Market Fragmentation and Inefficiencies in Maritime Shipping

Kostas Bimpikis Giacomo Mantegazza

Salomón Wollenstein-Betech*

Abstract

Maritime transportation is essential for global supply chains, but ballasting—vessels traveling without cargo—imposes significant economic and environmental costs. This paper focuses on the oil transportation industry, where about half of the total traveled miles are sailed empty, and illustrates that market fragmentation is the most important cause to ballasting after demand imbalances, accounting for 17-20% of the total. We find that it is possible to reduce carbon emissions associated with ballasting by as much as 15% by consolidating the market into small shipping pools (sets of vessels controlled by a single operator), which avoids concerns about excessive market power. Consolidation improves utilization because larger pools coordinate their vessels more efficiently and diversify the set of ports they serve, shortening vessels' relocations. At a broader level, this work shows the extent of the sustainability gains that can be obtained solely by organizing more efficiently the resources available in today's supply chains.

Keywords: Transportation markets, Maritime shipping, Fragmentation, Ballasting, Resource pooling, Supply chain sustainability

1 Introduction

Maritime transportation is critical for the global economy and supply chains, with about 90% of the trade of raw materials and finished products occurring by sea: 11 billion tons of goods were shipped in 2021, with volumes projected to triple by 2050.¹ Despite the steady increase in demand, empty miles are endemic in maritime transportation, a phenomenon also known as *ballasting*; it is estimated that between 40% and 50% of all miles traveled by ships are empty.² Empty miles pose a significant economic and environmental cost especially for the dry bulk and oil transportation

^{*}Graduate School of Business, Stanford University. Email: {kostasb, giacomom, salomonw}@stanford.edu.

We are grateful to Ludovico Crippa, Ömer Karaduman, Ilan Morgenstern, Daniela Saban, Stefan Wager for very helpful suggestions.

¹https://www.oecd.org/ocean/topics/ocean-shipping/

 $^{^{2}}$ Note that 50% empty miles resembles the least efficient system where trips are planned as out-and-back tours.

markets, because their specialized vessels are restricted to carry single cargoes from a one origin to a single destination, thus taking away the versatility that other shipping markets feaure.³

Ballasting has attracted considerable attention from both scholars and industry practitioners, mainly to study the role of and remedies to imbalanced demand patterns: if some locations have outflows larger than inflows, the system requires ballasting to be balanced. However, transportation markets usually feature many independent entities that control their transportation assets to provide their services. The interaction between these entities is potentially as important as demand patterns in shaping equilibrium flows, but it has not been as studied. In particular, the degree of market fragmentation affects individual incentives about strategic decisions such as the selection of routes and destinations to serve, which, in turn, affect the overall efficiency of the transportation market. The goal of this paper is to complement existing literature on transportation by focusing on the role of market fragmentation and its contribution to ballasting.

To study the role of market fragmentation, we adopt an applied perspective and partnered with a leading oil tanker operator and market analytics provider. We concentrate on the oil shipping market and ascertain that fragmentation is responsible for a considerable share of ballasting, on top of a "baseline" amount, to be attributed to trade imbalances. At the same time, our analysis shows that partial consolidation of the market can lead to a substantial reduction in ballasting, and therefore mitigate the economic and environmental impact associated with empty miles without excessively increasing the market power of market participants. Because our dataset provides details about the fleets managed by each shipowner, we further study the mechanisms whereby fragmentation exacerbates the challenges posed by structural imbalances. We observe a tendency among smaller shipowners to use their vessels on select routes, where they have previously found loads to serve. Conversely, larger operators exhibit a more flexible approach, diversifying their port destinations and allocating vessels accordingly. Thus, this work is among the first to (i) quantify the impact of fragmentation on ballasting, and (ii) identify the mechanisms through which fragmentation impairs industry efficiency.

We exploit a rich proprietary dataset containing the voyages of approximately six thousand oil tankers from 2018 to 2022. The data allows us to sketch a complete picture of demand, supply, and market structure, because they include information about the *Commercial Operator* of vessels, i.e., the entity that controls and determines the loads of each tanker. From the data we observe stark imbalances in trade flows: for example, oil imports to ports in China and Taiwan are three times larger than oil exports. This pronounced trade imbalance is depicted in the map of Figure 3. Consistent with anecdotal evidence and previous studies, we find that the vessels were empty for

 $^{^{3}}$ In contrast to bulk ships and oil tankers, container ships can be filled with different products, at different capacities, and can make intermediate stops.

42% of the total miles traveled. Moreover, from our data about commercial operators we conclude that fragmentation is extreme: no company controls more than 5% of the market capacity, with a very long tail of companies managing less than 1% of the world's tankers. The left panel of Figure 4 shows a histogram of the size of the commercial operators in the market.

We study the effect of *market fragmentation* on ballasting using applied and empirical methods. From an applied standpoint, we formulate and solve a set of integer and stochastic dynamic programs under different market structures and informational assumptions, and we use the solutions to estimate the share of ballasting associated to imbalanced demand, uncertainty about future loads, and market fragmentation. Intuitively, we find that trade imbalances explain the largest share of ballasting, followed by fragmentation, whose impact is twice as large as that of uncertainty. Finally, we show that even a slightly more concentrated market is substantially more efficient: consolidating oil tankers into shipping pools of 40 vessels each (equivalent to 5% of the global fleet) leads to a 15% reduction in costs and carbon emissions associated with empty miles.

From an empirical perspective, we analyze the strategies of existing operators to understand the mechanisms that underpin the positive effect of consolidation. We concentrate on *shipping pools*, a trend that has gained momentum in the oil shipping industry: with shipping pools multiple shipowners pool their vessels under a unique manager, who centralizes commercial operations. Our reduced-form specifications confirm that larger shipping pools enjoy higher utilization rates (lower fraction of empty miles) and explain that the efficiency gains of consolidation come mainly through two channels. First, we observe a coordination effect, such that vessels managed by the same entity can serve the same locations at a lower cost. Second, we find that increasing the pool size affects the incentives of the manager: larger pools serve a more complex network of locations, which in turn allows them to run more integrated operations, while smaller pools tend to concentrate on fewer routes. Moreover, while smaller pools are more likely to use the same vessels for the same routes, larger pools coordinate their tankers more efficiently and move them across the network.

1.1 Related Literature

Our work is related to a resurgent interest in analyzing and optimizing transportation systems sparked by the emergence of ride-sharing platforms. This rich line of work has concentrated on the design of pricing and relocation policies in response to long- and short-term demand patterns, and it has focused on a single decision maker that controls prices and/or relocations for the entire market. In the context of optimizing pricing decisions, Cachon, Daniels, and Lobel (2017) and Castillo, Knoepfle, and Weyl (2017) argue, theoretically and empirically, that surge pricing plays a role in rebalancing supply to respond to spikes in demand. Differently, Bimpikis, Candogan, and Saban (2019) focuses on the long-term demand imbalances to derive optimal pricing policies that encourage repositioning of drivers in the network.

In addition to pricing mechanisms, relocation (or repositioning) policies are another important lever to balance transportation systems, and many papers have investigated this direction: Braverman, Dai, Liu, and Ying (2019), Özkan and Ward (2020), and Banerjee, Freund, and Lykouris (2022) consider a central decision maker that can dispatch supply to different locations to better serve demand and derive static policies based on stochastic approximations, while Ata, Barjesteh, and Kumar (2020) and Banerjee, Kanoria, and Qian (2021) are the first to study theoretically and numerically state-dependent policies. Our work complements these lines of research, as we study a market with many independent decision makers and we focus on the relationship between market structure and ballasting. While we consider optimization problems with a central planner, we use their optimal values to estimate the impact of market fragmentation rather than to prescribe policies to be implemented in the real world.

Less attention has been devoted to studying the inefficiencies introduced by factors other than demand imbalances. Some recent works has taken the perspective of individual shipowners: for example, Prochazka, Adland, and Wallace (2019) and Adland and Prochazka (2021) use an approach similar to our model in Section 3 to estimate the value of foresight and contractual flexibility in the context of dry-bulk shipping. The key difference between our work and those papers is that they take the perspective of a single operator, and base their analysis on simulated data (calibrated on aggregate market data) instead of actual commercial decisions. On the other hand, Séjourné, Samaranayake, and Banerjee (2018) is among the first to adopt a market-wide view and study theoretically the inefficiencies generated by market fragmentation: they argue that fragmentation is detrimental because small fleets struggle at efficiently repositioning their assets in the network when demand is randomly split. We depart from their analysis along several dimensions: our datadriven approach allows us not only to quantify the inefficiencies caused by fragmentation, but also to identify as a new driver of ballasting the heterogeneity of networks served by large and small operators. Moreover, we study how increasing concentration affects efficiency, while Séjourné et al. (2018) take the level of fragmentation as given.

A growing emphasis on decentralized transportation markets comes from the Industrial Organization literature in Economics. For example, Frechette, Lizzeri, and Salz (2019) and Buchholz (2022) study the taxi market and show how different barriers to entry shape the market structure and welfare outcomes, while Harris and Nguyen (2022) analyzes how long-term relationships between truckers and brokers affect load assignments in the trucking industry. Particularly relevant for our setting are Brancaccio, Kalouptsidi, and Papageorgiou (2020) and Brancaccio, Kalouptsidi, Papageorgiou, and Rosaia (2023), who exploit voyage-level data in the dry-bulk industry to study the impact of search frictions on market efficiency. While they analyze a market similar to ours, their methods and focus are quite different: in a structural estimation set-up, their primary focus is to study how imperfect matching leads to inefficient equilibrium pricing, which in turn incentivizes inefficient relocation decisions. We complement this strand of literature by examining how different market structures influence resource relocation decisions, regardless of price formation mechanisms. In particular, since empty miles arise also under efficient pricing, market fragmentation has generally been overlooked by this literature: we innovate on this aspect by showing that fragmentation has a first-order effect and we provide a first measurement of its relative importance. Finally, in a market fragmentation set-up, we deal with the technical difficulty of studying decision makers that control multiple vessels and have pricing power, which departs from a common assumption used in the literature that each vessel represents an independent, non-atomic player.

At a higher level, our focus on shipping pools as a way to consolidate the market and achieve the benefits of centralization without its negative effects connects with the extensive literature on resource pooling in Operations Management. This idea has been formulated in different settings, ranging from inventory management (Eppen (1979), Benjaafar, Cooper, and Kim (2005), Corbett and Rajaram (2006), Bimpikis and Markakis (2016), Aflaki and Swinney (2021)), to manufacturing flexibility (Jordan and Graves (1995), Netessine, Dobson, and Shumsky (2002), Van Mieghem (2003), Simchi-Levi and Wei (2012), Moreno and Terwiesch (2015)) and stochastic processing networks (Bassamboo, Randhawa, and Mieghem (2012), Tsitsiklis and Xu (2013)). We innovate on this literature along several lines. First, we identify a new channel whereby pooling improves performance: not only it allows greater flexibility and coordination, but it also changes the incentives of participating decision makers, that start allocating their resources differently. Moreover, this is among the first papers to *measure in a real-world system* the extent to which pooling can benefit efficiency, does so at a *market-wide* level (as opposed to studying single decision-makers), and identifies what is the minimal amount of pooling that captures most of the gains from centralization.

Finally, our work also contributes to the growing literature concerned with improving the sustainability of firms' operations: see Lee and Tang (2018) for a comprehensive review on the subject. While indispensable to global trade, maritime shipping is also harmful to the environment in a variety of ways:⁴ we show that simple steps can substantially improve both the economic and environmental outcomes associated with this industry. This is particularly relevant in the context of ongoing discussions on new regulations about emissions associated with shipping, whose effectiveness can prove more elusive than expected as Hansen-Lewis and Marcus (2022) demonstrates.

⁴See Walker, Adebambo, Feijoo, Elhaimer, Hossain, Edwards, Morrison, Romo, Sharma, Taylor et al. (2019) for a comprehensive examination of the effects of maritime activity on the oceans.

2 Oil Shipping: Context and Data

Shipping activity can be broadly divided into three categories: raw materials (dry bulk, oil, and other liquid products), collectively constituting 85% of the global volume; containerized cargo, at approximately 11% of total volume; and other specialized categories, such as chemicals, for the residual portion.⁵ The oil transportation sector alone accounts for 30% of the seaborne trade volume, making it one of the most relevant markets in maritime transportation. While there exist extensive land infrastructures to move crude oil from extraction sites to refineries, it is estimated that 61% of the daily production volume of 90 million barrels of oil relies on maritime transport.^{6,7} A fleet of about 8,800 ocean-going oil tankers operates in this industry, divided into six (basically) independent segments based on the ships' displacement (i.e., cargo capacity), called *vessel classes* (see Table 5). The smallest vessels have a capacity of about 300,000 barrels of oil, while the largest can accommodate nearly ten times that volume. Ships within each class are relatively homogeneous in size, so that they can access the same ports and use the same canals.

While major oil producers and refiners maintain modest fleets of tankers, most players in this industry rely on a network of small shipowners, who make their tanker vessels available for leasing. Hiring of oil tankers occurs through a variety of contract types, characterized by two dimensions: the length of time for which a vessel is hired, and the financial obligations of the parties. Vessels may be chartered for individual journeys between a load port and a discharge port, or for longer periods of time (e.g., a year). In return, the hiring party pays a daily rate to the owner for the duration of the contract. Vessels are predominantly engaged on a per-voyage basis, on what is often referred to as a *spot* market. In fact, even vessels hired for longer periods of times are often let to other exporters to take advantage of rate fluctuations, so that virtually all loads are transported with spot contracts.

The spot market hinges on a network of intermediaries (brokers), who facilitate transactions between oil companies (exporters) and carriers. A typical transaction within this system unfolds through several distinct phases. Initially, an exporter needing transportation services contacts a broker and provides information regarding the cargo, including type, loading and unloading locations, and the desired departure schedule. In turn, the broker presents the cargo requirements to a selection of vessel operators. Only a subset of these will have tankers available to service the load, and the final choice of the vessel is determined through a negotiation process that includes bargaining on the charter rate. The contract agreed upon (known as the *fixture*) includes the

⁵https://unctad.org/system/files/official-document/rmt2022_en.pdf

⁶https://www.bp.com/content/dam/bp/business-sites/en/global/corporate/pdfs/energy-economics/sta tistical-review/bp-stats-review-2022-full-report.pdf

⁷https://www.eia.gov/international/analysis/special-topics/World_Oil_Transit_Chokepoints

precise load location and day, and the daily rate. Importantly, the daily rate is paid only for the duration of the voyage from the loading port to the final destination. Given the unpredictability of sea voyages, fixing of contracts is initiated shortly before the intended load time, so that only vessels close of the load area can be considered available to service the cargo. Thus, shipowners must make strategic decisions regarding the optimal positioning of their vessels to secure profitable contracts.

Consider a ship operator's perspective right after discharge: it is unlikely that the vessel has already been contracted for a new load.⁸ The operator must decide whether to wait for new offers in the same port (or neighboring anchorages), or to relocate somewhere else. It is during this phase that the shipowner typically receives proposals from brokers for new cargoes whose load area is a few days away from the current position of the vessel. If the offer is fixed, a new *voyage* begins, which ends at the time the new discharge port is reached. Thus, each voyage can be divided into an initial *ballast* leg, where the tanker is empty, and a final *laden* leg, where the tanker is full.

The market dynamics described above naturally lead to *empty miles*, a feature common to virtually all transportation markets. Empty miles substantially impact operating costs of oil shipping companies, and also impose a significant burden to the environment. Because of the economic and environmental costs associated with ballasting, understanding its causes and exploring potential remedies are subjects of primary importance to the industry.

2.1 Description of data

We partnered with a shipping and analytics company, who provided us with a list of 234,795 voyages completed from January 1st, 2018 to March 26th, 2022 by 5,886 vessels; Table 5 offers a breakdown by vessel class. For each voyage we have the identifier of the vessel, as well as its load and discharge ports with the respective arrival and departure dates; we are also provided with a *starting port*, that corresponds to the discharge location of the previously finished voyage. The ballast leg of the voyage is the trip from the starting port to the load port, and the laden leg goes from the load port to the discharge port. Thus, we know exactly how long each leg was for each voyage, and how much time each ship spent docked at a port, ballasting, or traveling with a load.

Geographical information is aggregated at various levels, from port level (the finest partition) to *wide areas* (the coarsest). We carry out our analyses at different levels, since the most appropriate depends on the specific application. By and large, we employ port level, *narrow area*, and *intermediate area* information: narrow areas partition the world into 51 regions, while intermediate areas into 28. A list of the intermediate areas can be found in Table 7.

⁸Exporters are reluctant to charter a tanker that has not yet discharged its previous cargo, because of potential delays associated with port congestion, discharging, and refurbishing of vessels.

In addition to geographical data, we have information about fuel consumption, pollutants emissions, speed, and the name of the *Commercial Operator* that managed the vessel for that voyage. The commercial operator is the agent that was in charge of the vessel's choice of which loads to fulfill as well as whether and where to ballast. The commercial operator may not correspond to the actual shipowner, since vessels with different owners may be part of the same shipping pools, operated by a single entity.⁹ We do not have access to the contracted rates; we discuss how this limits our analysis in Section 3.4.

Some features of this market reassure that potential demand should not be significantly different from realized loads: first, there is no outside option to sea transport once oil has arrived to the tankers terminal;¹⁰; second, since contracts are agreed upon at a short notice, oil exporters likely have their load ready for departure at the time they contract with brokers; third, significantly delaying the time of transportation is not economically viable, given the high cost of holding oil in the terminal. Given all this, the likelihood that exporters contact brokers for transportation and then decide to forgo the arranged contract is low, which implies that realized voyages are a good proxy for total demand.

To ensure consistency between the analysis in Section 3 and Section 4, we do not consider voyages shorter than 5 days (1.8% of observations) and for which we do not have information about the commercial operator (15% of voyages). Finally, we restrict the analysis to only those vessels for which we have a complete history, i.e., whose position can be accounted for from the first time they are assigned a load to the last time they discharge. Notice that this does not rule out vessels that undergo maintenance or become unavailable for other reasons, which we explicitly consider in Section 3, as long as we can make sure that we did not miss any loads they transported. This leaves 4,599 tankers and 146,745 voyages, approximately 63% of the full dataset. We now summarize some of the outstanding features of the oil shipping market from our dataset.

Trade imbalances Oil reserves are unevenly distributed in the world, which leads to substantial imbalances in the oil transportation patterns and, thus, are the main cause of ballasting. These imbalances are reflected in our dataset, as shown in Table 7 and Figure 3. For instance, the "China / Taiwan" region witnessed a threefold surplus of cargo arrivals over departures. Conversely, in other regions, there was a nearly five-fold surplus of cargo departures relative to arrivals. Given these patterns, it is unlikely for vessels to find a load to transport in the same location they discharge at.

⁹The same ships may, and in fact do, operate under different pools throughout the time span of our dataset.

¹⁰Contrast this with a passenger looking for a rideshare: if quoted prices are too high they might get a taxi or use public transportation instead.

Ballasting Ballasting is a common feature in the voyages we observe: about 95% of voyages begin in a port different from the one where the vessel had just discharged, and in 55% of the cases the tanker is required to change geographical area to reach the new load port. On average, 40% of the length of each voyage consists of the ballast leg, with minor differences across vessel classes. This translates to an annual consumption of approximately 10 million tons of fuel in ballasting, equivalent to 32 million tons of CO2 emissions (the emissions generated by 7 million passenger cars in a year) at an estimated cost of 7.2 billion USD (September 2023 prices). Table 6 summarizes this information. Ballasting usually occurs because vessel operators realize that the likelihood of finding new loads in their current location may be so low that it justifies paying the cost of relocating somewhere else.

Market fragmentation and pools Our dataset is noteworthy because we have access to information about the commercial operators of the tankers. We can trace a detailed picture of market concentration, its evolution over time, and of how the number of vessels under management correlates with the decisions of commercial operators. The control structure is extremely fragmented: as evidenced in the left panel of Figure 4, a substantial number of operators have fleet sizes ranging from 1 to 3 vessels; in contrast, very few control more than 10 vessels. The fleet size of the largest commercial operators exhibits notable variation across time, as shown in Figure 5, while smaller operators tend to have a more stable fleet. This difference can be attributed to the fact that large commercial operators manage vessels via *shipping pools*, i.e., they add to their own fleet vessels owned by others (usually, small shipowners themselves). A shipping pool is formed when a number of shipowners decide to "merge" their fleets under a unique manager. The pool manager assumes responsibility for all vessels and maximizes the collective pool income, while the respective owners remain entitled to a share of the earnings the pool generates.¹¹

Commercial operators can use shipping pools as a more flexible tool for expanding their fleets compared to formal mergers or acquisitions.¹² Pools have become increasingly popular thanks to the advancement in computational capabilities that allowed real-time optimization. Despite their advantages, shipping pools' market share remains low: Table 8 shows that for the Aframax vessel class the largest operator controls less than 4% of the total number of vessels, with the top 20 operators collectively controlling less than 40% of the asset share.¹³ Similar patterns prevail across other vessel classes, with a marginal uptick in concentration for larger tankers. Importantly, shipping pools exhibit substantial heterogeneity not only in terms of size, but also in the level of

¹¹In fact, optimally splitting a pool's income is itself an active area of research, see, e.g., Haralambides (1996).

¹²Vessels can ask to join a pool or withdraw at any time, although either process may take between one and six months before being finalized.

¹³One can compare this with the market structure in the market for container shipping, which is much more concentrated: in that case the four largest companies control 60% of capacity.

efficiency they achieve: the right panel of Figure 4 illustrates the distribution of quarterly utilization rates (i.e., the ratio of laden miles to total traveled miles). This could be attributed to many factors, and in this paper we concentrate on its relationship with pool size.

3 Determinants of ballasting

While ballasting is a well-known phenomenon within the transportation sector, its causes remain somewhat underexplored, and their relative significance is unknown. In the case of oil tankers, we have previously acknowledged the substantial role played by trade imbalances, but to what extent can we ascribe observed ballasting solely to disparities in demand? Answering this question is the necessary first step to understand whether the market is operating close to efficiency or if instead there exists untapped potential for improvement from other sources of inefficiency.

We classify the sources of ballasting inefficiencies in three categories: (i) trade imbalances; (ii) uncertainty, i.e., that time and locations of future loads are unknown;¹⁴; (iii) market fragmentation and, as a result, also of the flow of information.¹⁵ This form of informational asymmetry can be due to multiple features of the market, but anectdotal evidence suggests that there appears to be a relationship between the size of commercial operators and the quality of the information they receive about the market conditions.¹⁶

We first set to identify the share of empty miles that can be attributed to trade imbalances: it is a baseline level of ballasting intrinsic to oil trade patterns that cannot be avoided. Our strategy is based on the following intuition: if we can address the other two factors influencing ballasting, the intrinsic *baseline ballasting* can be characterized as the volume of empty miles incurred when we optimize in hindsight the assignment of all observed loads to tankers, with the objective of minimizing the ballasting cost. Optimizing in hindsight addresses any uncertainty, and assuming that all vessels are managed by a central operator guarantees no fragmentation in control and information. Using a similar analogy, we assess what would happen if many planners with perfect information (one for each shipping pool) were to make the tanker-load assignments. This allows us to obtain an estimate of the weight of market fragmentation and uncertainty as well.

Assumptions To perform these analyses, we use data from all four years of voyages to estimate costs and travel times between geographical areas, but we optimize only over the two years before the

 $^{^{14}}$ Therefore tankers may decide to ballast away from a port when in fact a load would have been offered shortly after.

¹⁵In fact, not all vessels in a given geographical area are always aware of all loads departing from there. This is because loads are not "advertised" on open platform, but orchestrated by brokers, who in turn work based on their relationships with individual managers.

¹⁶Size is a good proxy for complexity of the decision making structure and for the bargaining power of managers with brokers, both of which influence the ability to collect and process information.

outbreak of the COVID-19 pandemic for consistency with the uncertainty-aware method employed in Appendix A. We discretize time into five-day periods, so that e.g., voyages departing from January 1st, 2018 to January 5th, 2018 are imputed to period t = 0, use intermediate geographical areas described in Table 7; we impose that, if vessels reach a location at time t, they are also able to load cargoes at time t.¹⁷ Finally, we assume that (i) the laden leg of each voyage has the same duration irrespective of the ship assigned to it; (ii) loads are fungible within vessel class, so that each can be transported by any available vessel of the same class, but not across. The first assumption is justified by data and industry practice whereby all vessels sail at about the same speed.¹⁸ The second is less realistic because other idiosyncratic factors may also influence the choice of a vessel (e.g., energy certifications, nationality of operator); however, our data do not include additional information to impose further constraints.

3.1 Trade imbalances

Demand for oil transportation is influenced by many political, macro-economic, and other market factors. Hence, instead of fitting a necessarily imperfect model to observed demand, we accept demand as an exogenous input and force the transportation assets to deliver such demand. The share of ballasting exclusively attributable to demand imbalances can then be estimated as the ballasting cost that would obtain if transportation was organized by a benevolent and clairvoyant dictator that satisfies all demand, i.e., if it was perfectly efficient, relative to the observed total cost of empty travels. As we mentioned, the different vessel classes define virtually independent markets, and therefore we optimize separately within each class.

Formally, we cast each problem as an integer linear program whose objective is to minimize the total ballasting cost, defined in terms of CO₂ emissions. Since we impose the constraint that all observed voyages be transported, minimizing ballasting costs is equivalent to finding the minimum cost assignment that satisfies all demand. We choose CO₂ emissions as metric instead of travelled distance since emissions better capture important cost-generating events such as fuel consumption when waiting to cross a canal. We denote with $C_{l,o}^v$ the average CO₂ emissions for a ship in vessel class v of a trip between locations l and o belonging to the set of locations \mathcal{L} in the transportation network; let T be the number of 5-day periods from January 1st, 2018 to January 1st, 2020.

We introduce two types of decision variables: $X_{l,o,d,t}$ accounting for the number of vessels that start their voyage at time $t \in \{0, \ldots, T\}$ and location $l \in \mathscr{L}$ to serve a load picked up at origin $o \in \mathscr{L}$ and dropped at destination $d \in \mathscr{L}$; and $Y_{l,t}$ denoting the number of vessels in $l \in \mathscr{L}$ that remain unassigned in l at time t.

¹⁷This is equivalent to imposing that load and discharge times are less than five days. Figures 6 and 7 reassure that this is the case for most voyages.

¹⁸Research in maritime engineering set the most economical speed for tankers; e.g., Psaraftis and Kontovas (2013).

Parameter	Explanation
	Number of loads from area o to area d that depart in period t as observed in
$D_{o,d,t}$	the dataset.
	Indicator taking value 1 when a vessel departing location l at time τ reaches
$\chi_{l,\tau,o,t}$	location o at t (clearly $\chi_{l,\tau,o,t} = 0$ if $\tau > t$). It is obtained by estimating the
	average ballast time between the two areas.
	Number of vessels that become available in location l at time t , either for the
$I_{l,t}$	first time (i.e., the first they are observed in the data) or after a prolonged
	stop (i.e., longer that the 90% percentile of the distribution of port stops).
0.	Number of ships that cease to be available when in location l at time t ,
$O_{l,t}$	either because they exit the dataset or because of a prolonged pause.
$A_{d,t}$	Number of vessels that conclude a voyage in area d at time t .
C^v	Average CO_2 emissions in tons of a ballast travel between areas l and o
$\cup_{l,o}$	undertaken by a ship belonging to vessel class v .

Table 1: List of parameters estimated from the data and used in the optimization problems of Sections 3.1 and 3.2

Finally, the formulation uses the following parameters obtained from the data: $D_{o,d,t}$ represents the number of loads departing from location o towards location d at time t; $A_{l,t}$ represents the number of vessels that discharge at location l at time t;¹⁹ $I_{l,t}$ represents the number of vessels that first enter the data in location l at time t, and $O_{l,t}$ represents those vessels that leave the dataset after discharging in location l at time t; finally, $\chi_{l,\tau,o,t}$ indicates whether a vessel leaving location l at time τ arrives in location o by time t; as such, it encodes the average duration of a ballasting leg between the two locations. Further details about these parameters are in Table 1.

Using these parameters and variables, we formally state the integer program as:

$$\min \quad \sum_{l,o,d,t} C_{l,o}^v X_{l,o,d,t} \tag{1}$$

s.t.
$$\sum_{l,\tau \le t} X_{l,o,d,\tau} \chi_{l,\tau,o,t} = D_{o,d,t}, \qquad \forall (o,d,t) \in \mathscr{L} \times \mathscr{L} \times \{1,\dots,T\},$$
(2)

$$\sum_{o,d} X_{l,o,d,1} + Y_{l,1} = I_{l,1}, \qquad \forall l \in \mathscr{L},$$
(3)

$$\sum_{o,d} X_{l,o,d,t} + Y_{l,t} + O_{l,t} = A_{l,t} + I_{l,t} + Y_{l,t-1}, \qquad \forall l \in \mathscr{L} \ \forall t = 1, \dots, T,$$
(4)

$$X_{l,o,d,t}, Y_{l,t} \in \mathbb{N}.$$
(5)

Equation (1) accounts for the total balasting costs associated with the assignments in X. The

¹⁹Since we enforce that all loads be transported, and by definition the program is always feasible, arrivals can be taken as exogenous to the decision variables.

Vessel class	Weight of trade imbalances $(\%)$	No. vessels
MR1	77.14	357
MR2	74.49	1020
Panamax	73.76	308
Aframax	72.50	792
Suezmax	80.61	547
VLCC	88.34	756

Table 2: Share of ballasting attributed to intrinsic trade imbalances, obtained by comparing the current assignments against the assignments of a central planner matching tankers to loads in hindsight. Total number of vessels included in the optimization is 3,780; this figure is lower than the total number of vessels observed, because some entered the market after January 1st, 2020.

constraints in Equation (2) impose that all loads observed in the data are transported also under the optimal assignment solution, and that this solution is feasible in terms of travel times. Constraints in Equations (3) and (4) make sure that each vessel is used at most for one voyage at a time, and that the flow of ships in each location is conserved. On the right hand side of each constraint we have the total supply of vessels in location l prior to the decision at time t, composed by (i) the ships that remained in l from the previous period; (ii) the inflow of new vessels to the network; and (iii) the arrival of tankers from voyages. On the left hand side of Equation (3) and Equation (4), we model how supply is used, i.e., for assignment to loads ($\sum X_{l,o,d,t}$) and waiting ($Y_{l,t}$), or outflows ($O_{l,t}$). Finally, notice that without loss of generality our formulation does not allow for relocation of empty vessels ahead of their assignment to loads: since there is no randomness in the problem, for every optimal solution that would have repositioned a tanker beforehand, there exist another one that has the vessel ballasting to the origin of the same load just in time.

Let C_v^* be the optimal value of Equation (1) for vessel class v, and C_v^{obs} be the ballasting emissions coming from the observed real-world assignments for the same vessel class. We define the share of ballasting due to trade imbalances as the ratio between these two, $Share_v^{Trade} = C_v^*/C_v^{obs}$. Table 2 reports the results of the integer programs in terms of $Share_v^{Trade}$, and it shows that the share of emissions due to trade imbalances ranges from 70% for smaller vessel classes to almost 90% for larger ships. This trend is intuitive, since larger vessels serve more predictable routes: there exist relatively few ports able to accommodate these ships, which means that there is smaller scope for optimizing the network of locations served. For this reason, we expect those markets to operate closer to efficiency. Hence, factors beyond trade imbalances are relatively more important for smaller vessel classes in contrast to larger ones. Overall, these results confirm that trade imbalances are the single most important factor in determining the amount of ballasting we observe; however, they also highlight that for smaller vessel classes there is an ample margin of improvement.

3.2 Fragmentation and uncertainty

Suppose that, instead of a unique clairvoyant central planner, there was one clairvoyant planner for each commercial operator, and that they optimized their fleet independently of the others. The total ballasting cost incurred by this decentralized system is a lower bound to the true ballasting cost that would occur if the operators lived in a world with no uncertainty. It is a lower bound because it is obtained assuming that the planner minimizes ballasting costs, while in reality they maximize profits.²⁰ We compare this lower bound to the observed ballasting costs and the cost obtained from the central planner C_v^* to infer the weight of fragmentation on one side, and uncertainty on the other. As before, we apply this procedure to each vessel class separately.

For each commercial operator $i \in \mathscr{P}_v$, where \mathscr{P}_v denotes the set of operators in vessel class v, we solve the following problem.

$$\min \quad \sum_{l,o,d,t} C_{l,o}^v X_{l,o,d,t}^i \tag{6}$$

s.t.
$$\sum_{l,\tau \leq t} X_{l,o,d,\tau}^{i} \chi_{l,\tau,o,t} = D_{o,d,t}^{i}, \qquad \forall (o,d,t) \in \mathscr{L} \times \mathscr{L} \times \{1,\dots,T\}$$
(7)

$$\sum_{\substack{o,d}} X_{l,o,d,1}^i + Y_{l,1}^i = I_{l,1}, \qquad \forall l \in \mathscr{L},$$
(8)

$$\sum_{o,d} X^{i}_{l,o,d,t} + Y^{i}_{l,t} + O^{i}_{l,t} = A^{i}_{l,t} + I^{i}_{l,t} + Y^{i}_{l,t-1}, \qquad \forall l \in \mathscr{L}, \ \forall t$$
(9)

$$X_{l,o,d,t}^{i}, Y_{l,t}^{i} \in \mathbb{N}$$

$$\tag{10}$$

The parameters in the problem above have the same explanation as in in Section 3.1, with the only difference that they are computed at commercial operator level. Let C_*^i denote the optimal value of this program, and the system-wide ballasting cost of the decentralized system be $C_v^P = \sum_{i \in \mathscr{P}_v} C_*^i$. By definition, it must be that $C_v^* \leq C_v^P \leq C_v^{obs}$, so that

$$Share_v^{Trade} = \frac{C_v^*}{C_v^{obs}} \le \frac{C_v^P}{C_v^{obs}} \le 1$$

We now argue that the following two ratios estimate the share of empty miles due to fragmentation and uncertainty, respectively.

$$Share_{v}^{Frag} = \frac{C_{v}^{P} - C_{v}^{*}}{C_{v}^{obs}}$$
(11)

 $^{^{20}}$ In fact, with perfect information it is also possible that pool managers would have bid for different loads than the observed ones. We discuss limitations to our approach at greater length in Section 3.4.

Vessel class	Weight of uncertainty (%)	Weight of fragmentation $(\%)$
MR1	10.68	12.18
MR2	13.41	12.1
Panamax	10.75	15.49
Aframax	11.02	16.48
Suezmax	9.07	10.32
VLCC	4.71	6.95

Table 3: Share of ballasting attributed to uncertainty and fragmentation

$$Share_{v}^{Uncertainty} = 1 - \frac{C_{v}^{P}}{C_{v}^{obs}}$$
(12)

In the numerator of Equation (11) we compare the optimal cost of the assignments of a central planner with no uncertainty against the cost of multiple fleets without uncertainty, i.e., the only difference is the level at which assignments are made. Therefore, this difference can only be capturing the increase in cost due to fragmentation, akin to a "price of anarchy". Consider now Equation (12): the only difference between C_v^P and C_v^{obs} is that the former is calculated when vessel managers have perfect knowledge of the future, so that it can be thought of as capturing the effect of eliminating uncertainty; equivalently, the extent to which uncertainty affects ballasting. We report the results of these metrics for each vessel class in Table 3.

Notice that, with the exception of VLCCs, the weight of uncertainty is remarkably similar across vessel classes. This is consistent with our conjecture that the VLCCs follows a more predictable schedule. The other insight emerging from our results is that fragmentation always accounts for a larger share of ballasting than uncertainty. Unfortunately, we cannot further understand how much of the weight attributed from fragmentation derives from lack of coordination, and how much is associated with informational failures, because the data do not allow us to observe which vessels participated in the bargaining process that resulted in the assignment of a load.

These results indicate that the extreme fragmentation of the market is a major cause of waste, and that taking action to consolidate the control structure can yield sizable benefits, both to shipowners and to the environment. However, consolidating all tankers under a single entity is neither feasible nor desirable. The next section is devoted to understanding "how much" consolidation is sufficient to accomplish a substantial amount of efficiency.

3.3 Partial consolidation

Our data contains 169 Aframax commercial operators, and in Section 3.1 we considered a single entity controlling all the vessels to evaluate the inefficiency of the current transportation market. In order to measure the value of *partial* consolidation of the market, we create "synthetic" shipping pools that do not exist in reality. This allows us to test different degrees of consolidation by choosing the number of vessels belonging to each pool. Moreover, it serves us as a test bed to assess the trade-off between the efficiency gains of consolidation and market power.

To carry out this analysis we first define the efficiency gain due to centralization for vessel class v by Δ_v^* and given by $\Delta_v^* = (C_v^* - C_v^{obs}) / C_v^{obs}$. Then, we randomly split the tankers into shipping pools of identical size ϕ , where ϕ denotes the fraction of the total number of vessels that each pool manages. For each pool p, we observe the voyages performed by the vessels in p and solve an optimal assignment problem using Equations (6) to (10), which yields an optimal cost C_p^{ϕ} . Finally, we obtain the market-wide ballasting cost $C_v^{\phi} = \sum_p C_p^{\phi}$ and the gain from a ϕ -consolidation, denoted by Δ_v^{ϕ} .²¹ Our metric of interest is the ratio $\rho_v(\phi) = \Delta_v^*(\phi) / \Delta_v^*$, which reports the fraction of the central planner's fragmentation savings that can be achieved with pools of size ϕ , and can be essentially interpreted as a competitive ratio.

The left panel of Figure 1 plots $\rho_v(\phi)$ for the different vessel classes and for different pool sizes: it shows how much of the inefficiencies due to fragmentation can be eliminated with varying level of consolidation. Conversely, the right panel shows the ratio C_v^*/C_v^{ϕ} , that reports the improvement of the overall efficiency of the system as we vary the level of consolidation in the market. These figures establish that even with small pools, of approximately 5% of the market, it is possible to obtain ballasting savings that are close to the optimum of the central planner. This is particularly evident for smaller vessel classes such as Aframax or MR2: splitting all 792 Aframax ships into pools of 40 units each (about 5% of the fleet) yields between 80% and 90% of the savings that the central planner could have obtained. Using the weight of fragmentation estimated in Table 3, this corresponds to an overall decrease in ballasting emissions of about 15%, or 4.6 million tonnes of CO_2 not released in the environment (= 1 million fewer passenger cars every year). These results suggest that the benefit of centralized vessel-load assignments accrues with even little consolidation. at a level that is unlikely to cause significant threats to competition in the market. Moreover, the shape of the plots indicates that there are decreasing marginal returns to consolidation. In the next sections we investigate in greater detail the mechanisms whereby shipping pools reduce ballasting costs supported also by empirical analyses.

3.4 Discussion of the model

Our analysis so far suggests that the oil transportation market is operating far from its best level of efficiency, and that it stands to substantially benefit even from little consolidation. Moreover, despite being based on data from a specific market, our analysis can be applied with little modification to other decentralized transportation markets, such as dry bulk shipping and full-truckload

²¹By definition then $\Delta_v^* = \Delta_v^1$, because it is obtained with one fleet controlling 100% of the vessels.



Figure 1: Plots of the effects of partial consolidation. On the left, the share of fragmentation inefficiencies that can be recouped by pooling vessels in fleets of different fleet sizes. On the right, efficiency gap between the fully centralized and partially consolidated markets.

trucking, that are likely to suffer from the same inefficiencies we identify here. At the same time, it is important to recognize some limitations of our data and approach.

We stress that our estimates should not be interpreted as market counterfactuals for two reasons: (i) we do not consider the equilibrium response of demand for transportation to consolidation, and the latter's effect on the bargaining power of pool managers; (ii) we set up cost-minimizing problems, while in reality any decision maker attempts to maximize profits. Unfortunately, estimating a structural model to carry out counterfactual simulations is infeasible in our setting, because we do not have data on contracted rates. Moreover, any structural model must take into account that fleets dynamically optimize a set of discrete resources: this is a dynamic programming problem for which only approximate solutions exist in the monopolistic setting (e.g., Godfrey and Powell (2002a)), while the oligopoly model is intractable, and there is no algorithm guaranteed to find Markov Perfect/Oblivious Equilibria. For these reasons, papers in Industrial Organization that study decentralized transportation markets usually assume a collection of independent, nonatomic, agents (e.g., Brancaccio et al. (2020)). Notice, however, that the primary focus of this section is to *measure* the extent to which fragmentation impacts managerial efficiency. If we ran a counterfactual simulation and compared the outcomes, the difference in ballasting emissions would also capture efficiency gains from changes in served voyages (as we expect fulfilled demand to be lower when concentration increases) besides improved coordination. Instead, our approach keeps demand fixed. While the difference we calculate cannot be understood as "how outcomes change as concentration increases", we measure precisely the inefficiency due to lack of coordination, and how these inefficiencies become smaller with pooling.

As far as our choice of minimizing CO_2 emissions associated with ballasting is concerned, we remark that emissions calculations are based on fuel consumption data that vessels are obligated to log to comply with the International Maritime Organization (IMO) regulations. The mapping from fuel consumption to emissions is an industry standard and takes into account which type of fuel was used and for how long (cfr. Table 9, where we show that results are robust to using fuel consumption as metric of interest). Since vessels are crewed even when docked at port, fuel constitutes by far the largest component of variable costs associated with each voyage. Thus, minimizing ballasting emissions, taking served demand as given, is closely related to minimizing costs. Finally, notice that if a shipowner has little pricing power, as would appear natural in the case of Section 3.3 due to the level of market fragmentation, then cost minimization is a particularly good proxy for profit maximization. Thus, we conclude that our estimates, while rough, point in the right direction, and we believe it is unlikely that both the relative share of ballasting due to one factor compared to the others and the insights from Section 3.3 would be substantially different in a structural model with profit maximizing agents.

A final potential point of concern is that our treatment of uncertainty is "residual", in the sense that we have only analyzed optimization problems in hindsight, and obtained estimates for the share of ballasting due to uncertainty as the complement to 1 of all other factors together. A potential problem with this approach is that it might overstate the significance of uncertainty because it combines two forms of ignorance: not knowing where and when future cargo loads will be offered, and not knowing which operator will win them. In Section 3.1 this is irrelevant, because the central manager can allocate any load that appears in the data to any feasible vessel. However, in Section 3.2 we assume that each operator's clairvoyant planner knows exactly which loads it will win, while a more realistic assumption is that it knows which loads will be offered to its fleet, but not whether they will be won. We address this issue in Appendix A, where we consider a central planner that solves a stochastic dynamic program; we adopt the state representation of Godfrey and Powell (2002a,b), and obtain an approximate solution to the dynamic program using Reinforcement Learning tools. For the sake of simplicity we limit that analysis to the Aframax vessel class, and comparing the figures we obtain there with the figures in Table 3 we conclude that the share of ballasting due to uncertainty lies between 6% and 10% for that class, which implies that the estimate for the share due to fragmentation is in the neighborhood of 16% to 20%.

4 Observed shipping pools

In this section, we empirically investigate how consolidation affects the overall efficiency of pools and identify the main mechanisms that drive this improvement. We employ two main strategies that focus on understanding the connection between fragmentation and efficiency. First, we analyze the behavior of pools that grew and shrank over time to handle any non-observable disparities across operators. Second, we compare the differences between the behavior of large pools and sets of small fleets that operate at the same time and locations of the large pool.

Our estimates confirm that larger pools tend to achieve higher levels of efficiency, defined as the ratio of laden miles over the total miles traveled. Furthermore, we show that a fraction of the efficiency gain stems from improved coordination among ships belonging to the same fleet, while another portion can be ascribed to the fact that larger pools run more complex operations. Finally, we empirically validate that there are decreasing marginal returns to consolidation, suggesting that benefits from further consolidation become negligible when pools manage more than 20 tankers.

The central focus of this section is the pool size of each commercial operator. We track the pool size of each operator for the time it is featured in the dataset. In our analysis, we define the pool size for any given period, typically a month or quarter, as the weighted count of vessels operating under the operator's management during that same period. The pool size is calculated as

$$PoolSize_{i,t} = \sum_{s} \frac{T_{s,i,t}}{T_t},$$

where $T_{s,i,t}$ equals the number of days in period t that vessel s spent operating (ballast, laden, or at port) for commercial operator i, and T_t is the number of days in period t. For example, if operator X Shipping Co. had two ships during April 2018, one of which traveled for this operator for the whole month and one that joined its pool on April 16, then we say that X Shipping Co. managed a pool of size $1 + \frac{15}{30} = 1.5$ during April 2018.

4.1 Efficiency

We begin our analysis by considering again our first-order question: does market fragmentation increase ballasting? To answer this, we employ four specifications to argue that larger pools achieve higher utilization (higher proportion of laden miles on total miles travelled) and to measure the size of the attainable savings. The first two specifications establish that larger pools achieve higher levels of utilization using different assumptions, while the other two regressions confirm that improvements come from the ballast leg and allow to measure the size of the improvement.

We say that pools are more efficient the higher their *utilization* is, and we define utilization for

period t as the ratio between the laden and the total miles traveled by all vessels operated by a commercial operator with starting date in t. Formally, for period t and commercial operator i, we define

$$Utilization_{i,t} = \frac{\sum_{n} M_{n,i,t}^{L}}{\sum_{n} M_{n,i,t}^{T}}$$

where $M_{n,i,t}^{L}$ denotes the laden distance of the *n*-th voyage operated by commercial operator *i* whose start date is in period *t*, and $M_{n,i,t}^{T}$ denotes the total distance of the same voyage. Notice that with this definition we impute the utilization level based on the starting date of the voyage, even if, because of the long travel times, it may happen that by the time the tanker secures another contract the pool size has changed. Since we do not observe the fixing date of loads, we assume that the decision about ballasting destinations is made at the starting date of the voyage. Although we focus in this Section on our distance-based metric for Utilization, we run similar specifications to estimate the effect of Pool Size on: (i) the fraction of ballast CO₂ emissions²² (in line with Section 3); and (ii) a time-based Utilization definition to capture time delays. All our results consistently point that larger pools are more efficient than smaller ones for these outcomes.

A limitation of our data is that we do not observe the negotiated rate of a load, nor who competed in the process of getting a load, so we cannot ascertain whether larger pools have some systematic advantage over smaller ones.²³ To counter this issue, we exploit the variation in pool size of each commercial operator over time and observe the differences in efficiency. Figure 5 shows the variation in pool sizes across time for the seven largest commercial operators in each vessel class. Exploiting this variation addresses existing heterogeneity in bargaining power and negotiation skills across different operators, but does not resolve endogeneity in fleet growth, which we try to handle in our second specification.

Single pool utilization We try to understand if the same operator becomes more efficient when its pool size grows. Specifically, we estimate the following regression at a monthly level:

$$\log\left(Utilization_{i,t}\right) = \alpha \times \log\left(PoolSize_{i,t}\right) + \xi_t + \psi_{i,v}.$$
(13)

In Equation (13), ξ_t denotes a time fixed effect and $\psi_{i,v}$ an operator and vessel class-specific fixed effect. The former is necessary to account for seasonal changes in the demand for transportation, as well as the outbreak of COVID-19 and other unforeseen events that affect global trade;²⁴ the

 $[\]overline{{}^{22}\text{Calculated by } FracBallastCO2_{it} = \sum_{n} E^B_{n,i,t} / \sum_{n} E^T_{n,i,t} \text{ where } E^B_{n,i,t} \text{ denotes the CO}_2 \text{ emissions of the } n\text{-th voyage operated by } i \text{ whose start date is in period } t, \text{ and } E^T_{n,i,t} \text{ is the total CO}_2 \text{ emissions of the same voyage.}}$

²³Adland, Cariou, and Wolff (2016) emphasize that match quality between shipowner and charterer significantly influences freight rates.

²⁴E.g., the week-long obstruction of the Suez Canal (March 2021) that caused widespread delays in global shipping.

operator-class fixed effect is necessary to account for each vessel class market and any unobservable variables specific to commercial operators, such as their connections with brokers or age of the vessels they manage. The coefficient of interest is that associated with $PoolSize_{i,t}$: we expect the estimate of α to be positive, reflecting that a larger pool size corresponds to higher utilization.

Our results, reported in the first column of Table 4, confirm this intuition: larger pools achieve higher utilization levels. In particular, the estimated coefficient suggests that, on average, doubling the size of a fleet (e.g., from 5 to 10 vessels) is associated with a 4.4% increase in utilization. In addition to this result, which aggregates all vessel classes (markets), we estimate similar specifications by vessel class and show that the coefficients are statistically significant and in the same direction and size to the one described (see Table 10^{25}). The main assumption behind the validity of these estimates is that unobservable features of the commercial operators such as their bargaining abilities remain fixed over time, and in particular that they are not correlated with the pool size. This appears a tenable assumption for operators that experience only small variations in their pool size, but it is difficult to justify in the case of operators that grew/shrank substantially over time.

In fact, there are three types of endogeneity risks in our analysis. The first type arises from factors that correlate with both utilization and the size of a pool, i.e., time-varying unobservable characteristics of commercial operators that improve their utilization rate as the pool expands, e.g., an operator's capacity to secure more loads due to a stronger contracting team. We are not concerned about this confounding effect because we treat pool size exactly as a surrogate for an operator's "market weight", so that we expect that it will also capture the effect of other factors that correlate with size, but that we cannot observe.

The second type is reverse causality, i.e., that improved utilization might actually be driving a pool's expansion. For instance, an operator might first achieve higher utilization rates by investing in advanced analytical capabilities without expanding its pool. Other vessels may then decide to join the pool, attracted by the prospect of increased utilization for their tankers. To address this issue we would need a source of exogenous, yet protracted, variation in pool sizes. Exploiting natural experiments is particularly challenging in this context because of the long travel times. For example, the 2021 Suez canal blockage affected fleet availability for only six days, while the average voyage lasts 25 days. Exogenous supply shocks should persist for several months to impact pool-wide utilization.

Finally, the third type is the effect of unobserved demand or productivity shocks to the industry, e.g., a sudden change in trade flows. In turn, we may distinguish between shocks that affect the whole oil trade network, and those that influence only a subset of routes. As far as global shocks are considered, we capture their influence through time fixed effects. On the other hand, localized

²⁵Table 11 and Table 12 show analogous results for time-based utilization and emissions, respectively.

shocks are more challenging to deal with, because they may disproportionately affect smaller pools, that tend to provide service to fewer locations (cfr. Section 4.2). For this reason, we introduce in the next paragraph an approach to compare the efficiency of pools based on the locations that they serve; the results are consistent with the estimates from Equation (13).

Synthetic consolidation We run an alternative analysis to capture the effect of pool size in utilization with two main objectives. First, to address the issues of demand/productivity shocks and reverse causality in our previous regression; and second, to compare the effect of *consolidation* (a single larger entity vs. multiple smaller entities) rather than only comparing *size* (a small vs. a large entity). To do so, we look at the difference between large fleets and *sets* of small fleets, i.e., for every large pool we build a "synthetic" pool composed of smaller operators that serve similar geographical areas, and compare utilization between the groups.

The procedure works as follows. For every quarter we first identify large (size > 10 ships, corresponding to the 80% percentile in the distribution of pool sizes) and small pools, and then for each large pool we look at the subset of small pools that predominantly served the same locations that the large pool did.²⁶ We then create a "synthetic" fleet by randomly sampling pools from this subset, until their aggregate size is about the same as the large pool under consideration, and record the utilization level this synthetic fleet attains; we repeat this process 100 times. We then perform a difference-in-means t-test between the utilization levels of the large pools and the average utilization over the 100 synthetic fleets associated to each large pool.

Similar to our previous result, we obtain that the utilization of large fleets is higher than the synthetic counterpart, as displayed in the second column of Table 4 and more extensively for each vessel class in Table 13; the size of the difference in means is 8%.²⁷ Since the average pool size of the fleets composing the synthetic fleets and larger fleets are 3 and 22, respectively, using a linear transformation we get that doubling the fleet size is equivalent to a 2.18% increase in utilization, which corresponds to about half of our previous result of 4.4%. These results underscore the positive impact of centralizing decision-making. Since synthetic and large fleets serve the same locations, we can rule out that the observed positive difference is due to geographical factors or to structural differences in the network of voyages of the large pool. Moreover, by construction reverse causality is not an issue, because we do not compare the same entity over time. However, it is unclear whether centralization is beneficial primarily because it improves coordination or because larger pools, with more influence over brokers, secure loads more swiftly.

²⁶For this specification we use a partition of the globe into 51 "narrow geographical areas".

 $^{^{27}}$ Table 14 confirms that the same result holds for our definition of utilization in terms of time, and Table 15 shows the results of a similar analysis for emissions reporting a difference in means of -5.97%.

		Coordi	nation		Operations	s Complexity
Dependent Variable	Log Utilization	Δ Utilization	Ballast Delay	Laden Delay	Log Trips on a single OD	Log Ports visited by vessel
	[%]	[%]	[Days/month]	[Days/month]	[# trips]	[# ports]
Independent Variable						
Log Fleet Size	0.0439^{***}		-0.5283**	-0.0327	-0.0151***	0.2326***
	(0.006)		(0.240)	(0.125)	(0.001)	(0.005)
Fleet Size Large		8.0272***				
(Dummy)		(0.2848)				
Fixed Effect						
Time	Yes	No	Yes	Yes	Yes	Yes
Operator \times Vessel Class	Yes	No	Yes	Yes	No	Yes
OD Pair	No	No	No	No	Yes	No
Vessel ID	No		No	No	No	Yes
Fit Statistics						
R-squared Adj.	0.1531		0.0894	0.0598	0.2359	0.3715
N obs	$23,\!670$		47,597	47,597	119,957	66,158
Granularity						
Geographical	N/A	N/A	Port	Port	Narrow Area	Port
Time	Monthly	Quarterly	Monthly	Monthly	Quarterly	Quarterly

Standard errors in parentheses. Signif.Codes: ***:0.01, **:0.05, *:0.1

Table 4: Summary of coefficient estimates for the regression analyses

Ballast and laden delays We showed that large pools achieve higher utilization than smaller fleets, but we have not explained the mechanism underpinning the improvement. A first step is to confirm that improvements arise from more efficient ballast legs of the trip and not from the laden ones. It is reasonable to expect that both small and large fleets would not encounter delays during their *laden* legs. Since, once a vessel is loaded, it proceeds directly to its discharge port without interruption. Conversely, the ballast leg may be susceptible to delays for various reasons. For instance, a ballast delay occurs if a vessel incorrectly repositions to a region where no actual load materializes, so that it needs to sail empty to another location to secure cargo, which results in a longer observed time to ballast from the previous discharge port to the actual load port.

To perform this analysis, we calculate the *free flow* travel time of a route by taking the 20th percentile of the observed laden travel times for on the same route.²⁸ We then define the delay of a leg as the difference between the observed travel time for the trip and the free flow travel time on that route, and we estimate the following two regressions.

$$LadenDelay_{n,i,t} = \theta \times \log\left(PoolSize_{i,t}\right) + \xi_t + \psi_{i,v},\tag{14}$$

$$BallastDelay_{n,i,t} = \theta \times \log\left(PoolSize_{i,t}\right) + \xi_t + \psi_{i,v}.$$
(15)

The variable $LadenDelay_{n,i,t}$ takes the value of the laden delay of the *n*-th voyage of operator *i* in month t, and similarly for $BallastDelay_{n,i,t}$. The coefficient estimates are reported in the third and fourth columns of Table 4, and show that, while there is no correlation between fleet size and delays on laden legs, ships belonging to smaller pools take significantly longer time to ballast

 $^{^{28}}$ Taking the 20th percentile of the travel times is a common practice in the applied literature on transportation, as it allows to disregard delays due to, e.g., port congestion.

between two ports than their counterparts in large fleets (see results by vessel class in Table 16). Put into context, the estimated coefficient in Equation (15) implies that if two pools with five tankers each were merged, five days of ballasting every month would be saved; assuming average emissions of 60 tons of CO_2/day (as in our data for Aframax vessels), this amounts to 3600 tons $CO_2/year$, equivalent to 800 passenger cars. In contrast, no such difference is detected in the case of laden trips, as reported in the fourth column and expected by our intuition.

Since tankers start looking for new cargo only after discharging the previous load, these estimates suggest that fleet size affects the time it takes to fix a new contract. We cannot point to a precise explanation why smaller pools take longer to secure a load, but at least two conjectures are possible. It may be that vessels belonging to larger pools idle less after discharging because managers of large pools have better information about new offerings, and therefore can find a suitable load quicker.²⁹ An alternative explanation is that tankers of large pools sail directly to a new area upon indication of their manager, who is confident that a new load will be found in the new location, while in smaller pools the vessels take more detours before arriving to the port where they ultimately find a new cargo. Additional data will be needed to answer this question more definitely, but it appears that ballasting trips of large and small pools are substantially different.

4.2 Coordination and operations complexity

We established that there is a relationship between the size of an operator and the time its tankers take to reach the origin of the new cargo, but this does not fully explain the increase in utilization represented in the first column of Table 4. In this subsection we argue that larger fleets use their tankers in a more complex fashion, which points to better coordination of their fleet, that in turn helps them achieve higher efficiency; in particular, we observe that larger fleets serve a more diverse network of locations and do not repeat the same route as much as their smaller counterparts.

We believe this difference in behavior can be connected to the attitude towards risk of pool managers and, in turn, how this relates to the size of their fleet. Shipowners are generally considered risk averse agents;³⁰ in this context, large commercial operators have the physical and financial assets to diversify their portfolio of locations, while smaller operators are more constrained from both points of view, and therefore expected to make more conservative decisions. In turn, large pools' willingness to serve a more diversified network of locations and routes allows them to run more integrated and efficient operations.

It is natural that large commercial operator can serve more locations at the same time than

²⁹There is an ecdotal evidence that, after discharge, vessels immediately depart and reach "idling" positions at the intersection of important routes, where they wait to be contacted by brokers.

³⁰See, e.g., Albertijn, Bessler, and Drobetz (2011); Drobetz, Haller, and Meier (2016) for research on financing and risk attitudes of shipowners.

smaller operators; therefore, to obtain meaningful comparisons in our analysis we run regressions at vessel level (not pool level). We pose the question about the complexity along two complementary dimensions: the number of trips that a vessel belonging to an operator performs between the same origin-destination pair, and how many different ports a vessel visits. Formally, we estimate, at quarterly level, the following specifications:

$$\log\left(Loads_{s,i,t}^{o,d}\right) = \beta \times \log\left(PoolSize_{i,t}\right) + \xi_t + \phi_{o,d},\tag{16}$$

$$\log(No.Ports_{s,i,t}) = \eta \times \log(PoolSize_{i,t}) + \xi_t + \psi_{i,s}.$$
(17)

The variable $Loads_{s,i,t}^{o,d}$ counts how many loads with origin in area o and destination in area d the vessel s belonging to operator i served in quarter t, while the variable $No.Ports_{s,i,t}$ counts how many distinct ports the tanker s visited in quarter t when under management of operator i. Therefore, the purpose of Equation (16) is to understand whether pools tend to reserve a vessel to serve a small number of routes, while Equation (17) serves to establish if the same ship is used more extensively across the network when part of larger fleets.³¹

Summary results for these regressions are shown in the last two columns of Table 4: the estimate for β is negative, indicating that smaller fleets tend to travel more often the same route, while the estimate for η is positive, which shows that tankers managed by larger operators take on loads from a wider range of locations than their counterparts in the same time frame. Taken together, they imply that smaller pools tend to concentrate on fewer legs, where they travel the same route back and forth. In contrast, when the same vessels are used by large pools, they visit more ports, which supports our previous intuition that smaller pool take more conservative decisions by sticking to routes which they have previous experience of: they are willing to trade off less uncertainty with longer ballasting. At the same time, larger pools are willing (and able) to optimize the vessel assignments so that after discharge they ballast immediately to a closer location to get a new load, instead of traveling back towards their previous (loading) location.

4.3 Partial consolidation from data

All our insights so far demonstrated that larger fleets tend to be more efficient, and we have laid out some key differences in how large commercial operators conduct their operations. In Section 3.3 we argue that slight consolidation is sufficient to achieve most of the benefits of coordination, and we find that marginal benefits from consolidation effectively vanish as fleet size increases. With this section we aim to further validate that intuition and extend the results of Section 4.1, where

³¹Hence the addition of an (o, d)-pair level fixed effect in Equation (16) that takes into account factors such as the importance and regional considerations of route and compare the vessels that served that leg only. Notice that in Equation (17) we add instead a fixed effect for the vessel class each tanker belongs to.



Figure 2: Coefficients and 95% confidence interval of regression (18) with panel data at a monthlylevel. Results reveal monotonically increasing and concave functions that plateaus when fleet size is between 15 and 20, pointing to similar results as the ones provided by the optimization model. The baselines of the (0,2] excluded group are 54.1%, 58.5%, and 57.7% for utilization (distance), emissions, and utilization (time), respectively.

we only consider a linear relationship between size and utilization. This approach also serves to confirm that consolidation is beneficial even when pool managers are a profit-maximizing agents, and not just in the cost-minimizing perspective of Section 3.3.

To establish this result, we regress the monthly utilization against a categorical variable that identifies the pool size based on discrete ranges. In particular, we consider the variable $PoolSizeBin_{i,t}$ taking values in the set $\{(0,2], (2,3], (3,5], (5,10], (10,15], (15,20], (20,\infty)\}$ depending on the observed fleet size and, in the same way as the previous regressions, we include fixed effects for vessel class, commercial operator and time as follows:

$$Utilization_{i,t} = \beta \times PoolSizeBin_{i,t} + \xi_t + \psi_{i,v}.$$
(18)

Our results, shown in Figure 2, report the coefficients and 95% confidence intervals of the different categories of pool sizes against the excluded group of (0, 2]. As an example, the coefficient for the (10, 15] group in Figure 2a is approximately 0.02, indicating that, on average, the emissions generated by a fleet with a pool between 10 and 15 vessels has 2 more percentile points than the utilization of a fleet with one or two vessels. Specifically, the average utilization (distance) of fleets with pool sizes between 0 and 2 is 58.6%. Then the way to interpret Figure 2 is that the utilization for a fleet with 10-15 vessels is (58.6 + 2.0)% = 60.6%.

Figures 2a to 2c sketch increasing concave functions, which demonstrate both the benefit of larger fleet sizes and the diminishing marginal returns of pool size, a result that coincide with our results in Section 3.3. The similarity in shape and scale also suggests that the efficiency gains we observe in the data come mostly from operational efficiencies related to the better coordination of the pool rather than from increased bargaining power and other "unboservable features". In fact, we remind that the results of Section 3.3 arise solely from coordinating loads and vessels, since we cannot model nor estimate other factors. This leads us to believe that coordination is the prime factor driving higher utilization, and that other motivations may play a less prominent role. Unfortunately, because of the limited nature of our data, we cannot further back this conjecture.

5 Concluding remarks

In this paper we investigate the interplay between market fragmentation and ballasting in the context of decentralized transportation markets. Using a combination of numerical and empirical methods, we find that fragmentation exacerbates the adverse effects of demand imbalances. Our research further shows that even modest market consolidation brings significant benefits. Analyzing a dataset of six thousand oil tanker voyages, we determine that demand imbalances cause 70-75% of empty miles, while fragmentation accounts for 15-20% (with 5-10% attributed to unavoidable uncertainty). We propose two channels to explain the beneficial impact of consolidation on ballasting. First, centralizing decision-making improves coordination, so that vessels can serve the same locations more efficiently. Second, larger pools diversify their network of served ports, thereby optimizing the utilization of tankers and minimizing periods of ballasting.

Given the drawbacks associated with centralizing shipping operations under a single entity, we turn our attention to the emerging trend of shipping pools. In shipping pools, individual shipowners aggregate their resources and entrust a single decision-maker to carry out commercial operations. Our data suggest that modest shipping pools, unlikely to negatively affect competition, suffice to capture most of the benefits that would accrue with a central decision maker: for example, partitioning the fleet of oil tankers in pools of 30-40 units results in a 15% decreases in empty miles (i.e., 70% of the first-best with a central planner). These results underscore the growing relevance of this institutional arrangement, which has been proposed by industry players also to encourage the adoption of environmentally sustainable practices. Finally, at a higher level the paper displays the extent of sustainability gains that can be achieved by optimizing the use of current resources in transportation, and supply chains more generally.

We view our work as contributing to the literature on decentralized transportation markets and, in particular, to the growing strand that analyzes the effects of market structure and mechanisms on efficiency outcomes. We hope that this work will motivate further studies along many directions: for example, further analysis is needed to understand the effect of fragmentation on price formation, since we focused primarily on how size changes affects shipowners' incentives. Other directions include, e.g., studying the optimal implementation of fuel levies to discourage empty travels, theoretical analysis of the optimal policies of pool managers with finite resources, and mechanisms to price and assign loads to vessels.

References

- Adland, Roar, Pierre Cariou, Francois-Charles Wolff. 2016. The influence of charterers and owners on bulk shipping freight rates. *Transportation Research Part E: Logistics and Transportation Review* 86 69–82.
- Adland, Roar, Vit Prochazka. 2021. The value of timecharter optionality in the drybulk market. Transportation Research Part E: Logistics and Transportation Review 145 102185.
- Aflaki, Arian, Robert Swinney. 2021. Inventory integration with rational consumers. Operations Research 69(4) 1025–1043.
- Albertijn, Stefan, Wolfgang Bessler, Wolfgang Drobetz. 2011. Financing shipping companies and shipping operations: A risk-management perspective. Journal of Applied Corporate Finance 23(4) 70–82.
- Ata, Baris, Nasser Barjesteh, Sunil Kumar. 2020. Dynamic dispatch and centralized relocation of cars in ride-hailing platforms. *Available at SSRN 3675888*.
- Banerjee, Siddhartha, Daniel Freund, Thodoris Lykouris. 2022. Pricing and optimization in shared vehicle systems: An approximation framework. Operations Research 70(3) 1783–1805.
- Banerjee, Siddhartha, Yash Kanoria, Pengyu Qian. 2021. Large deviations optimal scheduling of closed queueing networks. arXiv preprint arXiv:1803.04959.
- Bassamboo, Achal, Ramandeep S Randhawa, Jan A Van Mieghem. 2012. A little flexibility is all you need: On the asymptotic value of flexible capacity in parallel queuing systems. Operations Research 60(6) 1423–1435.
- Benjaafar, Saif, William L Cooper, Joon-Seok Kim. 2005. On the benefits of pooling in productioninventory systems. *Management Science* 51(4) 548–565.
- Bertsekas, Dimitri P. 2018. *Dynamic Programming and Optimal Control*, vol. II. 4th ed. Athena Scientific, Belmont, MA, USA.

- Bimpikis, Kostas, Ozan Candogan, Daniela Saban. 2019. Spatial pricing in ride-sharing networks. Operations Research 67(3) 744–769.
- Bimpikis, Kostas, Mihalis G Markakis. 2016. Inventory pooling under heavy-tailed demand. *Management Science* **62**(6) 1800–1813.
- Boyd, Stephen P., Lieven Vandenberghe. 2004. Convex optimization. Cambridge university press.
- Brancaccio, Giulia, Myrto Kalouptsidi, Theodore Papageorgiou. 2020. Geography, transportation, and endogenous trade costs. *Econometrica* 88(2) 657–691.
- Brancaccio, Giulia, Myrto Kalouptsidi, Theodore Papageorgiou, Nicola Rosaia. 2023. Search Frictions and Efficiency in Decentralized Transport Markets. *The Quarterly Journal of Economics* doi:10.1093/qje/qjad023. URL https://doi.org/10.1093/qje/qjad023.
- Braverman, Anton, Jim G Dai, Xin Liu, Lei Ying. 2019. Empty-car routing in ridesharing systems. Operations Research 67(5) 1437–1452.
- Buchholz, Nicholas. 2022. Spatial equilibrium, search frictions, and dynamic efficiency in the taxi industry. *The Review of Economic Studies* **89**(2) 556–591.
- Cachon, Gerard P, Kaitlin M Daniels, Ruben Lobel. 2017. The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing & Service Operations Management* **19**(3) 368–384.
- Castillo, Juan Camilo, Dan Knoepfle, Glen Weyl. 2017. Surge pricing solves the wild goose chase. Proceedings of the 2017 ACM Conference on Economics and Computation. 241–242.
- Corbett, Charles J, Kumar Rajaram. 2006. A generalization of the inventory pooling effect to nonnormal dependent demand. *Manufacturing & Service Operations Management* 8(4) 351–358.
- Drobetz, Wolfgang, Rebekka Haller, Iwan Meier. 2016. Cash flow sensitivities during normal and crisis times: Evidence from shipping. *Transportation Research Part A: Policy and Practice* **90** 26–49.
- Eppen, Gary D. 1979. Note—effects of centralization on expected costs in a multi-location newsboy problem. *Management science* **25**(5) 498–501.
- Frechette, Guillaume R, Alessandro Lizzeri, Tobias Salz. 2019. Frictions in a competitive, regulated market: Evidence from taxis. American Economic Review 109(8) 2954–92.

- Godfrey, Gregory A, Warren B Powell. 2002a. An adaptive dynamic programming algorithm for dynamic fleet management, i: Single period travel times. *Transportation Science* **36**(1) 21–39.
- Godfrey, Gregory A, Warren B Powell. 2002b. An adaptive dynamic programming algorithm for dynamic fleet management, ii: Multiperiod travel times. *Transportation Science* **36**(1) 40–54.
- Gordon, Geoffrey J. 1995. Stable function approximation in dynamic programming. Machine Learning Proceedings 1995. Elsevier, 261–268.
- Hansen-Lewis, Jamie, Michelle M Marcus. 2022. Uncharted waters: Effects of maritime emission regulation. Working Paper 30181, National Bureau of Economic Research. doi:10.3386/w30181.
- Haralambides, Hercules E. 1996. The economics of bulk shipping pools. Maritime Policy and Management 23(3) 221–237.
- Harris, Adam, Thi Mai Anh Nguyen. 2022. Long-term relationships and the spot market: Evidence from us trucking .
- Jordan, William C, Stephen C Graves. 1995. Principles on the benefits of manufacturing process flexibility. *Management science* **41**(4) 577–594.
- Lee, Hau L, Christopher S Tang. 2018. Socially and environmentally responsible value chain innovations: New operations management research opportunities. *Management Science* **64**(3) 983–996.
- Moreno, Antonio, Christian Terwiesch. 2015. Pricing and production flexibility: An empirical analysis of the us automotive industry. *Manufacturing & Service Operations Management* **17**(4) 428–444.
- Munos, Rémi, Csaba Szepesvári. 2008. Finite-time bounds for fitted value iteration. Journal of Machine Learning Research 9(5).
- Netessine, Serguei, Gregory Dobson, Robert A Shumsky. 2002. Flexible service capacity: Optimal investment and the impact of demand correlation. *Operations Research* **50**(2) 375–388.
- Ozkan, Erhun, Amy R Ward. 2020. Dynamic matching for real-time ride sharing. *Stochastic Systems* **10**(1) 29–70.
- Prochazka, Vit, Roar Adland, Stein W Wallace. 2019. The value of foresight in the drybulk freight market. *Transportation Research Part A: Policy and Practice* **129** 232–245.
- Psaraftis, Harilaos N, Christos A Kontovas. 2013. Speed models for energy-efficient maritime transportation: A taxonomy and survey. *Transportation Research Part C: Emerging Technologies* 26 331–351.

- Séjourné, Thibault, Samitha Samaranayake, Siddhartha Banerjee. 2018. The price of fragmentation in mobility-on-demand services. Proceedings of the ACM on Measurement and Analysis of Computing Systems 2(2) 1–26.
- Simchi-Levi, David, Yehua Wei. 2012. Understanding the performance of the long chain and sparse designs in process flexibility. *Operations Research* **60**(5) 1125–1141.
- Tsitsiklis, John N, Kuang Xu. 2013. On the power of (even a little) resource pooling. *Stochastic Systems* 2(1) 1–66.
- Van Mieghem, Jan A. 2003. Commissioned paper: Capacity management, investment, and hedging: Review and recent developments. *Manufacturing & Service Operations Management* 5(4) 269– 302.
- Walker, Tony R, Olubukola Adebambo, Monica C Del Aguila Feijoo, Elias Elhaimer, Tahazzud Hossain, Stuart Johnston Edwards, Courtney E Morrison, Jessica Romo, Nameeta Sharma, Stephanie Taylor, et al. 2019. Environmental effects of marine transportation. World seas: an environmental evaluation. Elsevier, 505–530.

Appendices

A Uncertainty: approximate dynamic programming

Uncertainty is one of the main characterizing elements of the maritime world. In the case of oil transportation markets, shipowners and managers face uncertainty from three different sources: (i) from where new loads will be offered, (ii) when, and (iii) the level of competition for each of these (i.e., if they will outbid their opponents and thus win the load). The first two elements are unaffected by fragmentation in the market, but the third is: in a very fragmented market there are many competitors for each load, which decreases the likelihood that a given tanker will win it. With this appendix we aim at disentangling the share of ballasting due to "pure" uncertainty from the share due to uncertainty that can be mitigated by consolidating the market.

Our approach is similar to what we followed in Section 3: we consider a central planner whose objective is to minimize the ballasting cost associated with transporting the loads. The main difference is that now we assume that the central planner does not have perfect foresight; instead, it has a probabilistic assessment about the distribution of future loads, and makes decision based on this. Formally, the central planner solves a stochastic dynamic program. We then compare the optimal costs obtained in the DP with C_v^* from Section 3.1; since the two problems only differ in terms of uncertainty about the future, the difference in optimal cost can be taken as a measure of the share of ballasting due to it. Moreover, since we are comparing ballasting costs incurred by two central planners, in both cases there is no uncertainty about competition. It follows that the measure we derive considers only pure uncertainty as discussed above.

Following discussions with our industry partner, we focus on the Aframax market, which has a number of attractive features. Aframax vessels are of intermediate size, so they can dock at most ports in the world and can use both the Panama canal and Suez canal. Moreover, an increasing number of them are suitable for transporting both crude oil and refined products, which makes it a segment expected to grow in popularity as MR1 and MR2 decrease their market share.³²

We approximate the optimal value function of the stochastic DP, denoted by $V^*(s)$, where s is a state summarizing the present and future availability of tankers. We then find the ballasting decisions that the central planner would have made if faced with the demand realizations observed in the data. Because the central planner optimizes *online*, we cannot ensure that all loads observed are transported. In our simulation we find that the central planner can satisfy 75% of the observed loads. We take a conservative stance and compare the minimum cost obtained with the approximate value function with a benchmark calculated as follows: for every route (o, d) and every time period

³²https://splash247.com/lr1s-and-mr1s-becoming-niche-tankers/

t we sort the loads observed in reality in ascending order of ballasting emissions, and then consider only the $n_{o,d,t}$ least costly, where $n_{o,d,t}$ is the number of loads on (o, d) in period t that were served in the simulation; we then sum the ballasting emissions associated to these voyages. Compared against this benchmark the central planner can achieve ballasting costs that are 20% lower than the observed. This suggests that for the Aframax class uncertainty accounts for about 7.50% of the overall ballasting cost, with the remaining 20% imputable to operational inefficiencies and uncertainty regarding competition. The remainder of the section draws the formal arguments to compute an estimate of V^* .

A.1 Dynamic programming formulation

Consider a decision maker that minimizes ballasting costs over discrete time periods t = 0, 1, ...,that represent 5-day periods in the data. Demand from origin $o \in \mathscr{L}$ to destination $d \in \mathscr{L}$ at time t is denoted by $D_{o,d,t}$ and is drawn i.i.d. from a Poisson distribution with mean $\lambda_{o,d}$. As in the case of Section 3, \mathscr{L} is the set of geographical areas from Table 7; let $N_A = |\mathscr{L}| = 28$. We estimate $\lambda_{o,d}$ as the average number of loads observed on the (o, d) route in the data.³³ Travel times between locations are deterministic, denoted by $T_{o,d}$, and equal to the average travel time observed in the data; let T_{max} denote the maximum length of a voyage in this world, i.e.

$$T_{\max} = \max_{l,o,d} \quad T_{l,o} + T_{o,d}.$$

Given a total number of vessels equal to N, the state of the system at every time t is represented by $s^t \in S$. The state space S is finite, equal to the set of all $28 \times (T_{\text{max}} + 1)$ matrices with natural entries that sum up to N. Formally,

$$S = \left\{ s \in \mathbb{N}^{28 \times (T_{\max} + 1)} : \sum_{d,\tau} s_{d,\tau} = N \right\}.$$

Column τ of state s^t is the number of ships that will become available in $\tau - 1$ periods in the future as a result of voyages begun in all periods up to t - 1 (inclusive) and that have not reached yet their destination. So, the first column represents the number of vessels currently available in each location, the second column the number and destination of vessel that will terminate their voyage in the next period, and so on. Let $s_{,\tau}$ denote the τ -th column of state s. Therefore, s^t summarizes the future availability of tankers given the decision made until time t - 1.

At every time period t, the central planner observes the demand realization and decides how

 $^{^{33}}$ Here is clear why we need to consider only data before the outbreak of COVID-19. The pandemic substantially altered oil trade flows.

many loads to serve and how. Specifically, it acts on two decision variables: $X_{l,o,d}$ denotes the number of vessels available in l used to transport loads from demand $D_{o,d,t}$; $B_{l,d}$ represents the number of vessels ordered to ballast from l to d, with the convention that $B_{l,l}$ equals the number of vessels ordered to wait in l. We assume that the central planner can serve a load on route (o, d) at time t only with currently available tankers that can reach location o from their position by the same period t. Notice that in this formulation a vessel may be ordered to ballast with a cargo already secured $(X_{l,o,d} \text{ for } l \neq o)$, are in expectation of new loads in the future $(B_{l,d} \text{ for}$ $l \neq d$). Based on these decisions, the deterministically transitions to s^{t+1} : all travels scheduled to terminate in $\tau - 1$ periods in s^t will be in column $\tau - 1$ in s^{t+1} ; and travels that take time T will appear in column T - 1. In particular, s^{t+1} can be written as a linear function of s^t .

In Section 3.1 we imposed the constraints of Equation (2), that require that all load be transported. This is possible because problem is deterministic; for the stochastic DP at hand, we impose instead that the central planner suffers a penalty M > 0 for each load that remains unassigned. Together with the ballasting cost paid for relocations and assignments, we obtain a flow-payoff function

$$r_t(X, B|s, D) = \sum_{l,o,d} C_{l,o}^v X_{l,o,d} + \sum_{l,d} C_{l,d} B_{l,d} + M \sum_{o,d} \left(D_{o,d,t} - \sum_l X_{l,o,d} \right).$$
(A.1)

The central planner seeks minimizes $\sum_{t=0}^{\infty} \gamma^t r_t$, where γ is a discount factor.³⁴ It is well known that the optimal value starting from state s of this dynamic program, denoted $V^*(s)$, satisfies the Bellman equation, i.e.,

$$V^*(s) = \mathbb{E}_D\left[\min_{X,B} r(X, B|s, D) + \gamma V^*\left(s'\right)\right],\tag{A.2}$$

where s' is the state that obtains after decisions X and B, and $\mathbb{E}_D[\cdot]$ denotes expectation taken with respect to the realization of demand. While state and action spaces are finite, it is computationally intractable to find an exact solution:³⁵ we turn to approximate dynamic programming. In particular, we approximate the optimal value function using the Fitted Value Iteration approach (Bertsekas (2018), Munos and Szepesvári (2008)). Intuitively, in Fitted Value Iteration the classical value iteration procedure is performed only for a small subset of the states, and an estimate of the value function is computed by fitting an approximation to the values thus obtained. Following Godfrey and Powell (2002a,b), we choose a piece-wise linear, convex approximation; with this combination we can efficiently solve a sequence of linear and quadratic programs.

 $^{^{34}}$ We set $\gamma = 0.85$, which corresponds to an effective time horizon of approximately one month.

 $^{^{35}}$ There are about $10^{2,500}$ possible states.

A.2 Fitted Value Iteration

With this approach one first defines a set of candidate functions, and then looks for a function this set that is closest to the fixed point of Equation (A.2). From Godfrey and Powell (2002a,b) we know that V^* must be convex in s, and it is known that a convex function can be approximated arbitrarily well with a family of affine functions.³⁶ Thus, we restrict attention to a set of piecewise-linear, convex functions. Each function in the set is defined as the point-wise supremum of a family of affine *basis* functions, in turn obtained from a set of basis states. Formally, let the set of basis states be $S = \{s^i : i = 1, \ldots, N_A\}$. Each basis state s^i is defined as follows: if N is the total number of vessels in the environment, first $\lceil \frac{3}{4}N \rceil$ are allocated uniformly at random in the $N_A \times (T_{\max} + 1)$ matrix; then we modify each entry on the *i*-th row as

$$s_{i,\tau}^{i} \leftarrow s_{i,\tau}^{i} \left\lceil \frac{N}{4\left(T_{\max}+1\right)} \right\rceil.$$

Thus, s^i corresponds to a situation in which relatively more of the vessels become available in location *i* over time, so that we expect $V^*(s^i)$ to capture how "good" is having tankers in *i*. To each state $s^i \in S$ we associate an initial value V_0^i , defined as the γ -discounted value over 147 time periods earned by a myopic central planner.³⁷ Finally, we obtain a family of affine functions by solving the following problem for each *i*, where $\hat{V}_0^i \in \mathbb{R}$, $g_0^i \in \mathbb{R}^{N_A}$, and $x \cdot y$ is the usual dot product between *x* and *y*.

$$\min_{\hat{V}_{0}^{i},g_{0}^{i}} \sum_{i=1}^{N_{A}} \left(\hat{V}_{0}^{i} - V_{0}^{i} \right)^{2} \tag{A.3}$$
s.t. $\hat{V}_{0}^{j} \ge \hat{V}_{0}^{i} + \sum_{\tau=1}^{T_{\max}+1} \gamma^{\tau-1} \left[g_{0}^{i} \cdot \left(s_{\cdot,\tau}^{j} - s_{\cdot,\tau}^{i} \right) \right] \text{ for all } i,j$

This procedure yields a set $\mathcal{B}_0^S = \left\{ \left(s^i, \hat{V}_0^i, g_0^i\right) : i = 1..., N_A \right\}$. The set \mathcal{B}_0^S represents basis functions because for each *i* we can write the affine function

$$g^{i}(s) = \hat{V}_{0}^{i} + \sum_{\tau=1}^{T_{\max}+1} \gamma^{\tau-1} \left[g_{0}^{i} \cdot \left(s_{\cdot,\tau} - s_{\cdot,\tau}^{i} \right) \right) \right].$$

³⁶See, e.g., Boyd and Vandenberghe (2004), Chapter 3.

³⁷That is, we simulate 147 time periods and collect all the flow payoffs according to Equation (A.1) that a myopic planner would achieve. 147 five-day periods correspond to two years from January 1st, 2018 to January 1st, 2020.

We define our initial estimate for the value function of the DP as the pointwise supremum of these g^{i} 's:

$$\hat{V}_{0}(s) = \max_{i} \left\{ \hat{V}_{0}^{i} + \sum_{\tau=1}^{T_{\max}+1} \gamma^{\tau-1} \left[g_{0}^{i} \cdot \left(s_{\cdot,\tau} - s_{\cdot,\tau}^{i} \right) \right) \right] \right\}.$$
(A.4)

Notice that by construction \hat{V}_0 is convex. The Fitted Value Iteration procedure seeks a function in the form of Equation (A.4) that approximately solves the Bellman equation in (A.2).

Procedure The idea of the procedure is to iteratively define basis functions \mathcal{B}_k^S for k = 1, 2...whose pointwise supremum approximates V^* better and better. Towards this end, for each k we first perform one approximate Bellman step on each basis state, that yield new values V_k^i for $i \in \mathscr{L}$. Formally,

$$V_k^i = \frac{1}{N_S} \sum_p \left[\min_{X,B} r\left(X, B | s^i, D^p\right) + \gamma \hat{V}_{k-1}\left(s'\right) \right].$$
(A.5)

We approximate the expectation over the demand realization with a Monte Carlo method drawing N_S samples from D, independently for each state s^i . Since the new state s' can be written as a linear function of each s^i and following Equation (A.4) also $\hat{V}_0(s')$ has a linear representation in X and B, the Bellman step can be cast as a linear integer program. Then we obtain the new \mathcal{B}_k^S by solving the convex fitting problem.

$$\min_{\hat{V}_k^i, g_k^i} \qquad \sum_{i=1}^{N_A} \left(\hat{V}_k^i - V_k^i \right)^2 \tag{A.6}$$
s.t.
$$\hat{V}_k^j \ge \hat{V}_k^i + \sum_{\tau=1}^{T_{\max}+1} \gamma^{\tau-1} \left[g_k^i \cdot \left(s_{\cdot,\tau}^j - s_{\cdot,\tau}^i \right) \right] \text{ for all } i, j$$

The new basis functions are represented by $\mathcal{B}_{k}^{S} = \left\{ \left(s^{i}, \hat{V}_{k}^{i}, g_{k}^{i}\right) : i = 1 \dots, N_{A} \right\}$, where \hat{V}_{k}^{i} and g_{k}^{i} are the optimal solutions to Equation (A.6). In turn, this procedure generates a sequence $\left(\hat{V}_{k}\right)_{k=1}^{\infty}$ of approximate value functions. While this sequence cannot be guaranteed to converge to a limit,³⁸ it appears from Figure 8 that the Bellman error $e_{k} = ||\hat{V}_{k+1} - \hat{V}_{k}||_{2}$ quickly settles on small values, indicating that Equation (A.2) is approximately satisfied. Denote by \hat{V} the approximate value function obtained with this procedure.

 $^{^{38}}$ See the discussions in Gordon (1995) and Bertsekas (2018) for additional details on the reasons why Fitted Value Iteration may fail to converge.

A.3 Comparison with perfect information

We use \hat{V} as approximate value function to compute which decisions the central would have made when facing the demand realizations that we observe in the data. In practice, for every $t = 0, \ldots, T$ we solve

$$\min_{X,B} r\left(X, B|s^t, D^t\right) + \gamma \hat{V}(s')$$

where s^t is the state representation of the situation faced by the central planner as generated by its past decisions and D^t is the demand instance in period t. We collect all optimal decisions (X^t, B^t) and then we define the ballasting cost associated with them as

$$C^{DP} = \sum_{l,o,d,t} C^v_{l,o} \left(X^t_{l,o,d} + B^t_{l,o} \right).$$

As mentioned before, the central planner does not satisfy all the loads. Let \hat{C}^{obs} represent the benchmark ballasting emissions observed and computed as follows: for every route (o, d) and every time period t we sort the loads observed in reality in ascending order of ballasting emissions, and then consider only the $n_{o,d,t}$ least costly, where $n_{o,d,t}$ is the number of loads on (o, d) in period t that were served in the simulation; \hat{C}^{obs} is the sum the ballasting emissions associated to these voyages. Then we have that

$$\frac{C^{DP}}{\hat{C}^{obs}} \approx 80\%$$

Comparing this ratio with $Share^{Trade}$, we then conclude that the share of ballasting due to uncertainty is 7.50%. Because of our conservative way to compute \hat{C}^{obs} , it is likely that we are underestimating the share, which confirms our insight in Section 3 that uncertainty explains for a share between 7% and 11% of the ballasting costs.

B Tables and Figures

B.1 Tables

Vessel class	No. vessels	Min DWT	Max DWT
MR1	937	25,000	42,000
MR2	1,800	42,000	60,000
Panamax	475	60,000	80,000
Aframax	1,146	80,000	125,000
Suezmax	661	125,000	200,000
VLCC	867	200,000	∞

Table 5: Number of ships by vessel class in the dataset and their dimension. DWT refers to the *deadweight tonnage capacity*.

Vessel Class	Ballast distance (nm)	Ballast portion $(\%)$	\mathbf{CO}_2 emissions (tons)
MR1	1,054	40.2	218.8
MR2	1,411	38.6	293.8
Panamax	1,571	40.9	377.5
Aframax	1,424	42.1	443.8
Suezmax	2,611	43.7	933.1
VLCC	5,718	45.7	2943.71

Table 6: Average ballasting distance, average portion of ballasting on total voyage length and average emissions due to ballasting, broken down by vessel class.

Geographical Area	Avg. arrivals	Avg. departures
Red Sea	13.9	12.5
West Africa	21	26.9
Pacific Islands	2.1	0.2
Russian Pacific	0.1	7.6
Caribs	19.6	21.8
Baltic	11.9	31.8
South East Asia	60.5	51.8
US Gulf & Mainland	25.2	60.3
Korea / Japan	34.1	21.9
Black Sea / Sea Of Marmara	8	30.3
East Coast South America	21.3	20.5
West Coast Mexico	4.8	2.1
West Coast North America	13.1	8.2
West Coast South America	10.9	7.8
Australia / New Zealand	12.1	5.3
India / Pakistan	39.7	19.9
East Coast Canada	5.7	4.3
East Coast Central America	3.7	0.6
US Atlantic Coast	19.9	1.6
South East Africa	9.7	1.4
North Sea	1.2	5.4
UK Continent	59.2	31.3
West Coast Central America	4.2	2.3
East Coast Mexico	9.5	6.8
China / Taiwan	62.1	23.5
Mediterranean	68.1	56
Arabian Gulf	21.4	97.6
Arctic Ocean & Barents Sea	4	7.2

Table 7: Average number of ships arriving laden and departing laden from each geographical area. The table summarizes loads for all vessel classes, and the figures are obtained looking at the average number of loads arriving/departing in windows of 5 days.

Commercial Operator	Pool size	Share of $vessels(\%)$	Share of tonnage(%)
Teekay Corp	40	3.9	4.2
Sovcomflot	35	3.4	3.7
AET	33	3.2	3.4
Scorpio Commercial Management	27	2.6	2.8
Minerva Marine	23	2.2	2.4
ST Shipping & Transport	22	2.1	2.2
Cardiff Marine	21	2	2.3
Thenamaris	20	2	2.1
Shell	20	2	2.1
Navig8 group	20	2	2.1
Heidmar	18	1.8	1.9
Vitol	16	1.6	1.7
Trafigura	16	1.6	1.7
Equinor	15	1.5	1.7
Penfield Marine	15	1.5	1.6
Maersk	12	1.2	1.2
Signal Maritime	12	1.2	1.2
Zodiac Maritime	12	1.2	1.3
ExxonMobil	12	1.2	1.3
Frontline	12	1.2	1.2

Table 8: Average pool size for the 20 largest operators in the Aframax segment. The share of vessels indicates the percentage of active Aframax ships observed in the dataset controlled each commercial operator, and the share of tonnage indicates what share of total capacity of the segment is controlled by each commercial operator.

Vessel class	Trade imbalances (%)	Uncertainty $(\%)$	Fragmentation (%)
MR1	76.76	12.17	11.07
MR2	74.10	12.11	13.79
Panamax	73.41	15.43	11.16
Aframax	72.16	16.55	11.29
Suezmax	77.81	12.22	9.97
VLCC	87.06	7.79	5.15

Table 9: Share of ballasting attributed to trade imbalances, uncertainty and market fragmentation, using as objective function in Section 3 total fuel consumption associated to ballasting travels.

Dependent Variable	Log Utilization (Laden miles / Total miles)						
Depenaent Variable –	All	Aframax	Panamax	Suezmax	VLCC	MR1	MR2
Filtering: Fleet Size ≥ 1							
Log Fleet Size	0.0439^{***}	0.0655***	0.0745^{***}	0.0161	0.0229^{**}	0.0505^{***}	0.0347^{***}
	(0.006)	(0.014)	(0.028)	(0.019)	(0.012)	(0.013)	(0.012)
R-squared	0.1924	0.1836	0.1637	0.1606	0.2039	0.2176	0.1924
N obs	$23,\!670$	5,010	2,243	3,311	3,338	4,058	5,710
Filtering: Fleet Size ≥ 2							
Log Fleet Size	0.0521^{***}	0.0646***	0.0654^{*}	0.0203	0.0274^{*}	0.0571^{***}	0.0623^{***}
	(0.007)	(0.016)	(0.034)	(0.022)	(0.014)	(0.014)	(0.016)
R-squared	0.2219	0.1889	0.1609	0.1742	0.1926	0.3077	0.2268
N obs	17,924	4,081	1,580	2,728	2,826	2,495	4,214
Filtering: Fleet Size ≥ 3							
Log Fleet Size	0.0454^{***}	0.0430***	0.0747^{*}	0.0248	0.0300^{*}	0.0152	0.0807^{***}
	(0.008)	(0.017)	(0.039)	(0.023)	(0.017)	(0.018)	(0.016)
R-squared	0.2559	0.2004	0.1579	0.1957	0.1938	0.4019	0.2704
N obs	$14,\!624$	3,255	$1,\!304$	2,397	$2,\!450$	1,838	$3,\!380$
Controls							
Month-Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Operator	No	Yes	Yes	Yes	Yes	Yes	Yes
Operator * Vessel Class	Yes	No	No	No	No	No	No

Table 10: Effect of Fleet Size in Utilization (Distance)

Signif. Codes: ***:0.01, **:0.05, *:0.1

Notes: Table shows the results of the regression of fleet size and utilization at the Commercial Operator level as described in Eq. (13). Herein we extend our results and calculate the coefficient for the different markets (Vessel Classes) and we observe that all coefficients are significant besides the Aframax Class. Moreover, we filter the dataset to only observations when the average fleet size of a Commercial operator exceeds 1, 2, 3. This filtering is important given that Utilization is an aggregate metric and when few observations are available the metric could be biased. As an example, consider a fleet of 1 vessel with a Laden trip of over 30 days. Then, for the month of Laden the utilization of the fleet would be equal to 1 and for the ballasting trip the utilization would be equal to 0. Hence, this filtering helps on estimating the coefficient more accurately.

Dependent Variable	Log Utilization (Laden Days / Total Days)							
Dependent Variable –	All	Aframax	Panamax	Suezmax	VLCC	MR1	MR2	
Filtering: Fleet Size ≥ 1								
Log Fleet Size	0.0399^{***}	0.0426***	0.0848^{***}	0.0273	0.0125	0.0594^{***}	0.0249^{**}	
	-(0.006)	-(0.015)	-(0.029)	-(0.019)	-(0.013)	-(0.014)	-(0.012)	
R-squared	0.1991	0.1802	0.1475	0.1901	0.1851	0.2555	0.2095	
N obs	$23,\!804$	5,030	2,267	3,322	$3,\!354$	4,078	5,753	
Filtering: Fleet Size ≥ 2								
Log Fleet Size	0.0499^{***}	0.0610***	0.0812^{**}	0.0256	0.0187	0.0655^{***}	0.0421^{***}	
	-(0.007)	-(0.017)	-(0.035)	-(0.022)	-(0.017)	-(0.014)	-(0.016)	
R-squared	0.2203	0.2014	0.1575	0.1692	0.1681	0.3312	0.2396	
N obs	$17,\!986$	4,096	1,587	2,731	2,839	2,501	4,232	
Filtering: Fleet Size ≥ 3								
Log Fleet Size	0.0443^{***}	0.0507***	0.0821^{**}	0.0322	0.0251	0.0375^{*}	0.0440^{***}	
	-(0.008)	-(0.018)	-(0.039)	-(0.024)	-(0.020)	-(0.019)	-(0.015)	
R-squared	0.2421	0.2012	0.1573	0.1773	0.1857	0.3826	0.2809	
N obs	$14,\!659$	3,266	1,310	2,399	$2,\!458$	1,843	3,383	
Controls								
Month-Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Operator	No	Yes	Yes	Yes	Yes	Yes	Yes	
Operator * Vessel Class	Yes	No	No	No	No	No	No	

Table 11: Effect of Fleet Size in Utilization (Time)

Signif.Codes: ***:0.01, **:0.05, *:0.1

Notes: Table shows the results of the regression of fleet size and utilization at the Commercial Operator level as described in Eq. (13) but using the time definition of utilization instead of the distance one. Herein we extend our results and calculate the coefficient for the different markets (Vessel Classes) and we observe that all coefficients are significant besides the Aframax Class. Moreover, we filter the dataset to only observations when the average fleet size of a Commercial operator exceeds 1, 2, 3. This filtering is important given that Utilization is an aggregate metric and when few observations are available the metric could be biased. As an example, consider a fleet of 1 vessel with a Laden trip of over 30 days. Then, for the month of Laden the utilization of the fleet would be equal to 1 and for the ballasting trip the utilization would be equal to 0. Hence, this filtering helps on estimating the coefficient more accurately.

Dependent Variable	Fraction of	ballast em	issions (Ball	ast CO_2 Er	nissions /	Total CO_2 E	missions)
Dependent variable =	All	Aframax	Panamax	Suezmax	VLCC	MR1	MR2
Filtering: Fleet Size ≥ 1							
Log Fleet Size	-0.0078***	-0.0107	-0.0161	0.0010	0.0003	-0.0181^{***}	0.0012
	-(0.003)	(0.0071)	-(0.014)	(0.0095)	-(0.006)	(0.0065)	(0.0058)
R-squared	0.1909	0.1756	0.1326	0.1796	0.1938	0.3162	0.2237
N obs	$24,\!226$	5112	2,304	3374	$3,\!388$	4161	5887
Filtering: Fleet Size ≥ 2							
Log Fleet Size	-0.0128^{***}	-0.0152*	-0.0154	0.0021	0.0003	-0.0224^{***}	-0.0124
	-(0.004)	(0.0082)	-(0.017)	(0.0107)	-(0.008)	(0.0074)	(0.0078)
R-squared	0.2518	0.2106	0.143	0.1746	0.1785	0.4036	0.2579
N obs	$18,\!130$	4132	1,598	2745	2,861	2524	4270
Filtering: Fleet Size ≥ 3							
Log Fleet Size	-0.0118^{***}	-0.0149*	-0.0252	0.0007	-0.0015	-0.0103	-0.0143^{*}
	-(0.004)	(0.0089)	-(0.019)	(0.0117)	-(0.010)	(0.0105)	(0.0082)
R-squared	0.934	0.2210	0.1586	0.1739	0.1964	0.4493	0.3098
N obs	14,730	3283	$1,\!314$	2404	$2,\!475$	1853	3401
Controls							
Month-Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Operator	No	Yes	Yes	Yes	Yes	Yes	Yes
Operator * Vessel Class	Yes	No	No	No	No	No	No

Table 12: Effect of Fleet Size in CO₂ Emissions

Signif.Codes: ***:0.01, **:0.05, *:0.1

Notes: Table shows the results of the regression of fleet size and the fraction of CO_2 emissions at the Commercial Operator level as described in Eq. (13) but using the emissions definition of utilization instead of the distance one. Explicitly, the specification we run is $FracBallastCO2_{i,t} = \alpha \times \log(PoolSize_{i,t}) + \xi_t + \psi_{i,v}$, where $FracBallastCO2_{it}$ is the fraction of emissions stemming from the ballast legs over the total emissions of commercial operator *i* during time period *t*. It is calculated by: $FracBallastCO2_{it} = BallastEmissions_{i,t}/TotalEmissions_{i,t}$. In line with our Utilization specification, we filter the dataset to only observations when the average fleet size of a Commercial operator exceeds 1, 2, 3. This filtering is important given that the fraction of emissions is an aggregate metric and when few observations are available the metric could be biased. As an example, consider a fleet of 1 vessel with a Laden trip of over 30 days. Then, the fraction of ballast emissions would be 0. Hence, this filtering helps on estimating the coefficient more accurately.

	Aggregate	MR1	MR2	Panamax	Aframax	Suezmax	VLCC
A Utilization (07)	8.0272***	8.8444***	7.4764***	6.4532^{***}	12.5471^{***}	7.2892***	4.4762***
Δ Utilization (70)	(0.2848)	(1.0697)	(0.5185)	(0.9737)	(0.5567)	(0.8336)	(0.5508)

Table 13: Average percent difference in the utilization rate (distance based) between large pools and synthetic fleets that serve similar locations. Standard errors in parentheses.

	Aggregate	MR1	MR2	Panamax	Aframax	Suezmax	VLCC
Λ III the action (0%)	21.4096^{***}	23.4235^{***}	21.7873***	18.3534^{***}	33.9539^{***}	22.0736^{***}	7.0543***
Δ Utitzation (70)	(0.6271)	(2.1276)	(0.9647)	(2.0069)	(1.4369)	(1.9845)	(0.9684)

Table 14: Average percent difference in the utilization rate (time based) between large pools and synthetic fleets that serve similar locations. Standard errors in parentheses.

	Aggregate	MR1	MR2	Panamax	Aframax	Suezmax	VLCC
Λ Dallacting (07)	-5.9732***	-7.2260***	-5.0161***	-5.3114***	-8.2912***	-6.8617***	-3.7527***
Δ ballasting (70)	(0.4417)	(1.6402)	(0.9015)	(1.5876)	(0.9371)	(1.1360)	(0.8173)

Table 15: Average percent difference in the share of CO_2 emissions due to ballasting between large pools and synthetic pools that serve similar locations. Standard errors in parentheses.

			Dollo	ot Doloro						1 or	don Dolorro			
Danon dant Variahle			Dalla	se rreigios						IL2	an relays			
Toporeneries & no move	All	Aframax	Panamax	Suezmax	VLCC	MR1	MR2	All	Aframax	Panamax	Suezmax	VLCC	MR1	MR2
Geo Granularity: Port														
Log Fleet Size	-0.5283**	-1.2944^{**}	-1.0446	1.6364	-1.1144	-0.0305	-0.2886	-0.0327	0.2528	-0.0091	0.2792	0.5941	-0.341	-0.0073
	(0.240)	(0.573)	(0.816)	(1.179)	(1.492)	(0.482)	(0.363)	(0.125)	(0.304)	(0.356)	(0.542)	(0.808)	(0.235)	(0.207)
R-squared	0.1088	0.0959	0.0959	0.1629	0.2328	0.0946	0.065	0.081	0.0751	0.0518	0.1204	0.1458	0.1159	0.0623
R-squared Adj.	0.0886	0.0738	0.071	0.1176	0.1787	0.0664	0.0484	0.0601	0.0525	0.0257	0.0728	0.0856	0.0883	0.0456
N obs	47,602	10,246	5,637	3,351	2,734	7,867	17,767	47,602	10,246	5,637	3,351	2,734	7,867	17,767
Geo Granularity: Narrow Area														
Log Fleet Size	-0.5351^{***}	-0.7105**	+0.8990*	0.1663	-0.0785	-0.3706^{**}	-0.1867	-0.0608	-0.1403	0.3483	-0.0826	0.1107	-0.2666^{*}	0.1277
	(0.121)	(0.291)	(0.480)	(0.490)	(0.654)	(0.184)	(0.182)	(0.082)	(0.183)	(0.321)	(0.263)	(0.337)	(0.149)	(0.156)
R-squared	0.0852	0.0715	0.0919	0.0851	0.0914	0.088	0.0602	0.0659	0.066	0.045	0.0472	0.0707	0.0936	0.0503
R-squared Adj.	0.0774	0.0632	0.0797	0.0739	0.0771	0.0787	0.0529	0.058	0.0577	0.0322	0.0355	0.056	0.0843	0.043
N obs	140,286	30,171	12,055	15,599	13,966	24,709	43,786	140,286	30,171	12,055	15,599	13,966	24,709	43,786
Geo Granularity: Area														
Log Fleet Size	-0.5129^{***}	-0.6560**	-1.0145^{**}	0.3567	-0.2226	-0.3784^{**}	-0.1307	-0.043	-0.0728	0.289	0.0568	0.0068	-0.3123^{**}	0.2068
	(0.116)	(0.281)	(0.485)	(0.465)	(0.556)	(0.184)	(0.182)	(0.084)	(0.187)	(0.333)	(0.270)	(0.345)	(0.155)	(0.160)
R-squared	0.0964	0.0805	0.0881	0.0749	0.1108	0.0914	0.0605	0.0671	0.0699	0.0429	0.0442	0.0704	0.0956	0.0522
R-squared Adj.	0.0889	0.0726	0.0759	0.0638	0.0971	0.0822	0.0534	0.0593	0.0619	0.0301	0.0327	0.056	0.0864	0.0449
N obs	142,777	31,188	12, 171	15,979	14,241	24,907	44,291	142,777	31,188	12, 171	15,979	14,241	24,907	44,291
Controls														
Month-Year	Yes	Yes	Yes	Yes	Y_{es}	Yes	Yes	\mathbf{Yes}	Yes	γ_{es}	Yes	Yes	Yes	γ_{es}
Operator	No	Yes	γ_{es}	Yes	Yes	Yes	γ_{es}	No	γ_{es}	Yes	γ_{es}	Yes	Yes	γ_{es}
Operator * Vessel Class	Yes	No	N_{O}	No	No	N_{O}	N_{O}	Yes	N_{O}	No	N_{O}	No	N_{O}	No
Standard errors in varentheses.														
Cimit Codos: ***.0 01 **.0 05 *	*.0.1													
Digned Coucs:	1.0.1													

Table 16: Effect of Fleet Size in Delays (Idle time) for the different legs: Ballast and Laden Legs

Notes: This table reports the relationships between fleet size and the delays experienced in different sections of the voyages The main takeaway of this analysis is to validate that the gained efficiency of bigger fleet sizes comes from the Ballast portion of the trip. Specifically, we regress the fleet size with the delays in both the ballast and laden voyages. As expected, we observe that a bigger fleet experiences lower delays in ballast voyages but does not has an effect on improving the delays in laden voyages. Intuitively, laden voyages should not be improved because a vessel, independently of the management wants to minimize the time to its destination once it is loaded. To perform this regression, we calculate the delays as the difference between the travel time of the voyage and the *non-delayed* travel time of that leg which is estimated as the 20 percentile of the distribution of travel times from laden voyages.

Dependent Variable		L	og Number	of Trips in a	single OD pa	ir	
Dependent variable	All	Aframax	Panamax	Suezmax	VLCC	MR1	MR2
Geo Granularity: Area							
Log Fleet Size	-0.0187^{***}	-0.0171***	-0.0105	-0.0355^{***}	-0.0143^{***}	-0.0190^{***}	-0.0284^{***}
	(0.001)	(0.001)	(0.007)	(0.007)	(0.004)	(0.002)	(0.004)
R-squared	0.2578	0.2396	0.5601	0.6324	0.1856	0.328	0.3536
R-squared Adj.	0.2538	0.2337	0.4866	0.6007	0.1547	0.3138	0.3215
N obs	$111,\!475$	72,571	1,852	2,555	2,815	23,710	7,972
Geo Granularity: Narrow Area							
Log Fleet Size	-0.0151^{***}	-0.0126***	-0.0124*	-0.0312^{***}	-0.0136^{***}	-0.0148^{***}	-0.0207***
	(0.001)	(0.001)	(0.007)	(0.007)	(0.004)	(0.002)	(0.003)
R-squared	0.2462	0.2479	0.4552	0.48	0.2548	0.2724	0.3835
R-squared Adj.	0.2359	0.2333	0.3119	0.4126	0.2097	0.2383	0.3168
N obs	119,957	77,365	1,926	2,805	2,928	26,404	8,529
Fixed Effects							
Quarter-Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes
OD-pair	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Vessel Class	Yes	No	No	No	No	No	No

Table 17: Effect of the fleet size of a vessel and the number of trips per month made on an OD pair (Complexity of the operation)

_

Signif.Codes: ***:0.01, **:0.05, *:0.1

Notes: This table reports the relationships between fleet size and the intensity of trips that a vessel travel in the same lane. The main takeaway suggests that a vessel belonging to a larger fleet do not stick to a single lane and it is more likely to serve more routes than vessels belonging to smaller fleets. More generally, this result supports our findings that larger fleets have more complex operations via coordination and reach higher levels of utilization.

Demondont Variable		Log	Number of l	Ports per Ve	ssel in a Qua	arter	
Dependent Variable	All	Aframax	Panamax	Suezmax	VLCC	MR1	MR2
Number of ports							
Log Fleet Size	0.2326^{***}	0.2292^{***}	0.2707^{***}	0.2205^{***}	0.1783^{***}	0.2423^{***}	0.2113^{***}
	(0.005)	(0.014)	(0.019)	(0.017)	(0.012)	(0.012)	(0.010)
R-squared	0.3826	0.2625	0.2595	0.215	0.1919	0.4377	0.2613
R-squared Adj.	0.3715	0.2485	0.2417	0.2004	0.1791	0.4213	0.249
N obs	$66,\!158$	$13,\!113$	$5,\!475$	8,969	$11,\!991$	7,766	$18,\!844$
Fixed Effects							
Quarter-Year	Yes						
OD-pair	Yes						
Vessel Class	Yes	No	No	No	No	No	No

Table 18: Effect of the fleet size of a vessel and the number of different ports visited in a quarter (Complexity of the operation)

Signif. Codes: ***:0.01, **:0.05, *:0.1

Notes: This table reports the relationships between fleet size and the diversity of Ports, Narrow Areas and Areas that a vessel visit in a quarter. The main takeaway is to suggest that larger fleet typically have a better allocation of their vessels across the network of ports and that coordination allows them to have more complex operations and reach higher level of utilization. Specifically, we regress for each vessel its fleet size (the size of the fleet that the vessel belongs in a month) against the number of unique ports that the vessel visited in a month. We control by the trip itself such that the comparison is done across vessels that completed trips in the same month and route.

B.2 Figures



Figure 3: A map showing the demand imbalances across the network. The color of each country report the difference between the number discharge events and the number of loading events. A positive value is interpreted as a country which has more exports than imports while a negative value points to countries with higher imports than exports. Specifically, we observe that China is the larges importer while Russia and Saudi Arabia are the largest exporters.



Figure 4: Left histogram shows the variation in Fleet Size at a quarterly level, i.e., for each quarter and Commercial Operator we measure their average fleet size and generate the histogram using these values. The plot shows the market structure emphasizing its fragmentation (most of the observations happen on the left side). The blue dotted line report the median of the distribution which is equal to 3.3. Similarly, the plot on the right panel shows the utilization (laden miles/total miles) of each fleet at a quarterly level. The red solid line is at the 0.5 level, where all trips would be of the style *out and back* and the dotted blue the median of the distribution which equals 57.4%.



Figure 5: The figure shows the variation across time of the average fleet size of the seven largest Commercial Operators in each sub-market at a monthly level. This becomes relevant for our empirical analysis where we exploit the differences across time of the fleet size of an operator.



Figure 6: Distribution of the number of days between arrival of an empty vessel at port and its departure laden (i.e., loading time)



Figure 7: Distribution of the number of days between arrival of a laden vessel at port and its departure empty (i.e., unloading time)



Figure 8: L_2 distance between each successive set of basis values \hat{V}_k^i for the basis states of the Fitted Value Iteration procedure.