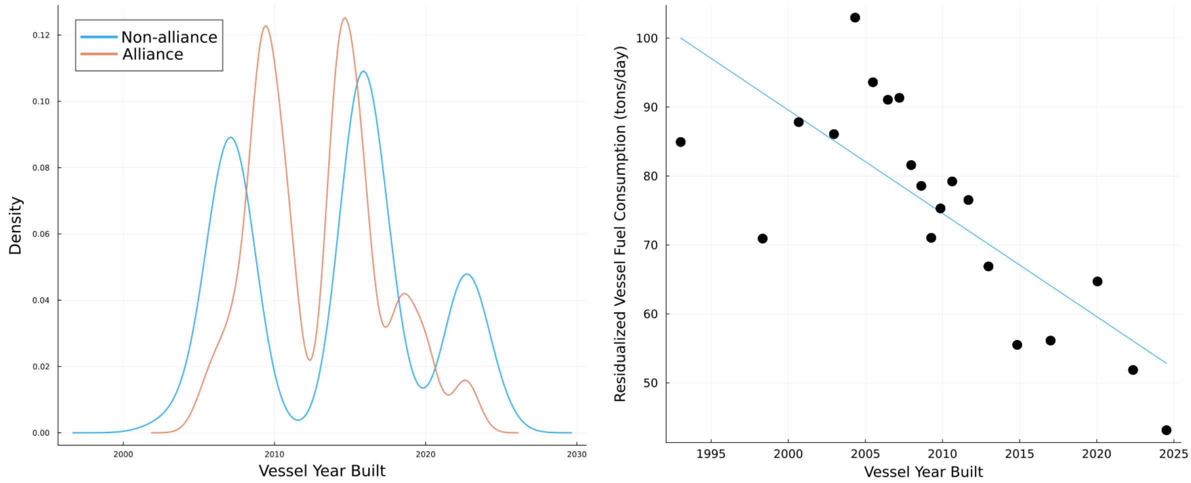

Figure 6: Distribution of vessel build date and its relationship with fuel consumption.

scale, those reductions may be offset by the larger capital cost they entail. For example, in Figure 7, fuel consumption rises with capacity, but unit (per container) consumption declines. Recall that the fuel consumption of younger ships is also lower conditional on capacity; since alliance services have newer, higher capacity vessels, holding all else fixed, they should have lower fuel costs. However, these vessel characteristics are directly associated with higher capital costs by construction. Figure 8 shows a mixed relationship between unit capital costs by service type and distance, reflecting the higher capital costs alliances face for the capacity they offer. Since vessels can be leased on a short term basis, this motivates the inclusion of both fuel and capital costs in the carriers' variable costs, allowing for a direct comparison of this tradeoff. This means that the cost curve, and the returns to scale, that can be estimated in the following sections are smooth over the quantities of containers sold.

3.3 Instruments

The instruments to be used in estimation (Z) are derived from cost variables and market size proxies for demand and cost estimation, respectively. Separating the instrument matrix $Z = [Z^\xi, Z^\omega]$ into cost and demand shifters creates two sets of instruments. First, on the demand side, Z^ξ is composed of i) a proxy for fuel costs is created from marine fuel prices and the distance between the origin and destination port, ii) the average

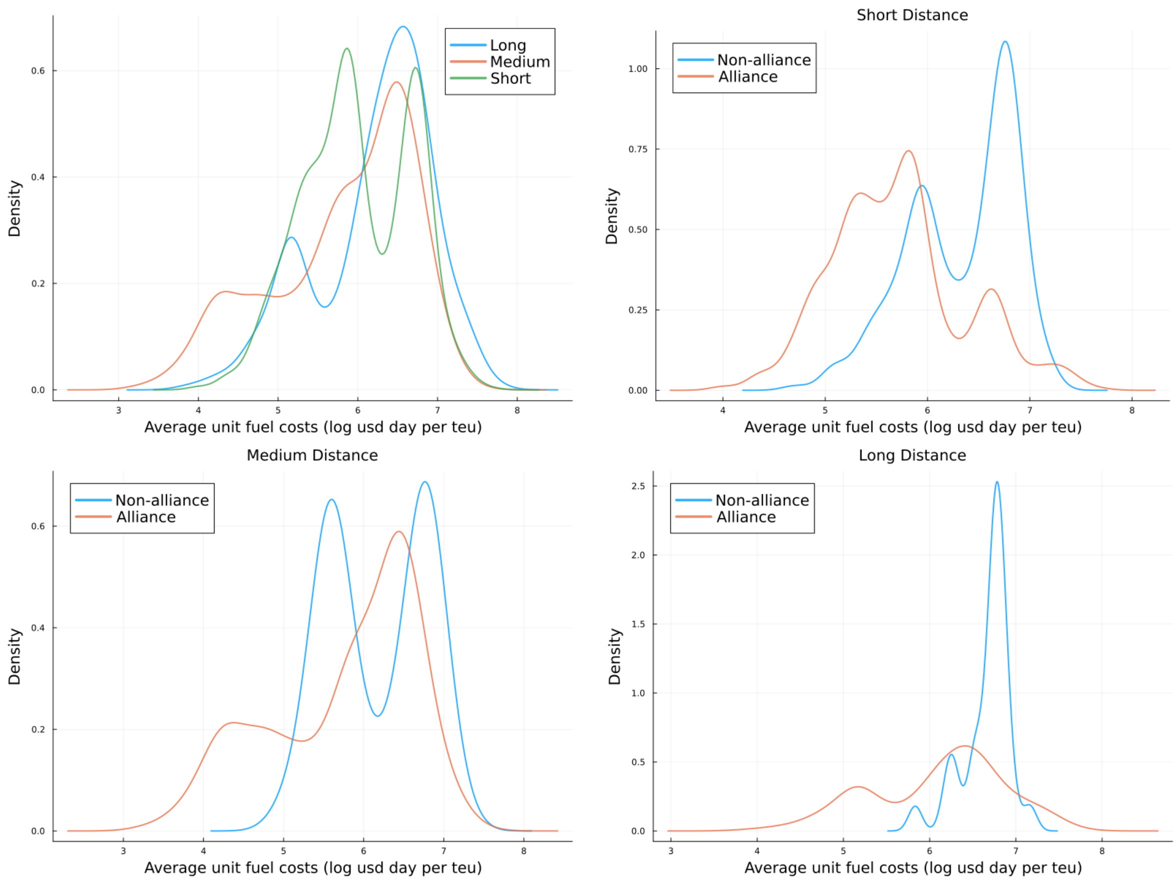

Figure 7: Distribution of average fuel costs by distance and alliance presence.

consumption profile and capital costs of the service vessel fleet, and iii) a proxy for carrier port fees constructed as the predicted importer port fees for a given carrier, route, quarter, and port-pair. Second, on the supply side, Z^{ω} is composed of i) predicted quantities by carrier, quarter, and port-pair, ii) predicted carrier prices by port-pair and quarter, and iii) two measures of marine fuel prices which are not included in costs. Summary statistics of these variables are reported in Table 3.

4 Demand for capacity

Importers and exporters i demand transportation services from origin ports to destination ports in any given month. A product j a slot for a standardized container (one teu) traveling from port o to port d in month t, which jointly define the market, $m := odt$. Consumers of the transportation service value two observed characteristics: price, p_m,

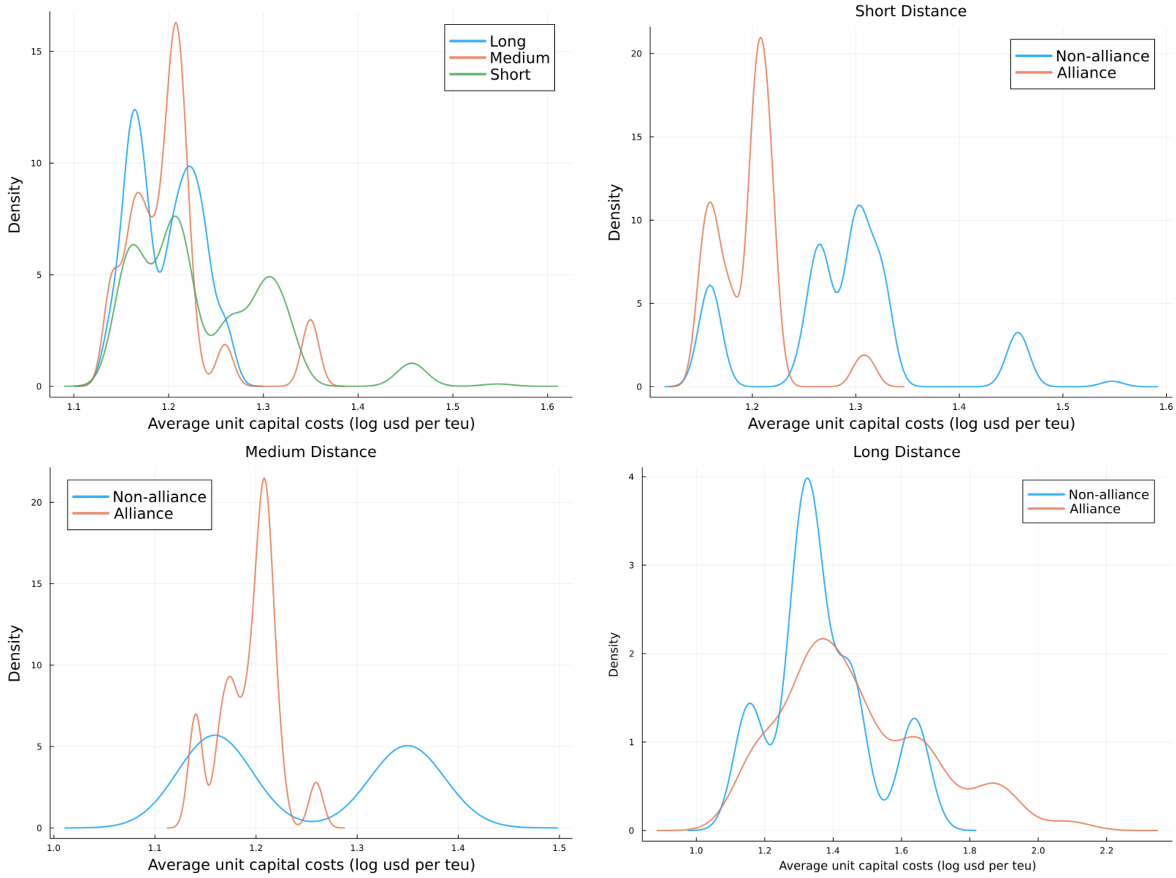

Figure 8: Distribution of average capital costs by distance and alliance presence.

Table 3: Summary statistics of constructed variables.

		Service Type	
	Total	Alliance	Non-alliance
Q_j^f mil teu	1.54 (1.61)	1.67 (1.64)	0.751 (1.11)
$\bar{Q}_j^f$ * mil teu	10.5 (19.8)	11.8 (21)	2.83 (7.09)
Q_{jm} mil teu	2.18 (1.84)	2.51 (1.78)	0.209 (0.318)
$\bar{Q}_{jm}$ * mil teu	14.9 (27.7)	17.2 (29.2)	0.732 (1.73)
Q_j mil teu	7.32 (6.06)	8.41 (5.85)	0.751 (1.11)
$\bar{Q}_j$ * mil teu	50.1 (94)	57.9 (99.3)	2.83 (7.09)

(∗) indicates constructed instrument.

and travel time duration, τ_j. They derive utility from their choice of service $j \in \mathcal{N}_m$ as a function of the observed characteristics and an unobserved quality term (ξ_{j_m}):

$$u_{ij_m} = \underbrace{\beta\tau_j - \alpha p_{j_m} + \xi_{j_m}}_{\delta_{j_m}} + \gamma_m + \epsilon_{ij_m}$$

The subscripts reflect the fact that the price data are aggregated at the market level, and that the duration of a transportation service is identical across all of the markets in which it is offered. Additionally, $\gamma_m = \gamma_{od} + \gamma_{ot} + \gamma_{dt}$ are market effects, further decomposed into distance and quarterly effects. The distributional assumption on the heterogeneity of consumer preferences, i.e. that $\epsilon \sim T1EV$, results in a closed form expression for product j's share in market m of the form:

$$s_{j_m} = \frac{\exp(\beta\tau_j - \alpha p_m + \gamma_m + \xi_{j_m})}{H_m^0 + \sum_{i \in \mathcal{N}_m} \exp(\beta\tau_i - \alpha p_m + \gamma_m + \xi_{i_m})}.$$

H_m^0 represents the value of the outside good, which is air freight in TEU.[4] After normalizing this value to 1, the equilibrium relationship between market shares and product characteristics can then be expressed linearly using the well known inversion from Berry 1994:

$$\log(s_{j_m}) - \log(s_{0_m}) = \beta\tau_j - \alpha p_m + \gamma_m + \xi_{j_m}. \tag{1}$$

4.1 Demand estimation

If market shares are observed with sufficient variation, this relationship can be treated as an estimating equation. The error term (ξ_{j_m}) is interpreted as an unobserved factor which rationalizes the residual variation in market shares not explained by observed

[4]Air cargo is traditionally measured in terms of volume, cubic meters. Panjiva, which provides the quantities data for this paper, estimates a TEU equivalent using the reported cargo weight in the raw data.

characteristics (τ_j and p_m) or fixed market effects (γ_m).

When characteristics are exogenous, an OLS regression of the difference in log market shares, the left hand side of Equation 1, on observed product and market characteristics will yield unbiased estimates of α and β. Results from such a procedure are reported in columns (1) and (3) of Table 4. If the characteristics p_m and τ_j are endogenous, instrumental variables can be used to correct the bias. Results from this procedure are reported in columns (2), (4) and (5) of Table 4. The specification in column (2) isolates price, instrumented for by the cost of marine fuel interacted with distance, as the only characteristic. It yields a high magnitude estimate of α, with an implied average own price elasticity of -4.91. The inclusion of duration does not affect the sign or magnitude of this estimate. The corresponding own duration elasticity is smaller (-0.16), suggesting that the choice of duration is a less flexible margin of adjustment for carriers in response to demand shocks. This is not unreasonable as τ_j is chosen for the full set of markets in which the service is offered whereas prices are set individually in each market.

The instrumental variables in this specification are the distance interacted fuel measure as in column (2), a second fuel measure again interacted with distance, and the average rental rate of the vessels in the service. The fuel based instruments are clearly relevant since they are a direct variable cost faced by carriers. They are plausibly excluded conditional on time fixed effects, since those would absorb macroeconomic shocks such as oil price fluctuations which in turn may be correlated with demand shocks.

4.2 Robustness to definition of market size

So far, market size has been defined as the total quantities observed in the data plus the total quantity of air cargo, approximated into equivalent units (TEU). This may be unsatisfactory for two reasons: i) air freight is not the only substitute for maritime goods transportation and ii) the completeness and quality of the air cargo data is not ideal in this setting. On the first point, truck and rail transportation are plausible substitutes for container shipping, though unlikely for the port-pairs in this sample which are all separated by sizeable bodies of water. On the second point, bills-of-lading derived data, such as those used in this paper, are documented by Flaaen et al. 2022 to have poor coverage of air freight.

Since the price data are aggregated at the market level, identification of α primarily relies on substitution patterns to the outside option. For these reasons, I test the robustness of the results to alternate definitions of market size. On one extreme, I define time-varying shipped goods plus air freight, and on the other, I consider a fixed market size for a given port-pair for the duration of study period. Following Huang and Rojas 2014, I restate the estimating equation such that the outside option quantity is a market-level unobservable on the right hand side which can be controlled for using traditional methods. Defining Q_m as the total size of market m,

$$\log(s_{j_m}) - \log(s_{0_m}) = \log\left(\frac{q_{j_m}}{Q_m}\right) - \log\left(\frac{q_{0_m}}{Q_m}\right) = \beta\tau_j - \alpha p_m + \gamma_M + \xi_{j_m} \tag{2}$$

$$\log(q_{j_m}) = \beta\tau_j - \alpha p_m + \gamma_M + \xi_{j_m} + \log\left(q_{0_m}\right) = \beta\tau_j - \alpha p_m + \xi_{j_m} + \gamma I_m \tag{3}$$

the new estimating equation 3 includes a vector of origin-destination and quarterly market effects which serve as a proxy for the unobserved outside good term, $\log\left(q_{0_m}\right)$. This term can be interpreted as the total quantity of all other modes of goods transport, converted into equivalent units. Columns (2) and (4) report estimates from two-stage least squares regressions analogous to columns (2) and (5) in Table 4, now with additional instruments. The sign of α in column (3), and to a lesser extent column (5), is consistent with the estimates obtained from the logit regression using air freight as the sole outside option, though the magnitude is attenuated.

5 Supply of shipping services

5.1 Price setting

Carrier f sets prices to maximize profits across its services $j \in \mathcal{J}^f$ and the markets m which they serve. Its profits are:

Table 4: Results from logit demand with air freight outside option.

	(1)	(2)	(3)	(4)	(5)
Price	-0.00	-0.85***	-0.00	-0.85***	-0.08***
	(0.00)	(0.07)	(0.00)	(0.07)	(0.01)
Duration			-0.00	-0.00	-0.00
			(0.00)	(0.01)	(0.00)
O X D FE	Yes	Yes	Yes	Yes	Yes
O X QY FE	Yes	Yes	Yes	Yes	Yes
D X QY FE	Yes	Yes	Yes	Yes	Yes
Estimator	OLS	IV	OLS	IV	IV
N	46,703	46,703	46,703	46,703	46,703
R^2	0.64		0.64		
First-stage F statistic		191.87		95.95	1,422.82
Endogenous variables	None	Price	None	Price	Price, Duration
Instruments	None	F	None	F, C	F, C, Po

Notes: Price (p_m) is average freight rate in thousands of USD per TEU for a port-pair-month. Duration (τ_j) is the length in months of the total journey of the service. Instruments used are fuel price × distance (F), service fleet capital cost (C), and predicted port fees (Po) for the average OD pair per carrier. Results are reported from 1,436 unique markets, with a median of 10 carriers per market. Robust standard errors are reported in parentheses.

$$\Pi^f = \sum_{j \in \mathcal{J}^f} p_{j_m} Q_j - VC_j \tag{4}$$

where VC_j is the variable cost associated with service j. It is parametrized as:

$$VC_j = F\left(\bar{\varphi}_j, \tau_j, p^{fuel}\right)^{\theta_1} C\left(\bar{K}_j, \bar{X}_j\right)^{\theta_2} \frac{Q_j^{\theta_3+1}}{\theta_3+1} \exp(\omega_{j_m}) \tag{5}$$

$$\tag{6}$$

where F_j represents the fuel costs for a service with an average fuel consumption of $\bar{\varphi}_j$ tons per day, a duration of τ_j days, facing a marine fuel price of p^{fuel} dollars, and C_j represents the cost of renting an average vessel—as a function of its capacity K_j and characteristics X_j—on the fleet of the service. I suppress firm superscripts—f for the single carrier, F for the service—on Q_j and VC_j because they are defined both ways in the analysis that follows. Marginal costs will then reflect the level at which quantity is defined.

5.2 Estimation of variable costs

I begin by recovering implied marginal costs using the pricing assumption and the estimated elasticity of demand. If variable profits take the form described in the previous section, then a Nash-in-prices equilibrium will satisfy:

$$(p_m - mc^f_{j_m})\frac{\partial s(\boldsymbol{p})}{\partial p} + s^f_{j_m}(\boldsymbol{p}) = 0 \tag{7}$$

$$mc^f_{j_m} = p_m + s^f_{j_m}(\boldsymbol{p})\frac{\partial s(\boldsymbol{p})}{\partial p_m} \tag{8}$$

Equation 8 relates observed market prices and unobserved marginal costs through the markup term $\left(\frac{\partial s^f_{j_m}(\boldsymbol{p})}{\partial p_m}\right)$ which has a simple analytic form under logit demand. I can then take these recovered costs to observed cost shifters to estimate the supply side of the model. Recall that variable costs are parametrized as the product of the monthly fuel costs faced by the service, F, its capital costs C, an aggregate measure of total quantities at the service level Q_j, and an idiosyncratic shock ω_{j_m}, such that marginal costs are:

$$mc_j = F\left(\bar{\varphi}_j, \tau_j, p^{fuel}\right)^{\theta_1} C\left(\bar{K}_j, \bar{X}_j\right)^{\theta_2} Q_j^{\theta_3}\exp(\omega_{j_m}) \tag{9}$$

$$\log(mc_{j_m}) = \theta_1\log(F_j) + \theta_2\log(C_j) + \theta_3\log(Q_j) + \omega_{j_m}. \tag{10}$$

With implied costs in hand, equation 10 can be taken directly to the data. The right hand side variables are constructed as follows. Fuel costs are defined as the cost of fuel for the consumption profile of an average ship in the service's fleet. They are a function of the price of a commonly used fuel source for ships $\left(p^{fuel}\right)$, the average fuel consumption

(φ), and the duration of the service in months (τ_j). Capital costs are predicted vessel lease rates from a model fit on the full history of lease contracts in the data. They are a function of average vessel capacity $\left(\bar{K}_j\right)$, average characteristics $\left(\bar{X}_j\right)$—age, speed, and service duration—as well as the lease year.

To capture many potential mechanisms through which scale economies may operate, Q_j is defined in four ways, two at the service level and two at the carrier level. At the carrier level, it is the firm's own quantity sold in the corresponding market (q_{j_m}) and the firm's own quantity sold at the route level Q_j^f. At the service level, it is the total quantity sold by the service in the corresponding market (Q_{j_m}) and the total quantity sold by the service in the trade lane (Q_j). Each measure corresponds to scale economies operating at the market (port-pair), or route (trade lane) level and within or across the component firms offering a service. I then consider two assumptions about firm conduct: first that firms jointly operating a service are also pricing jointly, and second that they are pricing individually.

Table 5 contains results from the preferred cost specifications across models of conduct. Overall, costs are decreasing with quantities, suggesting the presence of economies of scale. Under the assumption of individual carrier pricing, increasing returns appear most pronounced in capital costs, which I interpret as investment in the form of committed capacity. Conditional on market quantities, there are also returns to scale at the route level, supporting this interpretation. These results hold for implied marginal costs from a joint pricing assumption as shown in columns (4) - (6). Tables 6 and 7 report results for a larger range of cost specifications using service-level costs.[5] The first six columns correspond to Q_j defined as the carrier (q_j) and service (Q_j) quantities sold at the route level, and the last four allow costs to vary with two of the Q_j measures. The coefficient θ_2 is negative and statistically significant across pricing assumptions. This reflects the high fixed capital costs carriers incur when committing capacity to a given service. Fuel costs are always associated with higher marginal costs, as would be expected given the higher vessel count and average consumption profile of smaller services. In all of the reported specifications, both fuel and capital costs are assumed endogenous and instrumented for

[5]Results from a goodness-of-fit test across these are reported in Section 6.1.

by the set of cost shifters Z. The coefficient θ_3 is negative most consistently for route level quantities sold by the service. When inferring costs from an individual pricing assumption, service quantities at the market level increase costs, potentially because of higher handling costs for fuller vessels. These results hold when measuring costs at the container instead of the service level.

6 Testing

6.1 Test of conduct

To distinguish between individual carrier pricing and joint pricing at the service level, I perform a pair-wise test in the spirit of Rivers and Vuong 2002. The non-nested framework of the test allows comparison in the fit of two models without taking either as the default. The intuition behind the test is that the fit of a pricing model as measured by a criterion function can be compared in a two-sided sense against the null hypothesis that both fit equally well. A good starting point for this function is the objective used to estimate the supply parameters ($E[\boldsymbol{\omega}|Z] = 0$). To get more power out of this test, we can use insights from the optimal instruments literature and transform Z to isolate residual variaiton due to the predictions of each pricing assumption. Because the relevant difference between the two is in their predicted markups, the objective should include instruments which best capture the component of unexplained variation in prices due to this difference. Following Backus, Conlon, and Sinkinson 2021, I take as the expected markup difference conditional on the full set of instruments Z.

More specifically, given a set of supply instruments Z and a known cost specification $c(\Theta; C, F)$, I test the null hypothesis that both assumptions on conduct (h_1 and h_2) are equally far from the true model h_0, under which $E[\omega^{h_0}|Z] = 0$. For model h_i, the supply relationship between prices and costs is characterized by the firms' first order conditions:

$$\underbrace{p_m - \eta^{h_i}_{j_m}}_{mc_{j_m}} = \underbrace{F\left(\bar{\varphi}_j, \tau_j, p^{fuel}\right)^{\theta_1} C\left(\bar{K}_j, \bar{X}_j\right)^{\theta_2} Q_j^{\theta_3} \exp \omega^{h_i}_{j_m}}_{c(\Theta;C,F)}$$

$$\log(mc_{j_m}) = \theta_1 \log(F_j) + \theta_2 \log(C_j) + \theta_3 \log(Q_j) + \omega^{h_i}_{j_m}.$$

When estimating the parameters of this supply relationship, (Θ), instrumental variables are used to address bias arising from the endogeneity of cost variables C, F, Q, as discussed in subsection 5.2. The challenge in formulating instruments in this context is different. I need a set of strong instruments for the residual variation in ω_{j_m} conditional on $c(\Theta)$ due to the markup difference between the two conduct assumptions is in η^{h_i}. The set of instruments Z used in subsection 5.2 are valid, so the remaining task is to find a transformation, $A(Z)$, such that the test is sufficiently powered to target this difference. I follow Backus, Conlon, and Sinkinson 2021 and define this function as $A(Z) = E[\Delta\eta|Z]$, the predicted difference in markups $\Delta\eta = \eta^{h_1}_{j_m} - \eta^{h_2}_{j_m}$ conditional on Z. In practice, I take the predicted markup difference, $\Delta\eta := \widehat{\Delta\eta}$, from both a log-linear and a flexible function of the instruments. I then evaluate the criterion function under each model according to equation 11.[6] The difference in the value of this objective for h_1 and h_2 captures the degree to which either model fits the data relative to the other. Under the null, it can be summarized by the statistic in equation 13.

$$Q(\eta^{h_i}) = \left(\frac{1}{n}\sum_{j,m} \widehat{\Delta\eta} \cdot \omega^{h_i}_{j_m}\right)^2 \tag{11}$$

$$\tag{12}$$

$$T = \frac{\sqrt{n}Q(\eta^{h_1}) - Q(\eta^{h_2})}{\sigma} \sim \mathcal{N}(0,1) \tag{13}$$

Positive (negative) values greater (lesser) than (-)1.96 would imply rejection of the null in favor of model 2 (1). I bootstrap this procedure $N = 1000$ times to obtain standard

[6] Because this is a scalar-valued function, it circumvents the challenges concerning the choice of a weighting matrix highlighted by Hall and Pelletier 2011.

errors. Results are presented in Table 8 and Table 9. Across cost specifications, individual firm pricing is favored over joint pricing at the service level. The flexible functional forms for both $A(.)$ and $c(.)$ provide the strongest evidence in support of this. While the specifications in 5.2 allow for interpretable relationships between fuel and capital costs and implied marginal costs, the functional form is restrictive.

The testing framework also allows for comparison between specifications of costs. Taking $A()$ to be the expected price conditional on the instruments Z, I follow the same procedure as described above. Table 10 reports results. The specification corresponding to column (8) in Tables 6 and 7 is favored in most tests, so I use it as the primary cost specification going forward.

7 Welfare

Taking the demand and cost parameters are given, I investigate the effect of removing alliance membership benefits on marginal costs and prices. As a benchmark case, I take the fitted values from the preferred specification plus the joint pricing markup as counterfactual prices under a collusive regime. Figure 9 shows the distribution of welfare by pricing assumption on the left panel, and across markets served by alliances against those which are not on the right panel. Under the current baseline, alliance markets have higher prices and consequently consumers are worse off on average in those markets. This difference is exacerbated when alliance carriers price jointly. Because this model is rejected by the testing results, these changes are used as a benchmark case in evaluating the magnitude of the following counterfactual exercises.

The first thought experiment involves eliminating the premium alliances enjoy in vessel rental costs. For a given level of capacity, age, speed, and design of vessel rented in the same year, alliances face lower lease costs than non-alliance carriers. In the counterfactual scenario, I remove this premium, raising the capital costs alliances must pay for the same vessel fleet. Because higher capital costs imply lower marginal costs, alliance services lower their prices and welfare rises in the markets they serve. Figure 10 shows the change in average welfare for markets which are served by at least one alliance and markets which

Figure 9: Consumer welfare by pricing assumption and alliance presence.

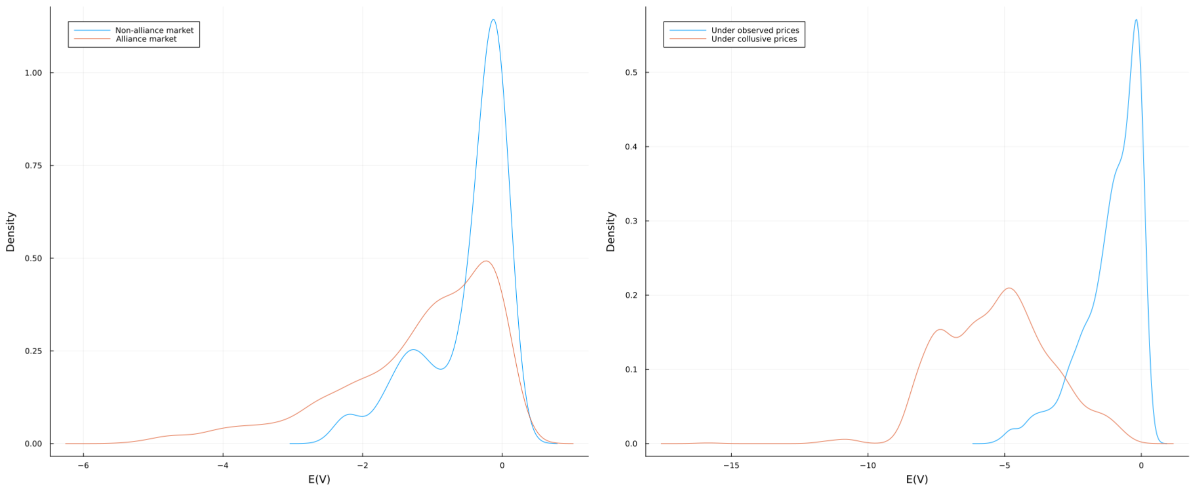

are not.

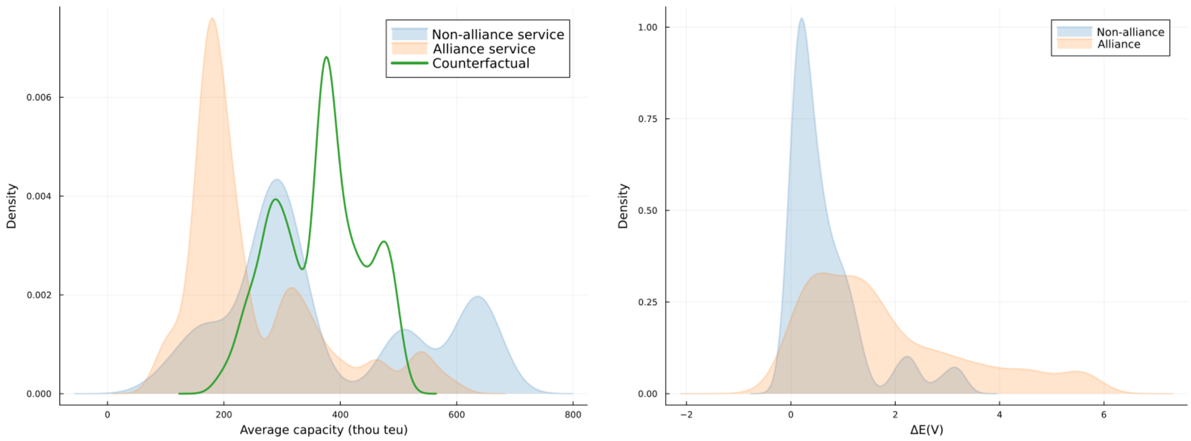

Figure 10: Change in welfare due to loss of alliance premium in capital costs.

The main feature associated with alliances in regulatory discussions assessing block exemptions is the higher capacity they can offer. In the second scenario, I bring closer together capacity levels observed in alliance and non-alliance services. Assuming the dissolution of alliances would imply a reduction in capacity levels, I compute counterfactual prices if the current alliance services set their capacity to the level of the median non-alliance service in the same route and going through the same origin port.[7] Since capital costs are primarily a function of capacity, the reduction in capacity leads to a decline in K_j and marginal costs rise. In alliance-served markets, prices fall and consequently the

[7]Non-alliance services are also assigned this median capacity, so this scenario can also be interpreted as a reduction in capacity for the majority of services in the sample.

change in welfare under this scenario is positive and largest in alliance served markets. This can be seen in Figure 11: the left panel shows the change in capacity in the counterfactual scenario, while the right panel shows the resulting change in welfare as a function of alliance status.

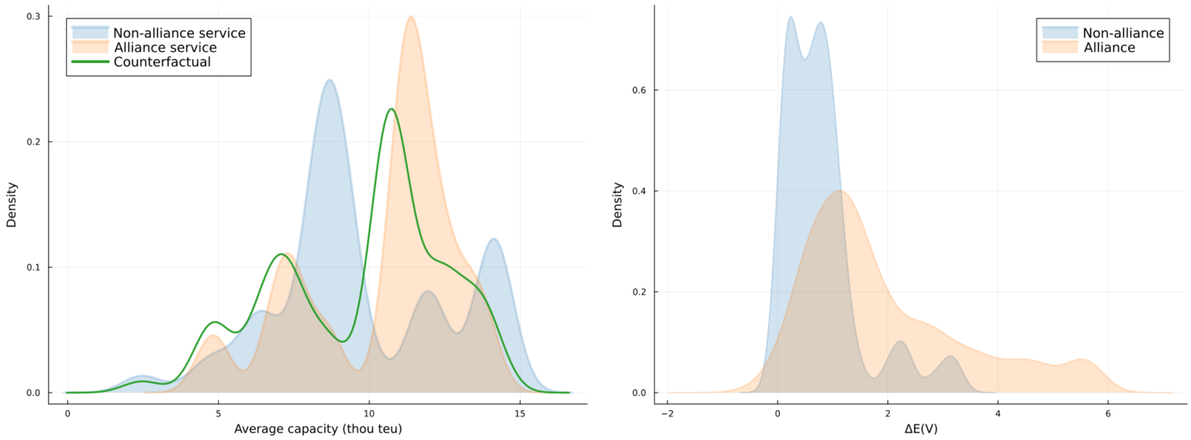

Figure 11: Change in welfare due to reduction in service capacity by alliance status.

Finally, I allow capital costs to depend on vessel flag, simulating the effect of flag based port fees proposed in 2025. Relative to Bermudan flagged vessels, all else equal, an American flagged vessel costs 65 times more to lease. For markets including an American port, I set the flag-based capital costs of all vessels serving that market to that of the U.S. flagged vessels. In most cases, this significantly increases the capital costs associated with each service, as seen in the left panel of Figure 12. Because marginal costs are decreasing in capital costs, prices decrease in the markets affected by this and the average consumer in those markets is left better off. The right panel of the same figure shows this change for markets falling under short, medium, and long distance bins.

8 Conclusion

To conclude, this paper presents novel results on firm conduct in the container shipping industry. Regulators who have grown increasingly concerned about capacity consolidation among the top liner companies have few empirical studies to inform their judgements. In particular, a countervailing force against welfare loss due to anti-competitive pricing could be cost synergies arising out of increasing returns to scale. The scheme that is

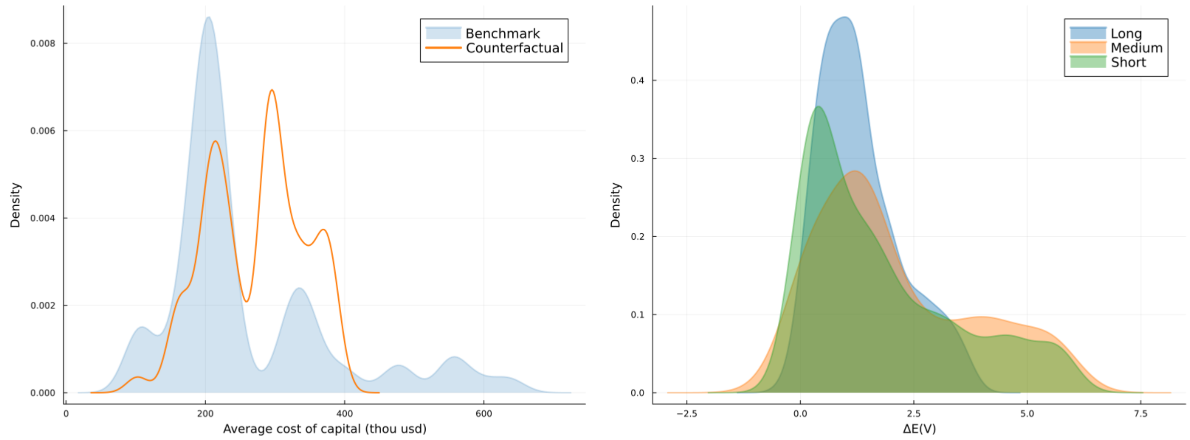

Figure 12: Change in welfare due to increase in regulation induced changes in costs.

currently under consideration—service alliances—allows carriers to jointly set capacity but not prices. However, it may still dampen price competition, harming consumers. It is crucial to understand the extent of economies of scale in the industry in order to assess the real effect of such a scheme on welfare.

To fill this gap, I build a dataset of shipping transactions and carrier costs using a variety of sources. These data are especially suited for the estimation of economies of scale because they contain detailed information on vessel and service level costs. Using these data, I estimate the contribution of fuel and capital expenses, as well as total number of containers shipped, to implied marginal costs. I repeat this exercise for a variety of assumptions on firm pricing behavior, most notably individual carrier pricing and joint venture (service level) pricing. The results strongly support increasing returns to scale, especially through capital cost investments. I then test across the assumptions on carrier pricing behavior, adapting the pair-wise testing approach of Rivers and Vuong (2002) and Backus, Conlon, and Sinkinson (2021). Individual carrier pricing is favored over alternative assumptions across the majority of specifications.

A limitation of the approach presented in this paper is that it assumes away selection into markets by alliance membership. This would be resolved by allowing carriers to make an entry decision then estimating the cost parameters conditional on that decision. Because market entry is correlated through service route design, comprehensively accounting for this choice is beyond the scope of this paper and is left for future work. However, the robustness of these results to alternative exlusion restrictions alleviates some of the concern about selection. As they are, the results in this paper provide crucial evidence regulators need to inform antitrust policies which are currently under consideration.

References

Backus, Matthew, Christopher Conlon, and Michael Sinkinson (2021). "Common Ownership in the Ready-to-Eat Cereal Industry". In: *NBER Working Paper*.

Berry, Steven (22 1994). "Estimating Discrete-Choice Models of Product Differentiation". In: *The RAND Journal of Economics* 25.2, p. 242. ISSN: 07416261. DOI: 10.2307/2555829.

Berry, Steven, Michael Carnall, and Pablo T. Spiller (May 2006). "Airline Hubs: Costs, Markups and the Implications of Customer Heterogeneity". In: *Advances in Airline Economics: Competition Policy and Antitrust* 1, pp. 183–214. DOI: 10.3386/W5561.

Berry, Steven and Philip Haile (Sept. 2014). "Identification in Differentiated Products Markets Using Market Level Data". In: *Econometrica* 82.5, pp. 1749–1797. ISSN: 1468-0262. DOI: 10.3982/ECTA9027.

Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papageorgiou (Mar. 2020). "Geography, Transportation, and Endogenous Trade Costs". In: *Econometrica* 88.2, pp. 657–691. ISSN: 1468-0262. DOI: 10.3982/ECTA15455.

– (May 2024). "Investment in Infrastructure and Trade: The Case of Ports". Cambridge, MA. DOI: 10.3386/W32503.

Brancaccio, Giulia, Myrto Kalouptsidi, Theodore Papageorgiou, and Nicola Rosaia (June 2023). "Search Frictions and Efficiency in Decentralized Transportation Markets". In: *Quarterly Journal of Economics*. DOI: 10.3386/W27300.

Bresnahan, Timothy F. (Jan. 1982). "The oligopoly solution concept is identified". In: *Economics Letters* 10.1-2, pp. 87–92. ISSN: 0165-1765. DOI: 10.1016/0165-1765(82)90121-5.

Fan, Ying (2013). "Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market". In: *American Economic Review* 103.5, pp. 1598–1628. DOI: 10.1257/aer.103.5.1598.

Flaaen, Aaron et al. (2022). "Bill of Lading Data in International Trade Research with an Application to the COVID-19 Pandemic * †". In.

Ganapati, Sharat, Woan Foong Wong, and Oren Ziv (2024). "Entrepôt: Hubs, Scale, and Trade Costs". In: *American Economic Journal: Macroeconomics* 16.4, pp. 239–78. ISSN: 1945-7707. DOI: 10.1257/MAC.20220250.

Hall, Alastair R. and Denis Pelletier (2011). "Nonnested Testing In Models Estimated Via Generalized Method Of Moments". In: *Econometric Theory* 27 (2), pp. 443–456.

Huang, Dongling and Christian Rojas (Jan. 2014). "Eliminating the outside good bias in logit models of demand with aggregate data". In: *Review of Marketing Science* 12.1, pp. 1–36. ISSN: 15465616.

Hummels, David et al. (2009). "The trade reducing effects of market power in international shipping". In: *Journal of Development Economics* 89.1, pp. 84–97. ISSN: 0304-3878.

Kalouptsidi, Myrto (Feb. 2014). "Time to Build and Fluctuations in Bulk Shipping". In: *American Economic Review* 104.2, pp. 564–608. ISSN: 0002-8282. DOI: 10.1257/AER.104.2.564.

Notteboom, Theo, Athanasios Pallis, and Jean-Paul Rodrigue (2022). *Chapter 7.1 – Ports, Policies, and Politics*, pp. 1–35. ISBN: 9780367331559.

Rivers, Douglas and Quang Vuong (June 2002). "Model selection tests for nonlinear dynamic models". In: *The Econometrics Journal* 5.1, pp. 1–39. ISSN: 1368-423X. DOI: 10.1111/1368-423X.T01-1-00071.

Ryan, Stephen P (2012). "The Costs of Environmental Regulation in a Concentrated Industry". In: *Econometrica* 80.3, pp. 1019–1061. DOI: 10.3982/ECTA6750.

Shapiro, Joseph S (2016). "Trade Costs, CO 2 , and the Environment". In: *American Economic Journal: Economic Policy* 8.4, pp. 220–254. DOI: 10.1257/pol.20150168.

9 Appendix

Table 5: Results from preferred cost specification across candidate models of conduct.

	Log MC (f)			Log MC (F)		
	(1)	(2)	(3)	(4)	(5)	(6)
Log fuel cost	0.13***	0.15***	0.30***	0.65***	0.27***	0.56***
	(0.01)	(0.01)	(0.01)	(0.03)	(0.03)	(0.03)
Log capital cost	-0.17***	-0.18***	-0.43***	-1.18***	-0.96***	-0.95***
	(0.02)	(0.02)	(0.02)	(0.05)	(0.05)	(0.05)
Service-Route Q	0.00	0.01	-0.14***	-0.08***	-0.29***	-0.04
	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)	(0.02)
Carrier-Route Q		-0.02**			0.36***	
		(0.01)			(0.02)	
Service-Market Q			0.24***			-0.15***
			(0.01)			(0.03)
O X D FE	Yes	Yes	Yes	Yes	Yes	Yes
O X QY FE	Yes	Yes	Yes	Yes	Yes	Yes
D X QY FE	Yes	Yes	Yes	Yes	Yes	Yes
N	46,703	46,703	46,703	46,703	46,703	46,703
First-stage F statistic	391.43	82.87	86.62	87.11	82.87	0.10

Notes: Fuel and capital costs are in log thousands of USD per container-month. Fuel is constructed as the product of i) the average consumption profile of the service vessel fleet, ii) the duration of the service, and iii) a primary measure of ship fuel prices. Capital costs are predicted vessel lease rates for an average vessel in the service fleet. Service-Market Q is the total capacity sold by the service for a given market (in logs) and Service-Route Q is the total capacity sold by the service across all markets in the trade lane (in logs).

Table 6: Results from service-level variable cost estimation assuming individual pricing for firms within a service.

	Log marginal cost									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Log fuel cost	0.21***			0.19***			0.23***	0.22***	0.25***	0.19***
	(0.02)			(0.02)			(0.02)	(0.02)	(0.02)	(0.01)
Log capital cost	-0.24***			-0.25***			-0.24***	-0.24***	-0.26***	-0.23***
	(0.03)			(0.03)			(0.03)	(0.03)	(0.03)	(0.03)
Carrier-Route Q	-0.04***	0.00	-0.00				-0.06***	-0.04***	-0.06***	
	(0.01)	(0.01)	(0.00)				(0.01)	(0.01)	(0.01)	
Service-Route Q				-0.01*	0.00	0.00		0.01		-0.09***
				(0.01)	(0.01)	(0.01)		(0.01)		(0.01)
Carrier-Market Q							0.09***			
							(0.02)			
Service-Market Q									0.03***	0.08***
									(0.01)	(0.01)
O X D FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
O X QY FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
D X QY FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Estimator	IV	IV	OLS	IV	IV	OLS	OLS	IV	IV	IV
N	46,703	46,703	46,703	46,703	46,703	46,703	46,703	46,703	46,703	46,703
R^2			0.64			0.63	0.63			
First-stage F statistic	256.68	3,310.82		10.68	4,247.97		0.27	8.58	122.14	255.61
Endogenous variables	F, C, Q	Q	None	F, C, Q	Q	None	None	F, C, Q, q	F, C, Q, q	F, C, Q, q
Instruments	All	All	None	All	All	None	None	All	All	All

Notes: Fuel and capital costs are in log thousands of USD per product-month. Fuel is constructed as the product of i) the average consumption profile of the service vessel fleet, ii) the duration of the service, and iii) a primary measure of ship fuel prices. Capital costs are predicted vessel lease rates for an average vessel in the service fleet. Service-Market Q is the total capacity sold by the service for a given market (in logs), Carrier-Market Q is the same measure but for individual firms in the service, and Service-Route Q is the total capacity sold by the service across all markets in the trade lane (in logs). Instruments denoted (All) are defined as follows: $\bar{\xi}^f_{od}$ is the constructed port fees instrument used in subsection 4.1, F_2 is the second measure of marine fuel price interacted with distance, r^f_t are route by carrier dummies and $\bar{Q}^f_j, \bar{Q}j_m, \bar{Q}_j$ are the corresponding quantity instrument for each definition of Q employed across all columns. Their construction is described in the text. Robust standard errors are reported in parentheses.

Table 7: Results from service-level variable cost estimation allowing for joint price setting across firms within a service.

	Log marginal cost									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Log fuel cost	0.51***			0.51***			0.52***	0.25***	0.53***	0.43***
	(0.03)			(0.03)			(0.04)	(0.04)	(0.04)	(0.03)
Log capital cost	-0.83***			-0.43***			-0.83***	-0.70***	-0.69***	-0.64***
	(0.06)			(0.06)			(0.06)	(0.06)	(0.06)	(0.05)
Carrier-Route Q	-0.11***	-0.09***	-0.09***				-0.11***	0.24***	-0.07***	
	(0.02)	(0.01)	(0.01)				(0.02)	(0.03)	(0.02)	
Service-Route Q				-0.37***	-0.26***	-0.24***		-0.38***		-0.10***
				(0.01)	(0.01)	(0.01)		(0.02)		(0.02)
Carrier-Market Q							0.03			
							(0.05)			
Service-Market Q									-0.28***	-0.21***
									(0.01)	(0.02)
O X D FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
O X QY FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
D X QY FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Estimator	IV	IV	OLS	IV	IV	OLS	OLS	IV	IV	IV
N	46,703	46,703	46,703	46,703	46,703	46,703	46,703	46,703	46,703	46,703
R^2			0.55			0.56	0.56			
First-stage F statistic	256.68	3,310.82		10.68	4,247.97		0.27	8.58	122.14	255.61
Endogenous variables	F, C, Q	Q	None	F, C, Q	Q	None	None	F, C, Q, q	F, C, Q, q	F, C, Q, q
Instruments	All	All	None	All	All	None	None	All	All	All

Notes: Fuel and capital costs are in log thousands of USD per product-month. Fuel is constructed as the product of i) the average consumption profile of the service vessel fleet, ii) the duration of the service, and iii) a primary measure of ship fuel prices. Capital costs are predicted vessel lease rates for an average vessel in the service fleet. Service-Market Q is the total capacity sold by the service for a given market (in logs), Carrier-Market Q is the same measure but for individual firms in the service, and Service-Route Q is the total capacity sold by the service across all markets in the trade lane (in logs). Instruments are defined as follows: $\bar{\varsigma}^f_{od}$ is the constructed port fees instrument used in subsection 4.1, F_2 is the second measure of marine fuel price interacted with distance, r^f_t are route by carrier dummies and $\bar{Q}^f_j, \bar{Q}j_m, \bar{Q}_j$ are the corresponding quantity instrument for each definition of Q employed across all columns. Their construction is described in the text. Robust standard errors are reported in parentheses.

Table 8: T statistics for pair-wise conduct test across three assumptions of firm pricing.

m^1	NB^f			NB^F		NB^A
m^2	NB^F	NB^M	NB^A	NB^M	NB^A	NB^M
$c(\Theta)1$	-3.94	-0.83	-6.35	2.42	-3.01	0.05
$c(\Theta)_4$	-5.96	-2.99	-5.16	0.24	-7.26	-1.53
$c(\Theta)_7$	-5.04	-2.34	-7.74	1.23	-8.36	-1.14
$c(\Theta)_8$	3.62	-2.89	-2.55	-0.35	-5.16	-1.66
$c(\Theta)_{10}$	0.14	-1.16	-5.58	1.24	-6.41	-0.98

Notes: Cells represent a pair-wise test between the model listed in the first row and that listed in the second row. The values listed are T statistics as described in the text: positive (negative) values greater than (-)1.96 imply a rejection of the null in favor of the second (first) model. NB^f, NB^F, NB^A, and NB^M denote own pricing, joint pricing, alliance pricing, and monopoly pricing, in that order.

Table 9: T statistics for pair-wise conduct test using flexible h().

m^1	NB^f			NB^F		NB^A
m^2	NB^F	NB^M	NB^A	NB^M	NB^A	NB^M
$c(\Theta)_1$	-12.15	-6.39	-6.31	-14.17	-15.28	-4.27
$c(\Theta)_4$	-13.13	-18.11	-17.16	-16.78	-16.69	-15.48
$c(\Theta)_7$	-13.41	-16.66	-15.73	-17.33	-17.11	-14.54
$c(\Theta)_8$	-12.32	-16.93	-16.29	-15.21	-15.8	-15.39
$c(\Theta)_9$	-12.39	-17.7	-16.5	-16.04	-15.68	-14.73
$c(\Theta)_{10}$	-13.45	-17.19	-16.1	-17.79	-16.42	-14.62
c_{flex}	-1.11	-1.32	-1.22	-1.18	-1.31	-0.22

Notes: Cells represent a pair-wise test between the model listed in the first row and that listed in the second row. The values listed are T statistics as described in the text: positive (negative) values greater than (-)1.96 imply a rejection of the null in favor of the second (first) model. NB^f, NB^F, NB^A, and NB^M denote own pricing, joint pricing, alliance pricing, and monopoly pricing, in that order.

Table 10: T statistics for pair-wise test across specifications of firm costs.

c_1	(10)					(9)				(8)			(7)		(4)
c_2	-(9)	-(8)	-(7)	-(4)	-(1)	-(8)	-(7)	-(4)	-(1)	-(7)	-(4)	-(1)	-(4)	-(1)	-(1)
NB^f	-0.06	-2.12	-1.52	-3.38	4.33	-3.47	-2.69	-3.98	4.92	1.32	-2.6	4.76	-2.94	5.23	4.68
NB^F	-3.39	2.84	7.68	7.45	5.35	5.4	7.93	6.52	6.12	5.0	3.98	5.53	-2.8	4.51	4.03

Notes: Cells represent a pair-wise test between the specification listed in the first row and that listed in the second row. The values listed are T statistics as described in the text: positive (negative) values greater than (-)1.96 imply a rejection of the null in favor of the second (first) specification. S and U denote capital and fuel costs at the service and unit level, respectively.

Table 11: Results from logit demand with proxy outside option.

	Log Quantity				
	(1)	(2)	(3)	(4)	(5)
Price	0.00	-0.03***	-0.03***	-0.01	-0.03***
	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)
Duration			-0.00	-0.00	-0.00
			(0.00)	(0.00)	(0.00)
O X D FE	Yes	Yes	Yes	Yes	Yes
O X QY FE	Yes	Yes	Yes	Yes	Yes
D X QY FE	Yes	Yes	Yes	Yes	Yes
Estimator	OLS	IV	IV	IV	IV
N	46,703	46,703	46,703	46,703	46,703
R^2	0.41				
First-stage F statistic		4,033.12	4,034.72	2,134.05	2,017.55
Endogenous variables	None	Price	Price	Price	Price, Duration
Instruments	None	Po	Po	F, Po	Po, C

Notes: Log Quantity is the logged total number of containers sold in a port-pair-month by an individual carrier. Price (p_m) is average freight rate in thousands of USD per TEU for a port-pair-month. Duration (τ_j) is length of total journey of the service in months. Instruments used are the primary marine fuel price × distance (F), service average capital cost (C), and predicted port fees (Po) for the OD pair by carrier as well as their interaction with F. Results are reported from $1,436$ unique markets, with a median of 10 carriers per market. Robust standard errors are reported in parentheses.

Table 12: Results from unit (per TEU) variable cost estimation assuming individual price setting across firms within a service.

	Log marginal cost									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Log unit fuel cost	0.22***			0.19***			0.23***	0.22***	0.27***	0.18***
	(0.02)			(0.02)			(0.02)	(0.02)	(0.02)	(0.02)
Log unit capital cost	-0.27***			-0.24***			-0.24***	-0.28***	-0.24***	-0.28***
	(0.02)			(0.02)			(0.03)	(0.02)	(0.03)	(0.03)
Carrier-Route Q	-0.03***	0.00	-0.00				-0.06***	-0.04***	-0.07***	
	(0.01)	(0.01)	(0.00)				(0.01)	(0.01)	(0.01)	
Service-Route Q				-0.02**	0.00	0.00		0.01		-0.01
				(0.01)	(0.01)	(0.01)		(0.01)		(0.01)
Carrier-Market Q							0.09**		0.12***	-0.04*
							(0.03)		(0.03)	(0.02)
O X D FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
O X QY FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
D X QY FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Estimator	IV	IV	OLS	IV	IV	OLS	OLS	IV	IV	IV
N	46,703	46,703	46,703	46,703	46,703	46,703	46,703	46,703	46,703	46,703
R^2			0.76			0.76	0.76			
First-stage F statistic	60.50	3,310.82		8.05	4,247.97		23.45	90.83	100.77	79.87
Endogenous variables	F, C, Q	Q	None	F, C, Q	Q	None	None	F, C, Q, q	F, C, Q, q	F, C, Q, q
Instruments	All	All	None	All	All	None	None	All	All	All

Notes: Fuel and capital costs are in log thousands of USD per product-month. Fuel is constructed as the product of i) the average consumption profile of the service vessel fleet, ii) the duration of the service, and iii) a primary measure of ship fuel prices. Capital costs are predicted vessel lease rates for an average vessel in the service fleet. Service-Market Q is the total capacity sold by the service for a given market (in logs), Carrier-Market Q is the same measure but for individual firms in the service, and Service-Route Q is the total capacity sold by the service across all markets in the trade lane (in logs). Instruments are defined as follows: $\bar{\varsigma}_{od}^{f}$ is the constructed port fees instrument used in subsection 4.1, F_2 is the second measure of marine fuel price interacted with distance, r_t^f are route by carrier dummies and $\bar{Q}_j^f, \bar{Q}j_m, \bar{Q}_j$ are the corresponding quantity instrument for each definition of Q employed across all columns. Their construction is described in the text. Robust standard errors are reported in parentheses.

Table 13: Results from unit (per TEU) variable cost estimation allowing for joint pricing for firms within a service.

	Log marginal cost									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Log unit fuel cost	0.57***			0.52***			0.52***	0.33***	0.56***	0.46***
	(0.04)			(0.03)			(0.04)	(0.04)	(0.04)	(0.03)
Log unit capital cost	-0.72***			-0.51***			-0.83***	-0.50***	-0.83***	-0.56***
	(0.05)			(0.05)			(0.06)	(0.05)	(0.06)	(0.06)
Carrier-Route Q	-0.18***	-0.09***	-0.09***				-0.11***	0.11***	-0.14***	
	(0.02)	(0.01)	(0.01)				(0.02)	(0.02)	(0.02)	
Service-Route Q				-0.35***	-0.26***	-0.24***		-0.37***		-0.32***
				(0.01)	(0.01)	(0.01)		(0.02)		(0.01)
Carrier-Market Q							-0.28***		-0.24**	0.00
							(0.08)		(0.08)	(0.05)
O X D FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
O X QY FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
D X QY FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Estimator	IV	IV	OLS	IV	IV	OLS	OLS	IV	IV	IV
N	46,703	46,703	46,703	46,703	46,703	46,703	46,703	46,703	46,703	46,703
R^2			0.55			0.56	0.56			
First-stage F statistic	60.50	3,310.82		8.05	4,247.97		23.45	90.83	100.77	79.87
Endogenous variables	F, C, Q	Q	None	F, C, Q	Q	None	None	F, C, Q, q	F, C, Q, q	F, C, Q, q
Instruments	All	All	None	All	All	None	None	All	All	All

Notes: Fuel and capital costs are in log thousands of USD per product-month. Fuel is constructed as the product of i) the average consumption profile of the service vessel fleet, ii) the duration of the service, and iii) a primary measure of ship fuel prices. Capital costs are predicted vessel lease rates for an average vessel in the service fleet. Service-Market Q is the total capacity sold by the service for a given market (in logs), Carrier-Market Q is the same measure but for individual firms in the service, and Service-Route Q is the total capacity sold by the service across all markets in the trade lane (in logs). Instruments are defined as follows: $\bar{\varsigma}^f_{od}$ is the constructed port fees instrument used in subsection 4.1, F_2 is the second measure of marine fuel price interacted with distance, r^f_t are route by carrier dummies and $\bar{Q}^f_j, \bar{Q}j_m, \bar{Q}_j$ are the corresponding quantity instrument for each definition of Q employed across all columns. Their construction is described in the text. Robust standard errors are reported in parentheses.