

# Beached Assets? Capital Turnover and Emissions in Shipping

Allen Peters\*

University of British Columbia

Job Market Paper

[CLICK HERE FOR THE MOST RECENT VERSION](#)

January 6, 2024

## Abstract

This paper examines the equilibrium impacts of emissions regulations on the time path of CO<sub>2</sub> emissions for the maritime shipping industry. Notably, it explores the interactions between travel speed, price, and capital turnover: Fleet fuel efficiency improves when larger ships replace existing ones, but the long lifespan of ships makes turnover slow. Regulations that reduce travel speeds lower emissions quickly, but also limit the supply and increase the price of shipping, thereby impacting shipbuilding and scrapping incentives. To quantitatively assess these mechanisms and their time horizons, I construct a dynamic model of the dry bulk shipping industry with endogenous entry, exit, and travel speed, as well as fleet heterogeneity across age and size. Using a rich dataset on the global fleet and its operation, I structurally estimate the model and use it to simulate the dynamic effects of a fuel tax, an efficiency standard that limits speeds, and an entry subsidy. I find that a fuel tax has a persistent impact, while the effect of a speed limit diminishes considerably over time due to induced ship building. Counterintuitively, rather than hastening the exit of older ships, both policies initially suppress exits, even while reducing emissions. An entry subsidy is more effective at removing old ships from service.

---

\*Email: [apeters@protonmail.com](mailto:apeters@protonmail.com). I thank Hiro Kasahara, Werner Antweiler, Henry Siu, and Patrick Baylis for their guidance and support. I am grateful to Elyse Adamic, Federico Ricca, Sebastian Gomez Cardona, Luis Alejandro Rojas-Bernal, and Mahdi Ebrahimi Kahou for many fruitful discussions and support. This paper benefited greatly from feedback from Paul Schrimpf, Katherine Wagner, Jesse Perla, Ron Yang, and participants of the environmental economics and macroeconomics discussion groups at VSE, the IO workshop at Sauder School of Business, a seminar at the University of Victoria, and the 57th Canadian Economics Association Conference. I further thank members of the Green Shipping Project at UBC and two industry experts for sharing their industry knowledge and for their assistance obtaining data. Data acquisition was funded by a generous grant from the Centre for Innovative Data in Economics Research.

# 1 Introduction

Climate change demands urgent action. To limit warming to 1.5 °C likely requires a 45% decrease in carbon dioxide emissions by 2030 (IPCC, 2022). A prevalent obstacle to rapid decarbonization is that much of capital is both difficult to upgrade and long-lived. This means that the rate of emissions reduction depends largely on the rate of turnover of existing capital, which can take multiple decades. I study this challenge in the context of maritime shipping, where upgrades are impractical and ships last over 25 years.

Maritime shipping is both indispensable to global trade and a significant contributor to global CO<sub>2</sub> emissions. The industry’s importance has recently been highlighted by both the temporary blockage of the Suez Canal and widespread COVID-related port slowdowns, which snarled global supply chains. Such severe impacts are perhaps unsurprising in light of the fact that ships carry around 70% (by value) of all traded goods (UNCTAD, 2017). At the same time, shipping contributes over a quarter of trade-related CO<sub>2</sub> emissions (Shapiro, 2016), or roughly 3% of global emissions (Faber et al., 2020), placing it on par with the total emissions of Germany. The industry has targeted net-zero emissions by 2050, and has recently imposed a fuel efficiency standard that effectively limits the speeds of existing ships as a step toward this goal.

I evaluate the effect of potential CO<sub>2</sub> emissions policies on the time path of maritime shipping emissions, focusing on the role of capital turnover. For this purpose, I construct and structurally estimate a dynamic model of emissions from the dry bulk shipping industry that incorporates endogeneity in both ship operation and ship entry and exit. I conduct counterfactual experiments to compare the dynamic effects of the policies to a baseline business-as-usual scenario and find three key results: (1) A speed limit loses efficacy over time due to the induced entry and exit response, (2) both a fuel tax and a speed limit initially reduce scrapping of older ships by reducing the supply of shipping, and (3) an entry subsidy is effective at accelerating capital turnover, but does not reduce emissions.

A key contribution of this paper is to account for the endogeneity of ship building and scrapping decisions in evaluating emissions policies. Almost all of the existing literature and industry analyses either take entry and exit rates as exogenous, based on expert estimates (e.g., Faber et al., 2020; Parry et al., 2022), or assume seamless replacement of capacity that is removed by intensive margin adjustments, especially travel speed (e.g., Rutherford et al., 2020).<sup>1</sup> These approaches either preclude counterfactual analyses regarding entry, or assume away an important channel driving both the level and dynamics of emissions. My model captures the fact that policies directed at the intensive margin of speed induce not only a short-run response along that margin, but also longer-run responses along the extensive margin of entry and exit. Further, taking a structural approach allows me to conduct counterfactual analyses and quantify the result of the dynamic *interaction* between the two margins.

In modeling the extensive entry/exit margin, I employ the game theoretic approach developed by Ericson and Pakes (1995) and Doraszelski and Satterthwaite (2010), specifically building on

---

<sup>1</sup>One exception to this is Sheng et al. (2018) who apply a calibrated computable general equilibrium model to study a fuel tax.

Kalouptsidi (2014). With respect to the latter, I add the dimension of ship size, which is important for emissions due to physical economies of scale in fuel efficiency. Correspondingly, I incorporate greater detail in modeling fuel efficiency and I estimate key parameters from data on ship characteristics, operation, and fuel consumption. Additionally, I allow for a richer exit process, with exit permitted over a range of ship ages.

A common factor influencing technology transitions is the persistence of technology choices made at the time of investment, commonly known as ‘lock-in.’ This a key feature in a wide range of industries, including manufacturing, electricity generation, housing, and aviation (Foxon, 2013; Hawkins-Pierot et al., 2023). Lock-in is particularly striking in maritime shipping, as ships cost tens of millions of dollars and operate for multiple decades. Moreover, there are currently few options for upgrading ships to reduce their CO<sub>2</sub> emission intensities,<sup>2</sup> as the sheer size of ship engines makes replacement impractical.<sup>3</sup>

I focus on two key implications of lock-in: First, the energy *efficiency* of the capital base adjusts through retiring old equipment and building new equipment. Following the shipping investment literature, I model these as firm entry and exit.<sup>4</sup> Aggregate emission intensity improvements therefore occur solely through firm entry and exit, which naturally occur slowly. Crucially however, these rates are determined endogenously and are influenced by prices and regulations.

Second, lock-in creates significant heterogeneity, because entering capital is not the same as retiring capital. Equipment may wear or age so that otherwise identical new equipment is more efficient. Additionally, new capital may differ across various technological dimensions. In shipping, a single fuel type is dominant, so technology choices that impact emission intensity consist almost entirely of choices regarding fuel efficiency.<sup>5</sup> I decompose fuel efficiency improvements in recent decades using novel ship-level data and find that increasing ship size has been the predominant source of improvements in recent decades. All other technological improvements contributed the equivalent of a 1.2% increase in ship size per year *on average*, with virtually no improvements for two decades prior to 2013. Due to physical scaling laws, larger ships consume considerably less fuel per quantity of goods transported, and ship sizes are primarily constrained by infrastructure rather than the technology frontier. Consequently, I focus on heterogeneity over age and size, and abstract from technological innovations in the model.

To estimate the model, I compile a rich dataset of the dry bulk market and construct a monthly record of the global fleet, entry, and exit, spanning 2010 to 2021. My estimation strategy follows the two-step estimation procedure of Kalouptsidi (2014), which centers around non-parametric estimation of the value function using second-hand ship sales. I adapt the entry and exit processes employed by Kalouptsidi to include the additional dimension of ship size and I allow for and estimate

---

<sup>2</sup>Emission intensity is the amount of emission per useful work or output.

<sup>3</sup>Minor improvements to the rotor and hull are possible, but provide marginal gains. Recently, wind propulsion technology has been developed that can be installed on existing ships, but gains are limited, it remains expensive, and it is therefore not widespread.

<sup>4</sup>Accordingly, I will use the terms ‘ship’ and ‘firm’ interchangeably throughout.

<sup>5</sup>Alternative fuel ships are starting to be built, but the lower energy density of alternative fuels decreases cargo space. Currently, these are primarily ferries and container ships, which operate on fixed routes, because fuelling infrastructure is costly and will there take time to become sufficiently widespread.

exit at a range of ages. I estimate the quantitative effect of size and age on fuel efficiency outside of the model, using a combination of data on ship characteristics, movements, and reported emissions.

Using my estimated model, I simulate the dynamic effects of a fuel tax, a speed limit, and an entry subsidy in comparison with a scenario with no emissions policy. A fuel tax is typically considered the first-best policy as it directly prices the emissions externality. It is useful to consider both as a benchmark and as a policy that has been proposed in the past by groups of nations and firms. The speed limit, while a potential policy itself, also approximates a recently implemented energy efficiency standard that effectively limits the maximum power of non-compliant ships. I explore an entry subsidy as a method to accelerate capital turnover that is politically feasible, as it would not require universal consensus for implementation and similar industrial policy is used by many governments with the goal of promoting economic growth.

I simulate the effects of these policies over the course of a decade, thereby highlighting the differences in mid-run dynamic impacts. Strikingly, I find that the emissions effect of a speed limitation is reduced by roughly half over the course of ten years. Slowing ships reduces the effective supply, which in turn promotes entry and dampens exit. Over time, this causes the supply of ships to grow, thereby decreasing the price of shipping. At a lower price, ships optimally travel at a lower speed, and the speed limit therefore becomes less binding. In contrast, the effect of a fuel tax is persistent in both reducing emissions and raising price. After the initial drop, emissions continue to drop very slightly as a result of the larger ships entering and the maintained downward pressure on speeds.

Additionally, I find that both the fuel tax and speed limit suppress the exit of old ships in the short-run. This is somewhat counterintuitive, given that efficiency standards often aim to hasten the replacement of old, inefficient capital. In this context however, the strong intensive margin adjustment, namely lower speeds, causes a price increase that increases the profitability of ships, even older ones. I find that an entry subsidy is more effective at hastening capital turnover, as it increases both entry and exit.

Lastly, with a fuel tax level of 20%, I find an immediate aggregate emissions reduction of 15.8% and a mere 0.1 percentage point additional reduction over the following ten years, with a similarly meager increase in shipping quantity. The corresponding speed limit is less effective than the tax, and an entry subsidy increases the quantity of shipping but does not reduce emissions. With each of these policies, attaining net-zero within another 20 years would require a drastic acceleration of emissions reductions.

The maritime shipping industry inhabits a complex policy space, for which my findings have interesting implications. First, the result that a fuel tax causes persistently higher shipping prices rationalizes its opposition from trade-dependent nations. Given that shipping emissions lie outside of national emissions commitments and policy implementation relies on international consensus, this poses a significant challenge. The speed limit, which resembles a policy that *has* been implemented, achieves emissions reductions at a lower price, but largely because these reductions diminish over time. Finally, the emissions reductions from these policies that are attributable to capital turnover

are small. This suggests that current policies will not create a significant number of stranded assets. However, to achieve sufficient reductions, it is necessary to implement a policy such as an entry subsidy that aggressively incentivizes both the retirement of old ships and the construction of new ships that emit at significantly lower rates than those using current fuel technology.

This paper contributes to three strands of literature. The first concerns shipping emissions, for which there exists industry calculations and projections of aggregate emissions that do not incorporate equilibrium effects (Faber et al., 2020; Jalkanen et al., 2009; Johansson et al., 2017; Olmer et al., 2017). One study that allows for some equilibrium effects is Sheng et al. (2018), which adapts a multi-country, multi-sectoral computable general equilibrium model and calibrates various elasticities to examine a shipping fuel tax. The second strand of literature models the shipping industry, with a principal focus on explaining the cyclical nature of investment through uncertainty and time-to-build (Campello et al., 2021; Gkochari, 2015; Kalouptsi, 2014; Kyriakou et al., 2018) and informational frictions (Greenwood & Hanson, 2015). Of particular note is Kalouptsi (2014) whose modeling and estimation framework I adapt and extend. Additionally, Brancaccio et al. (2020) study the impacts of endogenous trade costs with a search and matching model of shipping. Finally, this paper joins a growing branch of literature that applies structural modeling to study the impact of environmental regulation on industry (e.g., Cullen and Reynolds, 2017; Fowle et al., 2016; Ryan, 2012; Scott, 2014; Toyama, 2019).<sup>6</sup> In bridging this literature, this paper is the first to study shipping emissions using a dynamic entry/exit model and structural estimation, and is unique in its focus on emissions dynamics in the presence of important intensive and extensive adjustment margins.

The remainder of this paper proceeds as follows: Section 2 provides a background on the shipping industry and Section 3 describes the data employed. In Section 4, I first present empirical analyses that drive modeling choices before describing the model. The estimation procedure and results are provided in Section 5. Section 6 presents the results of counterfactual analyses and Section 7 concludes.

## 2 Industry Background

The shipping industry is divided into sectors according to the format of goods that are transported and the corresponding type of ships that are used. The three largest sectors are container ships, dry bulk ships, and oil tankers. I focus on the dry bulk sector, given both its individual importance to emissions and the generalizability of its market structure. While container ships individually contribute the most to shipping CO<sub>2</sub> emissions at 26%, the dry bulk sector is a close second at 22% (Faber et al., 2020), and a sizeable portion of oil tankers (which constitute another 18%) operate in a similar fashion.

Dry bulk goods are typically raw materials such as iron ore, coal, and grains. The dry bulk sector is highly competitive, as shown in Figure 1. Over three quarters of firms own five ships or

---

<sup>6</sup>See Kellogg and Reguant (2021) for a review.




Figure 1: Number and fleet share by firm for the dry bulk sector.

fewer, and the largest firm owns only 2.5% of all bulk ships. It is a global market—unlike container ships, dry bulk ships are not limited to specific regions or routes.<sup>7</sup> Instead, they operate in a fashion similar to taxi cabs, known as ‘tramping’: upon completion of a trip, a ship searches for a new contract originating in the vicinity, with little regard for where the shipment is destined.<sup>8</sup> Shipping prices are quite volatile due to a combination of inelastic demand and an inflexible supply of ships in the short-run. The short-run supply response is limited to adjusting ship speed. However, on longer time scales, older ships are sold as scrap and exit the market and new ships are built and enter, with a typical time-to-build between 2–5 years (Kalouptsidi, 2014).

Sectors are further divided into sub-sectors based on the size of ships, measured in deadweight tonnes (DWT). From smallest to largest, these are Handysize, Handymax, Panamax, and Capsize. As reflected in some of these names, these size-based sub-sectors are typically imposed by infrastructure capacity such as canal sizes, port depths, and loading equipment. Due to physical scaling laws, there are strong within-ship economies of scale, i.e., the fuel consumption per mass of payload for ships decreases significantly with increasing ship size. For this reason, ship sizes have increased significantly as global trade has expanded. Figure 2 illustrates the relationship between built year and size, as well as the approximate division of size categories for all except the largest category of ships. I focus on the intermediate size category with the largest number of ships, known as Handymax. This is primarily for reasons of data quality, and I expect my results to generalize to all except perhaps the very largest, most specialized ships. Being an intermediate size category, Handymax ship sizes are not limited by the frontier of ship-building technology, but rather by infrastructure, which has been enhanced over time in response to growing trade.

The predominant fuel used in maritime shipping is heavy fuel oil. While alternative fuels are beginning to emerge, the development and spread of fueling infrastructure will take time and their

<sup>7</sup>To precisely study the emissions for sectors with fixed routes would require spatial modeling, however ship size and age are similarly important for fuel efficiency and driving entry and exit.

<sup>8</sup>Brancaccio et al. (2020) explores the nuances of this, which are not of first order importance here.

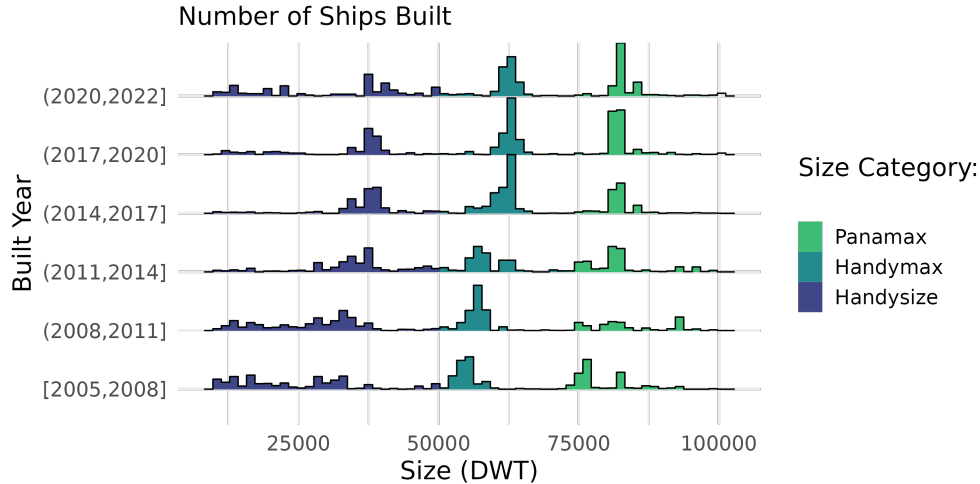


Figure 2: Number of new dry bulk ships by built year and size (excludes largest Capesize class).

use will likely be limited in the near future both geographically and to sectors operating on fixed routes.<sup>9</sup> For this reason, CO<sub>2</sub> emission intensity in the dry bulk sector is directly proportional to fuel efficiency, and I will use these terms interchangeably. Burning heavy fuel oil emits various gases and particulate matter that contribute to poor air quality and climate change, however my focus is specifically on CO<sub>2</sub> due to its contribution to global climate change.

Emissions from the shipping industry lie outside the scope of national emissions commitments and fall instead under the jurisdiction of the International Maritime Organization (IMO), which is a United Nations organization that facilitates international cooperation on safety and environmental standards in shipping. Its most recent emissions target calls for net-zero emissions by 2050, but the path to this remains uncertain. To date, the IMO has chosen to implement efficiency standards rather than a fuel tax. Despite support for a tax from a number of key stakeholders, some nations that rely heavily on low-cost trade oppose taxation due to a concern over its impact on the cost of shipping.

Three main CO<sub>2</sub> regulations are in place: In 2012, the IMO implemented minimum design efficiency standards for all new ships, with increasing stringency over time (Wan et al., 2018). In 2022, it introduced new requirements on operational efficiency and extended the design efficiency standards to existing ships. As a result of the limited technologies for upgrading, the majority of ships not meeting the design standard have opted to install power limiting devices on their engines. The result of this is akin to a speed limit, which I explore in a counterfactual scenario. Looking forward, the European Union (EU) is discussing applying some form of its existing emissions trading framework to the shipping industry, however the details are yet to be determined. This type of region-based policy is beyond the scope of the present model, however the tax counterfactual I study is equivalent to emissions trading to a first approximation.

<sup>9</sup>An extension of the current model could be used to study one important effect of an eventual fuel switch: a capacity decrease. All alternative fuels have a lower energy density and therefore take up more space on the ship, leaving less for cargo.

### 3 Data

I compile a rich dataset describing the dry bulk shipping industry from 2010 to 2021 and detailed ship operation from 2019 to 2021. This data includes a register of the global fleet, demolitions, second-hand sales, freight contracts, tracking data, and emissions reports. While much of the general industry data is commonly used in the literature, this is the first work to combine fleet data with detailed operational data for a complete picture of emissions. To link ship operation and emissions, I match fleet technical characteristics to tracking data and annual emissions reports to estimate key parameters related to fuel efficiency. Also of note, I obtain a large sample of freight contracts (important for demand estimation) from the Baltic Exchange, and am among the first to use ship-level emissions reports from the EU’s Monitoring, Reporting, and Validation (MRV) program. In this section I briefly describe the various data sources and the procedures used to compile the data used in estimations. For brevity, any data whose source is not explicitly stated was obtained from Clarksons Research. All monetary values are adjusted for inflation to 2015 levels. Further details are provided in Section A.1.

I first compute annual fuel efficiency in terms of fuel consumed per travel work<sup>10</sup> for individual ships that enter and exit the EU. This is used both for preliminary estimations that guide model choices, and for estimation of the relation of fuel efficiency with age and size used directly in the model. Fuel consumption is strongly dependent on operational speed (see Section 4.4), and therefore must be accounted for in determining a ship’s fuel efficiency. I obtain annual fuel consumption and emissions reports from the EU’s MRV emissions reporting program. This is matched by ship registration ID to a register of the global fleet from Clarksons Research that contains detailed technical characteristics as well as the ID for each ship’s automatic identification system (AIS), which reports each ship’s location and speed every few minutes. This AIS ID allows to further match to hourly AIS tracking records for the global bulk fleet from 2019 to 2021, obtained from Spire. However, because emissions are only reported for trips into and out of the EU, it is necessary to identify these trips within the raw tracking data. For this, I follow similar methodology to that developed by Brancaccio et al. (2020), Faber et al. (2020), Olmer et al. (2017), and Van et al. (2019)—details are outlined in Section A.1.<sup>11</sup>

Estimation of the demand for shipping requires the monthly quantity of shipping in tonne-miles for the Handymax sector. In order to calculate this, contract data is obtained from the Baltic Exchange, which gathers a record of contracts from a global network of ship brokers to establish industry standard price indices. While this data is extensive, it is only a sample and constitutes an unknown fraction of total contracts. Separately, I obtain a time series of *all* seaborne dry bulk trade. As shown in Figure 3, the fraction of contracts that are observed appears to be decreasing over time, as the mass of goods shipped observed in the contracts sample decreased while the total amount of trade increased steadily. In order to compute the total amount of Handymax shipping, I assume

---

<sup>10</sup>See Section 4.1 for a definition.

<sup>11</sup>To ensure accuracy, I select only those ship-year observations for which the detected distance traveled on EU trips matches the distance reported in the MRV data within  $\pm 10\%$ .






Figure 3: Monthly quantity observed from contracts sample, total sector quantity, computed observation probability, and final scaled quantity (left to right). Masses in million tonnes.

that the probability of observation is the same across sub-sectors and varies smoothly over time: I calculate probability of observation as the aggregate quantity from observed contracts for *all* size categories (sub-sectors), divided by the volume of dry bulk trade.<sup>12</sup> I then compute the total mass of goods shipped per month as the observed contract quantity divided by the probability of observation. To obtain quantity in mass-distance units, I further multiply this quantity by the average distance traveled per month by a ship.<sup>13</sup> Other data used in the demand estimation include commodity price indices obtained from UNCTAD and industrial production and GDP from the OECD.

To define the supply state and for exit regressions with heterogeneity over age and size, I construct a monthly panel of the global fleet of ships spanning 2010 to 2021. Annual fleet registers from Clarksons Research specify the ID number, name, size, month built, and other details for each ship in operation, as well as a list of those that were scrapped during the year, but not their scrap date. Scrap months were obtained from a separate weekly publication. Fleet registers were not available for some intermediate years, but given the long lifespan of ships, I am able to reconstruct the monthly fleet (Figure 4) and exits (Figure 5). Details are provided in Section A.1.

Central to the estimation strategy, a sample of 1336 reported second-hand ship prices is compiled from a weekly publication from Clarksons Research, with an average of 9.2 observations per month and an average price of \$11.8 million in 2015 USD. Finally, shipping price is calculated as the average monthly charter rate in the sample of observed contracts, with a sample average of \$7.86/tonne-month in 2015 USD.

<sup>12</sup>Monthly values are interpolated from an annual time series, as monthly data begins in 2015.

<sup>13</sup>This is calculated using the median ship speed over the sample period from a time series of average ship speed.

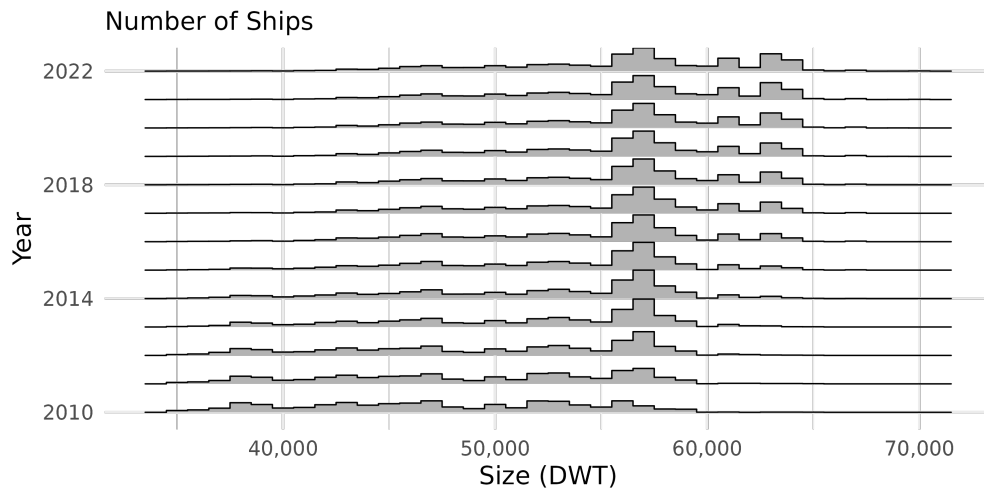


Figure 4: Yearly (January) evolution of the fleet size distribution.




Figure 5: Monthly ship scrapping by age group.

## 4 Model

I develop a dynamic model of the global bulk shipping industry that incorporates a key dimension of ship heterogeneity with regard to emissions: ship size. The model belongs to a large literature of dynamic games stemming from Ericson and Pakes (1995) and Doraszelski and Satterthwaite (2010). My model builds on that of Kalouptsi (2014), however I discard the endogenous time-to-build component, adding the size dimension, a richer exit process, and more detailed structure to the payoff function, which is necessary to accurately model fuel consumption. I begin by discussing key modeling choices regarding fuel efficiency and presenting empirical evidence that guides them before proceeding to the model description.

### 4.1 Modeling Fuel Efficiency

In order to construct a model of emissions that is both accurate and tractable, I first identify and characterize the essential factors that drive fuel efficiency. This includes both technical factors related to the ship and operational factors related to its use. *Technical efficiency* is primarily determined at time of build by a ship’s design. This includes a ship’s size, engine, propeller, etc., which I will discuss below. For reasons discussed previously, I model technical efficiency as locked-in for the life of the ship, *except* for an age-related deterioration, as fuel efficiency tends to worsen with age due to mechanical wear, etc. *Operational efficiency* includes speed of travel, loading, and routing, which determines the winds and currents it encounters. Of these, I model only speed. Fuel use is roughly cubic in speed and thereby constitutes a crucial short-run adjustment margin for supply. Ship loading has a relatively smaller effect on fuel use and varies little in practice.<sup>14</sup> Routing, and therefore weather, are largely random due to the tramping structure of the industry, so I do not include them explicitly.<sup>15</sup> Ship operator choices determining operational efficiency are driven by both market conditions and technical efficiency.

Which factors are most important for overall technical efficiency improvements? To answer this, I look at recent efficiency improvements and compare the contribution of size to that of technical innovations. Specifically, I regress annual fuel efficiency on size and built year, where built year proxies all other sources of fuel efficiency improvements over time:

$$\log(\text{fuel efficiency}_{it}) = \beta_0 + \beta_1 \text{built year}_i + \beta_2 \log(\text{size}_i) + \varepsilon_{it}, \quad (1)$$

where  $i$  indicates the ship and  $t$  indicates the year. I define fuel efficiency as the quantity of cargo shipped per fuel consumed and assume that ships are fully loaded. For robustness, I compare three different measures of fuel consumption: (i) *rated* fuel consumption *per day*, (ii) *reported* fuel consumption per reported *distance* traveled, and (iii) reported fuel consumption per *detected travel work*. Rated fuel consumption is a design specification and therefore not representative of actual

---

<sup>14</sup>This is according to consultations with industry experts and draught levels that I observe in tracking data. Adland et al. (2018) find only 2 percentage point change in loading between price extremes for the largest ships (Capesize sub-sector).

<sup>15</sup>Randomness in period profits is loaded into random scrap value draws.

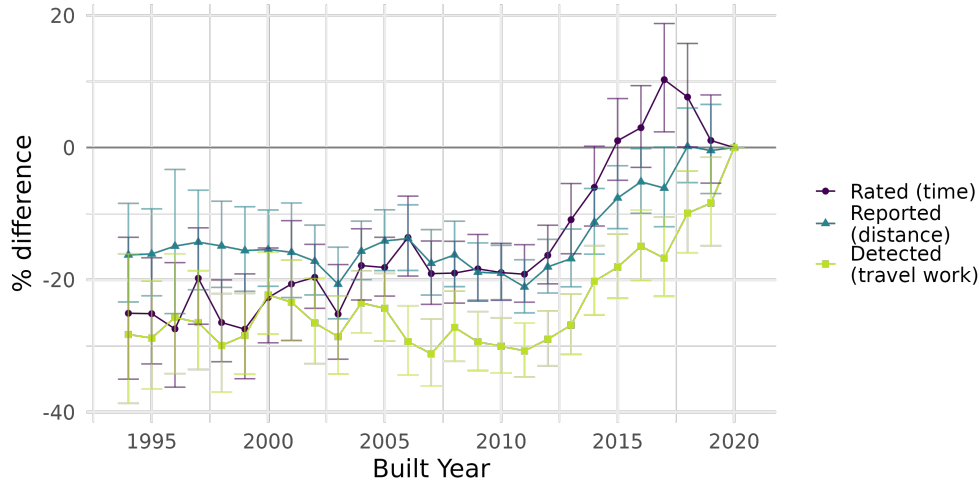


Figure 6: Built year fixed effects relative to 2020 from fuel efficiency regression controlling for size. Three alternative efficiency measures are shown. Error bars indicate 95% confidence intervals.

fuel consumed. It uses a rated speed and could reflect the use of lower powered engines that merely limit the maximum speed, rather than a true efficiency improvement. Reported fuel consumption better reflects true fuel use, but does not account for speed. To account for speed, I construct a measure of annual travel work from tracking data, detailed in Section A.2.1, which I combine with reported annual fuel consumption.<sup>16</sup>

The regression results in Table 5 suggest that increasing size is a more important source of fuel efficiency improvements than technological advances for the fleet of ships currently in use. For example, using the work-based efficiency measure, I estimate that a ship that is one year newer is equivalent to only a 1.2% increase in size, or an increase of 720 DWT for a typical 60,000 DWT ship.

Alternatively, I perform a similar regression, but replace the linear built-year term with more flexible built-year fixed effects. The fixed effects are plotted in Figure 6 and reveal that essentially all efficiency improvements corresponded to size increases from 1995 to 2013. The improvements over the subsequent years may be due to a new-build efficiency standard enacted in 2013, or potentially a lagged effect of high fuel prices a few years prior. There are two caveats to these results: First, there may be a survivor selection bias in this sample biasing the implied technical improvements downwards as less efficient ships may be more likely to be scrapped. However, very few ships under 20 years old are ever scrapped, so this would only affect fixed effects for the oldest ships.<sup>17</sup> Second, the measures using reported fuel consumption attribute any aging effects to built year.<sup>18</sup> However, this would bias the implied technical improvements upwards, further supporting my conclusion.

The relative importance of size in efficiency improvements is perhaps unsurprising in light of the maturity of shipping technology, which has not seen major disruptions for decades. Looking forward,

<sup>16</sup>This remains an imperfect measure due to a number of potential sources of measurement error.

<sup>17</sup>In robustness checks, I find that excluding the oldest ships from the regression with the linear trend does not qualitatively change the result.

<sup>18</sup>The panel is too short to distinguish the effects of aging from vintage.

most improvements, such as propeller modifications or hull coatings also promise relatively minor improvements. In light of this, I focus on size as the source of efficiency improvements in the model, rather than the other sources of technological change.

## 4.2 Environment

A firm consists of a single ship,<sup>19</sup> which operates in a discrete time environment with an infinite horizon. Firms are either incumbents or potential entrants. Incumbents face both operational (intratemporal) and exit (intertemporal) decisions, while potential entrants face only an entry decision. Entry occurs with a delay and potential entrants that do not enter perish. An incumbent ship is fully characterized by its:

1. age  $j \in \{0, \dots, J\} = \mathcal{G} : |\mathcal{G}| = J + 1$  and
2. size  $q \in \{q_1, \dots, q_K\} = \mathcal{Q} \subset \mathbb{R} : |\mathcal{Q}| = K$ .

The aggregate state of the market is characterized by:

1. demand for shipping services as captured by the intercept of the inverse demand curve  $d_t \in \mathcal{D} \subset \mathbb{R}$ , and
2. supply capacity as captured by the distribution of existing ships jointly over size and age  $\mathbf{n}_t \in \mathcal{N} \subset \mathbb{N}_0^{J+1 \times K}$ , with element  $n_t^{j,q}$  denoting the number of ships of age  $j$  and size  $q$  in period  $t$ .

The demand state variable  $d_t$  follows an exogenous first-order Markov process, while the fleet distribution evolves according to a combination of aging and endogenous entry and exit.

The problem of an incumbent is as follows. At the beginning of the period, firms observe the aggregate state  $\mathbf{x}_t = (d_t, \mathbf{n}_t)$ . Based upon this information, each ship chooses its operational speed, which will determine its per-period payoffs. It then draws a private per-size<sup>20</sup> scrap value  $\phi_t$ , which is independent and identically distributed (i.i.d.) across time and firms, from a distribution  $F_\phi$ . Based on this value, the firm decides to continue operating or to exit and receive scrap value  $q\phi_t$ . The exit decision is implemented at the end of the period, and the firm receives its per-period payoff (and scrap value if exiting).

In each period there is a large pool of identical potential entrants for each ship size  $q$  that may enter the market freely. Each observes the size-specific common entry cost  $\kappa_q(d_t, \mathbf{n}_t)$  and chooses whether or not to enter based upon their expected future profit. Entry decisions are realized with a one-period lag and are made simultaneously with incumbents' exit decisions.

---

<sup>19</sup>I will use the terms firm and ship interchangeably.

<sup>20</sup>The scrap value of a ship is comprised mainly of the value of the steel hull, which is proportional to its size.

### 4.3 Dynamic Equilibrium

Incumbents maximize the intratemporal speed choice and the intertemporal exit decision independently. I therefore define the dynamic equilibrium before elaborating on the details of the intratemporal choice. As such, let the optimal per-period payoffs be denoted simply as  $\pi^*(j, q; d_t, \mathbf{n}_t)$ . The value function  $V(\cdot)$  of a ship at the beginning of a period before it observes its scrap value can be written:

$$V(j, q; d_t, \mathbf{n}_t) = \pi^*(j, q; d_t, \mathbf{n}_t) + \beta E_{\phi_{it}} \max\{q\phi_{it}, VC(j, q, d_t, \mathbf{n}_t)\}, \quad (2)$$

with continuation value

$$VC(j, q; d_t, \mathbf{n}_t) \equiv E_{d_{t+1}, \mathbf{n}_{t+1}}[V(j+1, q; d_{t+1}, \mathbf{n}_{t+1}) | j, q, d_t, \mathbf{n}_t]. \quad (3)$$

Randomness in the fleet distribution stems from the number of entrants of each size in each period  $M_t^q$  and the number of exiting firms of each age and size in each period,  $Z_t^{j,q}$ . An incumbent firm will choose to exit if its scrap value  $q\phi_t$  is greater than its continuation value, which occurs with probability

$$\zeta(j, q; d_t, \mathbf{n}_t) \equiv Pr(q\phi_{it} > VC(j, q; d_t, \mathbf{n}_t)) = 1 - F_\phi \left( \frac{VC(j, q; d_t, \mathbf{n}_t)}{q} \right). \quad (4)$$

$Z_t^{j,q}$  thus follows a binomial distribution.

For potential entrants, the value of entry is:

$$VE(q, d_t, \mathbf{n}_t) \equiv \beta E_{d_{t+1}, \mathbf{n}_{t+1}}[V(0, q, d_{t+1}, \mathbf{n}_{t+1}) | d_t, \mathbf{n}_t]. \quad (5)$$

The expectation is due to the lag on ship delivery, and the ship is age zero upon delivery. Following Weintraub et al. (2008), I assume that each large pool of potential entrants plays a symmetric mixed entry strategy. They show that because the binomial distribution approaches a Poisson distribution in the limit, the number of entrants is well approximated by a Poisson distribution. The mean of each distribution,  $\lambda_q(d_t, \mathbf{n}_t)$ , is determined endogenously via the free entry condition. Because the value of entry is non-increasing in the number of entrants it must be that either the value of entry is equal to the entry cost or the entry value is below the entry cost and there is no entry.

I define a symmetric Markov Perfect Equilibrium in which all firms use a common stationary exit strategy. The cut-off exit rule described above is equivalent to a mixed exit strategy  $\zeta$  mapping every potential state to an exit probability (Doraszelski & Satterthwaite, 2010). With entry rates  $\lambda \equiv \{\lambda_q\}_{q \in \mathcal{Q}}$  defined for each state, we can rewrite the incumbent's value function in (2) in terms of mixed strategies as:

$$V(j, q, d, \mathbf{n}; \zeta', \zeta, \lambda) = \pi^*(j, q, d, \mathbf{n}) + \zeta' \beta E [q\phi | q\phi > VC(j, q, d, \mathbf{n}; \zeta, \lambda)] \\ + (1 - \zeta') \beta VC(j, q, d, \mathbf{n}; \zeta, \lambda), \quad (6)$$

where  $\zeta'$  is the firm's own strategy and  $\zeta$  is the common strategy of all other firms.

A symmetric equilibrium consists of an exit strategy  $\zeta$  and entry rate functions  $\lambda$  such that:

(i) firms choose exit strategies optimally in each state:

$$V(j, q, d, \mathbf{n}; \zeta, \zeta, \lambda) = \sup_{\zeta' \in [0,1]^{(j,q,d,\mathbf{n})}} V(j, q, d, \mathbf{n}; \zeta', \zeta, \lambda),$$

$$\forall j \in \mathcal{J}, q \in \mathcal{Q}, d \in \mathcal{D}, \mathbf{n} \in \mathcal{N}; \text{ and}$$

(ii) the free entry conditions are satisfied:

$$VE(q, d, \mathbf{n}) \leq \kappa_q(d, \mathbf{n}), \quad \forall q \in \mathcal{Q}, d \in \mathcal{D}, \mathbf{n} \in \mathcal{N}, \quad (7)$$

with equality if  $\lambda_q(d, \mathbf{n}) > 0$ .

A symmetric equilibrium exists, as shown by Doraszelski and Satterthwaite (2010), under the assumptions in this section on the scrap value density and the finiteness of the state space, as well as the following specification of the profit function (it is bounded) and the continuity of (6) in strategies.

#### 4.4 Static Equilibrium

Next I turn to the intratemporal problem and derive equilibrium price and profit. Within the dry bulk industry, there are three broad contract types: In a time charter, a charterer with goods to transport charters a ship and assumes responsibility for the operation of it for an extended period of time. In a voyage charter, the shipowner is responsible for operating costs to deliver a load for a charterer. The trip-time charter is a hybrid of these in which the charterer charters a ship for a single voyage and pays for voyage-specific fees including fuel and canal and port fees. The owner remains responsible for the ship maintenance and crew.

In my data, the trip-time charter is by far the most commonly observed contract type for Handymax ships, and I therefore model contracts as trip-time charters. Given the highly un-concentrated nature of the dry bulk industry, I assume perfect competition in the charter market.<sup>21</sup> The problem of a charterer desiring to ship a load a distance  $x$  is therefore to choose speed  $s$  in order to minimize cost for the trip:<sup>22</sup>

$$C^*(j, q, x, p, c) = \min_s pq \frac{x}{s} + c\eta(j, q)xs^2 + F_{j,q}^c, \quad (8)$$

In order, the additive terms of the objective function represent the rental cost, fuel cost, and fixed trip costs such as port and canal fees. The charter rate  $p$ , which I will also refer to as the time-based shipping price, is defined per-mass $\times$ time,<sup>23</sup>  $c$  is the per-unit fuel price,<sup>24</sup> and  $\eta(j, q)$  is

<sup>21</sup>See discussion in Section 2. Furthermore, Kalouptsi (2014) finds competitive pricing when allowing for Cournot competition.

<sup>22</sup>This specification assumes the charterer is not responsible for any ‘ballast’ leg in which the ship travels empty.

<sup>23</sup>I assume the charter rate scales linearly with size, which approximates the data.

<sup>24</sup>In order to maintain a tractable state space, I abstract from stochastic fluctuations in the fuel price. In estimation, variations in fuel price will be incorporated into the demand state variable.

the age and size dependent fuel consumption of the ship. The squared speed term is a result of the physical force of drag on the ship hull.<sup>25</sup>  $F_{j,q}^c$  are fixed trip-related fees.

I model fuel consumption as a power  $\alpha$  of ship size according to physical scaling laws,<sup>26</sup> and exponentially increasing with age, with parameter  $\delta$ :<sup>27</sup>

$$\eta(j, q) = \alpha_0 q^\alpha e^{\delta j}. \quad (9)$$

The optimal speed  $s^*$  is derived from the first order condition, and in the absence of constraints on speed is given by:

$$s_{j,q}^* = \left( \frac{pq^{1-\alpha}}{2c\alpha_0 e^{\delta j}} \right)^{1/3} \quad (10)$$

Ceteris paribus, optimal speed is increasing in the charter rate and size ( $\alpha < 1$ ), and decreasing in fuel cost and age ( $\delta > 0$ ). The corresponding elasticities of speed with respect to charter rate and fuel cost are, respectively:

$$\frac{\partial s^*/s^*}{\partial p/p} = \frac{1}{3},$$

$$\frac{\partial s^*/s^*}{\partial c/c} = -\frac{1}{3}.$$

Given the optimal speed, the equilibrium price follows from aggregating individual choices. I assume a log-linear demand system so that prices are given by the inverse demand curve:

$$p(d, Q) = \bar{x} e^d Q^\gamma, \quad (11)$$

where  $1/\gamma$  is the price elasticity of demand and  $Q$  is the total quantity (mass $\times$ distance) of shipping per period, obtained by aggregating over the fleet. The average distance travelled per-period  $\bar{x}$  is required to convert between the quantity in tonne-miles and the time-based price. With speed defined in terms of distance *per-period*, the average distance is equal to the average speed  $\bar{s}$ , given by:

$$\bar{s} = \frac{1}{n} \sum_{j',q'} n^{j',q'} s_{j',q'}^*, \quad (12)$$

where  $n$  is the total number of ships in the fleet.

The quantity of shipping supplied by a single ship is simply the amount of cargo transported multiplied by the distance traveled per period. The aggregate quantity of shipping supplied is therefore:

---

<sup>25</sup>This specification is commonly used in the literature (e.g. Kalouptsi (2014) and Ronen (1982)) and derives from a well-established engineering relationship for approximating fuel consumption around the design speed (MAN, 2018).

<sup>26</sup>This an approximation of engineering relationships (see e.g., MAN, 2018) that is widely used, for example by the IMO in defining the reference curve for fuel efficiency regulations (MEPC, 2012)

<sup>27</sup>Exponential aging is a reasonable assumption used in many contexts and is implied by regression specifications of a log dependent variable on age. As the aging parameter is small in practice, a linear function would provide similar results.



$$Q = \psi \sum_{j',q'} n_{j',q'} q s_{j',q'}^*, \quad (13)$$

where  $\psi$  accounts for the fraction of time a ship spends either in port or traveling without a load.<sup>28</sup>

Finally, the profit of the shipowner, who receives the charter rate and pays a per-period fixed cost  $F_{j,q}$  (including crew, insurance, maintenance, etc.) is given by:<sup>29</sup>

$$\pi^*(j, q; d, \mathbf{n}) = pq - F_{j,q}. \quad (14)$$

In the case where there are no constraints on speed, an analytical solution is available for the equilibrium price. The corresponding equations are used in estimation and are helpful to illustrate the intuition of the mechanisms. Denoting values corresponding to this case with a tilde, the supply curve is obtained from (10) and (13):

$$\tilde{Q} = \left( \frac{\tilde{p}}{2c\alpha_0} \right)^{1/3} \Sigma_Q(\mathbf{n}),$$

where I define  $\Sigma_Q(\mathbf{n}) \equiv \psi \sum_{j',q'} n_{j',q'} \left( \frac{q'^{4-\alpha}}{e^{\delta j'}} \right)^{1/3}$  for brevity. This term can loosely be interpreted as the aggregate supply capability of the fleet. Similarly, optimal speeds can be substituted into (12). Setting quantity supplied equal to quantity demanded in the inverse demand equation from (11), the equilibrium price is:

$$\tilde{p} = e^{\frac{3d}{2-\gamma}} (2c\alpha_0)^{\frac{\gamma+1}{\gamma-2}} \Sigma_s(\mathbf{n})^{\frac{3}{2-\gamma}} \left( \frac{1}{\Sigma_Q(\mathbf{n})} \right)^{\frac{3\gamma}{\gamma-2}}, \quad (15)$$

where  $\Sigma_s(\mathbf{n}) \equiv \frac{1}{n} \sum_{j',q'} n_{j',q'} \left( \frac{q'^{1-\alpha}}{e^{\delta j'}} \right)^{\frac{1}{3}}$  is related to the average ship fuel efficiency. This equation is written so that the exponents are positive with the demand elasticity that I estimate later. We see that the time-based charter rate will be increasing in demand and fuel price, and decreasing in the total shipping capability of the fleet. It is also increasing in the average ship efficiency, which can be rationalized by the fact that a more efficient fleet of ships provides travels faster and therefore provides more shipping services (distance and/or quantity) within a period.

The profit of the ship owner with no constraints on speed is then:

$$\widetilde{\pi}^*(j, q; d, \mathbf{n}) = e^{\frac{3d}{2-\gamma}} (2c\alpha_0)^{\frac{\gamma+1}{\gamma-2}} \Sigma_s(\mathbf{n})^{\frac{3}{2-\gamma}} \left( \frac{1}{\Sigma_Q(\mathbf{n})} \right)^{\frac{3\gamma}{\gamma-2}} q - F_{j,q}. \quad (16)$$

Given fixed costs  $F_{j,q}$  that are increasing in both size and age, profit will be strictly decreasing in age, but in general will have a maximum at an optimal size that depends on the charter rate.

Finally, it may be more intuitive to consider a quantity-based measure of the price of shipping

---

<sup>28</sup>Here, I assume  $\psi$  is constant across all ships and does not depend on the aggregate state. This assumption could potentially be relaxed in future work.

<sup>29</sup>This specification implicitly assumes ships are always under hire. This could be relaxed by multiplying profit by a demand-varying term.

$\tilde{p}_Q \equiv \frac{\tilde{p}}{\tilde{x}}$  in dollars per tonne-mile. The equilibrium equation for this is:

$$\tilde{p}_Q = e^{\frac{2d}{2-\gamma}} (2c\alpha_0)^{\frac{\gamma}{\gamma-2}} \left( \frac{1}{\Sigma_s(\mathbf{n})} \right)^{\frac{\gamma}{\gamma-2}} \left( \frac{1}{\Sigma_Q(\mathbf{n})} \right)^{\frac{2\gamma}{\gamma-2}}. \quad (17)$$

Again, this is written such that all exponents are positive with the estimated demand elasticity. As with the time-based charter rate, this is increasing in demand and fuel price and decreasing in the total shipping capability. However, it is decreasing in average fuel efficiency, so that a more efficient fleet results in a lower shipping price *per quantity*.

## 5 Model Estimation

My empirical strategy is based on that of Kalouptsi (2014), in which the value function is estimated non-parametrically directly from resale prices. Since firms are homogeneous and the market is competitive, the resale price must equal the value of owning the ship. Value functions and state transitions are used to recover the model primitives: the entry costs, the scrap value distribution, and the variable and fixed costs. The additional dimension of ship size (compared to Kalouptsi, 2014) adds a significant degree of complexity, especially to the entry and exit estimation. As well, I allow for and estimate exit for a range of ship ages. Fuel efficiency parameters are estimated exogenously from the model. I focus on the Handymax sub-sector, but the results likely generalize well to the other sizes of dry bulk ships.

The estimation procedure consists of three steps. First I estimate exogenous objects, namely the price elasticity of shipping demand, the demand state variable, and its transition process. Next, I estimate two equilibrium objects: the value function and the transition matrix. Finally, in the third step the model primitives are obtained. Optimal profits are calculated using the Bellman equations. The profit parameters (variable and fixed costs) are obtained by fitting a parametric profit function to the optimal profits.

### 5.1 Exogenous Object Estimation

Demand for shipping is taken to be exogenous. Given the global nature of the bulk shipping industry, estimation of the demand system relies on time-series variation at the monthly level, over a 12-year time frame.

#### Demand

As discussed above, I assume log-linear demand as specified in (11), for which the empirical counterpart is:

$$\log \left( \frac{p_t}{\tilde{x}_t} \right) = \gamma_0 + \gamma_1 \log(\mathbf{H}_t) + \gamma \log(Q_t) + \varepsilon_t,$$

where  $p_t$  is the monthly average charter rate in \$/tonne-month and  $\bar{x}_t$  is the average distance travelled by a ship in a month, calculated from the observed average speed.  $\mathbf{H}_t$  is a vector of demand shifters and  $Q_t$  is the combined amount of shipping realized per month in tonne-miles. I do not directly observe  $Q_t$ , but construct it as described in Section 3.<sup>30</sup> I instrument for  $Q_t$  with supply shifters: total fleet capacity and the standard deviation of ship ages. This serves two purposes: to account for endogeneity and measurement error in  $Q_t$ . The underlying identification assumption is that since the fleet is pre-determined, it is uncorrelated with contemporaneous demand shocks.<sup>31</sup> Demand shifters include the world volume of industrial production (WIP),<sup>32</sup> global price indices for food, agricultural raw materials, and minerals, ores and metals, and the global average price of ship fuel. I also include the fleet utilization of the previous (smaller) and next (larger) size categories, calculated as the mass quantity of realized contracts divided by the existing fleet capacity. The demand regression results are presented in Table 1 and the implied elasticity of demand is -0.85, suggesting that demand is rather inelastic to price. This corresponds with the fact that dry bulk ships carry goods that are either inputs to or are themselves staples. While there are no directly comparable estimates in the literature, the most similar estimates are commodity-specific, volume-based elasticities from Cariou et al. (2023). They estimate elasticities of -0.785 and -0.944 for grain and soybeans, respectively.

The demand state variable is defined as the intercept of the inverse demand curve, as given in (18), and the estimated series is plotted in Figure 7.

$$d_t = \hat{\gamma}_0 + \hat{\gamma}_1 \log(\mathbf{H}_t) + \hat{\epsilon}_t. \quad (18)$$

Lastly, I model the evolution of demand as a first-order autoregressive process with a normally distributed error.<sup>33</sup>

$$d_t = \rho d_{t-1} + \rho_0 + \epsilon_t \quad (19)$$

Linear regression estimates  $\rho = 0.863$  with a standard error of 0.045, suggesting that the demand state variable is stationary, as  $\rho < 1$ . While one might assume demand has grown over time, the observed fleet growth has been accompanied by a considerable decrease in freight price. The standard error of the residual is 0.056, or roughly 6% of the estimated steady-state demand level, which implies considerable volatility at a monthly frequency.

---

<sup>30</sup>I verify robustness to this process by varying the smoothing amount and the smoothing function used in calculating the quantity of shipping.

<sup>31</sup>While similar instruments were used by Kalouptsidi (2014), one challenge to the instrument validity is if demand is auto-correlated and past demand affects past entry and exit, which determine the present fleet. This concern is slightly, but not entirely, mitigated by the long lifespan of ships, combined with a low persistence and high volatility of demand. I have experimented with alternative instruments, such as weather events, but they did not have sufficient explanatory power.

<sup>32</sup>I calculate WIP as a GDP-weighted average of industrial production for the OECD, Russia, Brazil, and India.

<sup>33</sup>In contrast, Kalouptsidi (2014) uses a heavy-tailed distribution. The Jarque Bera test fails to reject normality of the residuals in my estimation and a fat-tailed distribution does not change my result significantly.

Table 1: Demand regression results.

	(1) 1st Stage	(2) 2nd Stage
	$\log(\text{Quantity})$	$\log(\text{Rate})$
log(SD Age)	0.551 (0.545)	
log(Fleet Capacity)	1.974*** (0.393)	
log(WIP)	-0.057 (0.572)	2.269*** (0.581)
log(Food Prices)	-0.865* (0.432)	-0.504 (0.532)
log(Ag. Raw Materials Prices)	1.451*** (0.374)	-0.492 (0.296)
log(Minerals, Ores, Metals Prices)	-0.355 (0.202)	1.591*** (0.265)
log(Ship Fuel Price)	-0.202 (0.121)	-0.132 (0.161)
log(Previous Size Utilization)	0.131** (0.046)	0.255** (0.097)
log(Next Size Utilization)	0.606*** (0.110)	0.555** (0.212)
log(Quantity)		-1.172*** (0.227)
Constant	-17.145 (9.495)	-4.421 (4.599)
Observations	144	144
F-test (IV only)	30.2	64.5

Signif. codes: \*\*\*=0.001, \*\*=0.01, \*=0.05, .=0.10. Newey-West (L=3) standard errors in parentheses. Fleet capacity and the standard deviation of ship ages instrument for Quantity.




Figure 7: Demand state variable.

## 5.2 Equilibrium Object Estimation

Equilibrium objects include the value function and the transition probability matrix. Both are defined on a discrete grid of states. I first describe the state space discretization before proceeding to the value function estimation. Computation of transition probabilities is straightforward, but to obtain a tractable state space, the aggregate state variable is aggregated, requiring a simulation procedure to map to the aggregated state space.

### State Space Discretization and Aggregation

Individual state variables age  $j$  and size  $q$  are defined on evenly spaced grids. Age is defined monthly, ranging from zero to a maximum absorbing age of 360 months (30 years). Size is divided into 6 bins, ranging from 38,000 DWT to 68,000 DWT. Taken together, these form 2160 individual states.

The aggregate state is composed of the demand state previously estimated and the supply state, i.e., the distribution of the existing fleet over age and size. In order to best represent the most frequently observed demand states  $d_t$ , I use K-means clustering on the observed demand states to define  $d$ , consisting of ten demand levels. This process results in a higher density of levels at intermediate demand, with low and high demand being more sparsely represented.

The supply state variable is naturally high-dimensional, consisting of the count of ships of every age and size. Defined on the same bins for age and size as the individual state variable, this is given by:

$$\mathbf{n}_t = \begin{bmatrix} n_t^{0,q_1} & \dots & n_t^{0,q_K} \\ n_t^{1,q_1} & \dots & n_t^{1,q_K} \\ \vdots & \ddots & \vdots \\ n_t^{J,q_1} & \dots & n_t^{J,q_K} \end{bmatrix} \in \mathbb{R}^{((J+1) \times K)}. \quad (20)$$

However to reduce dimensionality, I define an aggregated version of the supply variable  $\mathbf{N}$  as follows. I first define three age bins  $j \in \{[0, 119], [120, 239], [240, 360]\}$ <sup>34</sup> and three size bins  $q \in \{\{q_1, q_2\}, \{q_3, q_4\}, \{q_5, q_6\}\}$ , resulting in a nine dimensional supply variable of much lower dimension than the individual state. To discretize the continuous and infinite set of values that each variable can take on, I use multidimensional K-means clustering jointly on the elements of  $\mathbf{N}_t$  to select 30 levels that optimally cover the observed supply states. The disadvantage of this approach is that the resulting supply state variable levels are difficult to interpret and cannot be ranked, however alternative methods require more levels and also do not permit ranking.<sup>35</sup>

Finally, the state is taken as the Cartesian product of age, size, demand, and supply. For brevity, I denote the high-dimensional aggregate state by  $\mathbf{x}_t = (d_t, \mathbf{n}_t)$  and the low-dimensional version by  $\mathbf{X}_t = (d_t, \mathbf{N}_t) \in \mathbb{R}^{10}$ . This procedure results in 300 aggregate states  $\mathbf{X}$ , and 432,200 overall states  $(j, q, \mathbf{X})$ , which is large but tractable.

<sup>34</sup>Any ships over 360 months are attributed this maximum age.

<sup>35</sup>For example, defining levels on each of the bins separately and taking the Cartesian product of these levels would result in an unmanageably large supply state space.

## Value Function

The next step is to estimate the value function on this grid of states. The assumptions that shipowners are identical and the resale market is competitive are key to the estimation strategy. These imply that a seller is willing to sell a ship second hand only at a price  $p^{SH} \geq V(j, q; d, \mathbf{n})$ , and a buyer will buy only if  $p^{SH} \leq V(j, q; d, \mathbf{n})$ . Therefore, observed resale prices can be written as:

$$p^{SH} = V(j, q; d, \mathbf{n}) + \varepsilon,$$

where  $\varepsilon$  represents measurement error. It follows that the value function is the expectation of the secondhand price, conditional on the individual and aggregate states:

$$V(j, q; d, \mathbf{n}) = E[p^{SH} | j, q, d, \mathbf{n}]$$

I estimate the value function non-parametrically using local linear regression on grid points  $[j, q, \mathbf{X}]$ :

$$\begin{aligned} \min_{\beta_0, \beta_j, \beta_q, \beta_N, \beta_d} \sum_i \{ & p_i^{SH} - \beta_0(j, q, \mathbf{X}) - \beta_j(j, q, \mathbf{X})(j_i - j) - \\ & \beta_q(j, q, \mathbf{X})(q_i - q) - \beta_N(j, q, \mathbf{X})(\mathbf{N}_i - \mathbf{N}) - \\ & \beta_d(j, q, \mathbf{X})(d_i - d) \}^2 K_h([j_i, q_i, \mathbf{X}_i] - [j, q, \mathbf{X}]), \end{aligned}$$

where the  $i$  subscript denotes the 1325 observed secondhand sales data points and  $K_h$  is a multivariate normal kernel with bandwidth  $h$ . The bandwidth corresponding to each dimension is  $h$  scaled by the standard deviation of observations across that dimension. The optimal  $h$  is chosen using a k-fold method: The data are randomly partitioned into five equal-sized subsets. The above local linear regression is estimated repeatedly at the *data points* (as opposed to grid points) from each combination of four subsets (i.e., leaving out one subset), and the average root mean squared error (RMSE) is computed. This is performed for a grid of candidate bandwidth values, and the one resulting in the lowest RMSE is chosen. This was achieved at 2.0 standard deviations. The estimated value functions are shown in Figure 8. As would be expected, older and larger ships are typically more valuable, with age generally having a more pronounced effect than size.

## Transition Probabilities

State transitions involve deterministic aging, stochastic demand fluctuations, and stochastic entry and exit. Ship size is fixed. Given these underlying processes, transition probabilities can be computed directly for the high-dimensional state  $\mathbf{x}_t$ , however these probabilities cannot be mapped directly to transition probabilities for  $\mathbf{X}$ . Therefore, I simulate many transitions from  $\mathbf{x}_t$  states to  $\mathbf{x}_{t+1}$  states and map each to its nearest  $\mathbf{X}$  states. Transition probabilities for  $\mathbf{X}$  are then obtained by counting the transitions between the corresponding  $\mathbf{X}$  states. This procedure is detailed below.

Beginning in the high-dimensional state space, the supply state at time  $t$  given in (20) will evolve




Figure 8: Value functions for representative ages and sizes.

to:

$$\mathbf{n}_{t+1} = \begin{bmatrix} M_t^{q1} & \dots & M_t^{qK} \\ n_t^{0,qK} - Z_t^{0,qK} & \dots & n_t^{0,q1} - Z_t^{0,q1} \\ \vdots & \ddots & \vdots \\ n_t^{J-1,q1} - Z_t^{J-1,q1} + n_t^{J,q1} - Z_t^{J,q1} & \dots & n_t^{J-1,qK} - Z_t^{J-1,qK} + n_t^{J,qK} - Z_t^{J,qK} \end{bmatrix},$$

where  $M_t^q$  denotes the number of entrants of each size and  $Z_t^{j,q}$  denotes the number of exits for each age and size. Entry follows a Poisson distribution, while exit follows a binomial distribution, which I approximate as a Poisson distribution.<sup>36</sup> Note that ships do not change size and age increments deterministically until age  $J$ , which is an absorbing age. Demand evolves as per the stochastic autoregressive process specified in (19). The entry, exit, and demand transition processes are independent, conditional on the state. In order to simulate  $\mathbf{x}_t$  transitions, it remains to determine the entry and exit process parameters, which are endogenous to the model. In the first step, I estimate these ‘offline.’

**Offline Entry and Exit** I assume a simple functional form for the parameters of Poisson distributions describing entry and exit, and estimate them using observed entry and exit rates.<sup>37</sup> Due to the limited number of entries and exits that occur, it is impossible to estimate parameters separately for each age and size. Instead, I assign entry observations a size bin  $q$  (as defined for the supply state variable, described at the beginning of this section). I assume the following functional form and estimate via ordinary least squares:

<sup>36</sup>This is justified if the exit probability is small compared to the number of ships (which is the case here), as a binomial distribution approaches a Poisson distribution as the number of trials tends to infinity.

<sup>37</sup>While the exit process follows a binomial distribution, this can be approximated by a Poisson distribution because the number of exits is small compared to the existing fleet.

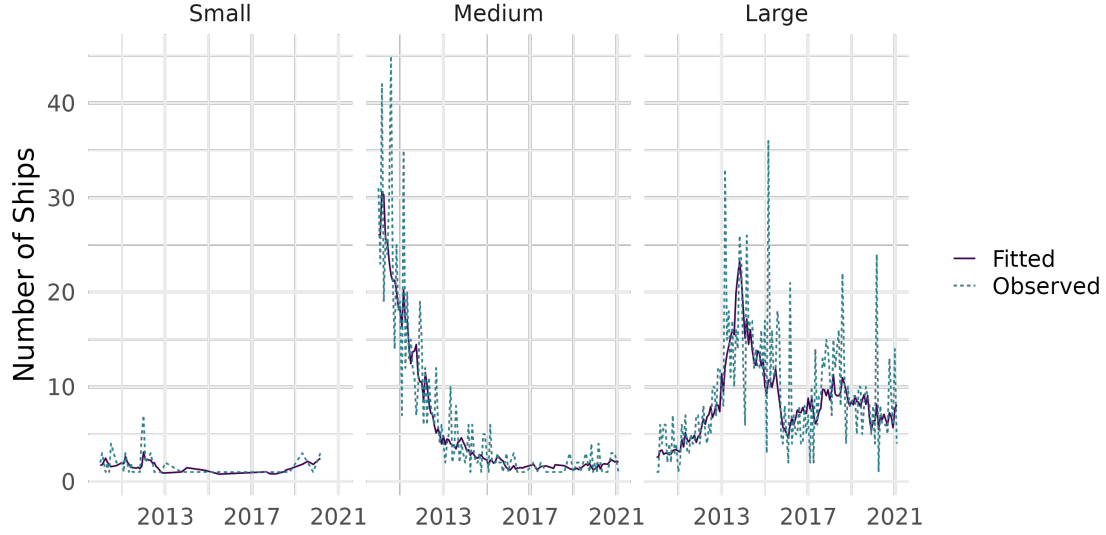


Figure 9: Number of observed entries and predicted entries using offline estimates by size group.

$$\lambda(q, d_t, \mathbf{n}_t) = \sinh \left( \gamma_q + \gamma_d d_t + \sum_{j', q'} \gamma_{q, n}^{j', q'} n_t^{j', q'} \right).$$

For exit, I distinguish ships by age ranges corresponding to the aggregate supply age bins  $j$ .<sup>38</sup> Given the lack of observed exits for ships younger than 10 years, I assume a zero exit probability for these ages. For the older two age groups, I assume the following functional form and estimate parameters separately via OLS:

$$\mu_j(d_t, \mathbf{n}_t) = \sinh \left( \delta_j + \delta_{j, d} d_t + \sum_{j', q'} \delta_{j, n}^{j', q'} n_t^{j', q'} \right), \quad j \neq [0, 119]$$

Regression results are provided in Table 6 and Table 7 and the fits are shown in Figure 9 and Figure 10.

**Transition Simulation** State transitions are simulated beginning from observed states as well as additional states obtained by drawing demand states and pairing them with observed fleet states.<sup>39</sup> For the three stochastic processes, random draws are taken from the appropriate distributions. The fleet is updated according to the simulated number of entry and exits for each age and size of ship, as well as the aging process. Exits for age bins used in the offline exit estimation are distributed by drawing randomly from the existing fleet of ships within the corresponding age range. For each transition simulated, the prior and post states are assigned to their nearest aggregate state

<sup>38</sup>I do not distinguish by size because it is not significant when included in regressions.

<sup>39</sup>This ensures that simulations cover the entire state space so that probabilities can be computed.






Figure 10: Number of observed exits and predicted exits using offline estimates by age group.

$\mathbf{X}$  representation using a  $k$ -dimensional tree algorithm to reduce computations. The probability matrix is then computed by counting the relative number of transitions between states. Finally, the transition probabilities from each state are smoothed using a multivariate normal kernel (i.e., the transition probability matrix is smoothed column-wise). The resulting transition probability matrix is depicted in Figure 11. Given the prominent diagonal, it is clear that states are very likely to remain unchanged from month to month. The diagonal bands correspond to nearby demand states—transitions between demand states are much more likely, while the supply state evolves much less frequently.

### 5.3 Model Primitives Estimation

Next, I estimate the model primitives: the entry costs, the scrap value distribution, and the cost parameters, which are obtained via the model implied profits. The discount factor is set at 0.9959, which corresponds to 5% per year, as is typically employed in the literature.

#### Entry Costs

The cost of entry for a new ship is obtained using the free entry condition in (7). Given strictly positive entry, the entry cost is equal to the expected value function one period (the time-to-build) ahead, when the ship will begin operation:

$$\kappa_q(\mathbf{X}) = \beta \mathbf{P}_X \mathbf{V}_{0,q}, \quad (21)$$

where  $\mathbf{V}_{0,q} \in \mathbb{R}^{L \times K}$  is the matrix of values of an age zero, size  $q$  ship for each aggregate state and size,  $\mathbf{P}_X$  is the row of  $\mathbf{P}$  corresponding to state  $\mathbf{X}$ . The estimated entry costs are shown in Figure 12




Figure 11: Aggregate state transition probability matrix.

along with an index of new Handymax ship prices (which was not used in estimation) as a point of comparison, since the ship cost is the primary cost of entry.

### Scrap Value Distribution

The scrap value distribution is estimated using the following model-implied relation between continuation values  $\nu$  and the observed exit frequencies:

$$\frac{Z_{j,q}}{n_{j,q}} = 1 - F_{j,\phi} \left( \frac{\nu}{q} \right) \quad (22)$$

Continuation values are straightforward to calculate from the value functions and transition matrix as:

$$\mathbf{VC}_{j,q} = \mathbf{P}\mathbf{V}_{j+1,q},$$

where  $\mathbf{VC}_{j,q}$  and  $\mathbf{V}_{j+1,q}$  are vectors over the aggregate state.

I assume distinct scrap value distributions for age bins  $j$  are normally distributed with mean  $\mu_{j,\phi}$  and standard deviation  $\sigma_{j,\phi}$ . Because exits follow a Binomial distribution, the corresponding maximum likelihood optimization for each age bin is:

$$\max_{\mu_{j,\phi}, \sigma_{j,\phi}} \sum_t \log \left[ F_{j,\phi} \left( \frac{\mathbf{VC}_j(\mathbf{X}_t)}{q}; \mu_{j,\phi}, \sigma_{j,\phi} \right)^{n_t^j - Z_t^j} \left( 1 - F_{j,\phi} \left( \frac{\mathbf{VC}_j(\mathbf{X}_t)}{q}; \mu_{j,\phi}, \sigma_{j,\phi} \right) \right)^{Z_t^j} \right]$$

The parameter estimates are provided in Table 2 and the number of exits implied by the estimated scrap value is compared to observed exits and the offline exit fit in Figure 13. The number of exits is generally similar, but the number implied by the scrap value distribution is less volatile. The



Figure 12: Model implied expected entry costs plotted with an index of new Handymax ship prices. Note that the ship size definition for the index changes from 56,000–58,000 DWT to 61,000–63,000 DWT in 2014.

Table 2: Scrap distribution estimation results.

Age	$\mu_\phi$	$\sigma_\phi$
10-20 years	-565	244
20+ years	-698	359

$\mu_\phi$  and  $\sigma_\phi$  are mean and standard deviation in \$/tonne.

model fails to capture a decrease in exits after 2015, which may be related to uncertainty regarding incoming environmental regulations, which is beyond the scope of the current work.

## Profits

Profits are directly computed from the Bellman equations in (2), so that profits are given by:

$$\boldsymbol{\pi}_{j,q}^* = \boldsymbol{V}_{j,q} - \beta(1 - \tilde{\boldsymbol{\zeta}}_{j,q})\boldsymbol{VC}_{j,q} - \beta\tilde{\boldsymbol{\zeta}}_{j,q}qE \left[ \phi \mid \phi > \frac{\boldsymbol{VC}_{j,q}}{q} \right],$$

where bold symbols denote vectors over aggregate states and the exit probabilities are computed from the scrap value distribution as:

$$\tilde{\boldsymbol{\zeta}}_{j,q} = 1 - F_{j,\phi} \left( \frac{\boldsymbol{VC}_{j,q}}{q} \right).$$

Figure 14 shows the estimated profits for select ages and sizes. Unfortunately there are no available data on profits of shipowners for comparison. However, given the price of a ship, these estimates seem reasonable, though high around the years 2011 and 2014. Model fit on this dimension

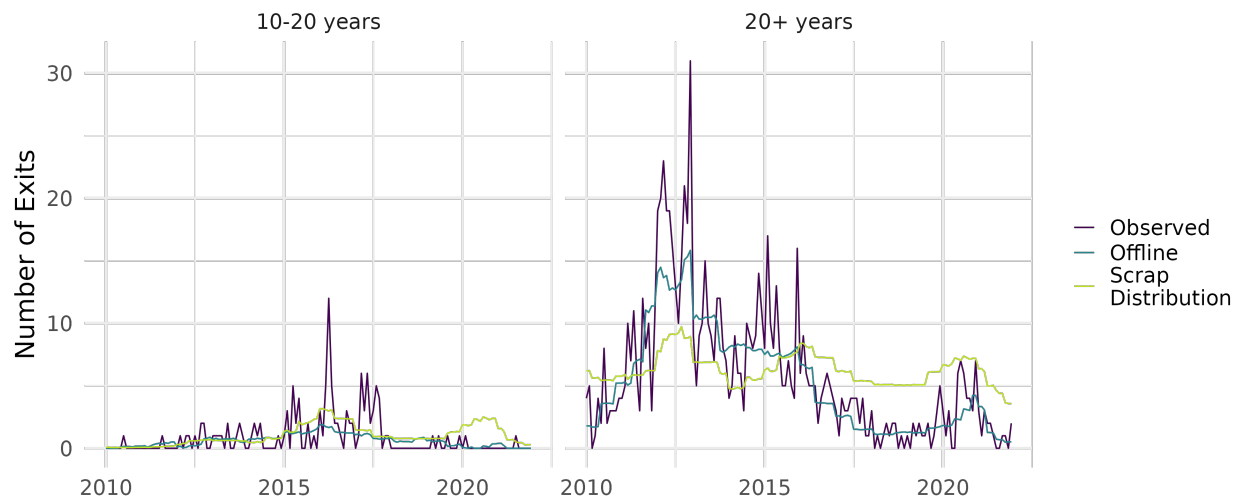


Figure 13: Observed exits, exits estimated offline, and exits implied by estimated scrap value distributions.

is not crucial for accuracy however, as entry and exit rates are calculated to be internally consistent. Importantly however, profits decrease in age and increase in size when prices are sufficiently high, as implied by the economies of scale in ship size.

### Profit Parameters

Lastly, I estimate the cost parameters for the profit function in (16) via nonlinear least squares. To do so, I first estimate the parameters relating fuel consumption to ship size and age from (9) offline. This provides greater accuracy for fuel consumption within the model and reduces the number of parameters to be estimated using the profit function.<sup>40</sup> I also calibrate the parameter  $\psi$ , which represents the fraction of time a ship is not laden and traveling, to 0.504.<sup>41</sup> This assumes that a ship is on hire roughly 60% of the time and spends 20% of that time in port, which are obtained from multiple sources.<sup>42</sup>

**Aging Parameter** The aging parameter  $\delta$  from (9) that relates fuel consumption to ship age is estimated using ordinary least squares, in a regression similar to that in Section 4.1. Unfortunately the available data consists of observations for only three years at an annual frequency, so that the effect of aging on fuel consumption cannot be clearly separated from ship vintage. As a next-best approach, I attempt to control for all technical factors that may be correlated with vintage and that

<sup>40</sup>I estimate marginal cost and fixed cost separately. While it is possible to simultaneously estimate these parameters, there is not a clear justification for how to weight the contribution from the observed prices versus the model profits.

<sup>41</sup>This value matches observed average speed data well. This parameter could potentially be estimated from the optimal speed equation, however, monthly data does not exist for the entire sample period.

<sup>42</sup>Brancaccio et al. (2020) report that 41% of the distance traveled by Handymax ships is unladen. Loading time is taken as an average of route-specific cost estimates in data from Clarksons Research.




Figure 14: Estimated profits for a 50,000 DWT ship of various ages (top) and a 10-year-old ship of various sizes (bottom).

could influence fuel consumption, including *rated* fuel consumption. However, the estimate of  $\delta$  will be biased away from zero by any unobserved efficiency improvements. The regression specification is:

$$\log(\text{fuel consumption}_{it}) = \tilde{\alpha}_0 + \delta \text{age}_{it} + \tilde{\alpha}_1 \log(\text{size}_i) + \boldsymbol{\alpha} \log(W_i) + \varepsilon_{it}$$

The dependent variable is the reported fuel *consumption* per detected travel work described in Section 4.1. I take the aging parameter  $\delta$  as the coefficient of age (in years).  $W_i$  is a vector of all available ship characteristics, including those related to the form factor of the ship (length, width, draught) and those related to its fuel consumption (rated speed, rated fuel consumption, engine power). To ensure a sufficient sample size, I use data not just from Handymax ships, but for ships from 10,000–100,000 DWT. Table 3 shows that all specifications suggest a very small aging effect. In model estimation I use the estimate with all controls, which is 0.031%/month. This value implies that a 20-year-old ship would consume around 8% more fuel than when it was new.

**Efficiency Scaling Parameter** To obtain the scaling parameter for fuel consumption, I regress log fuel consumption on log size. Here, I use *rated* fuel consumption to avoid any effects of aging, although I present a regression with *observed* work-based fuel consumption for comparison.<sup>43</sup> One problem with the rated fuel consumption measure is that it is time-based, and therefore is dependent on the rated speed, which also varies with ship size. I therefore control for the rated speed. All results are shown in Table 4, and the preferred estimate for  $\alpha$  is 0.441. For comparison, the IMO

<sup>43</sup>For the work-based regression, I take an average of fuel consumption observations for each ship to avoid over-weighting ships that were observed more than once.

Table 3: aging parameter regression results.

	log(Fuel Consumption per Travel Work)			
	(1)	(2)	(3)	(4)
Age	0.0021*	0.0013	0.0045***	0.0037**
	(0.0010)	(0.0010)	(0.0011)	(0.0012)
log(Size)	0.3436***	0.2554***	-0.6166**	-0.6467***
	(0.0147)	(0.0234)	(0.2204)	(0.1698)
Constant	-11.0938***	-9.0469***	-2.9366*	-2.6937
	(0.1631)	(0.3566)	(1.4390)	(1.3877)
Controls: consumption	No	Yes	No	Yes
Controls: form factor	No	No	Yes	Yes
Observations	1,563	1,563	1,563	1,563
Adjusted R <sup>2</sup>	0.36	0.39	0.40	0.42

Signif. codes: \*\*\*=0.001, \*\*=0.01, \*=0.05, .=0.10. Standard errors in parentheses, clustered at ship level. Consumption and form factor controls listed in main text.

uses a scaling value of 0.477 for its efficiency standards, however this is directly comparable to the specification *without* controlling for rated speed (MEPC, 2012).<sup>44</sup>

**Marginal Cost** I recover the marginal cost parameter by fitting the parametric prices  $\tilde{p}(c')$  from (15), where  $c' \equiv c\alpha_0$ , to observed charter rates  $p_t$  via nonlinear least squares. The efficiency parameter  $\alpha_0$  is not separately identified from  $c$ , so I calibrate to match the observed fuel price corresponding to the model steady-state demand level, which is \$452/tonne. To obtain  $c'$ , I minimize as follows:

$$\min_{c'} \sum_t (p_t - \tilde{p}(c'))^2.$$

The estimation fit is shown in Figure 15.<sup>45</sup>

**Fixed Cost** I obtain the fixed cost parameters via separate non-linear least squares regression:

$$\min_{\theta} \sum_i (\widehat{\pi}_i^* - \pi_i^*(\theta))^2$$

where  $i$  are model states  $\mathbf{X}$ ,  $\widehat{\pi}_i^*$  are the non-parametrically estimated profits, and  $\pi_i^*(\theta)$  is given in (16). I further assume that fixed costs scale linearly with age and parameterize  $F_q(j) = F_{q,0} + jF_{q,1}$ .<sup>46</sup> The parameters to estimate are therefore  $\theta = \{F_{q,0}, F_{q,1}\}_{\forall q}$ . Figure 16 shows the fit for a 50,000 DWT ship of various ages. As expected, the parametric profits are smoother, which mitigates

<sup>44</sup>The IMO reference line uses a calculated efficiency index that relies on rated speed and includes only ships built from 1999 to 2008.

<sup>45</sup>The specification I have used assumes that the main component of marginal cost, fuel price, is constant. This assumption could be relaxed by incorporating fuel price as an additional state variable.

<sup>46</sup>I have experimented with different specifications, including exponential cost in age, and the results are similar.

Table 4: Size parameter regression results.

	$\log(\text{FC}_{time})$ (1)	$\log(\text{FC}_{work})$ (2)	$\log(\text{FC}_{work})$ (3)
$\log(\text{Size})$	0.482*** (0.004)	0.441*** (0.004)	0.382*** (0.009)
$\log(\text{Rated Speed})$		1.261*** (0.047)	
Constant	-1.834*** (0.041)	-4.734*** (0.116)	-11.470*** (0.100)
Observations	5,447	5,447	1,362
Adjusted R <sup>2</sup>	0.75	0.78	0.56

Signif. codes: \*\*\*=0.001, \*\*=0.01, \*=0.05, .=0.10. Standard errors in parentheses. Dependent variable  $FC$  is fuel consumption, either rated time-based or detected work-based. Only a subset of ships are observed for the latter.




Figure 15: Estimated and observed monthly charter rate.




Figure 16: Profits implied by estimated cost parameters versus non-parametrically estimated profits for a 50,000 DWT ship.

the high non-parametric profit estimates. The estimated fixed profits are in Figure 17. There are no reported values for comparison, however the magnitude of these appears reasonable, except for negative values, however these can be considered as due to a level shift in profits that is not separately identified from fixed costs.






Figure 17: Estimated fixed costs by ship age and size.

## 6 Counterfactuals

This section explores how emissions policies that target intensive margins affect industry dynamics in the mid-run. I examine three emissions policies: a fuel tax, a speed limitation, and an entry subsidy, taking as a benchmark a business-as-usual baseline scenario with no emissions policy. I compute the equilibrium under each of these policy regimes and perform mid-run simulations. The simulation procedure is first described before presenting and discussing the results.

Equilibrium objects are computed from the model primitives in a somewhat reverse fashion to the estimation procedure. I first calculate optimal period profits in each state using the marginal and fixed costs, efficiency parameters, and the assumed functional forms for profit as per Section 4.4. For policy counterfactuals, the policy modifications occur at this stage. In the case of a fuel tax the marginal cost is adjusted, and in the case of a speed limit, a maximum speed constraint is applied to the charterer cost minimization problem in (8). Entry costs are decreased for the entry subsidy. I next compute equilibrium objects, taking as primitives the optimal profits, the scrap distribution, and the entry costs.<sup>47</sup> The equilibrium objects include entry and exit policies, state transition probabilities, and value functions. I use an iterative process to find a mutually consistent set of equilibrium objects, beginning with offline estimates for entry, exit, and state transitions. The Bellman equation is used to compute implied entries and exit probabilities,<sup>48</sup> which are in turn used to compute new implied transitions, and so on.

Simulations are computed in a process similar to that used to estimate state transition probabilities in Section 5.2, in which draws are taken for the three stochastic processes of demand, entry, and

<sup>47</sup>In a preliminary step, I first obtain model-consistent entry costs through fixed point iteration on the Bellman equation for the oldest, absorbing state, followed by backward calculation. See Section A.4 for details.

<sup>48</sup>In contrast to the estimation procedure, as well as Kalouptsi (2014), here I use exit *probabilities* rather than *numbers*, as this avoids an aggregation step and allows exit probabilities to vary over each age and size rather than by groups.

exit. Here however, the equilibrium entry and exit policies are used rather than the offline estimates. Linear regressions of entries and exit probabilities on the aggregate state are used to interpolate and extrapolate from the grid point values computed in the previous step.<sup>49</sup>

I begin by simulating a baseline business-as-usual scenario. All simulations begin from the steady-state demand level and the existing fleet as of January 2012.<sup>50</sup> The average aggregate values from 4800 simulated time paths are shown in Figures 18 and 19. While demand remains at its steady-state level, the fleet expands fairly consistently throughout the simulated period. This indicates that the initial fleet distribution is not the steady state distribution, and the slow fleet adjustment highlights the importance of employing a dynamic model to study the mid-run evolution of the industry. With the increased supply, the equilibrium price of shipping decreases and the quantity of shipping increases. Average speed declines gradually, while total emissions remain largely unchanged. Because the estimated demand was quite flat from 2010 to 2020, these simulations resemble the observed trends reasonably well. In particular, the fleet expansion is almost identical, rising from roughly 3000 to 4000 ships over the course of ten years. In reality, the price declined by about half from 2012 to 2020, though with a more rapid decline early on before flattening out. Average speed decreased from around 12 to 11.4, a similar decline to the simulation. While not an exact comparison in terms of ship sizes, reported emissions for ships between 35000 and 60000 DWT remained almost constant from 2012 to 2018. The entry time path in Figure 19 fluctuates similarly to the simulations in Kalouptside (2014), with the increase driven by the aging of the fleet.

I assess policy counterfactuals in relation to the baseline scenario. Each policy is implemented beginning at the thirteenth month, and the simulation is run for another ten years. I simulate a proportional fuel tax  $\tau$  by adjusting the marginal cost  $c$  by  $1 + \tau$ . A speed limitation is implemented by applying a maximum speed constraint in the charter optimization problem in (8). In order to compare the mid-run dynamics of the policies, I calibrate their stringency to induce the same reduction in emissions in the first period of implementation. I simulate a fuel tax of 20% and a corresponding speed limit of 12.4 knots, which is only partially binding across ships and states. For comparison, the average observed speed over the sample period is around 11.7 knots. To model a uniform entry subsidy, the entry costs  $\kappa$  are reduced by 5% for all sizes and aggregate states.

The immediate impact in the first year of a given policy is primarily due to adjustments on the intensive margin, and examining the initial response is therefore helpful to clarify and distinguish the intensive margin channels from the extensive margin channels. From Figure 20, we see that both the fuel tax and speed limit cause an average speed reduction in the neighborhood of 5% in the first month. With ships travelling more slowly (and before the number of ships changes significantly) emissions drop 15.8%. The reduced supply of shipping resulting from slower speeds causes the price to spike, both in terms of the time-based charter rate and the quantity-based price of shipping. The higher price reduces the equilibrium quantity of shipping and also translates into higher profits

---

<sup>49</sup>Details are provided in Section A.4.

<sup>50</sup>I begin with a historical date for two reasons: First, this makes it easier to judge the accuracy of the model by comparing with the actual industry evolution. Secondly, because ship sizes increased over the observation sample period, beginning with a historical date ensures that ship sizes remain with a state space that is comparable to that used for estimation.

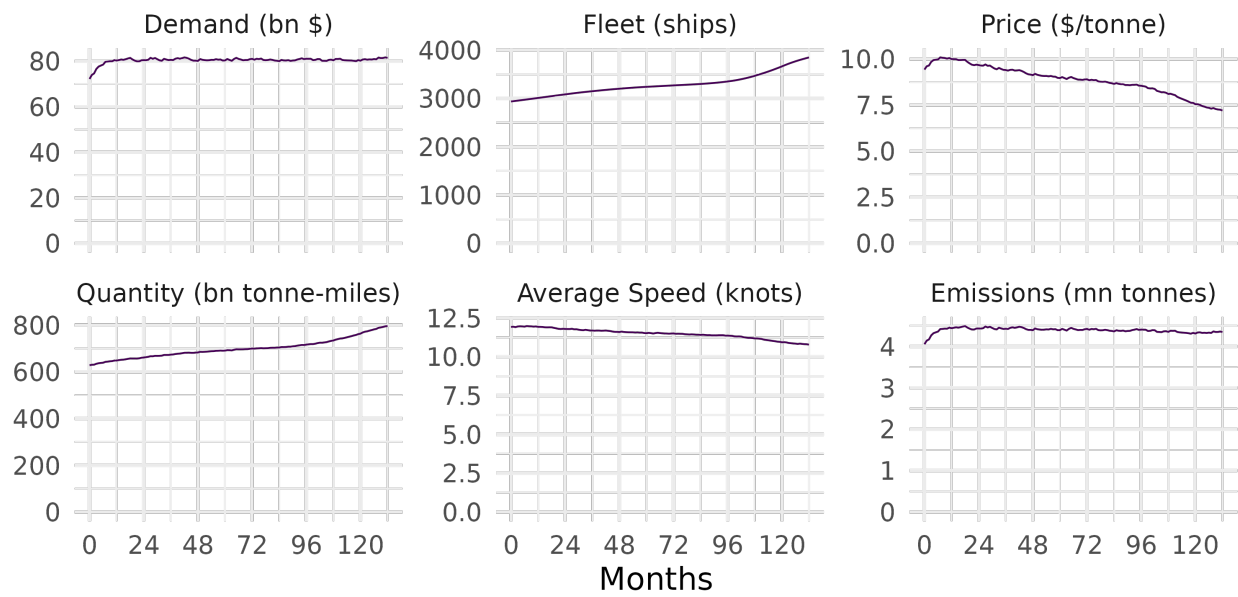


Figure 18: Time path of baseline counterfactual aggregate values.




Figure 19: Time path of baseline counterfactual entries and exits.

for shipowners. In contrast, the entry subsidy, which acts directly on the extensive margin, has essentially zero impact on speeds, price, and quantity in the short-run.

The longer-run, dynamic response incorporates the intensive margin channels as well as the effect of future expectations, which determine entry and exit choices. In Figure 21 we see that in the periods directly following the fuel tax or speed limit implementation, the higher expected profits incentivize more ships to enter the market, but also *decrease* the number of ships exiting. This is the first notable result of this model, as emissions policies are often promoted as a way to hasten the exit of older, less efficient capital. This counterintuitive result arises from the fact that the charter rate increases for all ships, thereby increasing current and expected future profits for both new and old ships. The magnitude of the profit increase varies across age and size, and this variation and the existing fleet distribution jointly determine the balance between entry and exit.

Looking further ahead, the importance of the extensive margin becomes clearer, with the fuel tax and speed limit counterfactuals diverging considerably. Whereas emissions under a fuel tax remain steady over time, under the speed limit they trend back toward the baseline level. This is the second striking result from this model and results from the interaction of the intensive and extensive margins. One way to understand this is that the fleet growth stimulated by the higher price acts counter to the speed reduction, increasing the supply of shipping and therefore lowering prices. The unconstrained optimal speed becomes lower over time and the speed limit is therefore less binding. As such, the initial strong effect of the speed limit diminishes over the course of ten years to less than half of the original emissions reduction relative to the baseline scenario. The tax, being proportional, maintains downward pressure on speeds and retains its effectiveness over time. Due to the small increase in the number and size of ships, there is an additional price decrease and corresponding quantity increase over time with the tax. While the tax is clearly superior in terms of emissions reduction, I find that it also reduces consumer surplus over the ten-year time frame less than the speed limit, by 5.2% compared to 5.6%. This is not immediately obvious looking at the average price and quantity, but results from the fact that a speed limit prevents the industry from adjusting to high demand shocks, while a fuel tax applies a more flexible downward pressure on speeds.

The entry subsidy has a very different effect from the two policies that operate on the intensive margin. It is effective at increasing both the entry rate of new ships *and* the exit rate of old ships, because it induces entry of new, more efficient ships without a price increase. However, it is ineffective at reducing emissions, with the fleet growth acting almost entirely to increase the quantity of shipping rather than reduce emissions.

A final characteristic of the response worth noting is that the induced entry for all policies is initially large, but diminishes over time. This is in part due to the drop in the charter price, however the evolution of the total number of ships, their ages, and their sizes all play a role in shaping the mid-run outcome of these policies. Nevertheless, the spike in entry has a persistent effect, either through decreasing emissions or increasing shipping quantity, which is due to the long lifespan of ships.

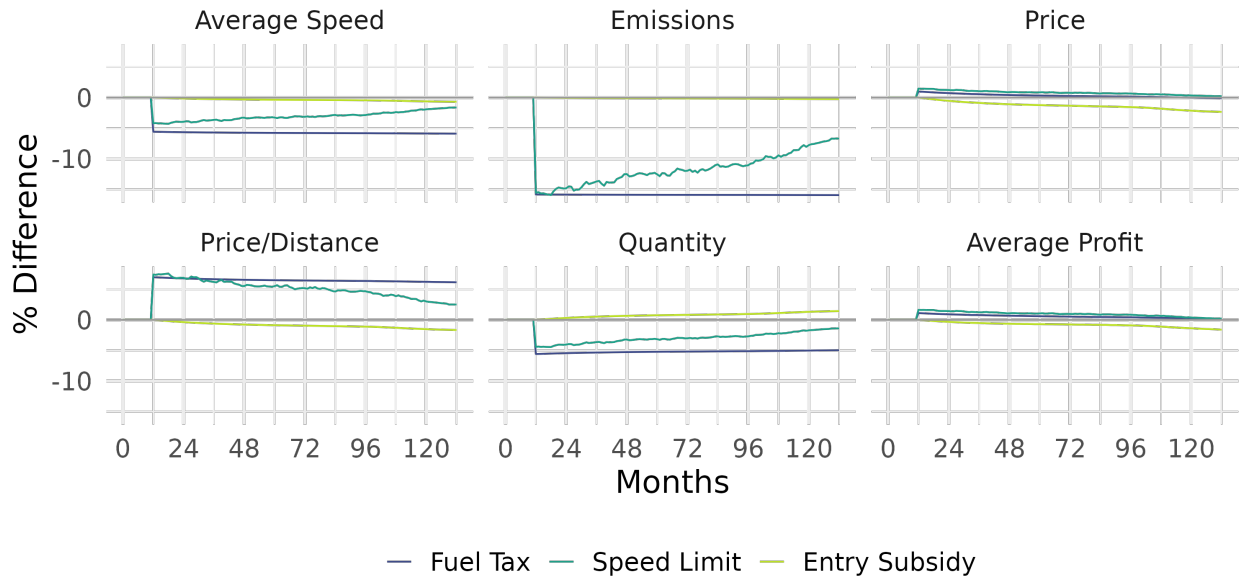


Figure 20: Time path of aggregate outcomes as a percentage difference from baseline scenario.




Figure 21: Time path of fleet evolution as percentage difference from baseline scenario (5-month moving average).

## 7 Conclusion

In this paper, I have constructed a dynamic model of the dry bulk shipping industry that takes seriously both the intertemporal and intratemporal firm choices to study how they jointly impact aggregate greenhouse gas emissions over time. I have adapted and extended an existing modeling and estimation framework from Kalouptsi (2014) by adding a critical dimension of heterogeneity over ship size and a detailed model of ship fuel efficiency and CO<sub>2</sub> emissions. Using this framework, I compared the mid-run dynamic effects of a fuel tax, a speed limiting regulation, and an entry subsidy to provide insight toward the development of effective shipping emissions policy.

Emissions reduction in shipping is particularly challenging, given the supranational nature of the industry, its importance to global trade, and the lack of technological solutions. Regulations have recently been enacted, but it is unclear if they will be sufficient to put the industry on a path toward its goal of net-zero by 2050. One of these regulations is in practice similar to a speed limit, which I find to have diminishing efficacy over time. While the stringency of this regulation will ratchet up over time, this may not fully mitigate the efficacy decrease if investors take actions based on long-run stringency levels, which is likely, given the long lifespan of ships. While a fuel tax is more effective at reducing emissions, I find that it comes at the cost of a persistently higher shipping price, rationalizing the opposition of some trade-reliant nations to this policy. My analysis suggests that regulation directly limiting speed is effective mainly in the short-run, and should be accompanied by other, more persistent measures. Furthermore, a policy such as an entry subsidy that acts directly on the extensive margin is more effective at promoting capital turnover.

The framework that I have advanced in this paper could be employed in future work to study alternative policy measures. For example, the existing efficiency standard for new ships more directly impacts the entry margin, but may deter investment and slow turnover if more efficient ships are more expensive to build. Further assumptions and data on the price of new ships in relation to their fuel efficiency would permit a straightforward application of my model to this policy. On a longer time horizon than I consider here, it is expected that the industry will transition to an alternative fuel, however the slow rate of capital turnover means that the fleet evolution up until that point is important. To that point, I have shown that emissions policies can both reduce emissions in the mid-run *and* maintain an older fleet that will be more conducive to rapid replacement when an alternative technology becomes viable. Additionally, an entry subsidy would be effective in accelerating this process. In the meantime, a more readily available alternative could be to leverage the economies of scale in ship size that I have highlighted above by enhancing infrastructure that constrains new ship sizes. This is yet another scenario to which this modeling framework could be adapted.

The slow rate of capital turnover is a key obstacle to rapid decarbonization. In this paper, I have quantitatively assessed how various emissions policies impact short- to mid-run industry dynamics in the presence of a strong intensive margin and important heterogeneity. I find that the dynamic interaction between the intensive margin and the extensive margin can greatly diminish the effect of an operational intensity standard—a problem that does not occur with a proportional fuel tax.

Additionally, in the presence of a strong intensive margin of adjustment, emissions policies that increase the price of shipping may simultaneously reduce emissions and reduce the exit rate of older, less efficient capital. My findings highlight the importance of accounting for equilibrium price effects in assessing how policies impact capital turnover and therefore aggregate emissions over time. These insights are relevant not only for the shipping industry, but for a range of industries that face this pervasive impediment to rapidly reducing their carbon footprints.

## References

- Adland, R., Jia, H., & Strandenes, S. P. (2018). The determinants of vessel capacity utilization: The case of brazilian iron ore exports. *Transportation Research Part A: Policy and Practice*, *110*, 191–201.
- Brancaccio, G., Kalouptsi, M., & Papageorgiou, T. (2020). Geography, transportation, and endogenous trade costs. *Econometrica*, *88*(2), 657–691.
- Campello, M., Kankanhalli, G., & Kim, H. (2021). *Delayed creative destruction: How uncertainty shapes corporate assets* (tech. rep.). National Bureau of Economic Research.
- Cariou, P., Halim, R. A., & Rickard, B. J. (2023). Ship-owner response to carbon taxes: Industry and environmental implications. *Ecological Economics*, *212*, 107917.
- Cullen, J. A., & Reynolds, S. S. (2017). *Market dynamics and investment in the electricity sector* (tech. rep.). Working paper.
- Doraszelski, U., & Satterthwaite, M. (2010). Computable markov-perfect industry dynamics. *The RAND Journal of Economics*, *41*(2), 215–243.
- Ericson, R., & Pakes, A. (1995). Markov-perfect industry dynamics: A framework for empirical work. *The Review of economic studies*, *62*(1), 53–82.
- Faber, J., Hanayama, S., Zhang, S., Pereda, P., Comer, B., Hauerhof, E., & Yuan, H. (2020). Fourth IMO greenhouse gas study [Online, accessed 11. Jul. 2021]. <https://www.imo.org/en/OurWork/Environment/Pages/Fourth-IMO-Greenhouse-Gas-Study-2020.aspx>
- Fowlie, M., Reguant, M., & Ryan, S. P. (2016). Market-based emissions regulation and industry dynamics. *Journal of Political Economy*, *124*(1), 249–302.
- Foxon, T. J. (2013). Technological lock-in. In J. Shogren (Ed.), *Encyclopedia of energy, natural resource, and environmental economics* (pp. 123–127). Elsevier Amsterdam, The Netherlands.
- Gkochari, C. C. (2015). Optimal investment timing in the dry bulk shipping sector. *Transportation Research Part E: Logistics and Transportation Review*, *79*, 102–109.
- Greenwood, R., & Hanson, S. G. (2015). Waves in ship prices and investment. *The Quarterly Journal of Economics*, *130*(1), 55–109.
- Hawkins-Pierot, J. T., Bureau, C. F. P., & Wagner, K. R. (2023). *Technology lock-in and costs of delayed climate policy* (tech. rep.).
- Intergovernmental Panel on Climate Change. (2022). Summary for policymakers. In *Global warming of 1.5°C: Ipcc special report on impacts of global warming of 1.5°C above pre-industrial levels in context of strengthening response to climate change, sustainable development, and efforts to eradicate poverty* (pp. 1–24). Cambridge University Press. <https://doi.org/10.1017/9781009157940.001>
- Jalkanen, J.-P., Brink, A., Kalli, J., Pettersson, H., Kukkonen, J., & Stipa, T. (2009). A modelling system for the exhaust emissions of marine traffic and its application in the baltic sea area. *Atmospheric Chemistry and Physics*, *9*(23), 9209–9223.
- Johansson, L., Jalkanen, J.-P., & Kukkonen, J. (2017). Global assessment of shipping emissions in 2015 on a high spatial and temporal resolution. *Atmospheric Environment*, *167*, 403–415.



- Kalouptsi, M. (2014). Time to build and fluctuations in bulk shipping. *American Economic Review*, *104*(2), 564–608.
- Kellogg, R., & Reguant, M. (2021). Energy and environmental markets, industrial organization, and regulation. In *Handbook of industrial organization* (pp. 615–742). Elsevier.
- Kyriakou, I., Pouliasis, P. K., Papapostolou, N. C., & Nomikos, N. K. (2018). Income uncertainty and the decision to invest in bulk shipping. *European Financial Management*, *24*(3), 387–417.
- MAN Energy Solutions. (2018). *Basic principles of ship propulsion* (tech. rep.).
- Olmer, N., Comer, B., Roy, B., Mao, X., & Rutherford, D. (2017). Greenhouse gas emissions from global shipping, 2013–2015 detailed methodology. *International Council on Clean Transportation: Washington, DC, USA*, 1–38.
- Parry, I., Heine, D., Kizzier, K., & Smith, T. (2022). A carbon levy for international maritime fuels. *Review of environmental economics and policy*, *16*(1), 25–41.
- Ronen, D. (1982). The effect of oil price on the optimal speed of ships. *Journal of the Operational Research Society*, *33*(11), 1035–1040.
- Rutherford, D., Mao, X., Osipova, L., & Comer, B. (2020). Limiting engine power to reduce co2 emissions from existing ships. *International Council on Clean Transportation*.
- Ryan, S. P. (2012). The costs of environmental regulation in a concentrated industry. *Econometrica*, *80*(3), 1019–1061.
- Scott, P. (2014). Dynamic discrete choice estimation of agricultural land use.
- Shapiro, J. S. (2016). Trade costs, co 2, and the environment. *American Economic Journal: Economic Policy*, *8*(4), 220–54.
- Sheng, Y., Shi, X., & Su, B. (2018). Re-analyzing the economic impact of a global bunker emissions charge. *Energy Economics*, *74*, 107–119.
- The Marine Environment Protection Committee. (2012). Resolution mepc.215(63): Guidelines for calculation of reference lines for use with the energy efficiency design index (eedi).
- Toyama, Y. (2019). *Dynamic incentives and permit market equilibrium in cap-and-trade regulation* (tech. rep.). Mimeo.
- United Nations Conference on Trade and Development. (2017). Review of maritime transport 2017.
- Van, T. C., Ramirez, J., Rainey, T., Ristovski, Z., & Brown, R. J. (2019). Global impacts of recent imo regulations on marine fuel oil refining processes and ship emissions. *Transportation Research Part D: Transport and Environment*, *70*, 123–134.
- Wan, Z., El Makhoulfi, A., Chen, Y., & Tang, J. (2018). Decarbonizing the international shipping industry: Solutions and policy recommendations. *Marine pollution bulletin*, *126*, 428–435.
- Weintraub, G. Y., Benkard, C. L., & Van Roy, B. (2008). Markov perfect industry dynamics with many firms. *Econometrica*, *76*(6), 1375–1411.

## A Appendix

### A.1 Data Details

#### Sub-Sector Definition

As seen in Figure 2, the size of ships within each sub-sector has evolved somewhat over time. As is standard in the industry, I define a ship’s market based on its size *and* its built year. For Handymax, the lower bound increases from 35,000 to 40,000 DWT in 1996, and again to 45,000 DWT in 2014. The upper bound increases from 60,000 to 65,000 DWT in 2000 and again to 70,000 DWT in 2013. This is roughly consistent with definitions used by Clarksons research and visually follows the troughs of Figure 2.

#### Matching Emissions to Tracking Data

Since 2018, the EU has collected and published reported annual fuel consumption, emissions, and distance traveled for each ship calling on an EU port. This constitutes roughly a quarter of all ships in any given year, which are representative of the total population in terms of size and age. A registration number (IMO number) is a unique identifier that allows to match the MRV data to ship characteristics from the Clarksons Research World Fleet Register. In addition to basic ship information, this contains detailed technical characteristics such as hull dimensions, engine power, propeller details, etc. as well as the MMSI number which identifies the ship’s AIS transceiver. Each ship is fitted with an AIS transceiver that reports location, speed, and draft (the vertical distance between the waterline and the bottom of the hull) every few minutes. I use hourly AIS data for 2019–2021, which is collected using both land-based and satellite systems. I first clean this data, for example filtering out clearly spurious observations such as those indicating impossible ship trajectories.<sup>51</sup> I then identify port calls as periods of 12 hours during which a ship does not exceed a speed of one knot and is within a nation’s economic exclusion zone (200 nautical miles from the coast), assigning the corresponding nation to the port call. To match with MRV data, I aggregate annual tracking data over all trips with either the origin or destination assigned to an EU country. Finally, to ensure accuracy, I select only those ship-year observations for which the AIS detected distance traveled on EU trips matches the distance reported in the MRV data within  $\pm 10\%$ , which constitutes 44% of MRV observations.

#### Fleet and Exits Reconstruction

A complete register including scrapped ships and their demolition date is only available back to 2018. Prior to that, I rely on annually published fleet registers, which I obtain for 2005 and 2012 through 2021. These specify the build month for each ship, but only the scrap year, not month for those that are demolished. Therefore, it is necessary to both reconstruct the 2006–2011 gap, as well as determine the scrap month of each vessel exiting during the entire sample period. The first is

---

<sup>51</sup>see Faber et al. (2020, sec. 2.2.3) for a discussion of AIS data errors and pre-processing

possible because the long lifespan of ships means that virtually no ship built after 2005 would exit before 2012. Therefore, I identify a set of about 1800 ships greater than 10,000 DWT present in 2005 but not in 2012 by matching on IMO number. I add this to the roughly 2600 exits from 2012 to 2021. Next, I compile a list of observed exits from a weekly publication that contains the ship name, built year, size, and scrap month, for which I have around 800 observations for 2006–2011 and 2400 for 2012–2021. I match these first on name, built year, and size (within a small margin of error), as well as scrap year when available. For those that do not match, I employ lists of ex-names for ships and perform a second round of matching on the same variables. For the remaining, I match on built year and size, as these two variables are the only that are critical to my analysis (I require only the age-size distribution). For the remaining unmatched observations, I expand the match criteria to built year within one year and DWT with 20%. Roughly half of potential exits are matched after this process. I search publicly available registers to find exit years for the remaining. Finally, I randomly draw exit months from the observed exit month distribution for 860 ships for which I am able to identify an exit year but not month. For 2018 onwards, I use the complete fleet register database as-is. I am able to validate the reconstructed data using time series of sub-sector fleet numbers, capacity and average age, as well as monthly publications of a slightly more granular binned fleet age-size distribution, both from the SIN. These agree well for the Handymax and Panamax size categories, but have some discontinuities and discrepancies for smaller and larger size categories.

## A.2 Modeling Choices Empirics

### A.2.1 Size vs. Technology

The rated fuel consumption is obtained from the Clarksons fleet register, while reported annual fuel consumption and the corresponding annual distance are from the MRV reports. Travel work is calculated from tracking data as follows: I define travel work for a trip segment of distance  $x$  at constant speed  $s$  as  $s^2 \cdot x$ , based on the physical force of drag on the ship. Travel work over a period of length  $t$  is aggregated by summing over the set of trip segments completed in that time,  $X_t$ :

$$W_t \equiv \sum_{x \in X_t} \cdot s_x^2 * x.$$

I calculate the travel work for each ship from hourly AIS data for trips into and out of the EU for each year, corresponding to annual reported fuel consumption in the MRV data. Due to the imperfect nature of the trip detection, I retain only ship-year observations for which the detected annual distance traveled matches the reported value with  $\pm 10\%$ . This remains an imperfect measure for a number of reasons, notably noise in the tracking data, use of speed over ground versus speed relative to water, and lack of data on how heavily laden the ship is.

Regression results are shown in Table 5. For consistency, I drop ships that are missing data for any of the three efficiency measures, however results are similar when retaining all observations.

Table 5: Fuel efficiency regression results.

	log FE <sub>time</sub> (1)	log FE <sub>distance</sub> (2)	log FE <sub>work</sub> (3)
log(Size)	0.509*** (0.010)	0.569*** (0.007)	0.612*** (0.010)
Built Year	0.012*** (0.001)	0.005*** (0.001)	0.007*** (0.001)
Constant	Yes	Yes	Yes
Size Change = 1 Year	2.5 %	0.8 %	1.2 %
Observations	2,291	2,291	2,291
Adjusted R <sup>2</sup>	0.79	0.82	0.75

Signif. codes: \*\*\*=0.001, \*\*=0.01, \*=0.05, .=0.10. Standard errors in parentheses, clustered at ship level. Regressions for three fuel efficiency (FE) measures: (1) rated, per *time*, (2) reported, per reported *distance*, and (3) reported, per detected travel *work*.

### A.3 Offline Entry/Exit Regressions

Table 6: Offline entry regression result.

	asinh(Entries) (1)
$q_1$	-19.977 (15.674)
$q_2$	-12.018 (6.811)
$q_3$	-5.625 (5.890)
d	0.501*** (0.133)
$n^{1,1}$	0.035 (0.029)
$n^{1,2}$	0.002 (0.002)
$n^{1,3}$	-0.006 (0.014)
$n^{2,1}$	-0.008 (0.029)
$n^{2,2}$	0.004 (0.008)
$n^{2,3}$	-0.002 (0.170)
$n^{3,1}$	0.007 (0.008)
$n^{3,2}$	0.039 (0.034)
$n^{3,3}$	0.123 (0.459)
$q_2 \times n^{1,1}$	-0.032 (0.031)
$q_3 \times n^{1,1}$	-0.039 (0.031)
$q_2 \times n^{1,2}$	-0.003 (0.002)
$q_3 \times n^{1,2}$	-0.002 (0.002)
$q_2 \times n^{1,3}$	0.007 (0.015)
$q_3 \times n^{1,3}$	-0.003 (0.015)
$q_2 \times n^{2,1}$	0.014 (0.031)
$q_3 \times n^{2,1}$	0.002 (0.030)
$q_2 \times n^{2,2}$	-0.003 (0.008)
$q_3 \times n^{2,2}$	0.000 (0.008)
$q_2 \times n^{2,3}$	-0.027 (0.181)
$q_3 \times n^{2,3}$	-0.066 (0.180)
$q_2 \times n^{3,1}$	-0.003 (0.008)
$q_3 \times n^{3,1}$	-0.015 (0.008)
$q_2 \times n^{3,2}$	-0.045 (0.037)
$q_3 \times n^{3,2}$	-0.005 (0.035)
$q_2 \times n^{3,3}$	-0.010 (0.544)
$q_3 \times n^{3,3}$	-0.416 (0.574)
Observations	283
R <sup>2</sup>	0.78

Signif. codes: \*\*\*=0.001, \*\*=0.01, \*=0.05, .=0.10. Bootstrapped standard errors on 1000 resamples in parentheses.

Table 7: Offline exit regression results.

	asinh(Exits)	
	(1)	(2)
	Age Bin [120,240)	Age Bin [240,Inf]
d	-0.2158 (0.2422)	-0.0985 (0.2216)
$n^{1,1}$	-0.0224*** (0.0000)	-0.0143*** (0.0000)
$n^{1,2}$	0.0003*** (0.0000)	0.0036*** (0.0000)
$n^{1,3}$	0.0050*** (0.0000)	0.0091*** (0.0000)
$n^{2,1}$	0.0033*** (0.0000)	0.0198*** (0.0000)
$n^{2,2}$	-0.0061*** (0.0000)	0.0033*** (0.0000)
$n^{2,3}$	0.1403*** (0.0000)	0.0710*** (0.0000)
$n^{3,1}$	-0.0024*** (0.0000)	0.0106*** (0.0000)
$n^{3,2}$	-0.0356*** (0.0000)	-0.0505*** (0.0000)
$n^{3,3}$	-0.0103*** (0.0000)	0.4257*** (0.0000)
Constant	10.0933 (9.7512)	-11.0437 (8.6610)
Observations	144	144
R <sup>2</sup>	0.30	0.70

Signif. codes: \*\*\*=0.001, \*\*=0.01, \*=0.05, .=0.10. Bootstrapped standard errors on 1000 resamples in parentheses.

#### A.4 Counterfactual Simulation Details

I present an alternative approach: I match individual contracts with ship tracking data to mitigate the measurement error in prices. The drawback of this approach is that there is the potential for endogeneity in the individual freight rate, for example from higher prices being offered for faster shipments. For this reason, I instrument for the individual freight rate using a global freight price index for the sub-sector. The identifying assumption is that the freight rate index affects the speed ships travel on a given voyage *only* through its effect on the freight rate. Measurement error in speed may arise in calculating an average speed from departure and arrival times divided by standard average route distances. This may not reflect actual distance traveled and may be affected by stopping or manoeuvring. Alternatively, I use the distance-weighted observed speed, which may more accurately reflect the speed choice made by the operator. My base regression specification is:

This section outlines the procedure used to compute counterfactual equilibria. I use bold letters to denote vectors over aggregate states  $\mathbf{X} = (d, \mathbf{n})$ .

To ensure model consistency, I first calculate the model implied value function to use as initial values in the entry/exit equilibrium computation. I obtain the value function for the oldest ships via

fixed point iteration on:

$$\mathbf{V}_{420,q} = \boldsymbol{\pi}_{420,q} + \beta(1 - \tilde{\zeta}_{420,q})\mathbf{P}_{\lambda,\mu}\mathbf{V}_{420,q} + \beta\tilde{\zeta}_{420,q}qE\left[\phi|\phi > \frac{\mathbf{P}_{\lambda,\mu}\mathbf{V}_{420,q}}{q}\right], \quad \forall q, \quad (23)$$

where  $\mathbf{P}_{\lambda,\mu}$  is the transition probability matrix computed using the offline entry and exit probability estimates, and the conditional expectation of the scrap value is given by:

$$E[\phi|\phi > A] = \frac{\int_A^\infty \phi f_\phi(\phi) d\phi}{1 - F_\phi(A)} \quad (24)$$

To avoid numerical integration, I use the fact that when  $\phi$  is normally distributed with mean  $\mu_\phi$  and standard deviation  $\sigma_\phi$ ,

$$\begin{aligned} E[\phi|\phi > A] &= \mu_\phi + \sigma_\phi \left( \frac{\phi\left(\frac{A-\mu_\phi}{\sigma_\phi}\right)}{1 - \Phi\left(\frac{A-\mu_\phi}{\sigma_\phi}\right)} \right) \\ &= \frac{\mu_\phi \left(1 - \Phi\left(\frac{A-\mu_\phi}{\sigma_\phi}\right)\right) + \sigma_\phi \phi\left(\frac{A-\mu_\phi}{\sigma_\phi}\right)}{1 - \Phi\left(\frac{A-\mu_\phi}{\sigma_\phi}\right)}, \end{aligned} \quad (25)$$

where  $\phi$  and  $\Phi$  denote the standard normal PDF and CDF.

The value function for the remaining ages is calculated sequentially for each  $q$  as follows:

$$\begin{aligned} \mathbf{V}_{419,q} &= \boldsymbol{\pi}_{419,q} + \beta(1 - \zeta_{419,q})\mathbf{P}_{\lambda,\mu}\mathbf{V}_{420,q} + \beta\zeta_{419,q}qE\left[\phi|\phi > \frac{\mathbf{P}_{\lambda,\mu}\mathbf{V}_{420,q}}{q}\right] \\ &\vdots \\ \mathbf{V}_{120,q} &= \boldsymbol{\pi}_{120,q} + \beta(1 - \zeta_{120,q})\mathbf{P}_{\lambda,\mu}\mathbf{V}_{121,q} + \beta\zeta_{120,q}qE\left[\phi|\phi > \frac{\mathbf{P}_{\lambda,\mu}\mathbf{V}_{121,q}}{q}\right] \\ \mathbf{V}_{119,q} &= \boldsymbol{\pi}_{119,q} + \beta\mathbf{P}_{\lambda,\mu}\mathbf{V}_{120,q} \\ &\vdots \\ \mathbf{V}_{0,q} &= \boldsymbol{\pi}_{0,q} + \beta\mathbf{P}_{\lambda,\mu}\mathbf{V}_{1,q} \end{aligned} \quad (26)$$

The model-implied entry costs are computed as in (21) using the baseline scenario with no emissions policy.

To find the equilibrium entry and exit policies, I iterate on the following procedure. I take the estimated scrap value distributions as primitives. For the initial entry and exit policies, I use  $\boldsymbol{\lambda}_q$  and  $(\boldsymbol{\mu}_{[120-239]}, \boldsymbol{\mu}_{240+})$  estimated offline through the linear regression of observed exits.

1. Compute  $\mathbf{P}_{\lambda_{k-1}, \mu_{k-1}}$  as described in Section 5.

2. Update value functions:

$$\begin{aligned}
V_{420,q,k} &= \pi_{420,q} + \beta(1 - \zeta_{240+})P_{\lambda_{k-1},\mu_{k-1}}V_{420,q,k-1} + \beta\zeta_{240+}qE \left[ \phi|\phi > \frac{P_{\lambda_{k-1},\mu_{k-1}}V_{420,q,k-1}}{q} \right] \\
V_{419,q,k} &= \pi_{419,q} + \beta(1 - \zeta_{240+})P_{\lambda_{k-1},\mu_{k-1}}V_{420,q,k} + \beta\zeta_{240+}qE \left[ \phi|\phi > \frac{P_{\lambda_{k-1},\mu_{k-1}}V_{420,q,k}}{q} \right] \\
&\vdots \\
V_{120,q,k} &= \pi_{120,q} + \beta(1 - \zeta_{120-239})P_{\lambda_{k-1},\mu_{k-1}}V_{121,q,k} + \beta\zeta_{120-239}qE \left[ \phi|\phi > \frac{P_{\lambda_{k-1},\mu_{k-1}}V_{121,q,k}}{q} \right] \\
V_{119,q,k} &= \pi_{119,q} + \beta P_{\lambda_{k-1},\mu_{k-1}}V_{120,q,k} \\
&\vdots \\
V_{0,q,k} &= \pi_{0,q} + \beta P_{\lambda_{k-1},\mu_{k-1}}V_{1,q,k}, \quad \forall q
\end{aligned} \tag{27}$$

3. Compute implied exit probabilities:

$$\zeta_{j,q}^* = \begin{cases} 0 & \text{for } j \in [0, 119] \\ 1 - F_{\phi,120-239} \left( \frac{P_{\lambda_{k-1},\mu_{k-1}}V_{j+1,q,k}}{q} \right) & \text{for } j \in [120, 239] \\ 1 - F_{\phi,240+} \left( \frac{P_{\lambda_{k-1},\mu_{k-1}}V_{j+1,q,k}}{q} \right) & \text{for } j \geq 240 \end{cases} \tag{28}$$

4. Compute new entry:

$$\lambda_q^* = \lambda_{q,k-1} \left( \frac{\beta^T (P_{\lambda_{k-1},\mu_{k-1},X})^T V_{0,q,k}}{\kappa_q} \right)^{1/2}, \quad \forall q,$$

where  $P_{\lambda_{k-1},\mu_{k-1},X}$  is the  $X^{th}$  row of  $P_{\lambda_{k-1},\mu_{k-1}}$ .

5. Update entry and exit parameters (only partially)

$$\begin{aligned}
\lambda_{q,k} &= \lambda_{q,k-1} + a_k(\lambda_q^* - \lambda_{q,k-1}) \\
\zeta_{j,q,k} &= \zeta_{j,q,k-1} + b_k(\zeta_{j,q,k}^* - \zeta_{j,q,k-1}), \quad \forall j, q
\end{aligned}$$