PRELIMINARY AND INCOMPLETE

PLEASE DO NOT CIRCULATE

Market Fragmentation and Inefficiencies in Maritime Shipping

Kostas Bimpikis   Giacomo Mantegazza   Salomón Wollenstein-Betech*

Abstract

Maritime transportation is essential for global supply chains, but the issue of 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 reveals that 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 13% by consolidating the market into small shipping pools, which avoids concerns about excessive market power. Consolidation improves utilization because larger pools better coordinate and diversify the set of ports they serve, which reduces the need for vessel relocations. At a higher 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 for sea transportation, this market suffers from the phenomenon of empty miles, also known as

*Graduate School of Business, Stanford University. Email: {kostasb, giacomom, salomonw}@stanford.edu.
We are grateful to Ludovico Crippa, Ilan Morgenstern, Daniela Saban, Stefan Wager for very helpful suggestions.

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

ballasting; it is estimated that between 40% and 50% of all miles traveled by ships are empty. Given the ever growing importance of sea transportation, empty miles pose a significant economic and environmental cost. The dry bulk and oil transportation markets are especially prone to ballasting, because their vessels usually transport only one load at a time from origin to destination.[2]

Ballasting has attracted considerable attention from both transportation scholars and industry practitioners. The attention has predominantly centered on the role of demand patterns: if some locations have outflows larger than inflows, the system requires ballasting to be balanced. However, transportation markets are usually characterized by many independent agents that compete to provide their services. Their interactions are as important as demand patterns in shaping equilibrium flows. 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, amplifies the effects of demand patterns. 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.

We adopt an applied perspective and partner 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 that cannot be eliminated, caused by trade imbalances. At the same time, our analysis shows that even partial consolidation of the market can lead to a substantial reduction in ballasting, and therefore mitigate the economic and environmental burdens associated with empty journeys without excessively increasing the market power of the shipowners. Because our dataset provides details about the fleets managed by individual shipowners, we can delve deeper and understand the mechanisms whereby fragmentation exacerbates the challenges posed by structural imbalances. Specifically, we observe a tendency among smaller shipowners to specialize 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 paper is among the first to (i) quantify the impact of fragmentation on ballasting and (ii) identify the channels channels through which fragmentation impairs industry efficiency.

Towards these ends, we exploit a rich proprietary dataset, that encompasses the voyages of approximately six thousand oil tankers throughout the four-year span from 2018 to 2022. Our data includes information about the *commercial operator* of vessels, i.e., the entity that controls and determines the loads of each tanker. Thus, we can sketch a complete picture of demand, supply, and the market structure. 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 on the map of Figure 3. Consistent with anectodal evidence and

[2] In contrast to container ships, that can be filled at various degrees and make intermediate stops.

previous studies, we find that the vessels were empty for 42% of the total miles traveled. Moreover, from our data about commercial operators we conclude that no company ever controls more than 5% of the market capacity, with a very long tail of companies managing less than 1% of the tankers: fragmentation is extreme. The left panel of Figure 4 sketches a complete picture.

Ballasting arises naturally when demand patterns are imbalanced. Moreover, the intrinsic uncertainty regarding the precise timing and locations of demand compounds the adverse effects of trade asymmetries. In this paper we focus instead on the effect of *market fragmentation* on ballasting, which is far less intuitive. We use the term "fragmentation" to describe both the situation where each operator has control over only a small number of vessels and the dispersion of information flows that arises from this market structure. Employing a set of integer and stochastic dynamic programs under different market structures and informational assumptions, we estimate the share of ballasting associated to imbalanced demand, uncertainty, and fragmentation. Consistent with intuition, we find that trade imbalances explain the largest share of ballasting; however, we also show that fragmentation is the second most important factor causing empty travels, with a share of ballasting twice as large as that of uncertainty. Moreover, we explore possibilities to reduce fragmentation and their impact on empty miles: our analysis shows that consolidating the market into shipping pools of 25 vessels each (equivalent to 3% of the global fleet) decreases carbon emissions associated to ballasting by 13%, about 70% of the savings that a central planner would achieve.

We then study the behavior of existing operators to understand the mechanisms that underpin the positive effect of consolidation. We concentrate on *shipping pools*, a trend that is gaining momentum in the oil shipping industry: with shipping pools multiple shipowners pool their vessels under a unique manager, that centralizes commercial operations. Our empirical analyses confirm that larger shipping pools have higher utilization (lower fraction of empty miles) rates. With reduced-form specifications, we show that consolidation works through two main 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 in a flexible way.

At a higher level, this paper also highlights the role played by market incentives in determining the efficiency levels of today's supply chains. Furthermore, it shows that economic and environmental costs associated with supply chains can be considerably reduced solely by organizing more efficiently the use of existing resources.

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. For example, Cachon, Daniels, and Lobel (2017) and Castillo, Knoepfle, and Weyl (2017) argue, theoretically and empirically, that surge pricing plays a role in rebalancing supply in response to spikes in demand. On the other hand, Bimpikis, Candogan, and Saban (2019) derives optimal pricing policies given long-term demand imbalances to encourage repositioning of supply in the network. Relocation policies are another important lever to balance transportation systems, and numerous 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 relocate supply in the network 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. This stream of literature has mainly concentrated on the problem faced by a single decision maker that can either controls prices or order relocations. Our study complements this work because we consider instead on a market with many independent decision makers and we focus on the relationship between market structure and ballasting. Furthermore, while we consider optimization problems with a central party, we use their optimal values to estimate the impact of demand imbalances and other factors on empty travels, rather than focusing on providing policies that can be implemented in the real world.

From an applied point of view, the optimization problems we solve are related to the vast literature on devising approximate solutions to the stochastic dynamic programming problem faced by fleet managers: see, e.g., Godfrey and Powell (2002a) and Godfrey and Powell (2002b). Also related to this paper are Prochazka, Adland, and Wallace (2019) and Adland and Prochazka (2021), who use an optimization-based approach similar to our to estimate the value of foresight and contractual flexibility in dry-bulk markets.

A growing emphasis on decentralized transportation markets comes from the Economics literature in Industrial Organization. For example, Frechette, Lizzeri, and Salz (2019) and Buchholz (2022) study the taxi market and show how different barriers to entry shape the marker 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 leads to 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 markets, the issue has generally been overlooked by the Industrial Organization literature: we innovate on this aspect by showing that fragmentation has a first order effect and we provide a first measurement of its importance.

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 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 by identifying 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.

Finally, our work also contributes to the growing literature concerned with improving the sustainability of firms' operations: see Lee and Tang (2018) for a review on the subject. While indispensable to global trade, maritime shipping is also harmful to the environment in a variety of ways:[3] we show that simple steps can substantially improve both the economic and environmental outcomes associated with the 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.

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, representing approximately 11% of total volumes; and other specialized categories, such as chemicals, accounting for the residual portion.[4] It is noteworthy that the oil transportation sector alone ac-

[3] See 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.

[4] https://unctad.org/system/files/official-document/rmt2022_en.pdf

counts for a substantial 30% share of the seaborne trade volume, making it one of the most relevant markets in maritime transportation.

Similar to raw materials and finished goods, oil and refined products often need to be transported over long distances. While there exist extensive land infrastructures to move crude 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.[5,6] A fleet of about 8,800 ocean-going oil tankers operates this industry, divided into six (basically) independent markets based on deadweight tonnage capacity (see Table 5).[7] For reference, the smallest of these tankers have a carrying capacity of approximately 300,000 barrels of oil, while the largest can accommodate nearly ten times that volume. For context, it estimated that the United States consume around 20M barrels of oil/day.[8]

In the context of oil product extraction, refining, and distribution, it is customary for companies not to possess a fleet of tankers for transportation purposes. While major industry players maintain modest tanker fleets, the majority of oil producers rely upon a network of heterogeneous shipowners, who make their tanker vessels available for leasing. Hiring of oil tankers occurs through a variety of contract types that cater to the specific needs of exporters. These contracts are primarily characterized by two fundamental aspects: 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, after which the shipowner is free to contract with a new exporter, or for longer periods of time (e.g., a year). The hiring party generally pays a daily rate to the owner for the duration of the contract, and the rate depends on whether the owner remains responsible for fuel and other variable costs associated to travel.[9] Vessels are predominantly engaged on a per-voyage basis, on what is often referred to as the *spot* market for charterers. In fact, even vessels hired for longer periods of times are often sub-let to other exporters on these markets to take advantage of rates fluctuations, so that virtually all loads are transported with spot contracts.

The operation of the spot market hinges on a network of intermediaries known as 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 providing detailed 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

[5] https://www.bp.com/content/dam/bp/business-sites/en/global/corporate/pdfs/energy-economics/statistical-review/bp-stats-review-2022-full-report.pdf

[6] https://www.eia.gov/international/analysis/special-topics/World_Oil_Transit_Chokepoints

[7] See also the UNCTAD Review of Maritime Transport for additional details on the maritime shipping activity https://unctad.org/publication/review-maritime-transport-2022.

[8] https://www.eia.gov/tools/faqs/faq.php?id=33&t=6

[9] For more information, the website of the Baltic Exchange provides excellent primers: https://www.balticexchange.com/en/who-we-are/guide-to-modern-shipping.html

available to service the load, and the final selection of the vessel is determined through a negotiation process that includes bargaining on the charter rate. The contract agreed upon (known as *fixture*) includes the precise load location and day, and the daily rate. Importantly, the daily rate is only payable for the duration of the voyage from the loading port to the final destination. Given the inherent unpredictability of sea voyages, fixing of the contract is initiated not long before the intended load time, so that only vessels sufficiently close of the load area can be considered available to service a load. It is clear then that shipowners must make strategic decisions regarding the optimal positioning of their vessels to secure lucrative contracts.

Let us examine the scenario from the perspective of a ship operator who has recently discharged its cargo. It is unlikely that the vessel has already been contracted for a new load.[10] The ship 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 their 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 will end 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.

This explanation of market dynamics sheds light on why oil transportation is especially susceptible to the phenomenon of *empty miles*, a feature common to virtually all transportation markets. Empty miles substantially impact operating costs of oil shipping companies, and they also impose a significant burden to the environment, it in terms of carbon emissions and, e.g., discharge of other pollutants at sea which disrupt fishing zones. Because of the economic and environmental costs associated with ballasting, understanding its causes and exploring potential remedies are subjects are primary importance to the industry.

2.1 Description of data

We partner with a shipping and analytics company, who provided us with data consisting of 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. Each voyage is characterized by the unique identifier of the vessel that performed it as well as its load and discharge ports with the respective arrival and departure dates; we are also provided 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

[10] Cargo assignments are typically confirmed with minimal lead time before their intended departure, because of the potential delays associated with port congestion, discharging, and refurbishing of vessels. Moreover, exporters are reluctant to charter a tanker that has not yet discharged its previous cargo.

exactly how long each leg was for each voyage, and how much time each ship spent docked at 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 ports level information, *narrow area*, and *intermediate area* information: narrow areas partition the world in 51 regions, while intermediate areas in 28.[11] A list of the intermediate areas can be found in Table 7.

In addition to geographical data, we also have information about fuel consumption, pollutants emissions, speed, and the name of the *commercial operator* that managed the vessel for that voyage. By commercial operator we mean the agent that was in charge of directing the vessel's ballasting decisions, contracting with the brokers, and that collects the earnings associated with travel. The commercial operator may or 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.[12] Shipping pools are one of the primary subjects of interest of the paper, and we discuss them in greater detail at the end of the section.

Voyages correspond to demand for transportation that is actually fulfilled, which is by definition lower than potential demand, and the loads that get transported may, in principle, differ in unobservable characteristics from demand that does not find a matching vessel. However, some features of this market reassure that potential demand should not be significantly different from realized loads: first, since contracts are agreed upon at short notice, oil exporters likely have their load (almost) ready for departure at the time they contract with brokers; second, there is no outside option to sea transport once oil has arrived to the tankers terminal;[13] third, significantly delaying the time of transportation is usually not an economical strategy, 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. Finally, unfortunately we do not have access to the contracted rates; the lack of prices somewhat limits the range of technical tools we can use for our analysis and therefore the scope of the paper, since we cannot run counterfactual estimations.

To ensure consistency between the analyses of Section 3 and Section 4, we drop 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., their sequence of voyages is a continuous path on the network of locations.

[11] We also conduct robustness checks at different granularity levels and the insights continue to hold.

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

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

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 The global distribution of oil reserves exhibits a pronounced asymmetry, leading to substantial trade imbalances in the oil transportation industry. This disparity is prominently reflected in our dataset, as shown in Table 7 and Figure 3. The table provides the average number of cargo shipments originating from and destined for various geographical regions during the period spanning January 1st, 2018, to January 1st, 2020, while the map visually portrays the same dataset.[14] Examining the data reveals marked disparities in shipment flows, with certain regions experiencing a significant influx of cargo arrivals in comparison to departing shipments. For instance, the "China / Taiwan" region witnessed a threefold surplus of cargo arrivals over departures. Conversely, in other regions such as the "Arabian Gulf" (i.e., Persian Gulf), there was a nearly five-fold surplus of cargo departures relative to arrivals. Finally, there are regions where shipment flows exhibit a greater equilibrium, as observed in the "Red Sea" and Mediterranean areas. Since this pattern makes is very unlikely that vessels always find a load to transport in the same location they discharge at, imbalanced demand for transportation is widely held to be the main reason why tankers sail empty.

Ballasting Owing to these trade imbalances, ballasting is very common 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. More precisely, 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 or an estimated cost of 7.2 billion USD based on September 2023 pricing. To contextualize this, it is tantamount to the emissions generated by 7 million passenger cars, a number approximating the total registered vehicles in the entirety of Greece.[15] 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. Throughout the paper, we argue that the very fragmented nature of asset control in the market exacerbates the issue of ballasting and that a marginal consolidation can substantially reduce the waste of resources.

Market fragmentation and pools Our dataset is noteworthy because we have access to information about the commercial operators of the tankers. This allows us to trace a detailed picture of

[14] For the purpose of Table 7 we only consider travels occurred before the outbreak of the COVID-19 pandemic.
[15] https://www.epa.gov/greenvehicles/tailpipe-greenhouse-gas-emissions-typical-passenger-vehicle

market concentration, its evolution over time, and of how the number of vessels under management influences the decisions of commercial operators. As mentioned above, commercial operators are responsible for the operational choices of tankers (which loads to bid for, where to ballast, etc...), so that they are the relevant players in the market. The control structure is extremely fragmented: as evidenced in the left panel of Figure 4, a substantial number of operators engage with a fleet size ranging from 1 to 3 vessels, with the median number of controlled vessels amounting to 2.66; 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 size. This difference can be attributed to the fact that large commercial operators manage vessels via *shipping pools*, i.e., they aggregate to the fleet of tankers they own vessels owned by other entities (usually, small shipowners themselves). A shipping pool is formed when a number of shipowners decide to "merge" their fleets under centralized management. 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.[16]

Shipping pools provide commercial operators with a more flexible avenue for expanding their asset base when compared to formal mergers or acquisitions, because vessels may join or withdraw at any time. Pools have become increasingly popular thanks to the advancement in computational capabilities that allowed real-time optimization.[17] Despite this increase in popularity, their market share remains low: Table 8 shows that for the Aframax 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.[18] Similar patterns prevail across other vessel classes, with a marginal uptick in concentration evident among larger tankers. Importantly, shipping pools exhibit substantial heterogeneity not only in terms of size, but also in the level of 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).

3 Determinants of ballasting

While the widespread occurrence of ballasting is a well-recognized phenomenon within the transportation sector, the factors underpinning its occurrence remain comparatively underexplored, and their relative significance is unknown. In the case oil tankers, we have previously acknowledged the substantial role played by geographical factors, but to what extent can we ascribe the ob-

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

[17] https://www.thesignalgroup.com/signal-maritime/pools

[18] One 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.

served prevalence of ballasting solely to the disparities in demand? Answering this question is the necessary first step in understanding whether the market is operating close to efficiency, or if instead if there exists untapped potential for improvement from other sources of inefficiency, such as ballasting decisions.

We classify the sources of ballasting inefficiencies in three main categories: (i) trade imbalances; (ii) uncertainty, i.e., that time and locations of future loads are unknown, and therefore tankers may decide to ballast away from a port when in fact a load would have been offered shortly after; (iii) fragmentation, i.e., that there is a large number of commercial operators with few vessels and that, as a result, also the flow of information is extremely fragmentary. In fact, not all vessels in a given geographical area are always aware of all loads departing from there.[19] 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.[20]

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 the transportation market 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. We then compare this baseline ballasting calculated from the model with the ballasting observed in the data. This approach estimates that trade imbalances explain a share from 70% (for smaller vessels) to 90% (for large vessels) of ballasting.

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 operational factors and uncertainty as well. We estimate that operational factors account for a share between 16 % (smaller vessel classes) and 7% (largest tankers), and uncertainty for the remaining 5-13%.

Optimization assumptions To perform the analysis described above, we use data from all four years of voyages to estimate costs and travel times between geographical areas; however, we only optimize only over the two-year period before the outbreak of the COVID-19 pandemic. While this

[19]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.

[20]Size 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.

is immaterial for hindsight optimization, restricting to pre-pandemic years is important to ensure consistency of this approach with the alternative one we propose in Appendix A to capture the role of uncertainty. We discretize time into five-day periods, so that voyages departing from January 1st, 2018 to January 5th, 2018 are imputed to period $t = 0$, and so forth. All optimization problems are solved at the level of intermediate geographical areas described in Table 7. Finally, we assume (i) that the laden leg of each voyage has the same duration irrespective of the ship assigned to it; (ii) that 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 of 12kn.[21] The second is less realistic because tankers are heterogeneous even within class; however, our data do not include type and quantity of the loads, so we cannot take further constraints into account.

3.1 Trade imbalances

Demand for oil transportation is influenced by many political, macro-economic, and other market factors (such as the price of transportation). Because of this complexity, instead of fitting a necessarily imperfect model to observed demand, we accept demand as an exogenous input. 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. Let T denote the number of 5-day period from January 1st, 2018 to January 1st, 2020, and let $\mathscr{L}$ denote the set of locations in the transportation problem. We introduce two types of decision variables: $X_{l,o,d,t}$ accounting for the number of vessels that start their voyage at time $t \in \{0, \dots, T\}$ 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. We use as proxy for the cost of ballasting between location l and location o the average CO_2 emissions, denoted by $C_{l,o}$. We choose this metric instead of miles since emissions better capture important cost-generating events such consumption of fuel when waiting to cross a canal.[22]

The program 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

[21] Extensive research in maritime engineering has set the most economical speed for tankers; see, e.g., Psaraftis and Kontovas (2013).

[22] CO_2 emissions are provided to us and are computed using a one-to-one mapping from fuel consumption. The mapping is an industry standard as required by IMO regulations.

Parameter	Explanation
$D_{o,d,t}$	Number of loads from area o to area d that depart in period t as observed in the dataset
$\chi_{l,\tau,o,t}$	Indicator taking value 1 if and only if a ship departing location l at time τ can reach area o at time t (clearly $\chi_{l,\tau,o,t} = 0$ if $\tau > t$). Obtained by estimating the average ballast time between the two areas. Vessel-class specific, since larger tankers cannot use the Suez and Panama canals, and therefore need to take longer routes.
$I_{l,t}$	Number of vessels that become available in location l at time t, either for the 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).
$O_{l,t}$	Number of ships that cease to be available when in location l at time 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
$C_{l,o}$	Average CO_2 emissions in tons of a ballast travel between areas l and o

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

that discharge al location l at time t;[23] $I_{l,t}$ 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, $\chi_{l,\tau,o,t}$ encodes whether a vessel leaving location l at time τ arrives in location o by time t. Further details about these parameters are in Table 1. Using these parameters and variables, we formally state the integer program as:

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

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

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

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

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

Equation (1) gives the total ballasting costs associated with the assignments in X. The constraints in Equation (2) impose that all loads observed in the data are transported also under the

[23] Since 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.

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 of 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 ($Y_{p,t}$), or outflows (O). 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 ballast 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:

$$Share_v^{Trade} = \frac{C_v^*}{C_v^{obs}}.$$

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 bigger classes. 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

Let us focus for a moment on the Aframax vessel class: Table 2 says that 72% of the ballasting we observe can be explained by trade imbalances. In other terms, there is potential to reduce the cost of empty miles by 28% by addressing fragmentation and uncertainty. Suppose that, instead of a unique clairvoyant central planner, there was one clairvoyant planner for each shipping pool, 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 shipping pools lived in a world with no uncertainty. It is a lower bound because it is obtained

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 exclusively 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 were built or entered the market after January 1st, 2020.

assuming that the pool planner minimizes ballasting costs, while in reality they maximize profits.[24] We can then 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 shipping pool $i \in \mathscr{P}_v$, where $\mathscr{P}_v$ denotes the set of shipping pools in vessel class v, we solve the following problem.

$$\begin{aligned}
\min \quad & \sum_{l,o,d,t} C_{l,o} X^i_{l,o,d,t} && (6)\\
\text{s.t.} \quad & \sum_{l,\tau} X^i_{l,o,d,\tau} \chi_{l,\tau,o,t} = D^i_{o,d,t}, \qquad \forall (o,d,t) \in \mathscr{L} \times \mathscr{L} \times \{1,\dots,T\} && (7)\\
& \sum_{o,d} X^i_{l,o,d,1} + Y^i_{l,1} = I_{l,1}, \qquad \forall l \in \mathscr{L}, && (8)\\
& \sum_{o,d} X^i_{l,o,d,t} + Y^i_{l,t} + O^i_{l,t} = A^i_{l,t-1} + I^i_{l,t} + Y^i_{l,t-1}, \qquad \forall l \in \mathscr{L},\ \forall t && (9)\\
& X^i_{l,o,d,t},\ Y^i_{l,t} \in \mathbb{N} && (10)
\end{aligned}$$

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 shipping pool level. Let C^i_* denote the optimal value of this program, and the system-wide ballasting cost of the decentralized system be

$$C_v^P = \sum_{i \in \mathscr{P}_v} C^i_*.$$

[24] 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.

By definition, it must be that $C_v^* \leq C_v^P \leq C_v^{obs}$, so that

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

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

$$Share_v^{Frag} = \frac{C_v^P - C_v^*}{C_v^{obs}} \tag{11}$$

$$Share_v^{Uncertainty} = 1 - \frac{C_v^P}{C_v^{obs}} \tag{12}$$

In the numerator of Equation (11) we compare the optimal cost of the assignments with no uncertainty and a central planner against the cost of no uncertainty with multiple decision makers, 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 computed when ship 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.

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

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 largest vessel class follows a somewhat more predictable schedule. The other striking feature that emerges 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 failure, because the data do not allow us to back out 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 (less CO_2 emissions). However, it is clear than putting all tankers in the world under a single entity is neither feasible, nor desirable, so the next section is devoted to understanding "how much" consolidation is sufficient.

3.3 Partial consolidation

Our data list 169 Aframax commercial operators, and in Section 3.1 we considered only one. In order to understand the value of *partial* consolidation, we need to create "synthetic" shipping pools, that do not exist in reality. This affords us the liberty of choosing the size of the pools; in other terms, we can test different degrees of consolidation and understand the minimal pool size required to achieve most of the gains that a central planner would get. We proceed as follows.

For a fixed vessel class v, we define Δ_v^* as

$$\Delta_v^* = \frac{C_v^* - C_v^{obs}}{C_v^{obs}}$$

which represents the gain from centralization. Then we randomly split the available tankers into shipping pools of identical size ϕ, where $\phi \in \{1.5\%, 2\%, 2.5\%, 3\%\}$ 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), that 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. Our metric of interest is the ratio $\rho_v(\phi) = \dfrac{\Delta_v^*(\phi)}{\Delta_v^*}$, which gives the fraction of the central planner's savings that can be achieved with pools of size ϕ, and can be essentially interpreted as a competitive ratio.

Figure 1 plots $\rho_v(\phi)$ for the different vessel classes and for different pool sizes. These figures show that even with small pools it is possible to obtain ballasting savings that are very close to the optimum of the central planner. This is particularly evident for smaller vessel classes such as Aframax: splitting all 792 ships in pools of 25 units each (about 3% of the fleet) yields between 70% and 80% of the savings that the central planner could have obtained. Using the figures estimated in Table 3, this corresponds to an overall cut in ballasting emissions of about 13%, or 4 million tonnes of CO_2 not released in the environment (= 1 million fewer passenger cars). These results suggest that the benefit of centralized vessel-load assignments accrue 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 section we investigate in greater detail the mechanisms whereby shipping pools reduce ballasting costs, and Figure 2 shows decreasing marginal returns are supported also by empirical analysis.

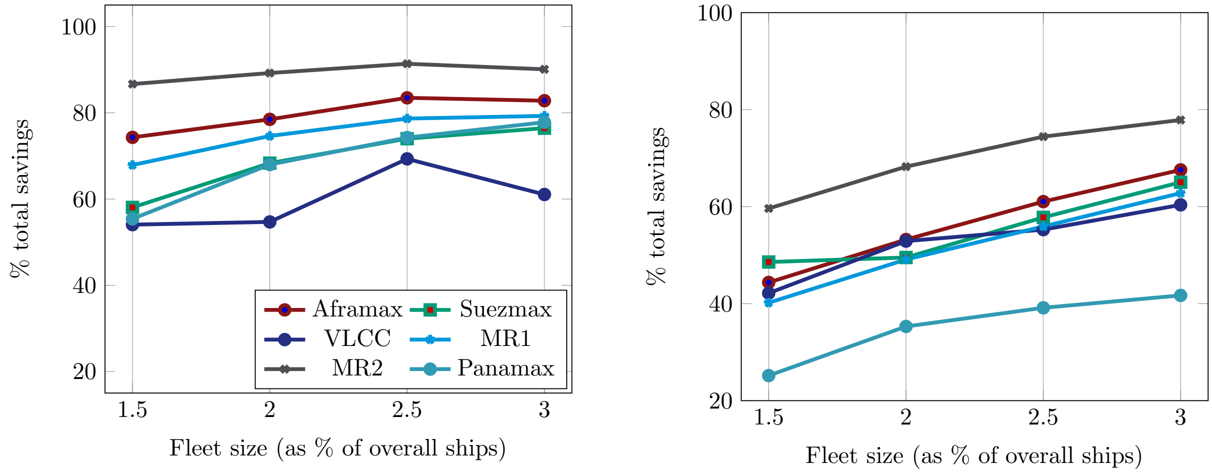

Figure 1: Plot of the fraction of savings that can be achieved by splitting all ships in a vessel class into small pools. The left plot is obtained by first splitting equally among pools the ships that start in the locations with highest demand, and the rest uniformly at random. The right plot is obtained by splitting all vessels uniformly at random.

3.4 Limitations

The estimated shares of ballasting attributed to various factors suggest that the market for oil transportation is operating quite far from efficiency. Moreover, despite being based on data from a specific market, the analysis is general enough that it can be applied with little modification to other decentralized transportation markets, such as dry bulk shipping and full-truckload trucking. At the same time, it is important to recognize some limitations of our data and approach.

First of all, we stress that our estimates should not be interpreted as counterfactuals in the sense of structural estimation, for two main reasons: (i) we set up cost-minimizing problems, while in reality any decision maker is a profit-maximizer; (ii) we do not consider the equilibrium response of demand for transportation to consolidation, and its effect on the bargaining power of pool managers. Because we do not have information on rates of the loads, we cannot estimate a structural model and obtain such counterfactuals. However, one should notice that all optimal assignments we compute *improve* over the total profit earned in the market, since revenues are unchanged and costs decrease. Thus, our estimates, while rough, point in the right direction, and it is unlikely that the relative share of one factor to the others would be substantially different in a structural model.

Second, 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 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 pool, but not whether they will secure them. We address this issue in Appendix A, where we consider a central planner that solves a stochastic dynamic program. 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 would imply that the estimate for the share of operational factors is closer to 20% than 16%.

4 Observed shipping pools

In this section, we empirically investigate how partial consolidation impacts the overall efficiency of pools and identify the main mechanisms that drive this improvement. To accomplish this, we employ two primary identification strategies that center 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, focusing on the relationship between fragmentation and efficiency.

The regression estimates reveal that larger pools tend to achieve higher levels of efficiency. A portion of this efficiency gain is attributed to enhanced coordination and more complex operations. We show that the relationship between efficiency and operations complexity is strong by the use of additional regressions. Finally, we empirically establish that there are decreasing marginal returns to consolidation, suggesting that the benefits from further consolidation become negligible when the pools manage more than 20 tankers. All results in this section validate and agree with our optimization results of Section 3.

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 15, then we say that *X Shipping Co.* managed a pool of size 1.5 during April 2018.

4.1 Efficiency

We begin our analysis by answering 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 allows 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 set

$$Utilization_{i,t} = \frac{\sum_n M^L_{n,i,t}}{\sum_n M^T_{n,i,t}},$$

where $M^L_{n,i,t}$ denotes the laden distance of the n-th voyage operated by commercial operator i whose start date is in period t, and $M^T_{n,i,t}$ 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.

A limitation of our data is that we do not observe the negotiated rate of a load, nor who was invited to submit offers, so we cannot ascertain whether larger pools have some systematic advantage over smaller ones.[25] To counter this issue, we exploit the variation in pool size of each commercial operator over time and observe the differences in efficiency when they change their size. Figure 5 shows the variation in pool sizes across time for the five largest commercial operators in each vessel class. Exploiting this variation enable us to alleviate the existing heterogeneity in bargaining power and negotiation skills across different operators, but does not solves the endogeneity of the intensive margin (growth/shrink of the same entity). We try to address this in our second specification.

[25] Adland, Cariou, and Wolff (2016) report that the quality of the match between shipowners and charterers is a first-order determinant of the freight rates.

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

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

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 and other unforeseen events that affect global trade;[26] the 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, thereby 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 run similar specifications for each vessel class and show that the coefficients are statistical significant and in the same direction and sizes to the one described (see Table 9).

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 mild variations in their pool size, but it can be more difficult to justify in the case of operators that grew/downsized substantially over time. In fact, there may well be factors that correlate with both utilization and the pool size (e.g., "market ability", because vessels may want to join the pool of an operator with fame of having high utilization rates).

There are two types of endogeneity risks in our analysis. The first type arises from time-varying unobservable characteristics of operators that improve their utilization as the pool expands, for example 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 as a surrogate for an operator's "market weight", so that we expect that it also captures the effect of other factors that correlate with size, but that we cannot observe. The second type is reverse causality, i.e., improved utilization drives a pool's expansion. For instance, an operator might achieve higher utilization rates by investing in advanced analytical capabilities without expanding its pool. New vessels may then decide to join the pool, attracted by the prospect of increased utilization for their tankers.

[26] E.g., the week-long obstruction of the Suez Canal that occurred in March 2021 that caused widespread delays in global shipping.

Synthetic consolidation. We run an alternative analysis to capture the effect of pool size in utilization with two main objectives. First, to address the issue of 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. In particular, for every large pool we build a "synthetic" pool composed of smaller operators that serve similar geographical areas, and look at the difference in utilization between the groups. Concretely, for any pool managing more than 10 tankers in a given quarter (which corresponds to the 80% percentile in the distribution of pool sizes) we generate a synthetic pool composed by smaller commercial operators. We do this by aggregating smaller pools that served loads between the same locations as the large fleet until their combined size equals that of the bigger one.[27] Then, we compare using a t-test the quarterly utilization of the large fleet and the "synthetic" fleet. Notice that the purpose of synthetic pools is to obtain a decentralized benchmark for the large pools; we do not change the vessel-load assignment observed.

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 10. The size of the difference in means is 8% when the average pool size of the fleets composing the synthetic fleets and larger fleets are 3 and 22, respectively. Hence, 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 provide evidence of a positive impact resulting from the centralization of decision-making. In fact, since our synthetic and large fleets server the same locations, we can rule out the possibility that the observed positive difference is attributable 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 are not comparing the same entity over time. However, it remains uncertain whether this centralization is beneficial primarily because it improves coordination or because a large pool, with more influence over brokers, is better at securing loads.

Ballast and laden delays. We have showed that large pools achieve higher overall utilization than smaller fleets, however, we have not disentangle the source of the effect. A first step is to confirm that such improvement in efficiency comes from the ballast leg of the trip and not from the laden one. It is reasonable to expect that both small and large fleets would not encounter delays during their *laden* legs. This is because once a vessel is loaded, it proceeds directly to its discharge port without interruption. Conversely, the ballast leg may be susceptible to delays for

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

Dependent Variable	Coordination				Operations Complexity	
	Log Utilization [%]	Δ Utilization [%]	Ballast Delay [Days/month]	Laden Delay [Days/month]	Log Trips on a single OD [# trips]	Log Ports visited by Vessel [# 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 × 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

various reasons. For instance, delays could occur if a vessel awaits a load outside the port; or if it incorrectly repositions to a region where no actual load materializes, so that the vessel 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 20% 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.

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

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

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 ships belonging to smaller pools take significantly longer time to ballast between two ports than their counterparts in large fleets (see results by vessel class in Table 11). 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. These exercises

[28] Taking the 20% 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.

also serves as a sanity check that confirms the fact that delays in laden legs would be rare for both large and small fleets but delays in ballast are expected to be larger for smaller fleets.

Since tankers start looking for new cargo only after discharging the previous load, these estimates suggest that size affects the time it takes to fix a new contract.[29] 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.[30] 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 seems certain that the 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. Although it is hard to disentangle all the mechanisms, 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.

A legitimate question is then why this happens: we believe the reason connects to the level of risk aversion of pools and how this relates to their size. In fact, shipowners are generally considered risk averse agents, especially when in making ballasting decisions, because it entails a cost in the present for future uncertain profits.[31] 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 can be expected to make more conservative decisions. This then results in more seamless and efficient operations for large pools.

It is natural that large commercial operator can serve more locations *at the same time* than smaller operators; therefore, to obtain meaningful comparisons in our analysis we run regressions at the vessel level (not pool level). We pose the question about the complexity along two complemen-

[29] It is implicitly assumed that, once under way, vessels sail at the same speed; this is supported by data that show that average speed is remarkably steady across voyages, as mentioned in Section 2.1.

[30] There is anecdotal 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.

[31] See, e.g., Albertijn, Bessler, and Drobetz (2011); Drobetz, Haller, and Meier (2016) for a description of the financing problems of shipowners and how they respond to risks.

tary 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:

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

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

The variable $Loads^{o,d}_{s,i,t}$ 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.[32]

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 have 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 partial consolidation is already sufficient to achieve most of the benefits of coordination, and we find that marginal benefits from consolidation effectively vanish beyond a size threshold. With this section we aim to further validate that intuition and extend the results of Section 4.1, where we only consider a linear relationship between size and utilization. This approach also serves to confirm that consolidation is helpful also when the pool manager is a

[32] Hence 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.

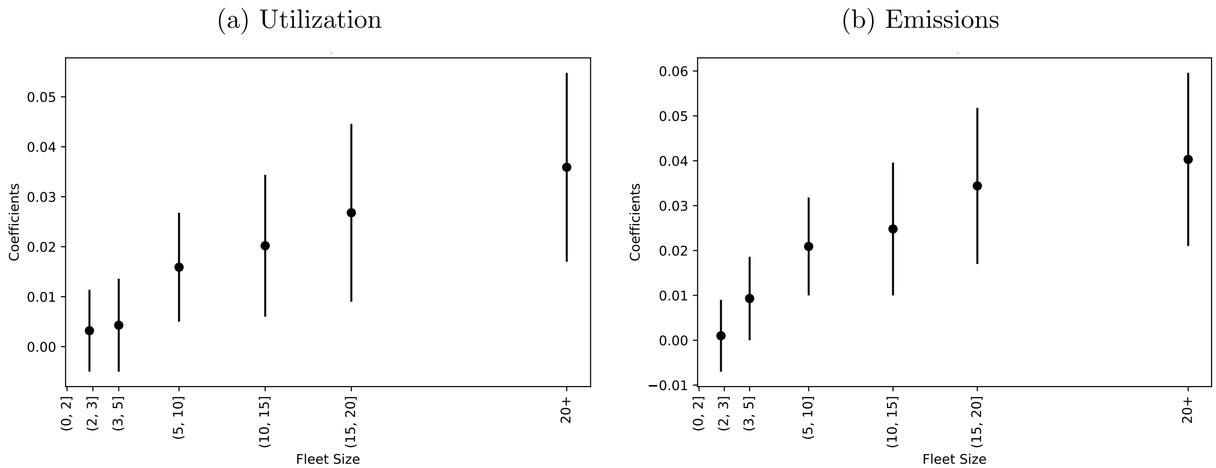

Figure 2: Coefficients and 95% confidence interval of regression (18) with panel data at a monthly-level. Results reveal monotonically increasing and concave functions that plateaus when fleet size is between 15 and 20, pointing to similar results as the ones provided by the optimization model. The baselines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

profit-maximizing agent, 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:

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

Our results, shown in Figure 2, report the coefficients and 95% confidence intervals of the different bins of pool sizes against the excluded group of $(0,2]$. As an example, the coefficient for the $(5,10]$ group in Figure 2b is ≈ 0.02, indicating that, on average, the emissions generated by a fleet with pool size between 5 and 10 has 2 more percentile points than the utilization of the $(0,2]$ group. Specifically, the average utilization 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 pool size between 5 and 10 is (58.6 + 2.0)% = 60.6%.

Figures 2a and 2b 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 smaller 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 trade imbalances. Notably, our research shows that even a modest degree of market consolidation can yield substantial advantages. Drawing on a dataset composed of the voyages of six thousand oil tankers, we assess that trade imbalances are responsible for 70-75% of empty miles, while fragmentation emerges as the second most significant contributor, accounting for 15-20% (with the remaining 5-10% attributed to structural uncertainty). We propose two channels to explain the beneficial impact of consolidation on ballasting. First, centralizing decision-making generates a coordination effect, enabling vessels to efficiently service the same locations with less reliance on ballasting. Second, larger pools diversify their network of served ports, optimizing the utilization of tankers and thereby minimizing protracted 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.[33] With shipping pools, individual shipowners aggregate their resources and entrust a single decision-maker to carry out commercial operations. Our data suggest that small shipping pools, unlikely to exert detrimental effects on 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 25-30 units results in a 13% 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.[34,35] Finally, at a higher level the paper displays the extent of sustainability gains that can be achieved by optimizing the

[33] Brancaccio et al. (2023) argues that, in a slightly different context, an Uber-like central planning platform would decrease social welfare, and it has been argued that extreme consolidation is detrimental to supply chain flexibility: https://www.freightwaves.com/news/shippings-extreme-consolidation-could-prolong-supply-chain-pain.

[34] https://www.spglobal.com/commodityinsights/en/market-insights/latest-news/shipping/061721-consolidation-key-to-tankers-meeting-climate-goals-international-seaways-ceo

[35] https://www.reuters.com/article/frontline-earning-ma/frontline-calls-for-consolidation-among-oil-tanker-firms-idUSL8N1NS52D

use of current resources in transportation, and supply chains more generally.

We view our paper 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 mechanisms, since we focused primarily on how size changes affects the incentives of shipowner. 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 production-inventory 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.

Özkan, 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.

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.[36]

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

[36] https://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 $t = 0, 1, \ldots,$ 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.[37] 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 $28 \times (T_{\max} + 1)$ 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 $s_{\cdot,\tau}$ 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

[37] 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 ($X_{l,o,d}$ for $l \neq o$), are in expectation of new loads in the future ($B_{l,d}$ 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

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

The central planner seeks minimizes $\sum_{t=0}^{\infty} \gamma^t r_t$, where γ is a discount factor.[38] 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:[39] 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.

[38] We set $\gamma = 0.85$, which corresponds to an effective time horizon of approximately one month.
[39] 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.[40] 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 $\mathcal{S} = \left\{s^i : i = 1, \ldots, N_A\right\}$. Each basis state s^i is defined as follows: if N is the total number of vessels in the environment, first $\left\lceil \frac{3}{4}N \right\rceil$ 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

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

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 $s^i \in \mathcal{S}$ we associate an initial value V^i_0, defined as the γ-discounted value over 147 time periods earned by a myopic central planner.[41] Finally, we obtain a family of affine functions by solving the following problem for each i, where $\hat{V}^i_0 \in \mathbb{R}$, $g^i_0 \in \mathbb{R}^{N_A}$, and $x \cdot y$ is the usual dot product between x and y.

$$\begin{aligned} \min_{\hat{V}^i_0, g^i_0} \quad & \sum_{i=1}^{N_A} \left(\hat{V}^i_0 - V^i_0\right)^2 \\ \text{s.t.} \quad & \hat{V}^j_0 \geq \hat{V}^i_0 + \sum_{\tau=1}^{T_{\max}+1} \gamma^{\tau-1} \left[g^i_0 \cdot \left(s^j_{\cdot,\tau} - s^i_{\cdot,\tau}\right)\right] \text{ for all } i, j \end{aligned} \tag{A.3}$$

This procedure yields a set $\mathcal{B}^S_0 = \left\{\left(s^i, \hat{V}^i_0, g^i_0\right) : i = 1 \ldots, N_A\right\}$. The set $\mathcal{B}^S_0$ represents basis functions because for each i we can write the affine function

$$g^i(s) = \hat{V}^i_0 + \sum_{\tau=1}^{T_{\max}+1} \gamma^{\tau-1} \left[g^i_0 \cdot \left(s_{\cdot,\tau} - s^i_{\cdot,\tau})\right)\right] .$$

[40] See, e.g., Boyd and Vandenberghe (2004), Chapter 3.

[41] That 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 Janurary 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:

$$\hat{V}_0(s) = \max_i \left\{ \hat{V}_0^i + \sum_{\tau=1}^{T_{\max}+1} \gamma^{\tau-1} \left[g_0^i \cdot \left(s_{\cdot,\tau} - s_{\cdot,\tau}^i \right) \right] \right\}. \tag{A.4}$$

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 $k = 1, 2 \dots$ 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,

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

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.

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

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,[42] it appears from Figure 6 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.

[42] 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 $t = 0, \dots, T$ 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

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

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 (%)	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.

Table 9: Effect of Fleet Size in Utilization

Dependent Variable	Log Utilization (Laden miles / Total miles)						
	All	Aframax	Panamax	Suezmax	VLCC	MR1	MR2
Filtering: Fleet Size >=0							
Log Fleet Size	0.0360***	0.0572***	0.0643**	-0.0047	0.0240**	0.0411***	0.0279**
	(0.006)	(0.013)	(0.027)	(0.018)	(0.011)	(0.013)	(0.011)
R-squared	0.1889	0.1807	0.1612	0.153	0.2145	0.211	0.1897
R-squared Adj.	0.149	0.1366	0.1002	0.1037	0.1642	0.1617	0.1423
N obs	24,132	5,090	2,288	3,367	3,359	4,190	5,838
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
R-squared Adj.	0.1531	0.14	0.1024	0.1129	0.1534	0.1687	0.1452
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
R-squared Adj.	0.1906	0.1503	0.1017	0.1315	0.146	0.2643	0.1869
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
R-squared Adj.	0.2276	0.1617	0.0992	0.1529	0.1468	0.3608	0.2332
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

Standard errors in parentheses.
*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 moth 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.

	Aggregate	MR1	MR2	Panamax	Aframax	Suezmax	VLCC
Δ *Utilization* (%)	8.0272***	8.8444***	7.4764***	6.4532***	12.5471***	7.2892***	4.4762***
	(0.2848)	(1.0697)	(0.5185)	(0.9737)	(0.5567)	(0.8336)	(0.5508)

Table 10: Average percent difference in the utilization rate between large pools and synthetic fleets that serve similar locations. Standard errors in parentheses.

Table 11: Effect of Fleet Size in Delays (Idle time) for the different legs: Ballast and Laden Legs

Dependent Variable	Ballast Delays							Laden Delays						
	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	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Operator	No	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes
Operator * Vessel Class	Yes	No	No	No	No	No	No	Yes	No	No	No	No	No	No

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

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.

Table 12: Effect of the fleet size of a vessel and the number of trips per month made on an OD pair (Complexity of the operation)

Dependent Variable	Log Number of Trips in a single OD pair						
	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

Standard errors in parentheses.
*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.

Table 13: Effect of the fleet size of a vessel and the number of different ports visited in a quarter (Complexity of the operation)

Dependent Variable	Log Number of Ports per Vessel in a Quarter						
	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	Yes	Yes	Yes	Yes	Yes	Yes
OD-pair	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Vessel Class	Yes	No	No	No	No	No	No

Standard errors in parentheses.
*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

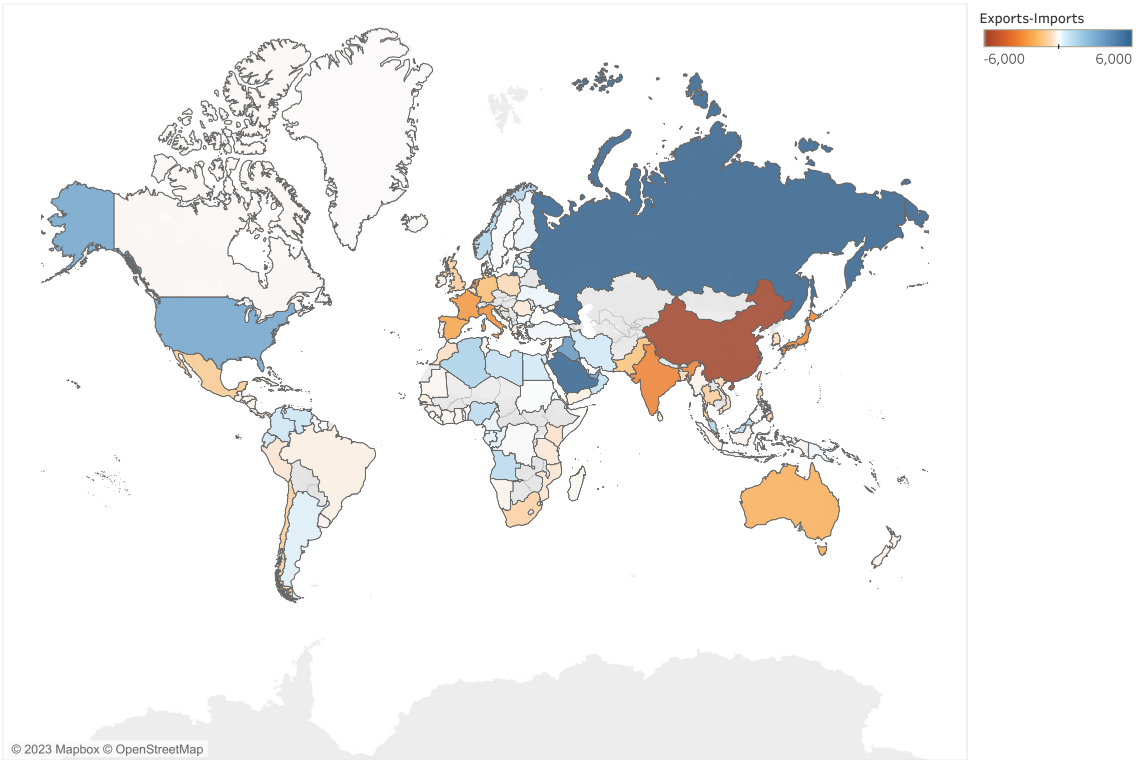

Map based on Longitude (generated) and Latitude (generated). Color shows sum of Export-Import. Details are shown for Country.

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.

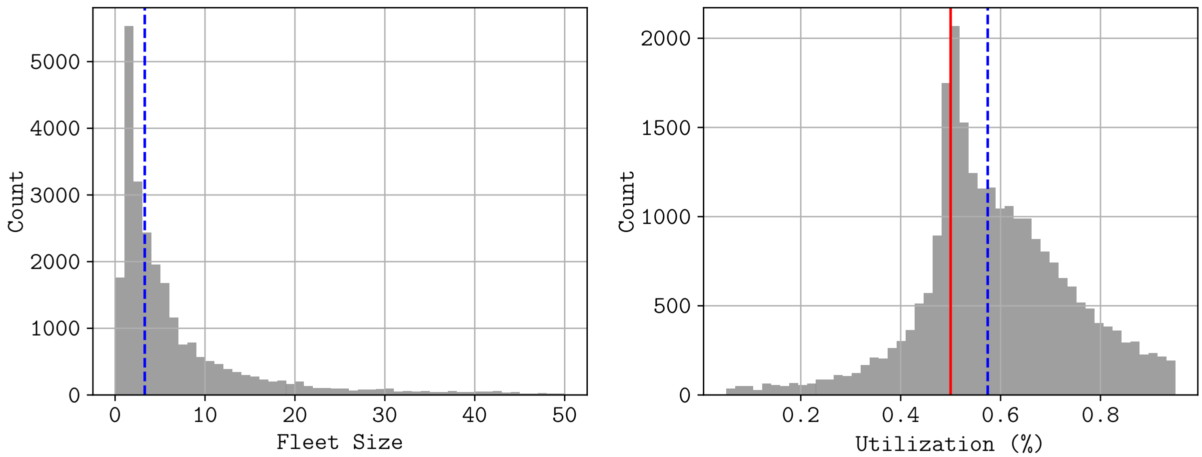

Figure 4: Left histogram shows the variation in Fleet Size at a quarterly level, i.e., for each quarter and Commercial Operator we measure their average fleet size and generate the histogram using these values. The plot shows the market structure emphasizing its fragmentation (most of the observations 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%.

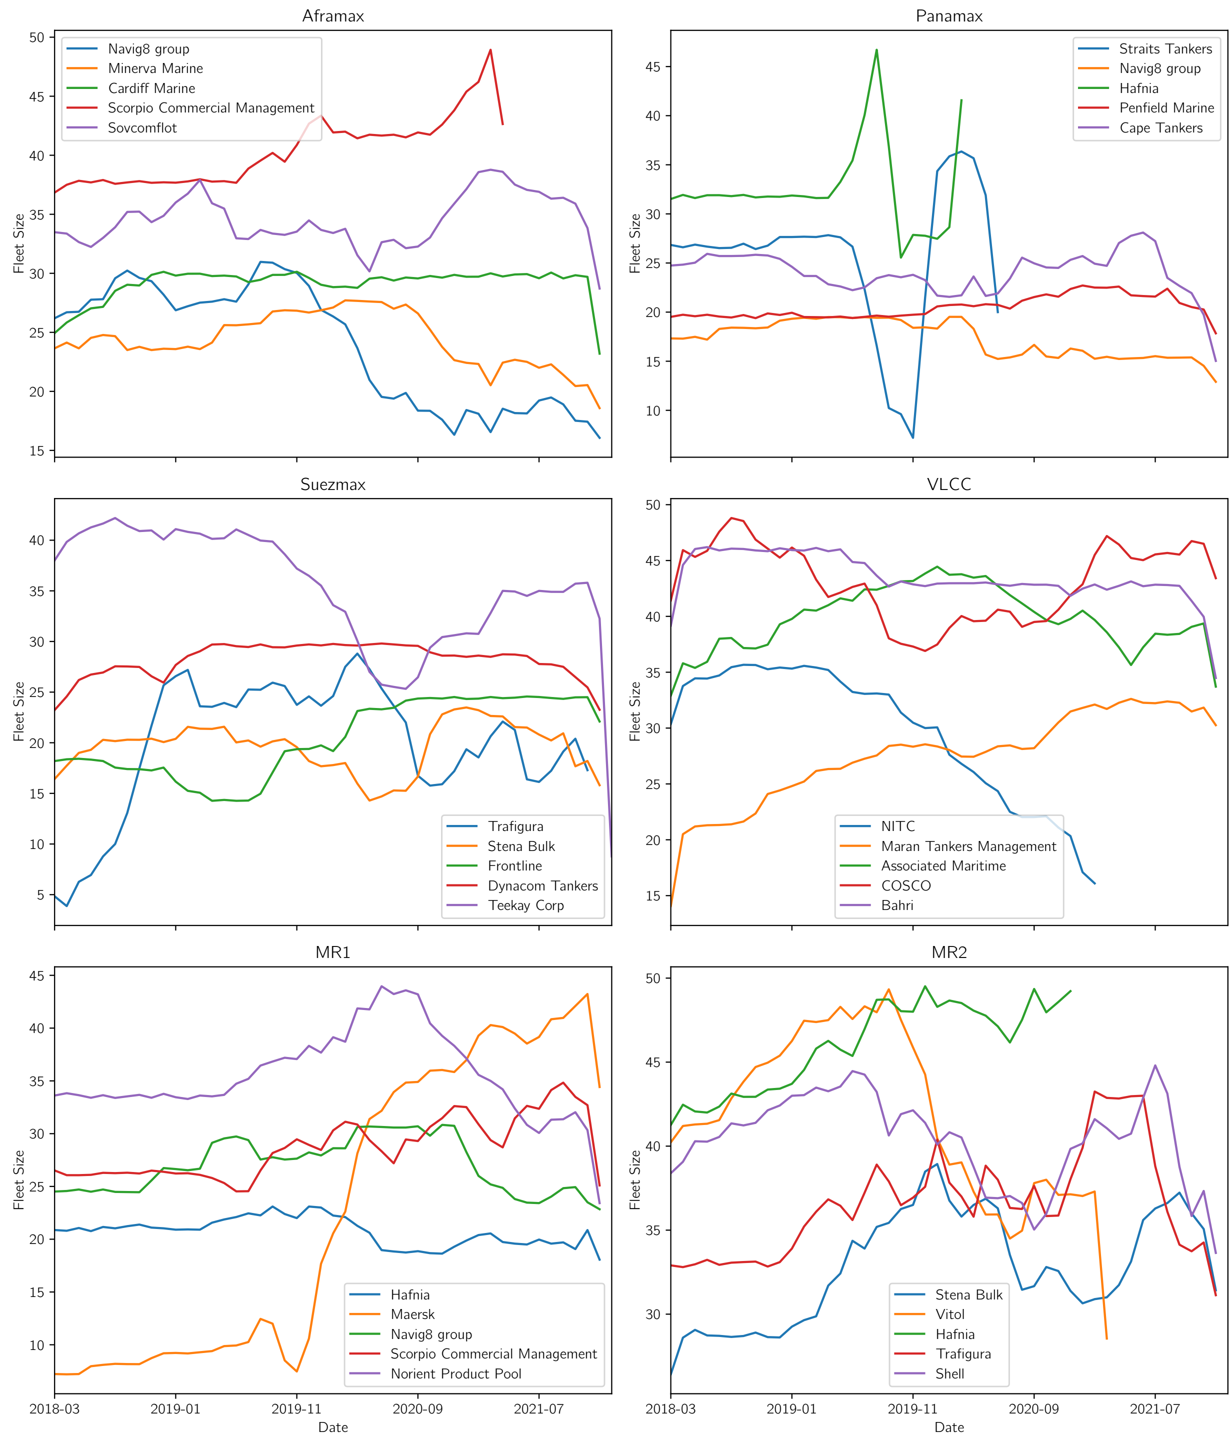

Figure 5: The figure shows the variation across time of the average fleet size of the five largest Commercial Operators in each sub-market at a monthly level. This becomes relevant for our empirical analysis where we exploit the differences across time of the fleet size of an operator.

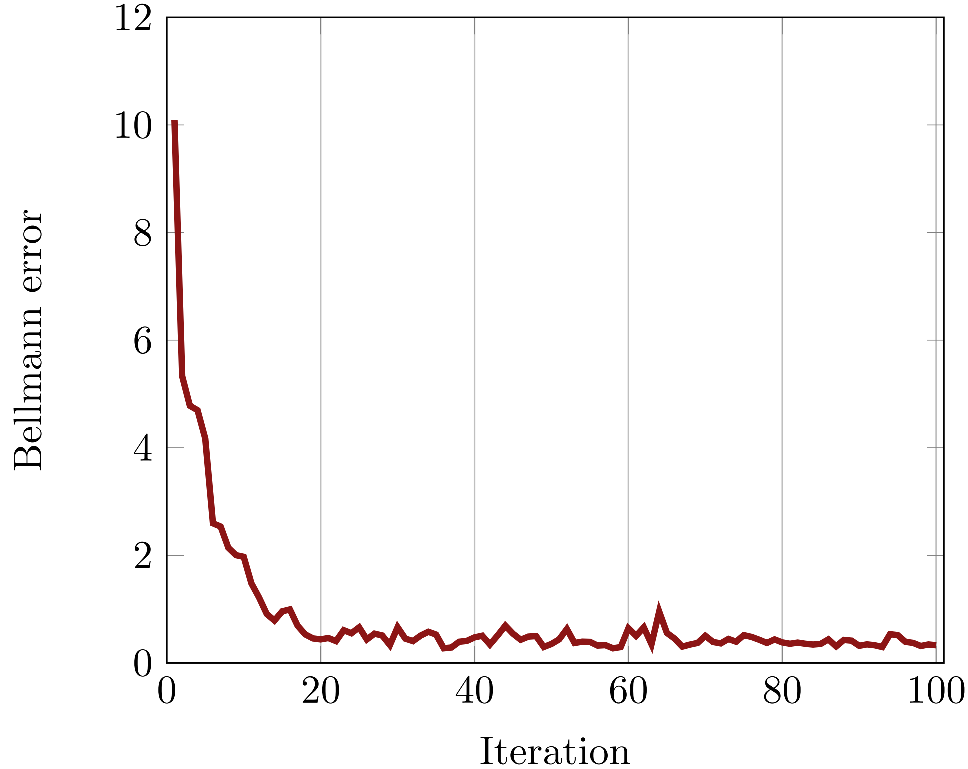

Figure 6: L_2 distance between each successive set of basis values $\hat{V}_k^i$ for the basis states of the Fitted Value Iteration procedure.