Market Fragmentation and Inefficiencies in Maritime Shipping

Kostas Bimpikis Giacomo Mantegazza Salomón Wollenstein-Betech*

Abstract

Maritime transportation accounts for 90% of global trade, but ballasting—vessels traveling without cargo—imposes substantial economic and environmental costs. This paper examines the oil transportation industry, where approximately half of all miles traveled are sailed empty. While some ballasting is necessary due to inherent supply-demand imbalances in oil markets, our analysis demonstrates that market structure, specifically the fragmentation of vessel ownership, is also a primary driver, accounting for 10-20% of the total empty miles traveled depending on the market segment. In addition, we show that consolidating vessels into small shipping pools—sets of vessels operated under unified management—can reduce ballasting-related carbon emissions by up to 15%. This market-driven approach, which is gaining industry adoption, maintains competitive dynamics, given the limited scale of consolidation, while significantly improving efficiency. The gains arise from enhanced coordination within larger pools and expanded port coverage, reducing unnecessary vessel repositioning. More broadly, our findings quantitatively demonstrate that organizational changes alone—specifically, the consolidation of vessel operations—can generate significant environmental improvements by reducing empty miles. This provides a practical path toward sustainability that can complement and amplify the benefits of technological innovation.

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

1 Introduction

Maritime transportation, while critical to global commerce, faces a significant economic and environmental challenge: vessels frequently travel empty, a phenomenon referred to as *ballasting*. With 11B tons of goods shipped in 2021 and volumes projected to triple by 2050, about 40-50% of all miles traveled are without cargo. This inefficiency is particularly acute in the dry bulk and oil transportation

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

markets, which employ specialized vessels restricted to single-cargo, point-to-point journeys, unlike container ships that can accommodate diverse cargo types and multiple stops.

While existing research has extensively examined how imbalanced demand and supply patterns necessitate ballasting, the role of market structure—particularly the interaction among competing service providers—remains understudied. This paper demonstrates that market fragmentation significantly impacts transportation efficiency through its effect on route selection and destination choices. By analyzing the oil shipping market, we show that fragmentation accounts for a sizable fraction of ballasting inefficiencies beyond those attributable to inherent trade imbalances.

Through a partnership with a leading oil tanker operator and market analytics provider, we analyze approximately six thousand oil tankers from 2018 to 2022, with comprehensive data on vessel movements, cargo status, and, crucially *Commercial Operators* controlling each vessel's deployment. The data reveals two critical patterns: first, significant regional trade imbalances exist, with East Asian ports importing more than three times their export volume (Figure 4), creating fundamental inefficiencies. Second, extreme market fragmentation prevails, with no single company controlling more than 5% of global load capacity, and the majority of companies managing less than 1% of global tankers (Figure 5). These patterns provide the foundation for our analysis of how market structure affects ballasting behavior.

To isolate the impact of market fragmentation on ballasting, we employ both optimization models and empirical analyses. Our optimization approach uses integer and stochastic dynamic programming models, taking into account variables such as vessel location, cargo capacity, and fuel consumption, to distinguish between three sources of inefficiency: trade imbalances, future demand uncertainty, and market fragmentation. While trade imbalances explain the largest share of ballasting (65-70%), our analysis shows that fragmentation accounts for 10-20% of inefficiencies depending on the market segment—twice the impact of uncertainty (5-8%). We quantify efficiency gains by measuring reduction in total empty miles traveled and associated fuel consumption costs.

Importantly, our analysis shows that consolidating vessels into modest-sized shipping pools (approximately 40 vessels, or 5% of the global fleet) can achieve 80% of the efficiency gains possible under theoretical centralized control. Our empirical examination of existing shipping pools—an emerging industry trend where multiple shipowners consolidate operations under single management—reveals two primary mechanisms driving these efficiency gains. First, unified management enables better vessel coordination within service regions, reducing redundant empty journeys. Second, larger pools develop more complex port networks and route diversification, leading to more integrated operations compared to smaller operators' route-specific focus.

These findings have significant implications both for industry practice and for environmental policy. They suggest that moderate market consolidation through shipping pools offers a practical

path to reducing empty miles while maintaining market competition. This approach could substantially reduce fuel consumption (and therefore raise industry profits) and emissions without requiring technological breakthroughs or extensive regulatory intervention.

1.1 Related Literature

Our research contributes to the growing literature on analyzing and optimizing transportation systems sparked by the emergence of ride-sharing platforms. This body of research has focused on the design of pricing and relocation policies in response to long- and short-term demand patterns, and it has mainly focused on a single decision maker that centrally manages pricing and fleet allocation across the entire market.

Regarding pricing optimization, some early work, e.g., Cachon et al. (2017), demonstrates that surge pricing serves as an effective mechanism for rebalancing supply to respond to spikes in demand. This is further reinforced by Besbes et al. (2021) and Ma et al. (2022), who derive optimal pricing policies that stress the importance of considering strategic behavior on both the demand and the supply side. Finally, Bimpikis et al. (2019) focuses on the long-term demand imbalances to derive optimal pricing policies that encourage repositioning of drivers in the network.

In addition to pricing mechanisms, relocation (or repositioning) policies represent another critical mechanism for balancing transportation systems, and a number of papers have investigated this direction. For example, Braverman et al. (2019) and Özkan and Ward (2020) consider a central decision maker that can dispatch supply to different locations to better serve demand, and derive static policies based on stochastic approximations. Our work complements these lines of research, as we examine a fragmented market with multiple independent operators, and we focus on the relationship between market structure and (in)efficiency as captured by empty miles. While we consider optimization problems with a central planner, we use their optimal values as benchmarks to estimate the impact of market fragmentation.

Recent research has focused on the matching quality between demand and supply of transportation services in environments with spatial and temporal imbalances. Azagirre et al. (2024) shows that profitability in ride-sharing systems can be substantially improved when matching algorithms takes into account the drivers' dynamic incentives, while Castro et al. (2023) studies how matching policies in the presence of autonomous vehicles need to account for the strategic response of human drivers to avoid inefficient service levels in equilibrium. While the optimization problems we solve can be recast as optimal dynamic matching problems (which can be particularly challenging to solve, especially when taking into account the agents' private information about the other side of the market, e.g., Manshadi et al. (2024)), we aim not to develop approximately optimal policies for shipping markets, but rather to quantify the level of inefficiencies due to their decentralized nature.

Although the transportation literature has extensively examined demand-related inefficiencies, our work addresses a gap by focusing on structural market inefficiencies that persist regardless of demand patterns. For example, Arlotto et al. (2019) shows that rational agents competing in an open-route queuing network demonstrate herding behavior in equilibrium, similar to what we observe in some ports in our dataset; and Acemoglu et al. (2018) demonstrates that congestion levels can actually worsen when agents receive information about routes they did not know. Recent work has also taken the perspective of individual shipowners: for example, Prochazka et al. (2019) and Adland and Prochazka (2021) use an approach similar to ours in Section 3 to estimate the value of foresight and contractual flexibility in the context of dry-bulk shipping. The key difference between our work and those papers is that they take the perspective of a single operator, and base their analysis on simulated data (calibrated on aggregate market data) instead of actual commercial decisions. On the other hand, Séjourné et al. (2018) adopts a market-wide view and studies theoretically the inefficiencies generated by market fragmentation: they argue that fragmentation is detrimental because small fleets struggle at efficiently repositioning their assets in the network when demand is randomly split among them. We depart from their analysis along several dimensions. Our research makes a distinctive data-driven contribution by leveraging granular commercial data to quantify the efficiency losses attributable to market fragmentation—moving beyond theoretical predictions to establish the actual magnitude of these effects in real-world maritime operations. In addition, we are able to identify a new driver of ballasting, i.e., the difference in the network of locations served by large versus small operators. Finally, we study how increasing market concentration may improve efficiency, while Séjourné et al. (2018) takes the level of fragmentation as given.

Our study also contributes to the expanding literature on decentralized transportation markets. For example, Frechette et al. (2019) and Buchholz (2022) study the taxi market and show how different barriers to entry shape the market structure and welfare outcomes, while Harris and Nguyen (2022) analyzes how long-term relationships between truckers and brokers affect load assignments in the trucking industry. Particularly relevant for our setting are Brancaccio et al. (2020) and Brancaccio et al. (2023), who exploit voyage-level data in the dry-bulk industry to study the impact of search inefficiencies 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 yields inefficient relocation decisions. We complement this strand of literature by examining how different market structures influence resource relocation decisions, independent of how prices are determined. In particular, since empty miles arise also under efficient pricing, market fragmentation has been largely overlooked by this literature: we make a significant contribution by empirically

demonstrating that fragmentation has a first-order effect on market efficiency and providing the first quantitative measurement of its relative importance in maritime shipping operations.

Our examination of shipping pools as a market consolidation mechanism builds upon the rich literature on resource pooling in Operations Management. This idea has been applied across diverse contexts, ranging from inventory management (Eppen (1979), Bimpikis and Markakis (2016)), to manufacturing flexibility (Jordan and Graves (1995), Netessine et al. (2002), Liu et al. (2016)), and the study of alliances in other transportation markets (Wright et al. (2010)). More recently, the value of sharing resources has also been investigated in emerging domains such as Robots-as-a-Service, where Jacquillat et al. (2024) finds that to reap the full benefits of pooling it is often necessary to combine two complementary dimensions (e.g, sharing robots across users and sharing workload across robots). We substantially extend this literature by uncovering a previously unidentified mechanism through which resource pooling drives performance improvements: beyond merely enabling greater flexibility and coordination, pooling fundamentally restructures incentives for participating decision makers, leading them to allocate their resources in a systematically different way. More importantly, this is among the first papers to *quantify in a real-world system* the extent to which pooling can benefit efficiency, does so at a *market-wide* level (as opposed to studying single decision-makers), and identifies in an applied setting what is the least amount of pooling that captures most of the gains from consolidation.

Additionally, our research advances the emerging literature concerned with improving the sustainability of firms' operations— see, e.g., the review article by Lee and Tang (2018), and applications that range from developing supply chain-aware ESG criteria in Dai et al. (2024) to studying decentralized forms of green certifications in Feldman et al. (2025). Transportation systems, both cargo and passenger, are among the largest generators of pollutants. Granular solutions have been proposed to curb pollution from cars (e.g., Arora et al. (2024)) but only higher-level policies, such as Emission Control Areas (Zhuge et al. (2024)), have been considered for maritime shipping. Our analysis reveals that policies aimed at encouraging some degree of market consolidation can simultaneously deliver substantial economic benefits and meaningful environmental improvements—a critical finding for an industry that has traditionally faced a perceived trade-off between profitability and sustainability. This insight becomes particularly relevant in the context of ongoing discussions on new regulations, 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, at approximately 11% of total volume; and other specialized categories, such as chemicals, for the remaining portion.[1] The oil transportation sector alone accounts for 30% of the seaborne trade volume, making it one of the most relevant markets in maritime transportation. While there exists extensive land infrastructure to move crude oil from extraction sites to refineries, it is estimated that 61% of the daily production volume of 90 million barrels of oil relies on maritime transport.[2,3] A fleet of about 8,800 ocean-going oil tankers operates in this industry, divided into six independent segments based on the ships' displacement (i.e., cargo capacity), called *vessel classes* (see Table 1). The smallest vessels have a capacity of about 300,000 barrels of oil, while the largest can accommodate nearly ten times that volume. Ships within each class are relatively homogeneous in size, so that they can access the same ports and use the same artificial waterways.

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	N/A

Table 1: Number of ships by vessel class in the dataset and their dimension. DWT refers to the *deadweight tonnage capacity*.

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

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

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

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

transported with spot contracts.

The spot market for maritime shipping operates through a network of brokers who facilitate transactions between oil companies (exporters) and carriers. A typical transaction begins when an exporter needs transportation and contacts a broker, providing essential information about the cargo: type, loading and unloading locations, and the desired departure schedule. The broker then presents these cargo requirements to select vessel operators whose ships are positioned near the departure area, and works to finalize an agreement between the parties. The resulting contract (known as a *fixture*) specifies the exact location and day of loading, along with the daily vessel rental rate. This rate is paid only for the duration of the trip from the loading port to the final destination. Due to the inherently unpredictable nature of sea voyages, this process typically begins shortly before the intended loading time. Consequently, only vessels located close to the loading area can be considered available to service a given cargo. For this reason, operators' decisions about where to position their vessels are critical to securing profitable contracts, as strategic positioning directly impacts their ability to compete for available cargoes in a dynamic market.

Consider a ship operator's perspective right after discharge: it is unlikely that the vessel has already been contracted for a new load.[4] The operator must decide whether to wait for new offers in the same port (or neighboring anchorages), or to relocate somewhere else. It is during this phase that the shipowner typically receives proposals from brokers for new cargoes whose load area is a few days away from the current position of the vessel. If the offer is fixed, a new *voyage* begins, which ends at the time the new discharge port is reached. Thus, each voyage can be divided into an initial *ballast* leg, where the tanker is empty, and a final *laden* leg, where the tanker is full.

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

2.1 Data Description

We partnered with a shipping and analytics company, which provided us with a list of 234,795 voyages completed from January 1st, 2018 to March 26th, 2022 by 5,886 vessels; Table 1 offers a breakdown by vessel class. For each voyage we have the identifier of the vessel, as well as its load and discharge ports with the respective arrival and departure dates; we are also provided with a *starting port* that corresponds to the discharge location of the previously finished voyage. The

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

ballast leg of the voyage is the trip from the starting port to the load port, and the laden leg goes from the load port to the discharge port. Thus, we know exactly how long each leg was for each voyage, and how much time each ship spent docked at a port, ballasting, or traveling with a load.

Geographical information is aggregated at various levels, from port level (the finest partition) to *wide areas* (the coarsest). We carry out our analysis at different levels of granularity, since the most appropriate one depends on the specific application. By and large, we employ port level, *narrow area*, and *intermediate area* information: narrow areas partition the world into 51 regions, while intermediate areas into 28. A list of the intermediate areas can be found in Table 9.

In addition to geographical data, we have information about fuel consumption, pollutants emissions, speed, and the name of the *Commercial Operator* that managed the vessel for that voyage. The commercial operator is the agent that was in charge of a vessel's choice of which loads to carry as well as whether and where to ballast. The commercial operator may not correspond to the actual shipowner, since vessels with different owners may be part of the same shipping pools, operated by a single entity.[5] We do not have access to the contracted rates; we discuss how this limits our analysis in Section 3.4.

Our dataset includes only completed voyages—we do not have information on potential demand that was not fulfilled. However, certain features of this market imply that potential demand closely aligns with realized loads: first, there is no outside option to sea transport once oil has arrived to the tankers terminal;[6] second, since contracts are fixed at a short notice, oil exporters have their load ready for departure at the time they contract with brokers; third, significantly delaying the time of transportation incurs high costs for storing oil in the terminal. Considering these factors, 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, we do not consider voyages shorter than 5 days (1.8% of observations) or voyages for which we do not have information about the commercial operator (15% of voyages). We restrict the analysis to only those vessels for which we have a complete history, i.e., whose position can be accounted for from the first time they are assigned a load to the last time they discharge. Note that this does not rule out vessels that undergo maintenance or become unavailable for other reasons, which we explicitly consider in Section 3, as long as we can make sure that we did not miss any loads they fulfilled. This leaves 4,599 tankers and 146,745 voyages (approximately 63% of the dataset); we now summarize some of the outstanding features of the oil shipping market from our dataset.

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

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

Trade imbalances. Oil reserves are unevenly distributed in the world, which leads to substantial imbalances in the oil transportation patterns and, thus, are the main cause of ballasting. These imbalances are reflected in our dataset, as shown in Table 9 and Figure 4. 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. Given these patterns, it is unlikely for vessels to always find a load to transport in the same location they discharge at.

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

Market fragmentation and pools. Our dataset is noteworthy because we have access to information about the commercial operators of the tankers. We can trace a detailed picture of market concentration, its evolution over time, and of how the number of vessels under management correlates with the decisions of commercial operators. The control structure is extremely fragmented: as evidenced in the left panel of Figure 5, a substantial number of operators have fleet sizes ranging from 1 to 3 vessels; in contrast, very few control more than 10 vessels. The fleet size of the largest commercial operators exhibits notable variation across time, as shown in Figure 6, while smaller operators tend to have a more stable fleet. This difference can be attributed to the fact that large commercial operators manage vessels via *shipping pools*, i.e., they add to their own fleet vessels owned by others (usually, small shipowners themselves). A shipping pool is formed when a number of shipowners decide to "merge" their fleets under a unique manager. The pool manager assumes responsibility for all vessels and maximizes the collective pool revenues, while the respective owners remain entitled to a share of the earnings the pool generates.

Commercial operators can use shipping pools as a more flexible tool for expanding their fleets compared to formal mergers or acquisitions.[7] Pools have become increasingly popular thanks to the advancement in computational capabilities that allow real-time optimization. Despite their

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

advantages, shipping pools' market share remains low: Table 10 shows that for the Aframax vessel class the largest operator controls less than 4% of the total number of vessels, with the top 20 operators collectively controlling less than 40% of the asset share.[8] Similar patterns prevail across other vessel classes, with a marginal uptick in concentration for larger tankers. Importantly, shipping pools exhibit substantial heterogeneity not only in terms of size, but also in the level of efficiency they achieve: the right panel of Figure 5 illustrates the distribution of quarterly utilization rates (i.e., the ratio of laden miles to total traveled miles). This could be attributed to many factors, and in this paper we concentrate on its relationship with pool size.

3 Determinants of Ballasting

While ballasting is a well-known phenomenon within the transportation sector, its causes remain somewhat under-explored, and their relative magnitude is largely unknown. In the case of oil tankers, the substantial role played by trade imbalances is quite intuitive, but the extent to which observed ballasting can be attributed solely to disparities in supply and demand is challenging to quantify. Answering this question is the necessary first step to understand whether the market operates efficiently or there exists untapped potential for improvement. To this end, we classify the sources of ballasting inefficiencies in three categories: (i) trade imbalances, (ii) uncertainty with respect to the time and locations of future loads;[9] and (iii) market fragmentation.[10]

We first quantify the share of empty miles that can be attributed to trade imbalances: it is a baseline level of ballasting intrinsic to oil trade patterns. Our strategy is based on the following intuition: if we can address the other two factors influencing ballasting, the intrinsic *baseline ballasting* would be given 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 asset control and information flows. Using a similar reasoning, we assess what would happen if several 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 relative share of ballasting due to market fragmentation and uncertainty as well.

[8] 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.

[9] E.g, tankers may decide to ballast away from a port even if a load would have been offered shortly after.

[10] Market fragmentation influences ballasting decisions in several significant ways. First, commercial operators often have access to varying information due to their diverse broker networks, creating information asymmetries in the market. Second, operators managing larger fleets typically have greater flexibility to position their vessels strategically, allowing them to be more opportunistic in pursuing profitable contracts. In our analysis, rather than attempting to isolate these individual factors, we collectively categorize them under the term *market fragmentation*, which serves as their common underlying cause. This approach acknowledges that these various effects stem from the same fundamental market structure.

Assumptions. To perform this analysis, we use data from all four years of voyages in our dataset to estimate costs and travel times between geographical areas. Given that the COVID-19 pandemic disrupted oil trade flows, we optimize only over the two year period preceding the onset of the pandemic, i.e., 2018-2020.[11] We discretize time into five-day periods, so that, for example, voyages departing from January 1st, 2018 to January 5th, 2018 are imputed to period $t = 0$. Also, we use the intermediate geographical areas described in Table 9. We impose that, if vessels reach a location at time t, they are also able to load cargoes at time t.[12] Finally, we assume that (i) the laden leg of each voyage has the same duration irrespective of the ship assigned to it; (ii) loads are fungible within a vessel class, so that each can be transported by any available vessel in the class, but not across classes.

3.1 Trade Imbalances

Demand for oil transportation is influenced by a number of factors, including geopolitics and macroeconomic trends. Hence, instead of fitting a necessarily imperfect model to observed demand, we accept demand as an exogenous input and impose the constraint that the available transportation assets have to deliver such demand. The share of ballasting exclusively attributable to demand imbalances can then be estimated as the ballasting cost that would obtain if transportation was organized by a benevolent and clairvoyant entity that satisfies all demand present in the data, relative to the total cost of empty travels measured within the dataset. Since the different vessel classes define virtually independent markets, we optimize separately within each class.

Formally, we cast each problem as an integer linear program whose objective is to minimize the total ballasting cost, defined in terms of CO_2 emissions. Since we impose the constraint that all observed voyages have to be fulfilled, minimizing ballasting costs is equivalent to finding the minimum cost assignment that satisfies all demand. We choose CO_2 emissions as the metric to optimize for instead of traveled distance since emissions better capture costs as they are directly linked to fuel consumption. We denote with $C^v_{l,o}$ the average CO_2 emissions for a ship in vessel class v of a trip between locations l and o belonging to the set of locations $\mathscr{L}$ in the network.[13] Also, we let T be the number of 5-day periods from January 1st, 2018 to January 1st, 2020.

We introduce two types of decision variables: $X_{l,o,d,t}$ that accounts 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

[11] This also ensures consistency with the analysis of Appendix A, which is based on reinforcement learning using synthetic data calibrated to the pre-pandemic trends in the market.

[12] This is equivalent to imposing that load and discharge times are less than five days. Figures 7 and 8 illustrate that this is the case for most voyages.

[13] Emissions estimates are based on fuel consumption data that vessels are obligated to log to comply with the International Maritime Organization (IMO) regulations. The mapping from fuel consumption to emissions is an industry standard and takes into account which type of fuel was used and for how long.

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 when a vessel departing location l at time τ reaches location o at t (clearly $\chi_{l,\tau,o,t} = 0$ if $\tau > t$). It is obtained by estimating the average ballast time between the two areas.
$I_{l,t}$	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.
$C^v_{l,o}$	Average CO_2 emissions in tons of a ballast travel between areas l and o undertaken by a ship belonging to vessel class v.

Table 2: List of parameters estimated from the data and used in the optimization problems of Sections 3.1 and 3.2

$o \in \mathscr{L}$ and dropped at destination $d \in \mathscr{L}$, and $Y_{l,t}$ that denotes the number of vessels in $l \in \mathscr{L}$ that remain unassigned in l at time t.

Finally, the formulation uses the following parameters obtained from the data: $D_{o,d,t}$ represents the number of loads departing from location o towards location d at time t; $A_{l,t}$ represents the number of vessels that discharge at location l at time t;[14] $I_{l,t}$ represents the number of vessels that first enter the data in location l at time t, and $O_{l,t}$ represents those vessels that leave the dataset after discharging in location l at time t. Finally, $\chi_{l,\tau,o,t}$ indicates whether a vessel leaving location l at time τ arrives in location o by time t. As such, this encodes the average duration of a ballasting leg between the two locations. Further details about these parameters can be found in Table 2.

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

$$\min \quad \sum_{l,o,d,t} C^v_{l,o} X_{l,o,d,t} \tag{1a}$$

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

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

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

$$X_{l,o,d,t},\ Y_{l,t} \in \mathbb{N}. \tag{1e}$$

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

Equation (1a) accounts for the total ballasting costs associated with the assignments in X. The constraints in Equation (1b) ensure that all loads observed in the data are fulfilled under the optimal assignment solution, and that this solution is feasible in terms of travel times. Constraints in Equations (1c) and (1d) make sure that each vessel is used at most for one voyage at a time, and that the flow of ships in each location is conserved. On the right hand side of each constraint we have the total supply of vessels in location l prior to the decision at time t, composed by (i) the ships that remained in l from the previous period; (ii) the inflow of new vessels to the network; and (iii) the arrival of tankers from voyages ending at location l. On the left hand side of Equation (1c) and Equation (1d), we model how supply is used, i.e., for assignment to loads ($\sum X_{l,o,d,t}$) and waiting ($Y_{l,t}$), or outflows ($O_{l,t}$). Finally, without loss of generality our formulation does not allow for relocation of empty vessels ahead of their assignment to loads.[15]

Let C_v^* be the optimal value of Equation (1a) for vessel class v and C_v^{obs} be the ballasting emissions measured for the same vessel class within the dataset. We define the share of ballasting due to trade imbalances as the ratio between these two, i.e., $Share_v^{Trade} = C_v^*/C_v^{obs}$. Table 3 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 ones. This trend is intuitive, since larger vessels serve a limited set of routes: there exist fewer ports able to accommodate these ships, which implies that there is limited scope for optimizing the network of locations served. For this reason, we expect those markets to operate closer to efficiency. Hence, factors beyond trade imbalances are relatively more important for smaller vessel classes in contrast to larger ones. Overall, these results confirm that trade imbalances are the primary factor in determining the amount of ballasting we observe; however, they also highlight that there exists potential for significant efficiency gains at least for smaller vessel classes.

3.2 Fragmentation and Uncertainty

To isolate the effects of market fragmentation from uncertainty, we propose a comparative framework with three scenarios. At one extreme, we have the ideal case of a clairvoyant central planner who can optimize across all vessels with perfect foresight, representing the theoretical minimum ballasting cost C_v^*, obtained in the previous subsection. At the other extreme lies our observed data, reflecting the full impact of both fragmentation and uncertainty. Between these, we introduce an intermediate benchmark: a system of decentralized planners—one for each commercial operator—each with perfect information about their own fleet's future loads but acting independently. This intermediate case incorporates the inefficiencies of market fragmentation while removing the element of uncertainty.

[15] Specifically, since there is no randomness in the problem as we set it up, for every optimal solution that would have repositioned a tanker beforehand, there exists another one that has the vessel ballasting to the origin of the same load just before the allocation.

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

By comparing these three scenarios, we can quantitatively attribute inefficiencies to either market structure (fragmentation) or information limitations.

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

$$\min \quad \sum_{l,o,d,t} C^v_{l,o} X^i_{l,o,d,t} \tag{2a}$$

$$\text{s.t.} \quad \sum_{l,\tau \le t} X^i_{l,o,d,\tau} \chi_{l,\tau,o,t} = D^i_{o,d,t}, \qquad \forall (o,d,t) \in \mathscr{L} \times \mathscr{L} \times \{1,\dots,T\} \tag{2b}$$

$$\sum_{o,d} X^i_{l,o,d,1} + Y^i_{l,1} = I_{l,1}, \qquad \forall l \in \mathscr{L}, \tag{2c}$$

$$\sum_{o,d} X^i_{l,o,d,t} + Y^i_{l,t} + O^i_{l,t} = A^i_{l,t} + I^i_{l,t} + Y^i_{l,t-1}, \qquad \forall l \in \mathscr{L},\ \forall t \tag{2d}$$

$$X^i_{l,o,d,t},\ Y^i_{l,t} \in \mathbb{N} \tag{2e}$$

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

$$Share^{Trade}_v = \frac{C^*_v}{C^{obs}_v} \le \frac{C^P_v}{C^{obs}_v} \le 1.$$

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

$$Share^{Frag}_v = \frac{C^P_v - C^*_v}{C^{obs}_v} \tag{3}$$

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 4: Share of ballasting attributed to uncertainty and fragmentation

$$Share_v^{Uncertainty} = 1 - \frac{C_v^P}{C_v^{obs}}. \tag{4}$$

In the numerator of Equation (3) we compare the optimal cost of the assignments of a central planner with no uncertainty against the cost of multiple fleets without uncertainty, i.e., the only difference is the level at which assignments are made. Therefore, this difference can only be capturing the increase in cost due to fragmentation. Consider now Equation (4): the only difference between C_v^P and C_v^{obs} is that the former is calculated when vessel managers have perfect knowledge of the future, so that it can be thought of as capturing the effect of eliminating uncertainty (equivalently, the extent to which uncertainty affects ballasting). We report the results of these metrics for each vessel class in Table 4.

Note that, with the exception of VLCCs, the weight of uncertainty is remarkably similar across vessel classes. This is consistent with our conjecture that the VLCCs follow a more predictable schedule due to the limitations imposed by their size. The other insight emerging from these figures is that fragmentation always accounts for a larger share of ballasting than uncertainty. The results indicate that fragmentation in the market is a major cause of waste, and that taking action to consolidate the control structure can yield sizable benefits, both to shipowners and to the environment. However, consolidating all tankers under a single entity is neither feasible nor desirable. The next section is devoted to understanding "how much" consolidation is sufficient to reap a substantial amount of the potential gains.

3.3 Partial Consolidation

Our data contains 169 Aframax commercial operators, and in Section 3.1 we considered a single entity controlling all vessels to evaluate the inefficiency of the current transportation market. In order to measure the value of *partial* consolidation in the market, we create "synthetic" shipping pools. This allows us to test different degrees of consolidation by choosing the number of vessels belonging to each pool. Moreover, it serves as a test bed to assess the trade-off between the efficiency gains of consolidation and market power.

To carry out this analysis we first define the efficiency gain due to centralization for vessel class v by Δ_v^* and given by $\Delta_v^* = \left(C_v^* - C_v^{obs}\right)/C_v^{obs}$. Then, we randomly split the tankers into shipping pools of identical size ϕ, where ϕ denotes the fraction of the total number of vessels that each pool manages. For each pool p, we observe the voyages performed by the vessels in p and solve an optimal assignment problem using Equations (2a) to (2e), which yields an optimal cost $C_{p,v}^*(\phi)$. Finally, we obtain the market-wide ballasting cost $C_v^*(\phi) = \sum_p C_{p,v}^*(\phi)$ and the gain from a ϕ-consolidation, denoted by $\Delta_v^*(\phi)$.[16] Our metric of interest is the ratio $\rho_v(\phi) = \Delta_v^*(\phi)/\Delta_v^*$, which reports the fraction of the central planner's fragmentation savings that can be achieved when the market is consolidate into $\frac{1}{\phi}$ pools each controlling a share ϕ of vessels.

The left panel of Figure 1 plots $\rho_v(\phi)$ for the different vessel classes and for different pool sizes: it shows how much of the inefficiencies due to fragmentation can be eliminated with varying levels of consolidation. The right panel illustrates the decrease in the average ballast time as we vary the level of consolidation in the market. These figures establish that even with small pools, each being approximately 5% of the market, it is possible to obtain ballasting savings that are close to the first-best solution induced by a central planner. This is particularly evident for smaller vessel classes such as Aframax or MR2: splitting all 792 Aframax ships into pools of 40 units each (about 5% of the fleet) yields between 80% and 90% of the savings that the central planner could have obtained. Using the weight of fragmentation estimated in Table 4, this corresponds to an overall decrease in ballasting emissions of about 15%, or 4.6 million tonnes of CO_2 not released in the environment (equivalent to one million fewer passenger cars every year). These results suggest that the benefit of centralized vessel-load assignments accrues with even modest consolidation that would not raise concerns about the level of competition in the market. Moreover, the shape of the plots indicates that there are decreasing marginal returns to consolidation. In Section 4 we investigate in greater detail the mechanisms by which shipping pools reduce ballasting costs.

3.4 Model Discussion

Our analysis so far suggests that the oil transportation market does *not* operate at an efficient level and that it stands to substantially benefit from even modest consolidation. Moreover, despite being based on data from a specific market, our analysis can be applied with little modification to other decentralized transportation markets, such as dry bulk shipping and full-truckload trucking, that are likely to feature similar inefficiencies as the ones we identify here. At the same time, it is important to recognize some limitations of our data and approach.

We stress that our estimates should not be interpreted as market counterfactuals (in the sense of comparing ballasting costs under different market equilibria) for two reasons: (i) we do not consider

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

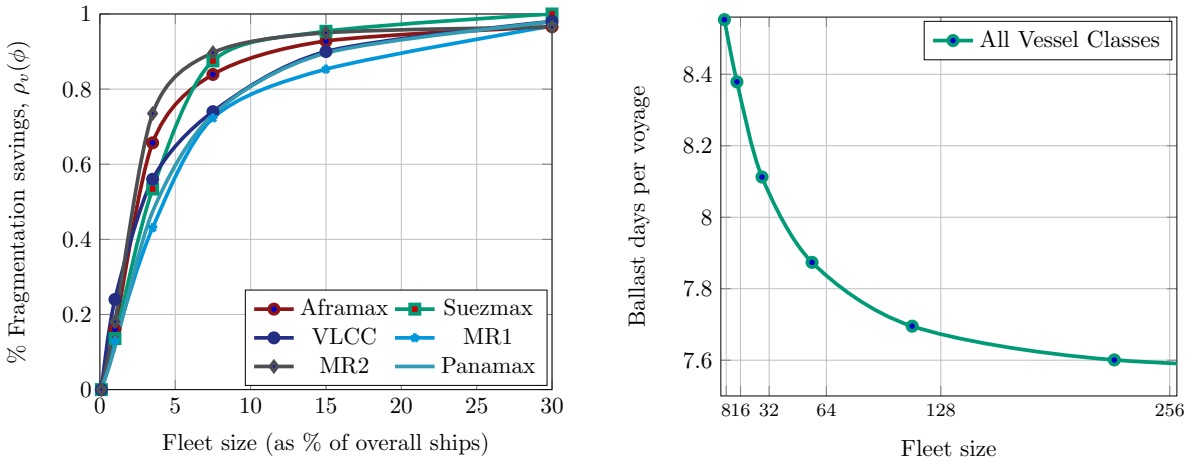

Figure 1: Plots of the effects of partial consolidation. On the left, the share of fragmentation inefficiencies that can be recouped by pooling vessels in fleets of different fleet sizes. On the right, the average length of the ballast leg (in days) of a voyage as the size of synthetic fleets increases (presented on a log axis); we show a weighted (by number of voyages) average of the curves by vessel classes for readability.

the equilibrium response of demand for transportation to consolidation, and consolidation's effect on the bargaining power of pool managers; (ii) we set up cost-minimizing problems, while in reality decision makers would maximize profits. Estimating a structural model to carry out counterfactual simulations is out of the scope of this paper given the lack of data on contracted rates.[17] Note, however, that the primary focus of this section is to *measure* the extent to which fragmentation impacts managerial efficiency. If we ran a counterfactual simulation and compared the outcomes, the difference in ballasting emissions would also capture efficiency gains from changes in served voyages (as we expect fulfilled demand to be lower when concentration increases) besides improved coordination. Instead, our approach keeps demand fixed. While the difference we calculate cannot be understood as "how outcomes in the real world change as concentration increases", we measure precisely the inefficiency due to the lack of coordination, and how this improves with pooling.

As far as our choice of minimizing CO_2 emissions associated with ballasting is concerned, we note that emissions are based on fuel consumption data that vessels are obligated to log to comply with the International Maritime Organization (IMO) regulations. The mapping from fuel consumption to emissions is an industry standard and takes into account which type of fuel was used and for

[17] Moreover, any structural model must take into account that fleets dynamically optimize a set of discrete resources: this is a dynamic programming problem for which only approximate solutions exist in the monopolistic setting (e.g., Godfrey and Powell (2002a)), while the oligopoly model is intractable, and there exists no algorithm guaranteed to find Markov Perfect/Oblivious equilibria. For these reasons, the literature that studies decentralized transportation markets usually assumes a collection of independent, non-atomic, agents (e.g., Brancaccio et al. (2020)).

how long (refer to Table 8, where we show that results are robust to using fuel consumption as the main metric of interest). Since vessels are crewed even when docked at port, fuel constitutes by far the largest component of variable costs associated with each voyage. Thus, minimizing ballasting emissions, taking fulfilled demand as given, is closely related to minimizing costs. Finally, note that if a shipowner has little pricing power, as would appear natural in the case of Section 3.3 due to the level of market fragmentation, then cost minimization is a good proxy for profit maximization. Thus, we conclude that our estimates point in the right direction, and we believe it is unlikely that both the relative share of ballasting due to one factor compared to the others and the insights from Section 3.3 would be substantially different in a structural model with profit maximizing agents.

Robustness check. Finally, we highlight that our treatment of uncertainty is "residual," in the sense that we have analyzed optimization problems in hindsight and obtained estimates for the share of ballasting due to uncertainty as the complement to the other factors we study. A potential limitation of this approach is that it might overstate the significance of uncertainty because it combines two types of random occurrence with different causes. The first type is intrinsic and unavoidable randomness about the timing and location of future loads: even if all the tankers were controlled by the same entity, the decision of oil exporters to seek transportation would still be subject to some idiosyncratic shocks. The second type of randomness has to do with vessel-load assignments: even if transportation requests arrived deterministically and in known locations, the outcome of the fixing process (which vessel eventually gets to transport a load) would remain uncertain. We argue that this latter form of uncertainty should be considered as a by-product of fragmentation, because under a centralized planner it would disappear, and hence its effect on empty miles should be included in our estimate of the share of ballasting due to fragmentation.

The distinction outlined above does not affect our calculation for the share of ballasting due to trade imbalances of Section 3.1, because the central manager can allocate any load that appears in the data to any feasible vessel. However, in Section 3.2 we assume that each operator's clairvoyant planner knows exactly which loads it will be assigned, while a more realistic assumption is that it only has information on which loads will be offered to its fleet. We address this issue in Appendix A, where we consider a central planner that solves a stochastic dynamic program. Specifically, we adopt the state representation of Godfrey and Powell (2002a,b), and obtain an approximate solution to the dynamic program using Reinforcement Learning tools. For the sake of simplicity we limit that analysis to the Aframax vessel class. With the approach of Appendix A we estimate that the share of ballasting due to uncertainty is 7% of observed costs, compared to 11% from Table 4. By construction the estimate of Appendix A considers only randomness about the timing and location of future loads, so it can be taken as lower bound on the share of ballasting due to uncertainty,

while the estimate of Table 4 can be considered an upper bound. We argue then that the share of ballasting due to fragmentation given in Table 4 for Aframax, equal to 16% of total costs, is a lower bound, while the share implied by the calculations of Appendix A, equal to approximately 20% of observed costs, is an upper bound. Because we do not observe the set of Commercial Operators to which a load was offered, we cannot further investigate the relationship between fragmentation and uncertainty, but the fact that the estimates we derive from two substantially different approaches are close to each other reassures us that our findings are robust.

4 Fleet Size and Operational Efficiency

In the previous section, we demonstrated the significant impact of market fragmentation on shipping efficiency and illustrated how the oil transportation market could substantially benefit from consolidation. This section aims to identify the key mechanisms driving these improvements through an empirical approach. We examine how changes in pool size affect efficiency measures and the overall operational complexity of shipping pools.

Our analysis reveals two primary findings. First, vessels managed by commercial operators with larger fleets experience substantially lower ballasting times, resulting in reduced costs and emissions. This confirms that the benefits of consolidation are realized almost exclusively during the ballast leg of voyages. Second, larger shipping pools operate differently compared to smaller fleets. Rather than focusing on select routes, larger pools utilize their vessels more extensively, i.e., serving a larger set of ports. Thus, this section transitions from a market-level analysis of fragmentation and ballasting to a micro-level examination based on individual commercial operators. This approach allows us to explore the specific factors that contribute to improved efficiency in larger fleets.

Pool size. Throughout this section, our primary independent variable is the *weighted pool size*, defined at the commercial operator level. For any given period (e.g., a quarter) we calculate the pool size as a weighted count of vessels under the operator's management during that period. This calculation takes into account the number of days each vessel spends operating (whether ballasting, laden, or at port) for the commercial operator, relative to the total number of days in the period:

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

where $T_{s,i,t}$ equals the number of days in period t that vessel s spent under commercial operator i and T_t is the total number of days in period t. This approach to measuring pool size allows for a more nuanced representation of an operator's fleet capacity over time, accounting for fleet changes occuring mid-period, which is important when we carry out our analysis at quarter/semester level.

For instance, if an operator managed two ships in April 2018, with one operating for the entire month and another joining mid-month, the calculated pool size would reflect this partial utilization.

Using pool size as a covariate in our analysis potentially exposes us to certain endogeneity risks. One such risk stems from time-varying unobservable characteristics of commercial operators that affect both the outcome variable and the pool size. However, we consider pool size as a surrogate for an operator's "market weight", exactly so that it encompasses these unobservable factors correlated with size. Another risk is that of reverse causality. For example, an operator might first improve efficiency through investments in advanced analytical capabilities, subsequently attracting more vessels to its pool due to increased utilization prospects. Addressing these issues would require a source of exogenous, yet prolonged, variation in pool sizes.

Exploiting natural experiments in this context is particularly challenging due to the long travel times in shipping. For instance, the 2021 Suez Canal blockage affected fleet availability for only six days, whereas the average voyage lasts 25 days. To have a detectable impact on pool-wide outcomes, exogenous supply shocks would need to persist for several weeks or months. Despite these limitations, our approach provides valuable insights into the relationship between pool size and operational efficiency in the shipping industry.

4.1 Fleet Size and Ballast Times

Building on our findings from Section 3, where we demonstrated that consolidation can reduce ballasting, we now empirically explore the mechanisms through which this reduction is achieved in practice. Our analysis focuses on how fleets of various sizes differ in their behavior across the two legs of a voyage, and our first objective is to show empirically that any efficiency improvements should primarily arise from the ballasting portion.

To conduct this analysis, we compare the travel times of each voyage leg with the "non-delayed" travel time on the same leg. We define this non-delayed time as the 20th percentile of all observed travel times for the same route, a concept similar to "free-flow" travel time in the transportation literature. This approach allows us to disregard delays due to factors such as port congestion. We introduce the following two variables for each voyage:

$$BallastDelay_{n,i,t} = TravelTimeBallast_{n,i,t} - NonDelayedTimeBallast_n$$
$$LadenDelay_{n,i,t} = TravelTimeLaden_{n,i,t} - NonDelayedTimeLaden_n,$$

where $BallastDelay$ ($LadenDelay$) captures the difference between actual ballast (laden) travel time and non-delayed ballast (laden) time. Using these as outcome variables helps us distinguish between direct travel between ports and legs involving inefficient detours. We expect to see no

substantial difference in laden delays across commercial operators, regardless of fleet size or the route. However, we anticipate that ballast delays will vary with fleet size due to differences in vessel positioning and load assignment strategies.

We identify two basic mechanisms by which a fleet can reduce empty vessel time. The first involves optimizing vessel positioning and load assignment to minimize average distance between ships and loads. The second focuses on minimizing the time it takes each vessel to secure cargo on a chosen ballast route. Section 3 concentrates on the first mechanism. The empirical approach we undertake in this section allows us to disentangle the relationship between fleet size and each of them. Toward this end, we estimate the following regression models, one for laden delays and one for ballast delays. These models include various fixed effects to account for factors such as seasonal changes in transportation demand, COVID-19 impacts, and operator-specific variables:

$$LadenDelay_{n,i,t} = \theta \times \log\left(PoolSize_{i,t}\right) + \phi_{o,d} + \xi_t + \psi_{i,v}, \tag{5}$$

$$BallastDelay_{n,i,t} = \theta \times \log\left(PoolSize_{i,t}\right) + \phi_{o,d} + \xi_t + \psi_{i,v}. \tag{6}$$

In the above, $\phi_{o,d}$ denotes an *origin-destination* fixed effect, ξ_t denotes a time fixed effect, and $\psi_{i,v}$ denotes a commercial operator and vessel class fixed effect. We are particularly interested in comparing the estimates for the coefficient θ with and without the origin-destination fixed effect $\phi_{o,d}$. Our results, presented in Table 5, confirm our expectations. The laden delay model shows no significant relationship with fleet size, supporting our approach in Section 3. The ballast delay regression, however, reveals that larger fleets experience shorter ballast legs. Specifically, doubling fleet size is associated with a half-day reduction in ballast time per voyage. To contextualize this result, a fleet of ten tankers, each completing one voyage per month, would save five days of ballasting every month compared to having two separate fleets each with five vessels; 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.

Further analysis using origin-destination fixed effects allows us to isolate the impact of factors such as a fleet's ability to quickly secure loads or avoid costly detours, conditioning on the choice of the route where to ballast. In other words, using origin-destination fixed effects amounts to analyzing differences between large and small fleets while taking the ballast route as given. It follows that the estimate of θ for Equation (6) with route fixed effects captures the importance of factors such as relationship with brokers. On the other hand, the difference between the resulting estimate and that without fixed effects captures the importance of coordination in improving ballasting efficiency, which is more easily achieved by larger fleets. In addition, Figure 2 illustrates the decreasing marginal returns from increasing pool size, confirming our intuition from earlier analyses.

Dependent Variable	Ballast Delay	Ballast Delay	Opt. Results	Laden Delay	Laden Delay
Independent Variable					
Log Fleet Size	-0.5006***	-0.3620**	-0.2210	0.0396	0.0085
	(0.1051)	(0.1192)	N/A	(0.0521)	(0.0621)
Fixed Effect					
OD Pair (Area)	No	Yes	N/A	No	Yes
Operator × Vessel Class	Yes	Yes	N/A	Yes	Yes
Time (Quarter)	Yes	Yes	N/A	Yes	Yes
Fit Statistics					
R-squared Adj.	0.106	0.332	N/A	0.041	0.253
N obs	141,778	141,778	N/A	141,778	141,778

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

Table 5: Summary of coefficient estimates for regressions (5) and (6).

Notes: For a ballast leg the origin is the starting port and the destination is the load port, while for a laden leg, the origin is the load port and the destination is the discharge port. Because imposing fixed effects for each port pair exponentially increases the number of variables to estimate, with resulting loss of statistical power, we estimate origin-destination fixed effects at area level. However, we compute travel times at port level (i.e., if two legs are between different ports that belong to the same origin/destination area, their travel times will be different). Moreover, we estimate time fixed effects at quarterly level and compute travel times at port level. Additional tables in Appendix B confirm that our estimates and insights are robust to using different geographical and time granularities.

We can further corroborate our findings by relating the estimates in Table 5 to the results presented in Section 3.3. In that section, we constructed synthetic shipping pools of uniform size and re-optimized the assignment of loads to vessels to minimize each fleet's ballasting cost. This approach allowed us to create pools of any size and compute the average length of a ballast leg for pools of size 2, 4, 8, and so on, as illustrated in the right panel of Figure 1. The third column of Table 5 reports the average reduction in ballast times as the fleet size doubles taken from that analysis. Notably, this value is of the same order of magnitude as the difference between the estimates of Equation (6) with and without fixed effects. It is important to note that in Figure 1, we only vary the size of the pools within which we re-optimize the load-vessel assignment, while keeping demand and available information constant. This approach allows us to isolate and measure how coordination efficiency improves as pool size increases.

The fact that the value in the third column of Table 5 is comparable to our regression estimates provides further validation for our analysis in Section 3. This congruence between our theoretical model and empirical findings strengthens our confidence in the robustness of our results and the accuracy of our analytical approach.

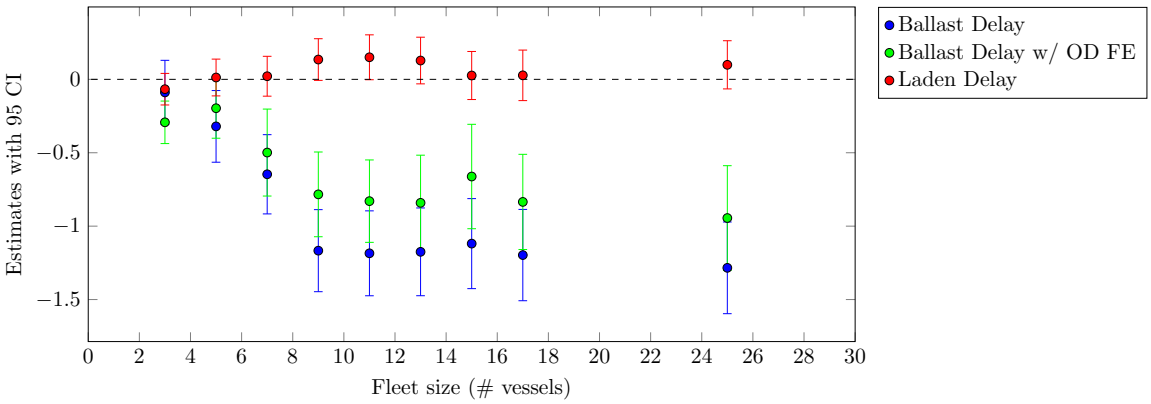

Figure 2: Difference in average length (in days) between the ballast legs of voyages performed by operators of increasing size. The figure is obtained by estimating the regression in Equation (6) using categorical variables corresponding to size bins (e.g., $[1, 2)$ vessels, $(2, 4]$ vessels, etc) instead of the continuous covariate $PoolSize$. The vertical bars represent the 95% confidence intervals, which increase in width as the the size of operators increases because there are fewer and fewer observations to estimate the coefficients with.

4.2 Operations Complexity and Coordination

Our analysis so far, employing two complementary approaches, demonstrates that the enhanced coordination facilitated by larger fleets significantly contributes to increased shipping efficiency. We now examine *how* this improved coordination is achieved, specifically investigating the impact of pool size on a fleet's strategic decisions. Because we cannot observe relational factors (such as a pool's relationship with brokers, which affects the ability to quickly secure new loads), we focus on the network of locations served by fleets and how vessels are utilized within this network.

When comparing fleets of different sizes, it is intuitive to assume that the larger fleet offers greater *flexibility*. For instance, a larger fleet is more likely to have a vessel nearby to serve an unexpected cargo opportunity or to reassign vessels efficiently if the operator does not win a load. This behavior is reminiscent of the flexibility designs described by Jordan and Graves (1995) in the context of production facilities. Consequently, we analyze our data for the maritime equivalent of a more connected manufacturing network: laden routes can be viewed as "product" whereas vessels act as "plants". For example, a "long chain" in this context would manifest as a vessel serving multiple locations intermittently rather than being tied to a single route. In essence, we seek evidence that larger fleets operate more complex shipping networks.

Because large commercial operators are expected to simultaneously serve more locations than smaller ones, we perform vessel-level regressions to ensure meaningful comparisons. We explore network complexity along two complementary dimensions: the number of laden trips a vessel

performs between the same origin-destination pair, and the number of different ports a vessel visits. Formally, we estimate the following specifications at a quarterly level:

$$\log\left(Loads_{s,i,t}^{o,d}\right) = \beta \times \log\left(PoolSize_{i,t}\right) + \xi_t + \phi_{o,d}, \tag{7}$$

$$\log\left(No.Ports_{s,i,t}\right) = \eta \times \log\left(PoolSize_{i,t}\right) + \xi_t + \psi_{i,s}. \tag{8}$$

Here, $Loads_{s,i,t}^{o,d}$ counts the number of loads with origin in area o and destination in area d that vessel s of operator i served in quarter t, while $No.Ports_{s,i,t}$ counts the distinct ports visited by tanker s in quarter t under operator i's management. Equation (7) aims to determine whether pools tend to assign vessels to a limited number of routes, while Equation (8) examines whether vessels in larger fleets are used more extensively across the network.[18] Continuing the manufacturing network analogy, Equation (7) estimates whether the strength of the link between a vessel and a specific laden leg depends on pool size, with a stronger link indicating less flexibility. Similarly, Equation (8) estimates whether the same vessel connects to more laden legs, collectively assessing how fleet size impacts the quantity and strength of connections between "products" and "plants".

Table 6 presents summary results for these regressions. The negative estimate for β indicates that smaller fleets tend to travel the same route more frequently, while the positive estimate for η shows that tankers managed by larger operators take on loads from a wider range of locations in the same time frame. Together, these results suggest that smaller pools concentrate on fewer legs, traveling the same route back and forth, while larger pools use the same vessels to visit more ports and take on more diverse loads. Essentially, larger fleets operate wider networks with weaker individual links, indicating a higher degree of fungibility between vessels and laden routes. These findings suggest that fleet size strongly correlates with an operator's ability to establish a flexible shipping network, which in turn affects efficiency in terms of reduced ballasting. This may be due to the costs associated with establishing links between vessels and laden legs: serving many locations requires not only connections with brokers offering loads but also arrangements for refueling, repairs, and other activities that smaller operators may struggle to manage. From this perspective, we highlight that consolidation through shipping pools may improve coordination not only through the simple effect of having more vessels, but also by placing more ships under the control of sophisticated operators capable of fully leveraging connected shipping networks.

[18] We include an (o, d)-pair level fixed effect in Equation (7) to account for factors such as route importance and regional considerations, comparing only vessels that served that specific leg. In Equation (8), we instead add a fixed effect for each tanker's vessel class.

Dependent Variable	Log Laden Trips in OD [#Days]	Log Laden Trips in OD [# trips]	Log Ports visited by vessel [# ports]
Independent Variable			
Log Fleet Size	-0.0074***	-0.0031**	0.2326***
	(0.0008)	(0.0012)	(0.005)
Fixed Effect			
Time (Quarterly)	Yes	Yes	Yes
Operator × Vessel Class	No	No	Yes
OD Pair × Vessel Class	Yes	No	No
OD Pair	No	Yes	No
Vessel ID	No	Yes	Yes
Fit Statistics			
R-squared Adj.	0.38520	0.39518	0.3715
N obs	135,853	135,853	66,158
Granularity			
Geographical	Port	Port	Port

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

Table 6: Summary of coefficient estimates for the regression analyses

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 methods, we find that fragmentation exacerbates the adverse effects of imbalances in trade flows. Our research further shows that even modest market consolidation can generate significant efficiency gains. Analyzing a dataset of six thousand oil tanker voyages, we determine that imbalances in trade flows account for 70-75% of empty miles, while fragmentation accounts for 15-20% (with 5-10% attributed to demand uncertainty inherently present in the market). This is among the very first studies that provides a quantitative assessment of the main factors contributing to inefficiencies in a transportation market vital for global supply chains.

In addition, we showcase empirically two mechanisms that explain the beneficial impact of consolidation on ballasting. First, centralizing decision-making improves asset coordination, so vessels can serve the same set of locations more efficiently. Second, larger pools diversify their network of served ports, thereby optimizing the utilization of tankers and minimizing ballasting.

Finally, we turn our attention to the emerging industry trend of shipping pools. Our data suggest that shipping pools of modest size suffice to capture most of the efficiency benefits that a centralized market would generate. For example, partitioning the fleet of oil tankers in pools of 30-40 vessels each results (which corresponds to fleets of approximately 5% of the global fleet of Aframax vessels) in a 15% decrease in empty miles (i.e., 70% of the first-best improvement associated with a central planner). Besides its practical value, this analysis provides a data-driven

illustration of the extent to which pooling resources improves operational efficiency.

We view our work as contributing to the literature on decentralized transportation markets and, in particular, to the growing strand that analyzes the effect of market structure on efficiency outcomes. We hope that this work will motivate further studies along many directions, including building an understanding of the effect of fragmentation on market prices and demand for transportation services, and exploring how carbon taxation mechanisms may be designed to discourage empty miles while taking into account the market structure.

References

Acemoglu, D., Makhdoumi, A., Malekian, A., and Ozdaglar, A. (2018). Informational braess' paradox: The effect of information on traffic congestion. *Operations Research*, 66(4):893–917.

Adland, R. and Prochazka, V. (2021). The value of timecharter optionality in the drybulk market. *Transportation Research Part E: Logistics and Transportation Review*, 145:102185.

Arlotto, A., Frazelle, A. E., and Wei, Y. (2019). Strategic open routing in service networks. *Management Science*, 65(2):735–750.

Arora, K., Zheng, F., and Girotra, K. (2024). Private vs. pooled transportation: Customer preference and design of green transport policy. *Manufacturing & Service Operations Management*, 26(2):594–611.

Azagirre, X., Balwally, A., Candeli, G., Chamandy, N., Han, B., King, A., Lee, H., Loncaric, M., Martin, S., Narasiman, V., et al. (2024). A better match for drivers and riders: Reinforcement learning at lyft. *INFORMS Journal on Applied Analytics*, 54(1):71–83.

Bertsekas, D. P. (2018). *Dynamic Programming and Optimal Control*, volume II. Athena Scientific, Belmont, MA, USA, 4th edition.

Besbes, O., Castro, F., and Lobel, I. (2021). Surge pricing and its spatial supply response. *Management Science*, 67(3):1350–1367.

Bimpikis, K., Candogan, O., and Saban, D. (2019). Spatial pricing in ride-sharing networks. *Operations Research*, 67(3):744–769.

Bimpikis, K. and Markakis, M. G. (2016). Inventory pooling under heavy-tailed demand. *Management Science*, 62(6):1800–1813.

Boyd, S. P. and Vandenberghe, L. (2004). *Convex optimization*. Cambridge university press.

Brancaccio, G., Kalouptsidi, M., and Papageorgiou, T. (2020). Geography, transportation, and endogenous trade costs. *Econometrica*, 88(2):657–691.

Brancaccio, G., Kalouptsidi, M., Papageorgiou, T., and Rosaia, N. (2023). Search Frictions and Efficiency in Decentralized Transport Markets. *The Quarterly Journal of Economics*.

Braverman, A., Dai, J. G., Liu, X., and Ying, L. (2019). Empty-car routing in ridesharing systems. *Operations Research*, 67(5):1437–1452.

Buchholz, N. (2022). Spatial equilibrium, search frictions, and dynamic efficiency in the taxi industry. *The Review of Economic Studies*, 89(2):556–591.

Cachon, G. P., Daniels, K. M., and Lobel, R. (2017). The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing & Service Operations Management*, 19(3):368–384.

Castro, F., Gao, J., and Martin, S. (2023). Autonomous vehicles in ride-hailing and the threat of spatial inequalities. *Available at SSRN 4332493*.

Dai, T., Lee, H. L., and Tang, C. S. (2024). Toward supply-chain-aware esg measures. In *Responsible and Sustainable Operations: The New Frontier*, pages 235–252. Springer.

Eppen, G. D. (1979). Note—effects of centralization on expected costs in a multi-location newsboy problem. *Management science*, 25(5):498–501.

Feldman, P., Li, Y., and Tsoukalas, G. (2025). Green incentives in decentralized consortia. *Available at SSRN*.

Frechette, G. R., Lizzeri, A., and Salz, T. (2019). Frictions in a competitive, regulated market: Evidence from taxis. *American Economic Review*, 109(8):2954–92.

Godfrey, G. A. and Powell, W. B. (2002a). An adaptive dynamic programming algorithm for dynamic fleet management, i: Single period travel times. *Transportation Science*, 36(1):21–39.

Godfrey, G. A. and Powell, W. B. (2002b). An adaptive dynamic programming algorithm for dynamic fleet management, ii: Multiperiod travel times. *Transportation Science*, 36(1):40–54.

Gordon, G. J. (1995). Stable function approximation in dynamic programming. In *Machine Learning Proceedings 1995*, pages 261–268. Elsevier.

Hansen-Lewis, J. and Marcus, M. M. (2022). Uncharted waters: Effects of maritime emission regulation. Working Paper 30181, National Bureau of Economic Research.

Harris, A. and Nguyen, T. M. A. (2022). Long-term relationships and the spot market: Evidence from us trucking. Technical report, Working Paper.

Jacquillat, A., Martin, S., and Wang, K. (2024). Value of sharing in robots-as-a-service operations. *Available at SSRN 4723289*.

Jordan, W. C. and Graves, S. C. (1995). Principles on the benefits of manufacturing process flexibility. *Management science*, 41(4):577–594.

Lee, H. L. and Tang, C. S. (2018). Socially and environmentally responsible value chain innovations: New operations management research opportunities. *Management Science*, 64(3):983–996.

Liu, F., Song, J.-S., and Tong, J. D. (2016). Building supply chain resilience through virtual stockpile pooling. *Production and Operations management*, 25(10):1745–1762.

Ma, H., Fang, F., and Parkes, D. C. (2022). Spatio-temporal pricing for ridesharing platforms. *Operations Research*, 70(2):1025–1041.

Manshadi, V., Rodilitz, S., Saban, D., and Suresh, A. (2024). Online algorithms for matching platforms with multichannel traffic. *Management Science*.

Munos, R. and Szepesvári, C. (2008). Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research*, 9(5).

Netessine, S., Dobson, G., and Shumsky, R. A. (2002). Flexible service capacity: Optimal investment and the impact of demand correlation. *Operations Research*, 50(2):375–388.

Özkan, E. and Ward, A. R. (2020). Dynamic matching for real-time ride sharing. *Stochastic Systems*, 10(1):29–70.

Prochazka, V., Adland, R., and Wallace, S. W. (2019). The value of foresight in the drybulk freight market. *Transportation Research Part A: Policy and Practice*, 129:232–245.

Séjourné, T., Samaranayake, S., and Banerjee, S. (2018). The price of fragmentation in mobility-on-demand services. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 2(2):1–26.

Wright, C. P., Groenevelt, H., and Shumsky, R. A. (2010). Dynamic revenue management in airline alliances. *Transportation Science*, 44(1):15–37.

Zhuge, D., Wang, S., and Zhen, L. (2024). Shipping emission control area optimization considering carbon emission reduction. *Operations Research*.

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.[19]

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

[19] 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 9; 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.[20] 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

[20] 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. Note 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 (1b), 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

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

The central planner seeks minimizes $\sum_{t=0}^{\infty} \gamma^t r_t$, where γ is a discount factor.[21] 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:[22] 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.

[21] We set $\gamma = 0.85$, which corresponds to an effective time horizon of approximately one month.
[22] 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.[23] 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

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

Thus, s^i corresponds to a situation in which relatively more of the vessels become available in location i over time, so that we expect $V^*(s^i)$ to capture how "good" is having tankers in i. To each state $s^i \in \mathcal{S}$ we associate an initial value V^i_0, defined as the γ-discounted value over 147 time periods earned by a myopic central planner.[24] 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

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

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

[24] 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:

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

Note 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 $k = 1, 2 \ldots$ 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 $\mathcal{B}_k^S = \left\{ \left(s^i, \hat{V}_k^i, g_k^i \right) : i = 1 \ldots, N_A \right\}$, 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,[25] it appears from Figure 3 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.

[25] See the discussions in Gordon (1995) and Bertsekas (2018) for additional details on the reasons why Fitted Value Iteration may fail to converge.

A.3 Comparison with Perfect Information

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

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

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

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

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

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

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

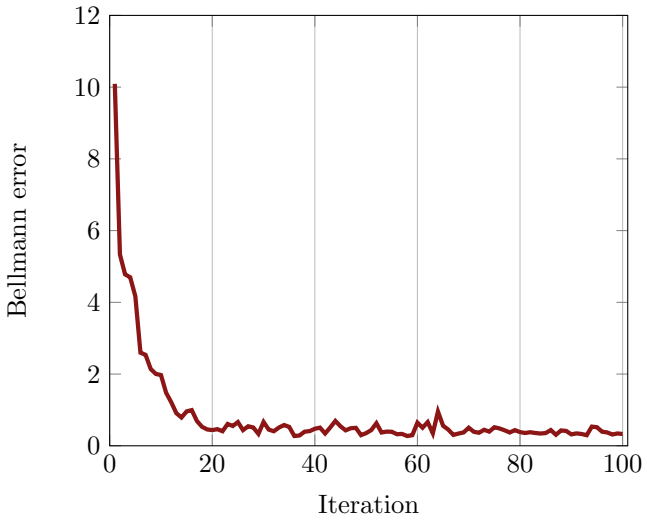

Figure 3: L_2 distance between each successive set of basis values $\hat{V}_k^i$ for the basis states of the Fitted Value Iteration procedure.

B Tables and Figures

B.1 Tables

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

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

Table 8: Share of ballasting attributed to trade imbalances, uncertainty and market fragmentation, using as objective function in Section 3 total fuel consumption associated to ballasting travels.

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 9: 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 10: 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.

The following tables (11 to 16) report an extensive set of robustness checks for Equations (5) and (6), the relationships between fleet size and the delays experienced in different sections of the voyages. The main purpose of this analysis is to validate the results in Section 4.1. 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.

Table 11: Ballast delay for Area granularity

Dependent Variable:	DelayBallast													
	All	All w/ OD	Aframax	Aframax w/ OD	Panamax	Panamax w/ OD	Suezmax	Suezmax w/ OD	VLCC	VLCC w/ OD	MR1	MR1 w/ OD	MR2	MR2 w/ OD
Variables														
log(fleet_size)	-0.5006***	-0.3620***	-0.3004	-0.3456	-0.6883	-0.4148	-0.7003	-0.4188	-1.291*	-1.271**	-0.2058	-0.0097	-0.1183	0.0841
	(0.1051)	(0.1192)	(0.2821)	(0.3450)	(0.9143)	(0.7926)	(0.6384)	(0.4458)	(0.7728)	(0.5726)	(0.1552)	(0.1350)	(0.1917)	(0.1329)
Fixed-effects														
Operator-VC	Yes													
Time	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Operator-VC-OD_Ballast_Area		Yes												
Operator			Yes		Yes		Yes		Yes		Yes		Yes	
Operator-OD				Yes		Yes		Yes		Yes		Yes		Yes
Fit statistics														
Observations	141,778	141,778	31,858	31,858	11,983	11,983	15,824	15,824	13,718	13,718	24,655	24,655	43,740	43,740
R^2	0.10561	0.33189	0.07454	0.32231	0.10514	0.32548	0.12901	0.40042	0.11945	0.35346	0.10761	0.28111	0.06909	0.26017
Within R^2	0.00026	0.00013	7.03×10^{-5}	9.05×10^{-5}	0.00041	0.00014	0.00027	9.57×10^{-5}	0.00053	0.00052	0.00011	2.45×10^{-7}	2.31×10^{-5}	1.07×10^{-5}

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

Table 12: Ballast delay for Narrow Area granularity

Dependent Variable:	DelayBallast													
	All	All w/ OD	Aframax	Aframax w/ OD	Panamax	Panamax w/ OD	Suezmax	Suezmax w/ OD	VLCC	VLCC w/ OD	MR1	MR1 w/ OD	MR2	MR2 w/ OD
Variables														
log(fleet_size)	-0.5006***	-0.4007***	-0.3004	-0.3390	-0.6883	-0.4473	-0.7003	-0.2691	-1.291*	-1.575**	-0.2058	-0.1012	-0.1183	0.0837
	(0.1051)	(0.1282)	(0.2821)	(0.3299)	(0.9143)	(0.8807)	(0.6384)	(0.4172)	(0.7728)	(0.7174)	(0.1552)	(0.1340)	(0.1917)	(0.1731)
Fixed-effects														
Operator-VC	Yes													
Time	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Operator-VC-OD		Yes												
Operator			Yes		Yes		Yes		Yes		Yes		Yes	
Operator-OD				Yes		Yes		Yes		Yes		Yes		Yes
Fit statistics														
Observations	141,778	141,778	31,858	31,858	11,983	11,983	15,824	15,824	13,718	13,718	24,655	24,655	43,740	43,740
R^2	0.10561	0.39705	0.07454	0.39357	0.10514	0.37527	0.12901	0.47046	0.11945	0.40286	0.10761	0.36972	0.06909	0.32983
Within R^2	0.00026	0.00015	7.03×10^{-5}	8.46×10^{-5}	0.00041	0.00015	0.00027	3.66×10^{-5}	0.00053	0.00074	0.00011	2.57×10^{-5}	2.31×10^{-5}	9.69×10^{-6}

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

Table 13: Ballast delay for Port granularity

Dependent Variable:	DelayBallast													
	All	All w/ OD	Aframax	Aframax w/ OD	Panamax	Panamax w/ OD	Suezmax	Suezmax w/ OD	VLCC	VLCC w/ OD	MR1	MR1 w/ OD	MR2	MR2 w/ OD
Variables														
log(fleet_size)	-0.5006***	-0.3725**	-0.3004	-0.4641	-0.6883	0.7269	-0.7003	-1.034	-1.291*	-1.767	-0.2058	-0.1734	-0.1183	0.2137
	(0.1051)	(0.1776)	(0.2821)	(0.3448)	(0.9143)	(1.043)	(0.6384)	(0.9508)	(0.7728)	(1.169)	(0.1552)	(0.2333)	(0.1917)	(0.2803)
Fixed-effects														
Operator-VC	Yes													
Time	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Operator-VC-OD		Yes												
Operator			Yes		Yes		Yes		Yes		Yes		Yes	
Operator-OD				Yes		Yes		Yes		Yes		Yes		Yes
Fit statistics														
Observations	141,778	141,778	31,858	31,858	11,983	11,983	15,824	15,824	13,718	13,718	24,655	24,655	43,740	43,740
R^2	0.10561	0.74298	0.07454	0.75281	0.10514	0.73245	0.12901	0.75488	0.11945	0.68694	0.10761	0.77361	0.06909	0.76933
Within R^2	0.00026	0.00012	7.03×10^{-5}	0.00016	0.00041	0.00037	0.00027	0.00045	0.00053	0.00077	0.00011	7.29×10^{-5}	2.31×10^{-5}	6.46×10^{-5}

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

Table 14: Laden delay for Area granularity

Dependent Variable:	DelayLaden													
	All	All w/ OD	Aframax	Aframax w/ OD	Panamax	Panamax w/ OD	Suezmax	Suezmax w/ OD	VLCC	VLCC w/ OD	MR1	MR1 w/ OD	MR2	MR2 w/ OD
Variables														
log(fleet_size)	0.0396	0.0085	0.1137	0.0892	0.1041	0.1883	0.0435	-0.0315	0.1265	0.0589	-0.1086	-0.0993	0.1094	0.0878
	(0.0521)	(0.0621)	(0.1186)	(0.1250)	(0.1621)	(0.1137)	(0.1527)	(0.1660)	(0.2614)	(0.2429)	(0.0920)	(0.1141)	(0.1387)	(0.1358)
Fixed-effects														
Operator-VC	Yes													
PeriodLaden	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Operator-VC-OD		Yes												
Operator			Yes		Yes		Yes		Yes		Yes		Yes	
Operator-OD				Yes		Yes		Yes		Yes		Yes		Yes
Fit statistics														
Observations	141,778	141,778	31,858	31,858	11,983	11,983	15,824	15,824	13,718	13,718	24,655	24,655	43,740	43,740
R^2	0.04122	0.25276	0.04360	0.27844	0.01942	0.23894	0.02771	0.32843	0.03894	0.24102	0.05497	0.22304	0.03005	0.22914
Within R^2	4.11×10^{-6}	1.66×10^{-7}	3.21×10^{-5}	1.76×10^{-5}	2.36×10^{-5}	6.33×10^{-5}	4.64×10^{-6}	2.21×10^{-6}	2.45×10^{-5}	4.95×10^{-6}	5.46×10^{-5}	4.08×10^{-5}	2.63×10^{-5}	1.42×10^{-5}

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

Table 15: Laden delay for Narrow Area granularity

Dependent Variable:	DelayLaden													
	All	All w/ OD	Aframax	Aframax w/ OD	Panamax	Panamax w/ OD	Suezmax	Suezmax w/ OD	VLCC	VLCC w/ OD	MR1	MR1 w/ OD	MR2	MR2 w/ OD
Variables														
log(fleet_size)	0.0396	-0.0146	0.1137	0.0046	0.1041	0.1572	0.0435	0.0168	0.1265	-0.0997	-0.1086	-0.0881	0.1094	0.0775
	(0.0521)	(0.0755)	(0.1186)	(0.1497)	(0.1621)	(0.1311)	(0.1527)	(0.1541)	(0.2614)	(0.2486)	(0.0920)	(0.1291)	(0.1387)	(0.1400)
Fixed-effects														
Operator-VC	Yes													
Time	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Operator-VC-OD		Yes												
Operator			Yes		Yes		Yes		Yes		Yes		Yes	
Operator-OD				Yes		Yes		Yes		Yes		Yes		Yes
Fit statistics														
Observations	141,778	141,778	31,858	31,858	11,983	11,983	15,824	15,824	13,718	13,718	24,655	24,655	43,740	43,740
R^2	0.04122	0.33123	0.04360	0.35477	0.01942	0.29774	0.02771	0.40804	0.03894	0.32345	0.05497	0.31086	0.03005	0.30867
Within R^2	4.11×10^{-6}	4.68×10^{-7}	3.21×10^{-5}	4.52×10^{-8}	2.36×10^{-5}	4.24×10^{-5}	4.64×10^{-6}	5.95×10^{-7}	2.45×10^{-5}	1.37×10^{-5}	5.46×10^{-5}	2.97×10^{-5}	2.63×10^{-5}	1.03×10^{-5}

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

Table 16: Laden delay for Port granularity

Dependent Variable:	DelayLaden													
	All	All w/ OD	Aframax	Aframax w/ OD	Panamax	Panamax w/ OD	Suezmax	Suezmax w/ OD	VLCC	VLCC w/ OD	MR1	MR1 w/ OD	MR2	MR2 w/ OD
Variables														
log(fleet_size)	0.0396	-0.2116**	0.1137	-0.2558	0.1041	-0.0780	0.0435	-0.0860	0.1265	-0.2130	-0.1086	-0.6752***	0.1094	0.0117
	(0.0521)	(0.1002)	(0.1186)	(0.1868)	(0.1621)	(0.2956)	(0.1527)	(0.2895)	(0.2614)	(0.4300)	(0.0920)	(0.2345)	(0.1387)	(0.2123)
Fixed-effects														
Operator-VC	Yes													
Time	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Operator-VC-OD		Yes												
Operator			Yes		Yes		Yes		Yes		Yes		Yes	
Operator-OD				Yes		Yes		Yes		Yes		Yes		Yes
Fit statistics														
Observations	141,778	141,778	31,858	31,858	11,983	11,983	15,824	15,824	13,718	13,718	24,655	24,655	43,740	43,740
R^2	0.04122	0.69016	0.04360	0.66651	0.01942	0.65347	0.02771	0.70379	0.03894	0.63828	0.05497	0.67236	0.03005	0.72800
Within R^2	4.11×10^{-6}	8.82×10^{-5}	3.21×10^{-5}	0.00013	2.36×10^{-5}	9.9×10^{-6}	4.64×10^{-6}	1.53×10^{-5}	2.45×10^{-5}	5.76×10^{-5}	5.46×10^{-5}	0.00127	2.63×10^{-5}	2.3×10^{-7}

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

Table 17: Effect of the fleet size of a vessel and the number of trips per quarter made on an OD pair **without** Vessel FEs (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	-0.0303***	-0.0327***	-0.0323***	-0.0261**	-0.0198**	-0.0232**	-0.0368***
	(0.0036)	(0.0102)	(0.0065)	(0.0105)	(0.0076)	(0.0101)	(0.0060)
Observations	111,709	23,821	10,004	13,064	12,362	15,936	36,522
R^2	0.30953	0.26923	0.34311	0.42845	0.19207	0.33385	0.24149
Within R^2	0.00873	0.00646	0.01362	0.00656	0.00556	0.00356	0.01770
Geo Granularity: Narrow Area							
Log Fleet Size	-0.0217***	-0.0238***	-0.0230***	-0.0199***	-0.0128***	-0.0168***	-0.0266***
	(0.0023)	(0.0076)	(0.0053)	(0.0068)	(0.0046)	(0.0050)	(0.0040)
Observations	120,104	25,712	10,320	13,725	12,845	19,193	38,309
R^2	0.29398	0.29011	0.33113	0.37202	0.24708	0.28900	0.25603
Within R^2	0.00582	0.00442	0.00766	0.00440	0.00294	0.00282	0.01212
Geo Granularity: Port							
Log Fleet Size	-0.0074***	-0.0078***	-0.0093***	-0.0140***	-0.0029**	-0.0047**	-0.0083***
	(0.0008)	(0.0021)	(0.0031)	(0.0038)	(0.0012)	(0.0021)	(0.0014)
Observations	135,853	29,730	11,281	14,997	14,002	23,148	42,695
R^2	0.38520	0.37150	0.43373	0.41945	0.31383	0.43001	0.33739
Within R^2	0.00144	0.00103	0.00198	0.00354	0.00031	0.00054	0.00279
Fixed Effects							
OD-pair × Vessel Class	Yes						
OD-pair		Yes	Yes	Yes	Yes	Yes	Yes
Quarter-year	Yes	Yes	Yes	Yes	Yes	Yes	Yes

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 18: Effect of the fleet size of a vessel and the number of trips per quarter made on an OD pair **with** Vessel FEs (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	-0.0046**	-0.0078*	-0.0164***	-0.0060	-0.0031	0.0159*	-0.0088**
	(0.0020)	(0.0046)	(0.0062)	(0.0061)	(0.0048)	(0.0081)	(0.0037)
Observations	111,709	23,821	10,004	13,064	12,362	15,936	36,522
R^2	0.41595	0.42402	0.45511	0.54708	0.34108	0.44294	0.36757
Within R^2	7.19×10^{-5}	0.00016	0.00105	0.00013	4.28×10^{-5}	0.00045	0.00039
Geo Granularity: Narrow Area							
Log Fleet Size	-0.0046**	-0.0071*	-0.0157***	-0.0119*	-0.0023	0.0067	-0.0065**
	(0.0019)	(0.0042)	(0.0058)	(0.0064)	(0.0040)	(0.0050)	(0.0026)
Observations	120,104	25,712	10,320	13,725	12,845	19,193	38,309
R^2	0.36968	0.41731	0.41765	0.45603	0.37409	0.36473	0.36381
Within R^2	9.02×10^{-5}	0.00016	0.00102	0.00054	3.04×10^{-5}	0.00011	0.00026
Geo Granularity: Port							
Log Fleet Size	-0.0031***	-0.0030	-0.0066	-0.0131**	-0.0018	0.0046	-0.0040**
	(0.0012)	(0.0028)	(0.0050)	(0.0056)	(0.0022)	(0.0029)	(0.0017)
Observations	135,853	29,730	11,281	14,997	14,002	23,148	42,695
R^2	0.39518	0.41642	0.48009	0.46515	0.38680	0.46096	0.39757
Within R^2	8.34×10^{-5}	5.79×10^{-5}	0.00028	0.00097	3.55×10^{-5}	0.00013	0.00023
Fixed Effects							
OD-pair	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Vessel ID	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Quarter-year	Yes	Yes	Yes	Yes	Yes	Yes	Yes

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 19: 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

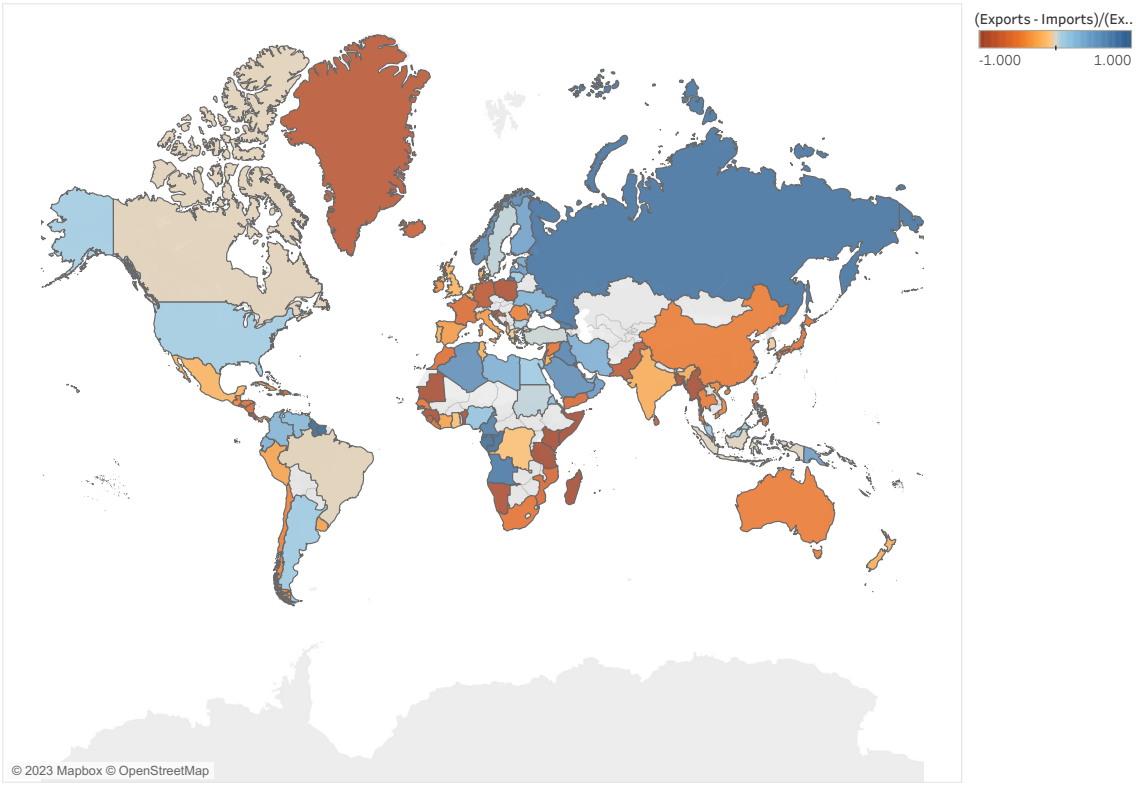

Figure 4: A map showing the demand imbalances across the network. The color of each country report the difference between the number discharge events and the number of loading events. A positive value is interpreted as a country which has more exports than imports while a negative value points to countries with higher imports than exports. Specifically, we observe that China is the larges importer while Russia and Saudi Arabia are the largest exporters.

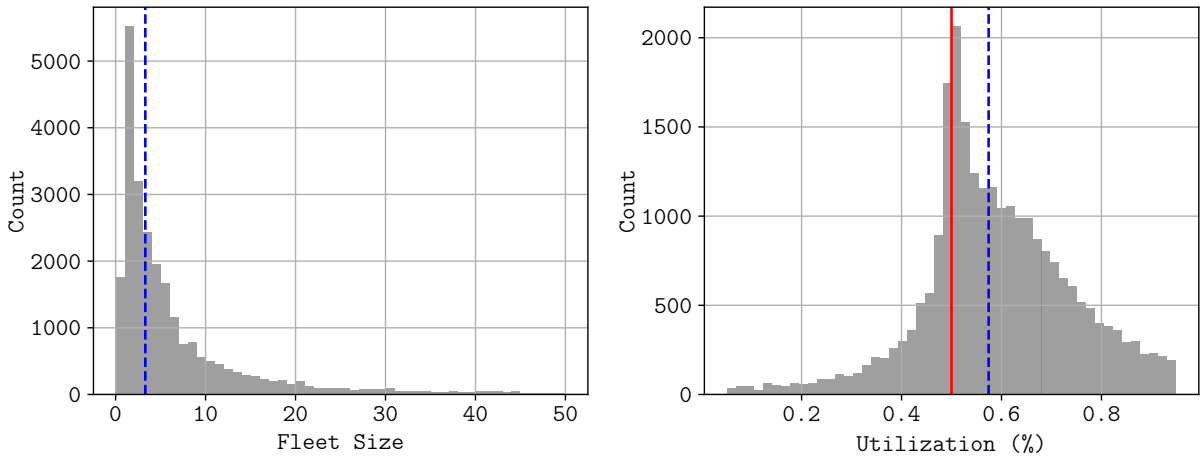

Figure 5: Left histogram shows the variation in Fleet Size at a quarterly level, i.e., for each quarter and Commercial Operator we measure their average fleet size and generate the histogram using these values. The plot shows the market structure emphasizing its fragmentation (most of the observations 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%.

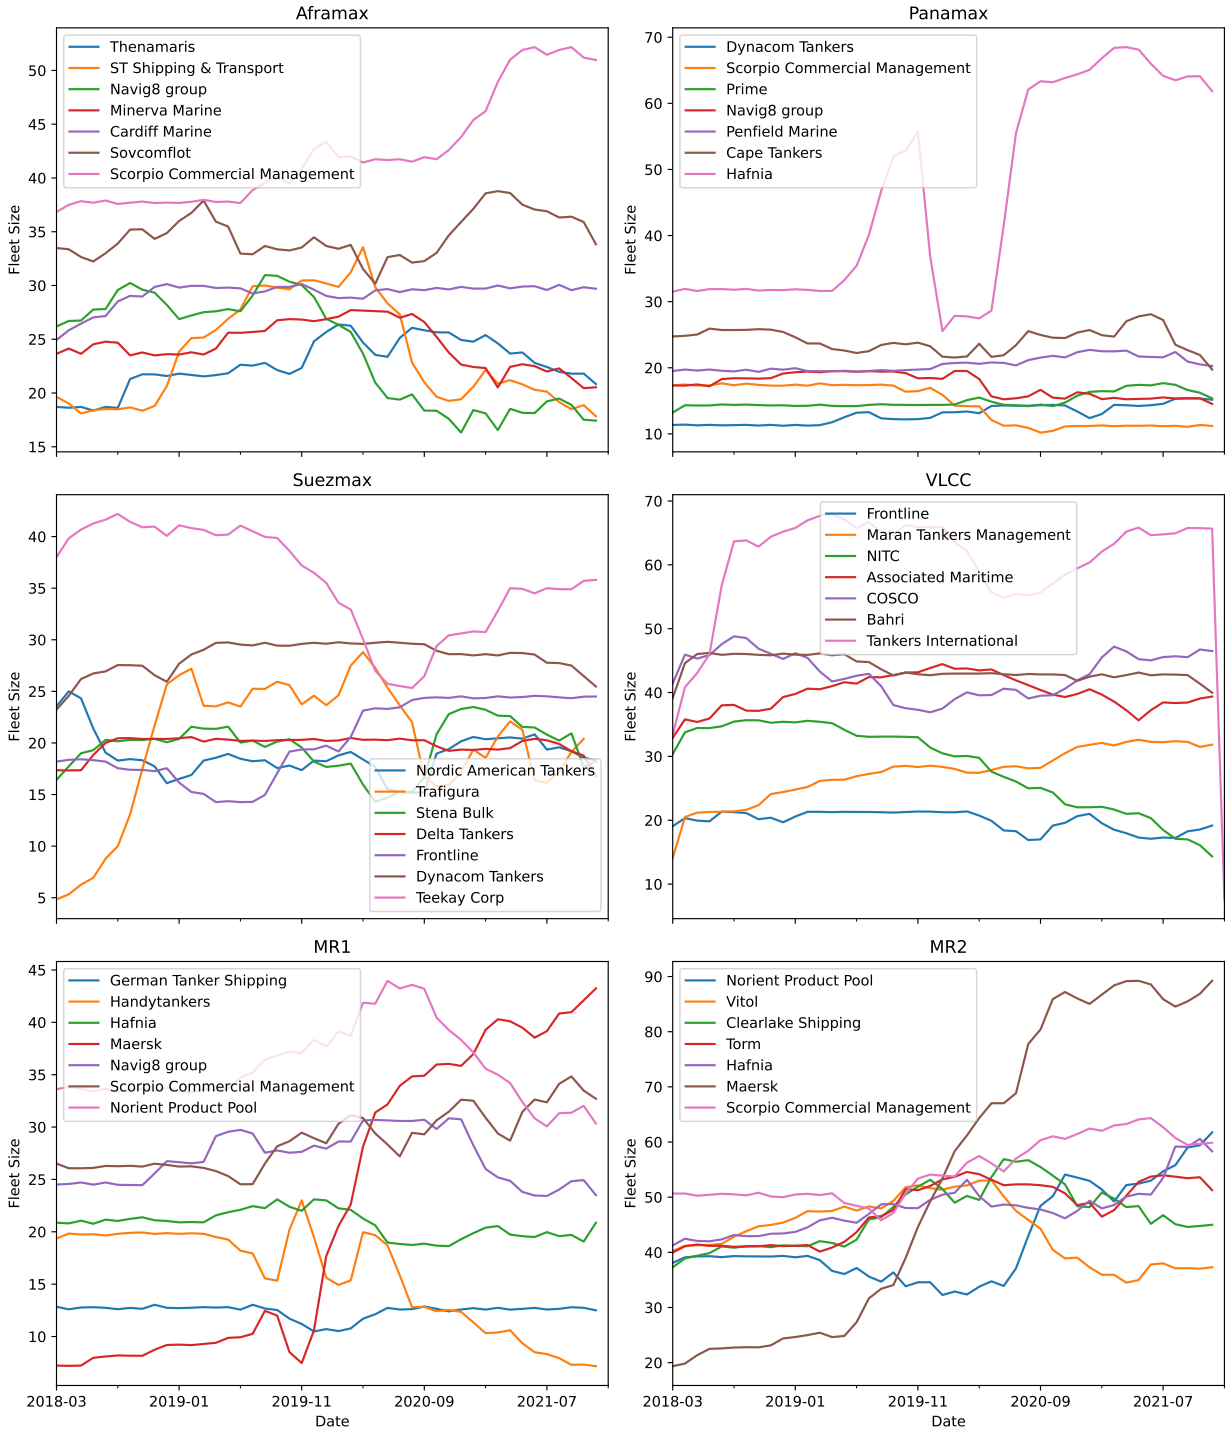

Figure 6: The figure shows the variation across time of the average fleet size of the seven largest Commercial Operators in each sub-market at a monthly level. This becomes relevant for our empirical analysis where we exploit the differences across time of the fleet size of an operator.

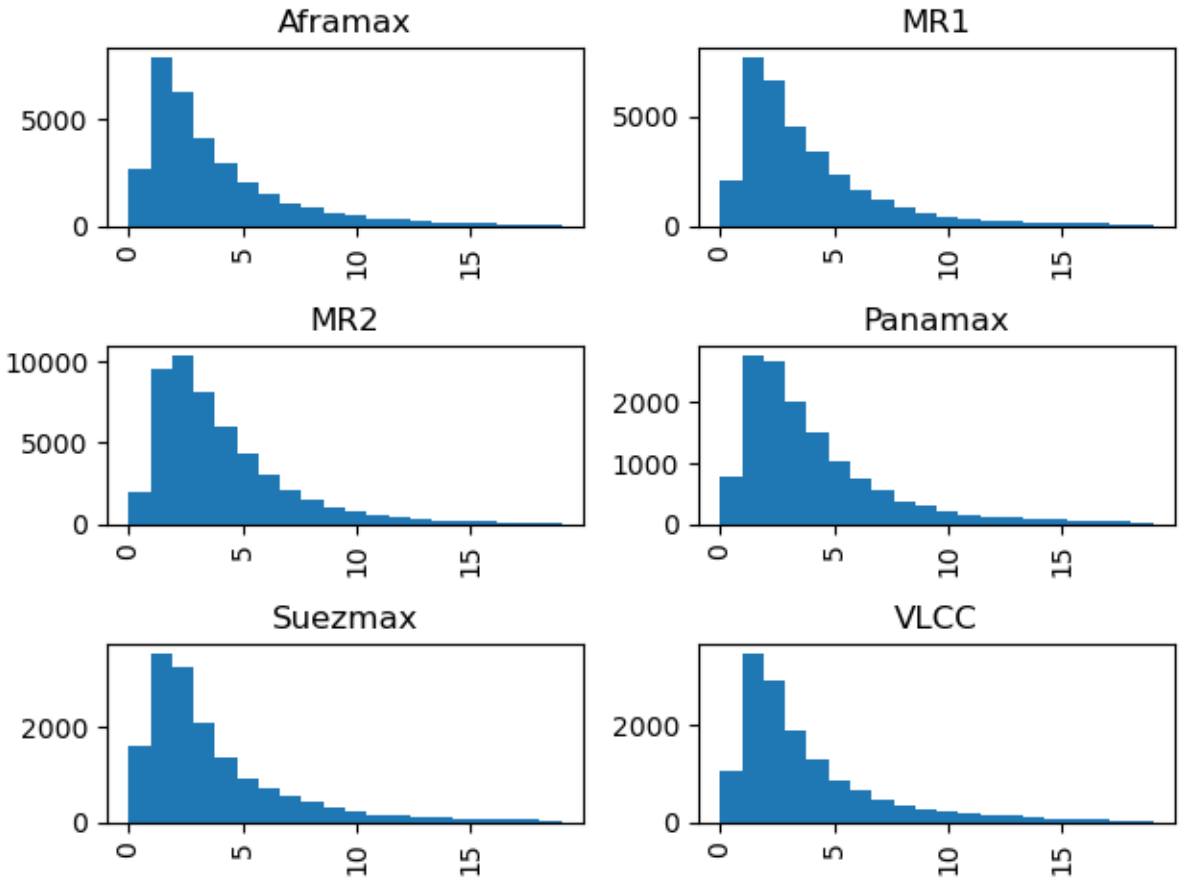

Figure 7: Distribution of the number of days between arrival of an empty vessel at port and its departure laden (i.e., loading time)

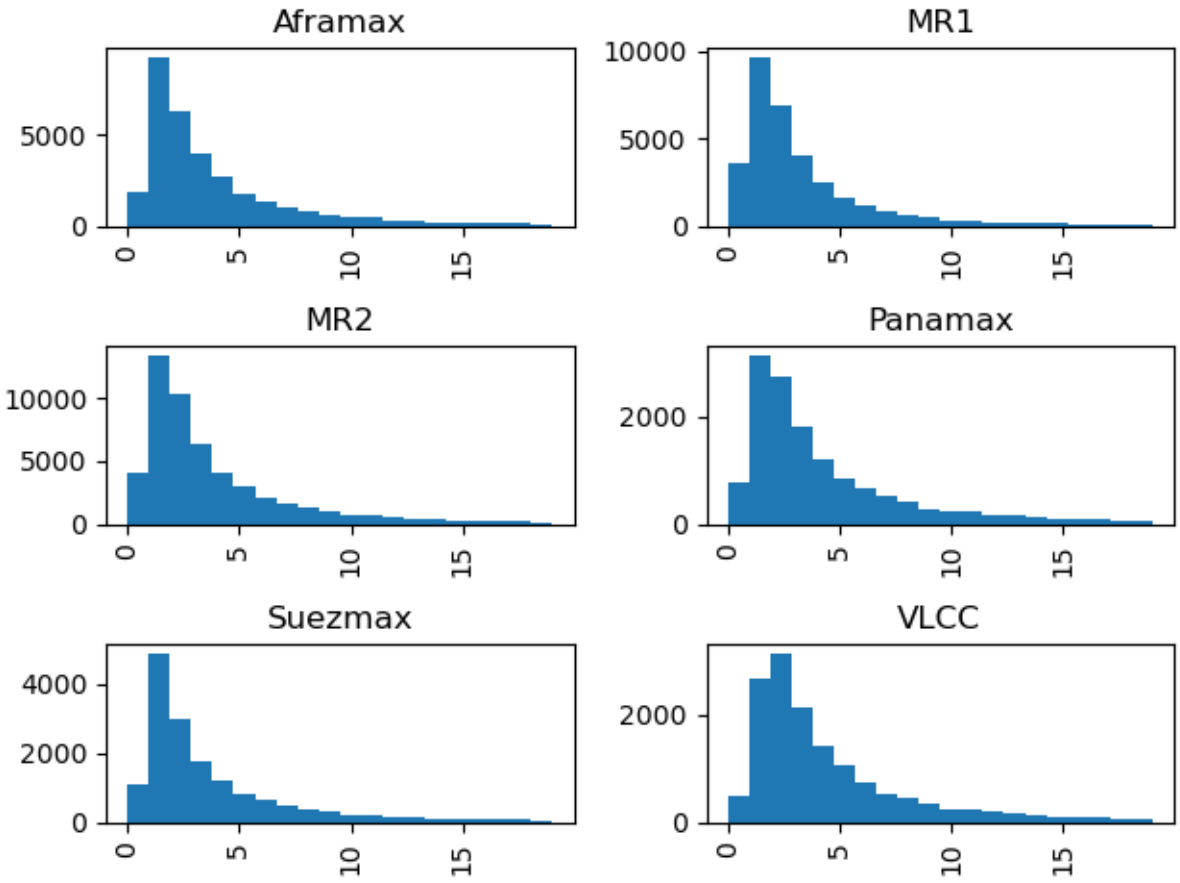

Figure 8: Distribution of the number of days between arrival of a laden vessel at port and its departure empty (i.e., unloading time)

C Analysis of utilization and fleet size

In Sections 3 and 4 we focused on ballast times and emissions to measure inefficiencies associated with fragmented shipping markets. Another commonly used measure of efficiency is *utilization*, computed as the ratio between how much a resource was used for productive purpose and the total usage. The aim of this appendix is to confirm that our insights from Section 4.1 are robust to using this alternative measure of efficiency.

We consider utilization expressed in terms of distance, time, and emissions. Specifically, for commercial operator i in time period t, we define:

(i) Distance-utilization as the ratio between the laden and the total miles traveled by all vessels operated by a commercial operator with starting date in t, i.e.,

$$Utilization^{D}_{i,t} = \frac{\sum_n M^{L}_{n,i,t}}{\sum_n M^{T}_{n,i,t}},$$

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.

(ii) Time-utilization as the ratio between the total time spent laden and the total duration of voyages beginning in time period t, i.e.,

$$Utilization^{T}_{i,t} = \frac{\sum_n T^{L}_{n,i,t}}{\sum_n T^{T}_{n,i,t}},$$

where $T^{L}_{n,i,t}$ denotes the duration of the laden leg of the n-th voyage operated by commercial operator i whose start date is in period t, and $T^{V}_{n,i,t}$ denotes the total distance of the same voyage.

(iii) (negative) Emissions-utilization as the ratio between the sum of emissions incurred during the ballast legs and the total emissions of voyages beginning in time period t, i.e.,

$$FracBallastCO2_{it} = \frac{\sum_n E^{B}_{n,i,t}}{\sum_n E^{T}_{n,i,t}}$$

where $E^{B}_{n,i,t}$ denotes the CO_2 emissions of the n-th voyage operated by i whose start date is in period t, and $E^{T}_{n,i,t}$ is equal to the total CO_2 emissions of the same voyage.

Note that with these definitions 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 the loads, we assume that the decision about ballasting destinations is made at the starting date of the voyage.

Single pool utilization. We seek to understand whether an operator becomes more efficient when its pool size grows. Specifically, we estimate the following regressions at a monthly level:

$$\log\left(Utilization_{i,t}\right) = \alpha \times \log\left(PoolSize_{i,t}\right) + \xi_t + \psi_{i,v} \tag{C.7a}$$

$$\log\left(FracBallastCO2_{i,t}\right) = \alpha \times \log\left(PoolSize_{i,t}\right) + \xi_t + \psi_{i,v} \tag{C.7b}$$

We estimate regressions for both distance- and time-based utilization. As in Section 4 ξ_t denotes a time fixed effect and $\psi_{i,v}$ an operator and vessel class-specific fixed effect. The coefficient of interest is the one associated with $PoolSize_{i,t}$: we expect the estimate of α to be positive for Equation (C.7a) and negative Equation (C.7b), reflecting that a larger pool size corresponds to higher efficiency.

We report the estimates for Equation (C.7a) in Table 20 and Table 21 for distance- and time-utilization respectively, and in Table 22 for Equation (C.7b). Our results confirm the intuition: larger pools achieve higher utilization levels. For example, the estimated coefficient in the first column of Table 20suggests 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. Naturally, these specifications are subject to the same endogeneity risks that we discuss in Section 4: unobserved variables that correlate with $PoolSize$ and $Utilization$ and reverse causality.

We can further connect the results of these regressions to Section 3.3 and Figure 2: 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 $\{(0,2],(2,3],(3,5],(5,10],(10,$ depending on the observed fleet size and, in the same way as the previous regressions, we include fixed effects for vessel class, commercial operator and time as follows:

$$Utilization_{i,t} = \beta \times PoolSizeBin_{i,t} + \xi_t + \psi_{i,v}. \tag{C.8}$$

Our results, shown in Figure 9, report the coefficients and 95% confidence intervals of the different categories of pool sizes against the excluded group of $(0,2]$. As an example, the coefficient for the $(10,15]$ group in Figure 9a is approximately 0.02, indicating that, on average, the utilization associated with a pool that has between 10 and 15 vessels is 2 percentage points higher than the utilization of a fleet with one or two vessels. Also in this case, we clearly the see that there are decreasing marginal returns from increasing fleet size, at least in terms of efficiency.

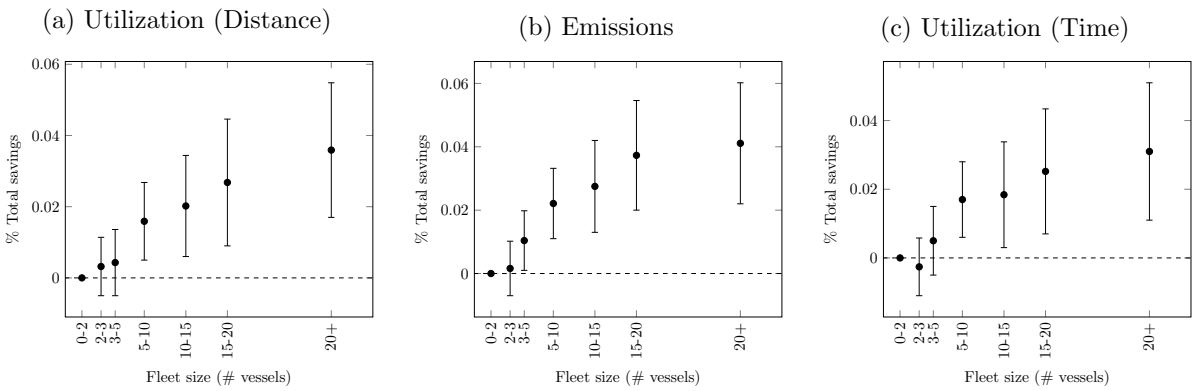

Figure 9: Coefficients and 95% confidence interval of regression (C.8) with panel data at a monthly-level. Results reveal monotonically increasing and concave functions that plateaus when fleet size is between 15 and 20, pointing to similar results as the ones provided by the optimization model. The baselines of the (0,2] excluded group are 54.1%, 58.5%, and 57.7% for utilization (distance), emissions, and utilization (time), respectively.

Coordination. We run a complementary analysis to study the relationship between pool size and utilization with two main objectives. First, to address the issue of reverse causality; second, to compare the effect of *consolidation* (a single larger entity vs. multiple smaller entities) rather than only comparing *size* (the same entity small and large). We analyze the difference between large fleets and *sets* of small fleets: for every large pool we build a "synthetic" pool composed of smaller operators that serve similar geographical areas, and compare utilization between the groups.

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

Similar to our previous result, we obtain that the utilization of large fleets is higher than the synthetic counterpart, as displayed in the first column of Table 23; the size of the difference in means is 8%.[27] Since the average pool sizes of the synthetic fleets and larger fleets are 3 and 22, respectively, using a linear transformation we get that doubling the fleet size is equivalent to

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

[27] Table 24 confirms that the same result holds for our definition of utilization in terms of time, and Table 25 shows the results of a similar analysis for emissions reporting a difference in means of -5.97%.

a 2.18% increase in utilization, which corresponds to about half of our previous result of 4.4%. These results underscore the positive impact of centralizing decision-making. Since synthetic and large fleets serve the same locations, we can rule out that the observed positive difference is due to geographical factors or to structural differences in the network of voyages of the large pool. Finally, by construction reverse causality is not an issue.

Table 20: Effect of Fleet Size in Utilization (Distance)

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

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. (C.7). Herein we extend our results and calculate the coefficient for the different markets (Vessel Classes) and we observe that all coefficients are significant besides the Aframax Class. Moreover, we filter the dataset to only observations when the average fleet size of a Commercial operator exceeds 1, 2, 3. This filtering is important given that Utilization is an aggregate metric and when few observations are available the metric could be biased. As an example, consider a fleet of 1 vessel with a Laden trip of over 30 days. Then, for the month of Laden the utilization of the fleet would be equal to 1 and for the ballasting trip the utilization would be equal to 0. Hence, this filtering helps on estimating the coefficient more accurately.

Table 21: Effect of Fleet Size in Utilization (Time)

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

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. (C.7) but using the time definition of utilization instead of the distance one. Herein we extend our results and calculate the coefficient for the different markets (Vessel Classes) and we observe that all coefficients are significant besides the Aframax Class. Moreover, we filter the dataset to only observations when the average fleet size of a Commercial operator exceeds 1, 2, 3. This filtering is important given that Utilization is an aggregate metric and when few observations are available the metric could be biased. As an example, consider a fleet of 1 vessel with a Laden trip of over 30 days. Then, for the month of Laden the utilization of the fleet would be equal to 1 and for the ballasting trip the utilization would be equal to 0. Hence, this filtering helps on estimating the coefficient more accurately.

Table 22: Effect of Fleet Size in CO_2 Emissions

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

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

Notes: Table shows the results of the regression of fleet size and the fraction of CO_2 emissions at the Commercial Operator level as described in Eq. (C.7) but using the emissions definition of utilization instead of the distance one. Explicitly, the specification we run is $FracBallastCO2_{i,t} = \alpha \times \log(PoolSize_{i,t}) + \xi_t + \psi_{i,v}$, where $FracBallastCO2_{it}$ is the fraction of emissions stemming from the ballast legs over the total emissions of commercial operator i during time period t. It is calculated by: $FracBallastCO2_{it} = BallastEmissions_{i,t}/TotalEmissions_{i,t}$. In line with our Utilizaiton specification, we filter the dataset to only observations when the average fleet size of a Commercial operator exceeds 1, 2, 3. This filtering is important given that the fraction of emissions is an aggregate metric and when few observations are available the metric could be biased. As an example, consider a fleet of 1 vessel with a Laden trip of over 30 days. Then, the fraction of ballast emissions would be 0. Hence, this filtering helps on estimating the coefficient more accurately.

	Aggregate	MR1	MR2	Panamax	Aframax	Suezmax	VLCC
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 23: Average percent difference in the utilization rate (distance based) between large pools and synthetic fleets that serve similar locations. Standard errors in parentheses.

	Aggregate	MR1	MR2	Panamax	Aframax	Suezmax	VLCC
Δ *Utlization* (%)	21.4096***	23.4235***	21.7873***	18.3534***	33.9539***	22.0736***	7.0543***
	(0.6271)	(2.1276)	(0.9647)	(2.0069)	(1.4369)	(1.9845)	(0.9684)

Table 24: Average percent difference in the utilization rate (time based) between large pools and synthetic fleets that serve similar locations. Standard errors in parentheses.

	Aggregate	MR1	MR2	Panamax	Aframax	Suezmax	VLCC
Δ *Ballasting* (%)	-5.9732***	-7.2260***	-5.0161***	-5.3114***	-8.2912***	-6.8617***	-3.7527***
	(0.4417)	(1.6402)	(0.9015)	(1.5876)	(0.9371)	(1.1360)	(0.8173)

Table 25: Average percent difference in the share of CO_2emissions due to ballasting between large pools and synthetic pools that serve similar locations. Standard errors in parentheses.