Endogenous Rigidities and Capital Misallocation: Evidence from Containerships

Maria Garcia-Osipenko Nicholas Vreugdenhil Nahim Bin Zahur*

November 8, 2024

Abstract

We investigate how endogenous rigidities inhibit efficient physical capital reallocation. We focus on the role of contract duration - a classic example of an adjustment rigidity. We argue that when agents choose to sign longer contracts in booms when asset markets are thin, they generate a contracting externality which further reduces available capacity and amplifies market thinness. This causes equilibrium contracts to be inefficiently long in booms and impedes the adjustment of these markets to shocks. We develop an equilibrium dynamic matching framework with booms and busts where agents search and choose how long to match for. We apply the framework to the market for containership leasing contracts, an important part of the supply chain. We (i) find substantial misallocation from the contracting externality, particularly in the transition after an aggregate shock (ii) quantify implications for the billions of dollars spent on industrial policy in the industry.

*Vreugdenhil: Department of Economics, Arizona State University (email: nvreugde@asu.edu); Garcia-Osipenko: Department of Economics, Arizona State University (email: mosipenk@asu.edu); Zahur: Department of Economics, Queen's University (email: nahim.zahur@queensu.ca). Thank you to Kevin Williams and Joris Hoste for discussing the paper, and audiences at conferences and seminars.

1 Introduction

Physical capital reallocation between firms is a central channel by which the aggregate economy adjusts to booms and busts.[1] For many forms of physical capital — like ships, drilling rigs, heavy equipment, and aircraft — capital reallocation occurs in decentralized leasing markets where agents need to search and match with trading partners. As a result, adjustment rigidities and frictions in these markets may be a cause of misallocation across the cycle. However, due to limited data, less is known about the exact process of capital reallocation in specific settings.[2]

In this paper we focus on rigidities in the form of fixed-term contracts in these leasing markets which prevent immediate reallocation of capital to its optimal use after a shock. We explore the mechanism that contract duration is an endogenous choice and can increase when markets get thinner in booms. An outcome of this mechanism is that it produces a contracting externality where even fewer assets are available, generating equilibrium contracts that are too long in booms, which leads to misallocation.[3]

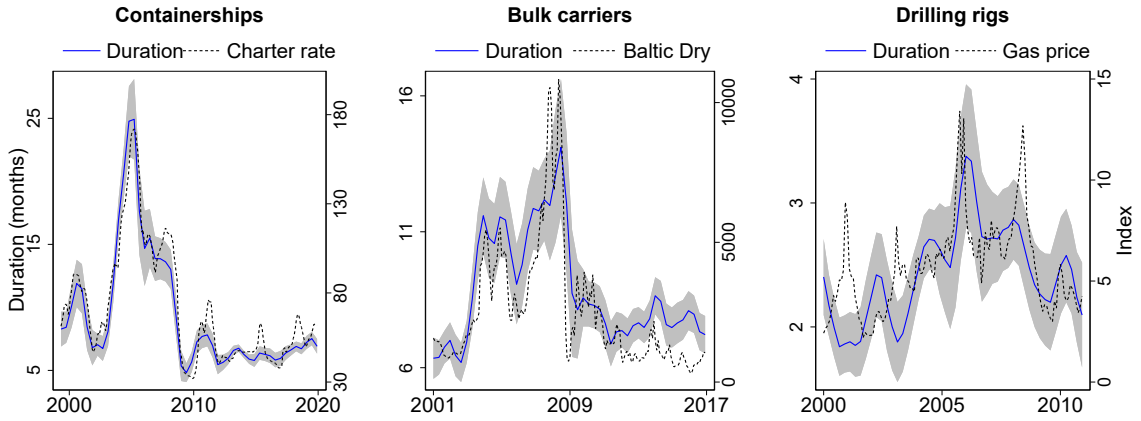
Overall, we answer the question: what are the implications of endogenous rigidities for capital misallocation as well as policy? Our starting point is a novel stylized fact about a selection of physical capital markets where firm-to-firm contract data are available: contract duration rises significantly during booms, as shown in Figure 1. Focusing on the leased containership market — an excellent example of a decentralized market with fluctuations that is also important in its own right in the supply chain — we make three empirical and conceptual contributions.

[1] For example, although the aggregate data typically do not contain information on capital leasing markets (the focus of this paper), Eisfeldt and Rampini (2006) document that reallocation through sales of physical capital alone accounts for one quarter of total capital reallocation, and is pro-cyclical.

[2] A blocker to empirical work has been the paucity of firm-to-firm data, with key information like contracts, allocations, and relationships, often confidential.

[3] Note that whether a longer duration is inefficient in booms is ultimately an empirical question and depends on the degree to which the decentralized market would internalize the contracting externality. Furthermore, (as we discuss further later) how firms need to allocate their physical capital could systematically change in a way that favors a longer duration in booms.

Figure 1: Contract duration over the business cycle for several physical capital markets



Note: See Appendix A.8 for data sources and details. We choose these markets because they have systematic time-series data on contracts (which are usually confidential in firm-to-firm markets) and because they are also important markets in their own right, comprising key parts of the supply chain. As well, we also provide anecdotal evidence from other markets (e.g. aircraft).

First, we construct a novel dataset of firm-to-firm contracts and information about the underlying capital allocations, and use these data to show descriptive evidence consistent with our main mechanism.[4] Second, based on this evidence, we develop a new empirical dynamic matching framework with booms and busts where agents choose *how long* to match for. Third, using the framework, we (i) quantify misallocation from the contracting externality (ii) quantify implications for the billions of dollars spent on industrial policy in this industry.

The majority of world trade in goods takes place via containerships and about half of these ships are rented under fixed-term contracts.[5] In this market charterers (like 'COSCO') need to lease physical capital — the containerships — from shipowners (like 'Seaspan'). Charterers then allocate ships to fixed scheduled routes and transact with downstream exporters for container slots on their routes.

[4]Unlike many firm-to-firm markets where individual contracts and other key data are often confidential (which has limited research progress), in our setting we have a rich dataset of contracts and allocations, including the exact location and amount of cargo that each ship is carrying.

[5]UNCTAD, 2021. Review of Maritime Transport 2021, United Nations Publications, New York. See https://unctad.org/system/files/official-document/rmt2021_en_0.pdf

Shipowners and charterers use time-charter contracts, which involve a negotiated 'day-rate' for a given duration. Charterers are specialized in their relationships with downstream exporters, and so face idiosyncratic demand shocks which create opportunities for reallocation. However, contract lock-in as well as search frictions (evidenced by the existence of brokers, the emergence of e-procurement, the fact that the market is fragmented on both sides, and other institutional details) make reallocation difficult. The industry is cyclical, reflecting the global business cycle.

We assemble a dataset of contracts and allocations from 2005-2015.[6] Using these data we show three key facts about the industry. First, we show that contract duration is pro-cyclical. This causes a substantial contract overhang after a crash, with charterers and ships locked into matches formed years previously in the boom.

Second, we show evidence consistent with this contract overhang preventing reallocation and generating misallocation. In our data we see observationally equivalent ships carrying systematically different amounts of cargo — measured both in volume and value — within a time period. This systematic dispersion in ship utilization suggests misallocation i.e. there are unexploited gains to reallocating ships to different firms.[7] We then document the role of contract rigidities in inhibiting efficient capital reallocation. Specifically, reallocations — defined as a ship moving to a new schedule where it is better utilized — tend to occur when a ship starts a new contract. The dispersion jumps significantly in the 2008-2009 crash, when the contract overhang is highest. Furthermore, matches in this period under longer contracts have relatively higher dispersion in utilization.

Finally, we show that contract duration responds to market thinness in the cross-section, consistent with our main mechanism: agents located in thinner markets (with fewer nearby

[6]The year 2005 is the first time that systematic satellite data are available on ship positions, while after 2015 there was consolidation of the industry into alliances; during 2005-2015 period the global market was unconcentrated on both sides. Furthermore, in this period we observe fluctuations in market conditions.

[7]As we discuss in our counterfactuals, however, we need to estimate a model to say so conclusively because some dispersion may be consistent with (constrained) efficient contracts in the presence of search frictions.

ships) choose longer contracts.

Based on the descriptive evidence, we estimate a model of the market. The model is dynamic and charterers (demand) need to search and match with ship-owners (supply). Charterers enter the market with a state-dependent valuation of a match, as well as a state-dependent probability that their value of a match will expire in each period. During booms there is a higher entry rate of charterers. Upon meeting,agents choose a contract length to maximize the total surplus of a match, given their types and the aggregate state. Longer contracts avoid agents having to search again (which may then result in agents not being matched) but may cause lock-in if the match value expires before the end of the contract. After signing a contract, prices are determined via Nash bargaining and the ship and exporter are removed from the market for the duration of the contract.

We estimate the model in two steps. In the first step we recover the demand process. In the second step we use simulated method of moments to estimate the structural parameters. The large state space (which includes the distribution of all matches under contract) in the second step results in a curse of dimensionality - we solve this by using a Moment-Based Markov Equilibrium solution concept (Ifrach and Weintraub, 2017). A key empirical challenge here that we solve is disentangling time-varying match-specific factors for why agents sign longer contracts from market thinness considerations.

We use the estimated model to perform two sets of counterfactuals. First, we quantify the inefficiency from the contracting externality. We solve for a constrained optimal social planner who can set the optimal contract length for the entire market, but is still subject to search frictions and cannot predict the exact future realizations of the demand process. Intuitively, the social planner fully internalizes the contracting externality that arises when agents sign a long contract.

We find that in booms, the planner would prefer to thicken the market with a shorter contract length than we see in the data. In fact, the optimal contract length is actually counter-

cyclical. Overall, misallocation in terms of lost output due to the contracting externality is 5.6%, but this total number masks substantial changes over the cycle. Misallocation rises in a boom and is highest (over 10%) at the start of the 2008 crash as the market transitions from a boom to a bust. but agents are locked in to long contracts.

Decomposing total misallocation, we find that the optimal contract generates higher quality matches at the expense of a smaller decrease in the total quantity of matches. However, due to search frictions, some dispersion in ship utilization from longer contracts is still efficient. We also investigate the effects of an intermediary who can eliminate search frictions and find that this would induce agents to sign much shorter contracts.

In the second set of counterfactuals we consider implications for the design of industrial policy, given the billions of dollars spent on subsidies (in the EU alone) in this industry. These subsidies target both costs typically borne by shipowners (e.g. subsidising the wages of the crews), and costs typically borne by charterers (e.g. fuel costs). In the presence of exogenous rigidities, the pass-through of these subsidies to industry profits is exactly 1.0 and a policymaker would be indifferent between subsidising either side of the market.

With endogenous rigidities, however, subsidies induce agents to sign longer contracts, which reduces the equilibrium market thinness and exacerbates search frictions. This worsens the inefficiency from the contracting externality, and reduces the pass-through of subsidies. We find counter-cyclical subsidies would be substantially more effective than a constant subsidy, since in the bust the contracting externality is lower. Furthermore, the effects are asymmetric depending on which side of the market is subsidized. The pass-through for a charterer subsidy is 0.89, but the pass-through for a ship subsidy is only 0.36 (meaning that $1 of subsidies only increases industry profits by $0.36).

1.1 Related literature

This paper is related to several strands of literature. The first is about the inner workings of decentralized asset markets. Some of these papers highlight the role of search frictions

and market thinness in determining efficient allocations e.g. Gavazza (2011a), Gavazza (2011b), Gavazza (2016). Other papers study how capital reallocation determines efficient allocations e.g. Lanteri and Rampini (2023), and Vreugdenhil (2023); and the role of adjustment costs in capital reallocation more generally e.g. Asker et al. (2014). Our contribution here is to shed light on the role of contract duration and the resulting externalities in causing misallocation.

The second strand is the literature that investigates the empirical determinants and effects of contract duration e.g. Hubbard (2001), MacKay (2022), Darmouni et al. (2023). These papers and the literature has primarily focused on a partial equilibrium analysis; our paper is focused on understanding the complete equilibrium effects of contract duration, which operates through a contracting externality. The main exception — to our knowledge – is Zahur (2024). But, while that paper focuses on a contracting externalities in a downstream spot market and the effects on investment through a hold-up channel, the economics of our paper are completely different. Concretely, we study a decentralized market for physical assets with search frictions that incorporates booms, busts, and transitions, and the contracting externality generates misallocation through distortions in the matching process.

More broadly, standard models of contract duration in search and matching frameworks in labor economics consider settings where there is one-sided commitment since employees cannot be forced to work.[8] Physical capital markets involve agreements between two firms, and so can operate quite differently with two-sided commitment and explicit agreements over duration, which is the central choice of agents in our framework. In terms of the equilibrium effects of organizational form outside labor economics, Harris and Nguyen (2023) explore the effects of relationships in the trucking industry. Our results are complementary, but formal contracts operate quite differently to informal relationships (which are zero-sided commitments). Amongst other things, long contracts formed in a

[8]Therefore, contracts must be carefully designed to retain workers with one-sided commitment in mind e.g. (Balke and Lamadon, 2022).

boom can persist well into a bust, inhibiting reallocation.

The third strand is the literature that studies shipping markets. These papers include Kalouptsidi (2014), Kalouptsidi (2018), Jeon (2022), Ganapati et al. (2023), Brancaccio et al. (2024), amongst others. Our paper incorporates key institutional details of these markets (like search frictions), but our focus is quite different in that we study trade in the fixed-term leasing market in booms and busts.

Finally, this paper is related to the broader literature in industrial organization that studies search-and-matching markets, particularly in the transportation sector. Recent examples include Fréchette et al. (2019), Brancaccio et al. (2020), Buchholz (2021), Gaineddenova (2022), Castillo (2023), Yang (2023), Rosaia (2023), and Brancaccio et al. (2023). While many of these papers are centered on markets like taxis or bulk shipping, where a match typically involves a single trip of a specific duration, in our market matches last longer with containerships operating on regular schedules involving many trips. As a result, agents use fixed-term contracts (as opposed to contracting over a single trip) and so choices about match duration are first-order.

2 Industry and data

2.1 Container shipping industry

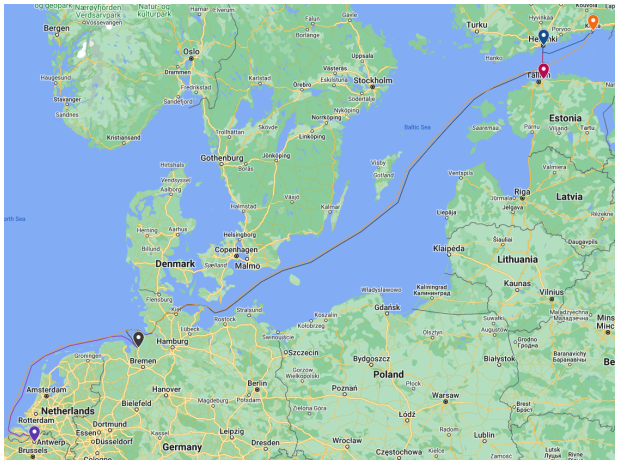
Our analysis focuses on the global container shipping industry. Containerships can be thought of as the 'buses of the ocean', typically operating on fixed schedules where they pass through a designated set of ports.[9] At each port a ship drops off and picks up a portion of its containerized cargo. Figure 2a demonstrates a map with the route for one of the containerships in our sample.

The value chain of the maritime shipping industry is visualized in Figure 2b. Cargo owners (or exporters) are the enterprises that use maritime transport providers and other suppliers

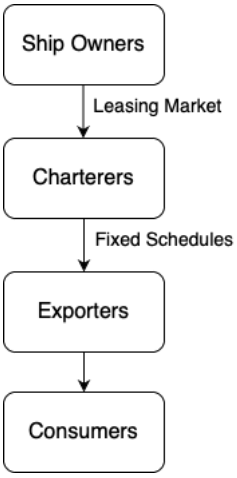
[9]This is in contrast to the dry bulk shipping industry, where ships operate more like taxi cabs and make voyage decisions on a trip-by-trip basis, often travelling without any cargo at all or in "ballast" (Brancaccio et al., 2020).

Figure 2: Container Shipping Industry Structure

(a) Container Ship's Course Visualization



(b) Value Chain



Note: Panel (a): This shows a geographical map of a ship in our data performing its scheduled stops. Here, the ship has the following schedule: Antwerp - Bremerhaven - Muuga - Helsinki - Kotka - Antwerp. Note that the ship uses the Kiel canal through Denmark. Panel (b): One component not pictured (but mentioned in the text) is that charterers can also own ships. These ships are not available on the leasing market, but can be allocated to the fixed schedules downstream.

of services to import/export their cargo. Containerships are operated by carriers (also known as liner companies) that specialize in the transport of containerized goods across the world (for example, COSCO). Carriers compete by setting up shipping itineraries or schedules. Each itinerary specifies a set of ports to be visited as well as a service frequency, which together determine the number of ships required to operate it (Stopford, 2009).[10] Running any such "liner service" poses complex operational problems, since the carrier has to address challenges posed by uncertain conditions at sea, port congestion and disruptions, seasonality of demand, and cargo imbalances.[11] As such, it is difficult for carriers to simultaneously maintain both high schedule reliability as well as high levels of capacity utilization (Stopford, 2009).[12]

[10]For example, if it takes 28 days for a ship to visit each of the ports on a given schedule, and the carrier wishes to run a weekly service, a minimum of four ships will be required.

[11]Cargo imbalances arise since demand may be higher on some legs of the itinerary than others, forcing the ship to sail with less cargo on other legs.

[12]For example, in 2019, average on-time arrival was only 47% (Zhang et al., 2022).

About half the containerships in the world are owned by the carriers themselves; these "owner-operated ships" are rarely leased out to other companies, and form the core of the fleet of carriers. The remaining "charter-operated ships" are owned by shipping companies (who we refer to as "shipowners") that do not themselves provide container shipping services, but instead specialize in leasing out these ships to the carriers (or "charterers"). It is this leasing market that we focus on.

Since containerships are a highly movable form of capital, practitioners treat the leasing market as a global market. The market is unconcentrated and fragmented, with a large number of agents searching on both sides of the market.[13] In our study period 2005-2015, the HHI for charterers is 415, and 136 for ship-owners who lease out their ships. Note that since 2015 — outside the period of our study — many charterers have consolidated into alliances. So, in order to focus on a period during which the market structure was relatively constant, we limit the scope of our analysis to the pre-2015 period.[14]

2.1.1 Leasing contracts

Shipowners lease to charterers using 'time-charter contracts'. These contracts are relatively simple and follow a standard template called "Boxtime 2004" (BIMCO, 2004). In this template, each contract specifies a 'dayrate' (the price per day to lease the ship that the charterer pays the shipowner) for a specified a duration. Shipowners pay the operating costs of the vessel (i.e. the crew, maintenance and repair), while charterers pay voyage expenses (i.e. bunker fuel, port charges, canal dues, and cargo-handling costs). Note that contracts do not restrict the route where charterers allocate the ship.

Contracts can be extended. However, contracts typically do not have explicit options for extensions: only 3.3% of contracts provide charterers a formal option to renew (Appendix

[13] This is contrast to downstream container shipping markets, which are regionally segmented and therefore more concentrated. In these markets, carriers are able to exercise significant market power when transacting with exporting firms, as shown by Hummels et al. (2009) and Ardelean and Lugovskyy (2023).

[14] There did exist alliances in the 2005-2015 period, but these were small. Re-computing the HHI for the charterers based on alliances only increases the index to 611 from 415. As a result, we abstract away from these alliances or market power considerations in our analysis.

A.7). Instead, these extensions are negotiated on an individual basis by the charterer if the ship does not have a subsequent contract. Similarly, contracts mostly do not include clauses for re-negotiation; re-negotiations of the contract (such as an early termination) must be agreed to by both the shipowner and charterer, and are reportedly rare.[15]

In theory charterers can sublet leased ships to other carriers. In practice subleases are rare and account for only 2.1% of all contracts (Appendix A.7). Two institutional factors suggest why subleasing is not widely used. First, charterers are generally unwilling to sublease to their competitors in downstream shipping markets.[16] Second, subleasing requires charterers to operate on the other side of the leasing market, with which they may not be experienced.

2.1.2 Reallocations

There are two types of reallocation in this industry. The first is the reallocation of a ship between different charterers which occurs across contracts. The second is the physical reallocation of a ship to different routes which, in theory, could happen within or across contracts. However, as we later show in Section 3.2, the physical reallocation of a ship typically occurs at the same time as a contractual reallocation; there is little within-contract reallocation. This is not surprising since charterers typically follow fixed shipping itineraries that they only infrequently adjust (Stopford, 2009).

What generates profitable opportunities for reallocation? Part of the answer is that charterers specialize geographically and are therefore differentially exposed to route-level demand shocks.[17] This also explains why ships tend to be reallocated across routes when there is a change in contract/charterer. But even charterers operating on the same route have heterogeneous relationships with downstream exporters (Ardelean and Lugovskyy,

[15] We are only aware of a handful of cases of early termination by mutual consent where the charterer was faced with financial difficulties: see https://www.seanews.com.tr/zim-downsizes-its-container-ship-fleet-as-demand-wanes/197918/.

[16] See, for example, Wackett (2021).

[17] To give one example, the Japanese carrier Mitsui OSK primarily operates services on the intra-Asia and trans-Pacific shipping routes.

2023). They also differ in how frequent the sailings are, which ports the ship stops along the way, and speed of transit (Stopford, 2009). Ex-post, this differentiation exposes charterers to idiosyncratic shocks: for example, port disruptions and port congestion (Brancaccio et al., 2023) will have differential effects on charterers depending on which ports are included in their itinerary. Likewise, when demand declines, a charterer offering weekly service may be hit harder than one offering monthly service on the same route, as filling the ship to capacity becomes more challenging. As a result, reallocation is valuable when a ship moves from a charterer with a bad shock who no longer needs a ship, to a charterer who does need a ship.

Reallocation comes with switching costs, in the form of lost time when a containership needs to move to a new contract. We include switching costs in the model, but these are small relative to the typical duration of a contract.[18]

2.1.3 Search frictions

Three features point to significant search frictions in the industry. The first is the presence of specialized ship brokers which, in itself, is indicative of search frictions (Brancaccio et al., 2020). Since the market is fragmented with a large number of agents searching on both sides, these brokers undertake matching on behalf of their shipowner and charterer clients. Common brokers are Bertling, Clarksons, and Maersk Broker.

A second feature is that the matching process is unstructured — that is, there is no centralized platform — leading to instances where matches do not form despite searching agents on both sides of the market. Information about available ships and searching charterers is usually shared via e-mail and brokers receive large number of e-mails each day.[19] As one manager in the industry observed: "Many shipping companies face e-mail

[18] For example, the standard transit time from the USA to Singapore is approximately 15 to 25 days, and from China to the UK between 35 and 45 days, whereas the typical leasing contract duration is 9 months. See https://www.maersk.com/logistics-explained/transportation-and-freight/2023/09/27/sea-freight-guide.

[19] This is not unique to the container shipping leasing market: Brancaccio et al. (2023) discuss how brokers in the dry bulk industry report receiving many thousands of emails every day.

overload - literally hundreds or thousands of e-mails each day. Failing to catch key operational information or an urgent e-mail from a broker can have a toll on a business."[20]

Finally, since the end of our study period in 2015, there have been attempts to use technology to improve matching. Concretely, participants have begun to create online centralized platforms (sometimes marketed as an 'Uber for ships') to better connect available ships and charterers.[21] The entry and adoption of these platforms suggests inefficiencies in the matching process in the study period.

2.2 Data

We use two main data sources. The first is data on containership time-charter contracts from Clarkson's. We provide more details about the dataset construction, and how we merge contracts with shipping movement data, in Appendix A.1. Our full dataset covers the period from 1999 to 2023, though the focus of our analysis is on the 2005 - 2015 period.

The other key dataset we utilize in our analysis is port call data from 2005-2015 provided by Lloyd's List Intelligence. This dataset provides information on the universe of port calls, including the dates of arrival and sailing, and the locations of the ports visited on each call. For a subset of port calls, we also observe the ship's "draught", or the vertical distance between the waterline and the bottom of the hull. A ship that is carrying more cargo will sink deeper into the water, causing its draught to increase; we can therefore use draught data to infer how much cargo the ship is carrying and measure capacity utilization (see Appendix A.2).

Our dataset includes a total of 1,655,140 port calls, with 299,903 of them matched with contracts data. In addition, as a proxy for aggregate demand, we use the container-ship time-charter rate index, a monthly index published by Clarksons based on their assessments

[20] Source: https://thedigitalship.com/news/maritime-software/item/3133-email-parser-for-charterers

[21] See https://www.pwc.com/gr/en/industries/Chartering_marketplaces.pdf for a description of these platforms and the connection with Uber's marketplace.

Table 1: Summary statistics for leasing contracts for 2005-2015

Variable	Obs	Mean	Std. dev.	Min	Max
Duration (months)	3,327	9.3	11.1	0.1	144.0
Contract dayrate ($/day)	3,327	10,233	6,138	2,600	47,000
Ship age (years)	3,327	8.6	5.7	0	29.0
Time-charter rate index	132	72.3	37.0	32.0	171.8
Capacity utilization	872,069	0.55	0.22	0	1
Reallocation	1,655,140	0.02	0.14	0	1

of the rates of newly negotiated time-charters. Table 1 provides descriptive statistics on the duration of the contracts, dayrates (contract price), age of the ships contracted, and capacity utilization in 2005 - 2015.

Reallocations We use the port call data to identify when a ship reallocates from one itinerary to another. This is challenging since we do not directly observe shipping itineraries in our raw data. Instead, what we observe are repeated sequences of port calls. We identify reallocations as large deviations across space from old port call sequences to a new set of port call sequences (where we use a threshold of 1000 km to define large deviations). We describe the algorithm for constructing reallocations in Appendix A.4. The threshold of 1000 km is not especially restrictive as ships are often moved across large geographical distances. The average distance when a ship is reallocated is 3100 km, and 25% of reallocations involve a transit of more than 4000 km from the original itinerary to the new itinerary.

2.2.1 Dispersion in capacity utilization and misallocation

We infer misallocation of ships from cross-sectional dispersion in capacity utilization, which we observe from the port call data. If a ship leased to charterer A is systematically under-utilized, whereas during the same period an otherwise equivalent ship leased to a different charterer B is fully utilized, we argue this is indicative of misallocation - industry output would increase if the first ship were reallocated from charterer A to B. This is consistent with the way that practitioners view fleet efficiency in the industry e.g., Adland

et al. (2018) and UNCTAD (2010).

To operationalize this argument empirically, we have to address a few issues. First, due to cargo imbalances, the amount of cargo a ship is carrying may differ across port calls depending on which direction of the route the ship is travelling on and which port it is visiting.[22] Second, ships may differ in capacity utilization not because of who they are allocated to, but because of underlying characteristics such as size and age (e.g., older ships may be less utilized as they need to undergo repair and maintenance more frequently). Third, capacity utilization also naturally varies over time with aggregate demand.

To deal with the first issue, we aggregate utilization to the ship-month level: while a ship's utilization on a given port call may vary due to cargo imbalances, these cargo imbalances are likely to cancel out when we aggregate over a sufficiently high number of port calls.[23] We then (i) residualize the capacity utilization of each ship by ship fixed effects, which controls for any heterogeneity in utilization arising from underlying ship characteristics (addressing the second issue) (ii) compute the standard deviation of residualized capacity utilization in each month. Note that cross-sectional dispersion in residualized utilization nets out the effect of any aggregate demand shifters that vary across time periods. As such, dispersion within a time period in this measure indicates that there are more productive matches available, but ships are not matched to them, which is indicative of misallocation.

Dispersion is large in magnitude, with the standard deviation in residualized capacity utilization at the monthly level ranging from 0.14 to 0.19 (relative to an average capacity utilization of 0.55). We also regress residualized capacity utilization on route-by-year month fixed effects, and find that the R^2 from that regression is only 0.025; thus, there is considerable cross-sectional dispersion in utilization even among ships operating on the

[22] For instance, on the Trans-Pacific trade route, there is typically considerably more demand when ships are travelling from Asia to North America than in the reverse direction.

[23] As a robustness check, we aggregate ship utilization to the ship-quarter level, since within a three month period almost every ship will have completed at least one round-trip. Results using this measure are very similar.

same trade route during the same time period.[24] Adding charterer fixed effects to this regression only marginally increases the R^2 to 0.082, implying dispersion is not driven by systematic differences across charterers. Instead, these patterns suggest that dispersion in utilization arises primarily due to idiosyncratic shocks. We investigate the link between boom-bust cycles, contracting and dispersion in utilization in Section 3.2.

We also obtain a revenue-based measure of dispersion in utilization because downstream exporter prices are heterogeneous across routes. Therefore, in theory, it may be efficient to deploy ships on routes where these prices are high, regardless of physical utilization. To construct our measure we multiply residualized capacity utilization by the freight rate of the itinerary the ship is currently operating on (Appendix Section A.5 provides more details). We then measure dispersion in the revenue across ships within the same time period. Note that, as we discuss further in Section 3.2, our results are similar regardless of whether we use physical or revenue-based dispersion in residualized capacity utilization.

3 Descriptive evidence

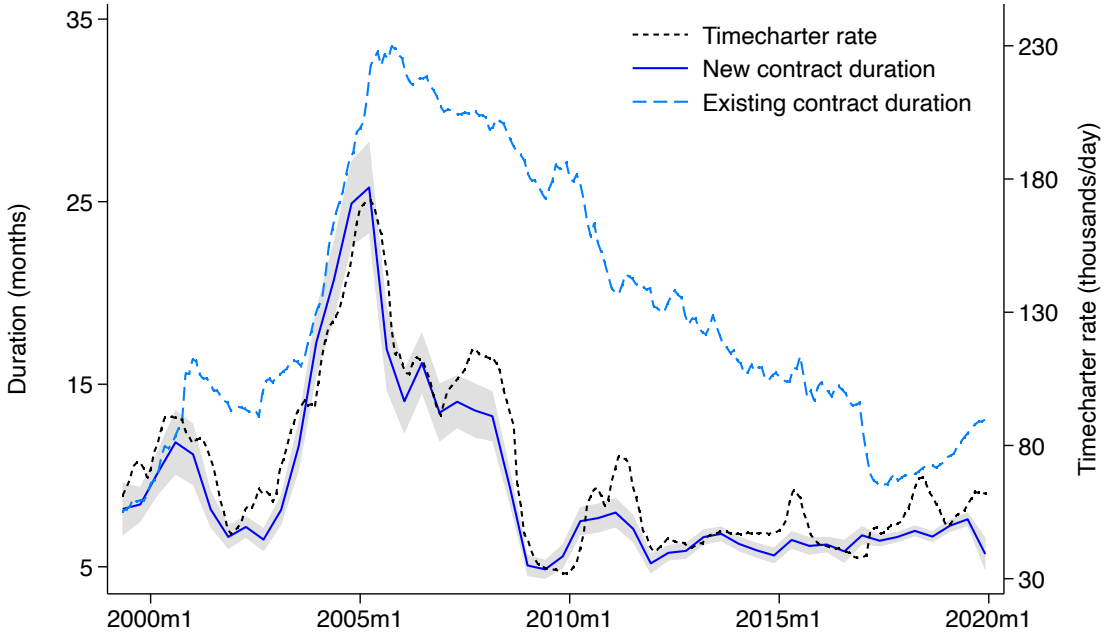
We describe three key empirical observations that underlie our analysis.

3.1 Observation 1: New contract duration increases in booms, leading to substantial contract overhang after a market crash.

Figure 3 shows how the average duration of newly signed leasing contracts changes over time. In the same figure, we also plot the containership 'timecharter rate index', which is a measure used within the industry to index whether the market is in a boom or a bust. The market is highly cyclical and firms sign significantly longer contracts during booms. This effect is especially pronounced during the boom in the mid-2000s, when the average duration of newly signed contracts increases from around 8 months to more than 24 months. More recently, during the upsurge in container shipping demand after the Covid-19 pandemic, there was a similar increase in the duration of new contracts (Appendix Figure

[24] See Appendix Section A.5 for how we identify which route a ship is on.

Figure 3: Average duration over the cycle: new vs existing contracts



12).[25] Furthermore, the correlation in Figure 3 is not driven by long contracts having systematically higher prices (Appendix Table 8).

Figure 3 also plots the duration of all existing contracts at each time period. This illustrates the long-run effects of a boom when there is pro-cyclical contract duration. Concretely, the boom causes an overhang of existing contracts (and the corresponding matches) that persists for years after the market crashes. The difference between new versus existing contracts is most pronounced during the "Great Trade Collapse" of 2008 - 2010. We now turn to how this contract overhang affected the reallocation of physical capital in the market.

[25]Note that later in the model we directly estimate the demand process rather than use the price index.

3.2 Observation 2: Evidence consistent with contract overhang inhibiting reallocation

3.2.1 Counter-cyclical dispersion in ship utilization

Figure 4a illustrates that dispersion in capacity utilization is counter-cyclical, rising significantly during the "Great Trade Collapse".[26] Table 2 documents that dispersion in revenue is similarly counter-cyclical in this period (Column (2)). The increase in dispersion is not simply explained by differential impacts of the recession on demand in different trade routes, since we find a similar increase in dispersion even if we control for route fixed effects when measuring dispersion (as shown in Columns (3) and (4)).[27]

We now explore two reasons that jointly explain why dispersion in utilization rises in the crash in this industry: (i) idiosyncratic effects of the crash across charterers (ii) contract lock-in as a factor that prevents the immediate reallocation of ships away from charterers with bad shocks to allocations where they would be better used.[28]

In terms of the first reason, the recession had heterogeneous impacts across charterers. The relationships between charterers and downstream exporters are governed by confidential contracts, resulting in significant price dispersion for downstream exporters, even on the same route (Ardelean and Lugovskyy, 2023); therefore, cargo demand faced by charterers responded heterogeneously to the recession. Moreover, charterers also differed in idiosyncratic ways in operational aspects, such as round-trip frequency and ports visited, and in their ability to respond to the recession by engaging in "slow steaming" or skipping some stops (Notteboom et al., 2021). As a consequence, the crash had very different financial consequences for charterers (Notteboom et al., 2021): for example, CMA CGM's

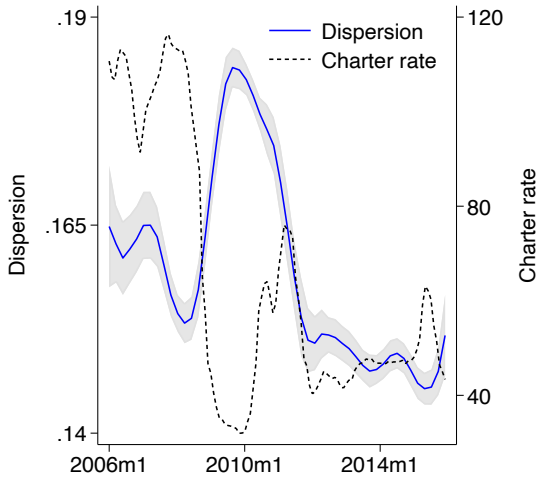
[26] Note that we see the same counter-cyclical pattern if we instead measure dispersion using the difference between the 10th and 90th percentile of utilization.

[27] One might wonder if, given that capacity utilization is bounded above at 1, dispersion would be mechanically lower during the boom if most ships were fully capacity utilized. However, capacity utilization equalled 1 for only 0.3% of ship-year observations during the boom, and was higher than 0.95 for 1.1% of observations.

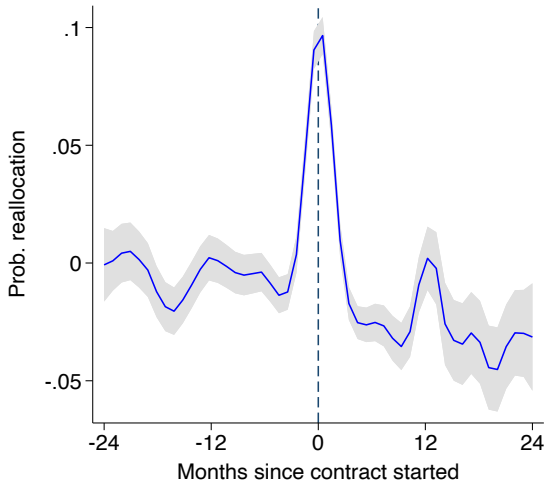
[28] Note that if (i) were the case, but ships could be costlessly redeployed between charterers, one expect dispersion in utilization to stay about the same following the recession.

Figure 4: Evidence consistent with contract overhang inhibiting reallocation

(a) Counter-cyclical dispersion in utilization



(b) Contract rigidities inhibit reallocation



Note: Panel (a): Plots dispersion in residualized utilization across ships every month, where residualized utilization is obtained by regressing utilization on ship fixed effects and a time trend. Panel (b): The figure plots the share of containerships reallocated every month at different in a 4 year window around the month when the ship was reallocated, after controlling for ship fixed effects and a time trend.

profits in 2008 dropped by 87%, whereas Yang Ming's profits decreased by 20%.

We next document evidence that contracting rigidities prevent the immediate reallocation of ships.

3.2.2 Contract rigidities inhibit reallocation

To examine whether being locked into a contract inhibits reallocation, we look at how the probability that a ship is reallocated to a different itinerary changes over the lifetime of a contract. If contracts did not inhibit reallocation, one would expect reallocations to happen more or less independently of which stage of the contract the ship happens to be in. Figure 4b shows, however, that the probability of reallocation jumps at the start of the contract (controlling for ship fixed effects and a time trend). We find little evidence of within-contract reallocation. This is likely because, as discussed earlier, charterers specialize by geography and it is costly for them to establish a new liner service (Haralambides, 2019).

Table 2: Dispersion in Utilization Rises Following the Bust

Panel A: Dispersion in Utilization, All Ships

	Baseline		Control for Route	
	(1)	(2)	(3)	(4)
	S.d., utilization	S.d., revenue	S.d., utilization	S.d., revenue
Great Trade Collapse	1.497	1.184	1.599	1.250
(July 2008 - End-2010)	(0.248)	(0.496)	(0.250)	(0.464)
N	120	120	120	120
R-squared	0.522	0.293	0.508	0.278

Panel B: Dispersion in Utilization: Ships Under Longer vs. Shorter Contracts

	Contracts $\geq$ 12 months		Contracts $<$ 12 months	
	(1)	(2)	(3)	(4)
	S.d., utilization	S.d., revenue	S.d., utilization	S.d., revenue
Great Trade Collapse	3.657	4.092	0.492	0.267
(July 2008 - End-2010)	(0.431)	(0.724)	(0.278)	(0.512)
N	120	120	120	120
R-squared	0.412	0.254	0.389	0.189

Note: Each observation is an year-month. The regressions also include a constant and a fixed effect for the post-2010 period. We consider two measures of dispersion in utilization. "S.d., utilization" is the standard deviation of residualized capacity utilization in each year-month (where utilization is normalized to range from 0 to 100), while "S.d., revenue" is the standard deviation of residualized revenue in each year-month, where revenue is the product of utilization and freight rate index.

In Panel A, we measure dispersion in these measures across all ships. In our baseline analysis (columns (1) and (2)), residualized utilization is obtained by regressing utilization on ship fixed effects and a time trend and recovering the residuals from that regression; residualized revenue is calculated in the same fashion. In columns (3) and (4), we also control for route fixed effects when residualizing utilization and revenue.
In Panel B, we measure dispersion separately for ships under longer contracts (Columns (1) and (2)), where the contract was signed at least 12 months prior, and for ships under shorter contracts (Columns (3) and (4)), where the contract was signed within the last 12 months. In Panel B, utilization and revenue are residualized by ship fixed effects and a time trend.

In Panel B of Table 2 we document additional descriptive evidence about the connection between contract overhang and counter-cyclical dispersion in utilization. We show that the rise in dispersion during the bust was concentrated among ships that were locked into longer contracts (at least 12 months old); among ships with more recently signed contracts, the increase in dispersion is much smaller and statistically insignificant.

3.3 Observation 3: Longer contracts are associated with thinner markets

How do firms choose their contract length? We present descriptive evidence that longer contracts are associated with thinner markets in Table 3. This consideration needs to be weighed against the possibility of lock-in which would favor a shorter contract duration, where conditions change and one party to the contract would prefer to break up the match.

In Table 3 we focus on isolating this relationship between thinness and duration in the cross-section, controlling for time-varying factors that might also affect contract duration. To do so we construct a measure for geographical market thinness for each match by counting the number of alternative ships nearby in time and space. We discuss the details about how we construct this measure in Appendix A.6. Intuitively, the thinness measure traces out different regions across the earth, with geographically isolated regions producing lower measures than trading hubs. We emphasize that this measure is not inconsistent with the view that the industry is ultimately a global market. However, since charterer needs may be time-specific and it can take a few weeks to move a ship across the world, the measure captures that it is more convenient to procure a ship nearby.

Overall, the results across the four specification in Table 3 show that thinner markets (with fewer nearby ships) are associated with longer contracts. As well as being statistically significant, the results are also economically significant: moving from the 1% quantile to the 99% quantile of market thinness results in about a 33% increase in contract duration. This increase in duration is robust across specifications. Although this number is relatively large, note that the effects of market thinness may be even larger across the cycle compared to within a time period, which is the primary focus of the paper. This is because the thinness

Table 3: Longer contracts are associated with thinner markets

	(1) Log(duration)	(2) Log(duration)	(3) Log(duration)	(4) Log(duration)
Log(# ships nearby, same type)		-0.10***		-0.11***
		(0.032)		(0.032)
Log(# ships nearby, all types)	-0.14***		-0.15***	
	(0.036)		(0.035)	
Log(time charter index)			0.85***	0.85***
			(0.040)	(0.040)
Year fixed effects	Yes	Yes	No	No
Observations	2,803	2,803	2,803	2,803
Adjusted R^2	0.137	0.136	0.158	0.157

Note: Standard errors in parentheses. We also include controls for ship heterogeneity (ship size fixed effects). We exclude very large ships (Post-Panamax or larger), newly built ships, as well as contracts longer than 8 years, to stay consistent with the estimation sample we use for the structural model. Our results are robust to these controls and sample restrictions, as shown in Appendix A.9.

of the regional markets is highly correlated across booms and busts.

3.4 Alternative explanations and factors for pro-cyclical contract duration

Insurance motives Firms may prefer longer lease contracts to manage risk. Longer contracts guard against the risk of not being able to find another match upon contract expiry, which is a core mechanism that we focus on and explicitly model. In addition, longer contracts may also be used to hedge against price risk, a channel not present in our model, since we assume risk-neutral firms. We make this assumption since shipping firms have ready access to financial derivatives (such as forward freight agreements) for price risk mitigation (Adland et al., 2019). Furthermore, if hedging against price risks were a central driver of pro-cyclical duration, we would expect prices of longer contracts to be less cyclical (and less volatile) than prices of shorter contracts; as Appendix Figure 14 shows, however, the prices of long and short contracts are similarly pro-cyclical.

Obtaining finance for new ships Some shipowners sign long-term leasing contracts with charterers when they acquire new ships, since the guaranteed contract revenue allows

them to obtain external finance at a cheaper cost.[29] This suggests a potential alternative explanation for pro-cyclical contract duration: during booms, demand for new ships rises, and lease contracts lengthen to facilitate financing. We find, however, that the pro-cyclical lengthening of contracts is of the same magnitude even if we control for ship age or drop new ships (Appendix Table 8). This is because the vast majority of our contracts are for existing ships; only 6.2% of leasing contracts are for brand-new ships, and we drop all these observations from the estimation sample (as we discuss later).

4 Model

4.1 Setup

Time is discrete with each period a month denoted by t. Agents are risk-neutral and forward-looking, and have discount factor β.

Agents There are two types of agents in our model, shipowners and charterers. We assume that each shipowner owns a single ship, and so refer interchangeably to ships and shipowners.[30] There are n_t homogeneous containerships. Shipowners lease out ships to charterers and we assume that each charterer requires a single ship.[31]

Payoffs Agents choose how long to match for; denote the contract length by τ. Every period the value of a match is given by $\pi_{i,t} - c$. Here, c is the cost of operating the ship borne by the shipowner. The component $\pi_{i,t}$ is charterer i's value of a match.[32] This $\pi_{i,t}$ is stochastic and subject to idiosyncratic shocks. With probability η_t the value is v_t and we say that the charterer is "alive". With probability $(1 - \eta_t)$ the value is 0 forever, and we say that the charterer is "dead". The parameters v_t and η_t are largely determined by the

[29] One example is described in https://splash247.com/ship-finance-international-acquires-three-mega-containerships/

[30] In 2015, the largest 25 shipowners only owned 41% of all chartered containerships (Clarksons, 2015). Thus shipowner market power is unlikely to be a first-order issue.

[31] Note that, in practice, charterers may own or lease other ships, and so one may be concerned that this could generate a more complicated portfolio problem. To keep the focus on the contract length decision we choose to not explicitly model this complex decision (if there is one). Rather, as we elaborate on later, the charterer's value of a match and idiosyncratic shock will embed such considerations.

[32] In practice, shipowners pay the operating costs of the vessel (e.g, crew), which are part of c. Charterers pay all voyage expenses (e.g., bunker fuel) and cargo-handling costs; these costs therefore affect $\pi_{i,t}$.

downstream market, where the charterer schedules a ship and sells the container slots to exporters. Note that v_t and η_t are allowed to be potentially state-dependent.

The fact that the $\pi_{i,t}$ shocks are idiosyncratic to each i is important because it generates profitable reallocation opportunities. Concretely, it generates the situation where a ship is contracted under a dead match but could be profitably re-matched to another charterer. What are these shocks? As we elaborated on in Section 2.1, charterers spatially specialize, have heterogeneous relationships with downstream exporters, and differ in schedule characteristics (service frequency, ports visited, and transit speeds) leaving them differentially exposed to demand shocks and other disruptions (such as the 2021 obstruction of the Suez Canal, or disruptions to port operations). Denote $\hat{\eta}_{t,k} = \prod_{l=0}^{k} \eta_{t+l}$. Then, the expected total value of a τ-length contract at time t is:

$$m_{t,\tau} = \mathbb{E}_t \sum_{k=0}^{\tau-1} \beta^k (\hat{\eta}_{t,k} v_{t+k} - c) \tag{1}$$

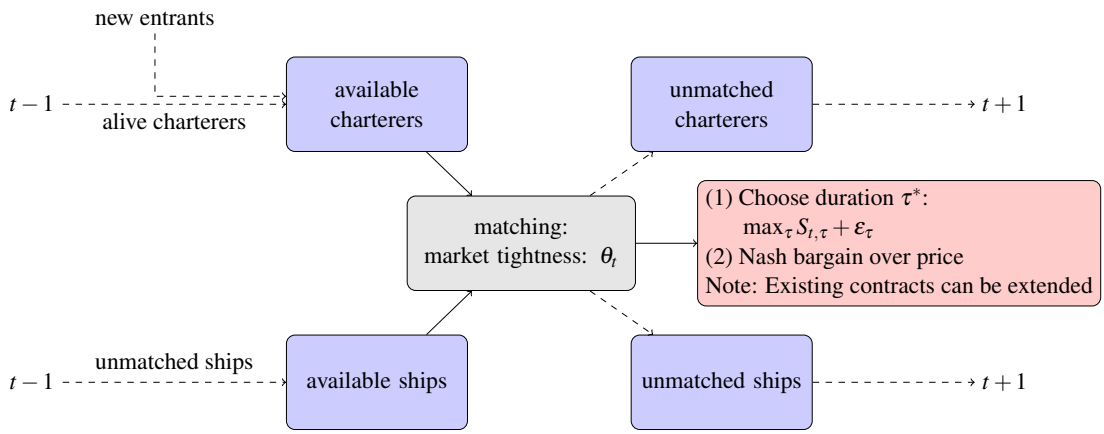
At the start of the contract, the ship has to be relocated from its previous location, so the charterer also incurs a one-time switching cost c_t^s on the initial contract.

Timing In each period, the timing is as follows. We also depict the timing visually in Figure 5.

1. *Match expiry and charterer exit:* The valuation v_t remains positive with probability η_t. If a charterer is not under contract and no longer alive, then they will exit.

2. *Contract extensions and entry*: Existing contracts that are ending are potentially extended with probability $\mathbb{P}_{\text{extend}}$ if the charterer is still alive. Otherwise the ship, and the charterer (if still alive), are added to the pools of available ships and available charterers. As well, e_t charterers enter the market, where e_t is a function of z_t and may vary over the boom-bust cycle (for instance, there may be more entry during a boom).

Figure 5: Timing within each period



3. *Search and matching:* Available charterers search and match with ships via random search.

4. *Choice of contract duration and price:* If agents are matched, they choose both how long to match for (the contract duration), and a fixed price paid by the charterer to the ship.[33] The ship and the charterer choose the contract duration to maximize their joint surplus from matching. Prices are then determined by Nash bargaining where $\delta \in [0,1]$ is the Nash bargaining weight of the ship.

5. *Output* The per-period output of the match (i.e., $(v_t - c)$ for alive matches minus switching costs) is realized.

Search and matching Matching takes place in a single global market.[34] The mass of available charterers equals $a_t^{charterer}$. The mass of available ships equals a_t^{ship}. We use a matching function to characterize the outcome of the matching process. The number of successful matches equals $m(a_t^{charterer}, a_t^{ship})$, where m is increasing in both of its arguments.

[33] This assumption is consistent with the fact that in practice, at the time of contract, the parties agree to a daily charter rate that is fixed over the duration of the contract.

[34] This is motivated by our finding that ships are often reallocated large distances (exceeding 1000 km) at the beginning of a new contract, so the set of possible matches is not necessarily constrained by the current location of the ship.

We assume there are constant returns to matching, consistent with prior literature (for example, Brancaccio et al., 2020). Let $\theta_t = a_t^{ship}/a_t^{charterer}$ denote the market thinness (the ratio of searching ships to searching charterers). Under the assumption of constant returns, the probability of finding a match is a function only of θ_t; denote these probabilities $q_t^{charterer}$ and q_t^{ship}.

4.2 Contract duration choice

Match surplus An important component to the contract duration choice is the match surplus $S_{t,\tau}$. This is the joint value to the charterer and ship from a τ-period contract at the time t state, minus their outside options. Note that a contract can be either an initial contract or an extension, and for the initial contract the total surplus also incorporates the switching cost i.e. it is $-c_t^s + S_{t,\tau}$. Here, $S_{t,\tau}$ is:

$$\begin{aligned} S_{t,\tau} = & \underbrace{m_{t,\tau}}_{\text{Value of } \tau\text{-period contract}} \\ & + \underbrace{\beta^{\tau}\mathbb{E}_t\left(\hat{\eta}_{t,\tau}(M_{t+\tau}^{ship} + M_{t+\tau}^{charterer}) + (1-\hat{\eta}_{t,\tau})U_{t+\tau}^{ship}\right)}_{\text{Continuation values after being matched}} \\ & - \underbrace{\beta\mathbb{E}_t\left(U_{t+1}^{ship} + \eta_t U_{t+1}^{charterer}\right)}_{\text{Outside options}} \end{aligned} \tag{2}$$

Here, M_t^{ship} and $M_t^{charterer}$ denote the value functions for ship and charterers at the end of the initial contract if the charterer is still alive.As we discuss further below, M_t^{ship}, $M_t^{charterer}$ embed that the match may be extended if the charterer is still alive. The components U_t^{ship} and $U_t^{charterer}$ denote the value functions for an unmatched ship and charterer. In the event of disagreement, both the charterer and ship need to wait until the next period, when they may enter the pool of searching agents.

Choice of contract duration After matching, agents choose the contract duration based on the match surplus and an idiosyncratic shock to the value of signing a τ-period contract, ε_τ, which is drawn from an i.i.d type-1 extreme value distribution with scale parameter

σ. We are dispensing with ship-specific and charterer-specific notation, but the ε_τ are idiosyncratic to each particular match, as well as to each contract length.

Denote W_t as the ex-ante surplus (the expected value of the surplus before the ε_τ are drawn):

$$W_t = \mathbb{E}_\varepsilon \left[\max_{\tau \in \{1,2,\ldots,\tau_{max}\}} \{S_{t,\tau} + \sigma\varepsilon_\tau\} \right] \tag{3}$$

where τ_{max} is the maximum possible contract duration (in practice, 48 months). For the initial contract, since Equation 3 is the inclusive value and the switching cost is an additive value, the ex-ante surplus is $-c_t^s + W_t$.

Let $P_{t,\tau}$ denote the probability that a matched charterer-ship pair chooses a contract of length τ:

$$P_{t,\tau} = \frac{exp(S_{t,\tau}/\sigma)}{\sum_{\tau'} exp(S_{t,\tau'}/\sigma)} \tag{4}$$

Note that it is individually rational for both ships and charterers to choose the contract duration that maximizes the surplus of a match. This is because Nash bargaining implies perfectly transferable utility. Furthermore, since the switching cost enters additively into all options for the initial contract, this choice probability is the same both for initial contracts and extensions.

Other value functions A searching charterer in period t will find a match with probability $q_t^{charterer}$, in which case they will split surplus via Nash bargaining. Otherwise they will wait until the next period. In Appendix B.1 we prove that the value of unmatched ships and charterers can be written as:

$$U_t^{charterer} = q_t^{charterer} \underbrace{\left((1-\delta)(W_t - c_t^s) + \beta\eta_t \mathbb{E}_t U_{t+1}^{charterer} \right)}_{\text{Charterer's expected payoff if matched}} + (1 - q_t^{charterer}) \beta\eta_t \mathbb{E}_t U_{t+1}^{charterer} \tag{5}$$

$$U_t^{ship} = q_t^{ship} \underbrace{\left(\delta(W_t - c_t^s) + \beta \mathbb{E}_t U_{t+1}^{ship}\right)}_{\text{Ship's expected payoff if matched}} + (1 - q_t^{ship})\beta \mathbb{E}_t U_{t+1}^{ship} \tag{6}$$

The functions M_t^{ship} and $M_t^{charterer}$ have a similar form, with a modification to the probability an agent is matched (explained more below) that incorporates extensions:

$$M_t^{charterer} = \hat{q}_t^{charterer} \underbrace{\left((1-\delta)W_t + \beta \eta_t \mathbb{E}_t U_{t+1}^{charterer}\right)}_{\text{(Expected) payoff if matched}} + (1 - \hat{q}_t^{charterer})\beta \eta_t \mathbb{E}_t U_{t+1}^{charterer} \tag{7}$$

$$M_t^{ship} = \hat{q}_t^{ship} \underbrace{\left(\delta W_t + \beta \mathbb{E}_t U_{t+1}^{ship}\right)}_{\text{(Expected) payoff if matched}} + (1 - \hat{q}_t^{ship})\beta \mathbb{E}_t U_{t+1}^{ship} \tag{8}$$

Here, we incorporate extensions in the objects $\hat{q}_t^{charterer}$, $\hat{q}_t^{ship}$ in the following way. We model the probability that a previously matched ship will continue to be matched as: $\hat{q}_t^{ship} = \mathbb{P}_{\text{extend}} + (1 - \mathbb{P}_{\text{extend}})q_t^{ship}$, where $\mathbb{P}_{\text{extend}}$ is a parameter to be estimated. Likewise, we model the probability that a previously matched (and still alive) charterer will continue to be matched as: $\hat{q}_t^{charterer} = \mathbb{P}_{\text{extend}} + (1 - \mathbb{P}_{\text{extend}})q_t^{charterer}$.

Our formulation for extensions allows for the possibility that alive matches are not always extended (i.e. $\mathbb{P}_{\text{extend}} < 1$).[35] Doing so is important in the context of our model: for example, consider the extreme case of $\mathbb{P}_{\text{extend}} = 1$ and $\sigma \to 0$. Here, it is optimal for agents to sign single period contracts and extend them month-to-month if and only if the match is alive each period. This is clearly rejected by the data because we observe very few single-period contracts (only 4.5% of all contracts).

A further justification for our formulation for extensions is that it can be viewed as an approximation of a continuous-time model where (i) unmatched charterers can contact ships who are finishing their contract (ii) information about whether charterers will be alive arises independently and continuously across charterers. Then, $(1 - \mathbb{P}_{\text{extend}})q_t^{ship}$

[35] Recall as well that options to extend are rarely used, as documented in Section 2.1.

approximates the probability that the ship is contacted by an alternative unmatched alive charterer, before information has arrived that the current matched charterer will be alive. In this case, the ship will switch and match with the new charterer. Similarly, $\mathbb{P}_{\text{extend}}$ is the probability that good information arrives that the current match will be alive before the ship is contacted by an alternative match.

4.3 States and computation

States The state for each agent consists of both its own state $\pi_{i,t}$, and a detailed industry state in period t, s_t. The detailed industry state consists of the distribution of searching agents and the distribution of current matches. Since the market is relatively fragmented, we assume that each agent takes the industry state as given.

The detailed industry state is high-dimensional.[36] As such computing value functions with such a state space is infeasible due to the curse of dimensionality. Instead we employ an approximation method similar to a Moment-Based Markov-Equilibrium (Ifrach and Weintraub, 2017). Specifically, we assume that agents approximate the industry state by only keeping track of the demand realization z_t and the market thinness θ_t, and that agents believe the transitions of these aggregate states follow AR(1) processes. We estimate the AR(1) process for the demand realization z_t "offline" (since it depends on global economic conditions and is therefore arguably exogenous to the containership leasing market)

Computation The computational algorithm involves an inner loop and an outer loop. In the inner loop we compute the value functions given the current iteration of the AR(1) parameters for θ_t. In an outer loop we iterate over the AR(1) parameters, updating them to be consistent with the detailed state evolution through forward simulation of the value functions in the inner loop. We provide more information about the computational algorithm in Appendix Section B.4.

[36] For example, the distribution of current matches includes, for each match, a state for the number of remaining periods and a state for whether the match is still alive.

4.4 Equilibrium

A dynamic equilibrium is characterized by a mass of searching agents ($a_t^{charterer}$, a_t^{ship}), a distribution of current matches (including how many periods remain on each contract, and whether or not each match is still "alive"), contract duration choice probabilities, prices, and agents beliefs about state transitions, such that at each state of the industry s_t, the following conditions are satisfied:

1. Agents optimally choose contract duration, according to equations (1)-(8).

2. Equilibrium prices are determined by Nash bargaining.

3. The mass of searching agents and charterers and the distribution of current matches follow the transition processes described in Section 4.1.

4. Agents have expectations over the evolution of the industry state governed by an AR(1) process that is consistent with the true industry state evolution.

4.5 Discussion

Main mechanisms behind contract duration choice Agents' contract duration choice within each period is determined through two opposing channels. On the one hand, there is the search frictions channel: signing a long contract is beneficial because it means the agents do not have to search again and then potentially fail to find a match. On the other hand, there is the lock-in channel: if the contract is too long then it may result in lock-in of a dead match. The optimal contract duration optimally balances these two channels and maximizes total surplus.

How then does the model generate overall contract duration changes in booms and busts? When a boom turns to a bust, there are two effects. The first is a 'match effect'. Here, the match itself could change through the parameters (ν_t, η_t). If, for example, the bust implies that downstream demand from exporters is less certain, this would increase the probability that the match dies $(1 - \eta_t)$, exacerbating the lock-in channel and favoring a

shorter contract. (Note that we estimate how these parameters change with the cycle, and so it is an empirical question what the direction of this effect is.)

The second effect is a 'market thickness effect'. Concretely, when a boom turns to a bust, market thickness increases because there are more unemployed ships and fewer searching charterers. This increases the probability of matching for charterers but decreases the probability of matching for ships. These probabilities affect total surplus through the outside option to search again. Since the market thickness effect goes in asymmetric directions for ships and charterers, the overall effect on total surplus is theoretically ambiguous. Ultimately, the bargaining parameter δ governs this asymmetry. For example, in the most extreme example where ships have no bargaining power, they receive no surplus in future matches, and so their outside option is 0 and fixed over time. As a result, the market thickness effect in this example would favor a longer contract in booms.

Contracting externalities When a ship and a charterer choose to sign a longer contract, this causes the market to become thinner, potentially exacerbating search frictions and imposing a negative externality on the rest of the market. Due to these contracting externalities, the decentralized choice of duration is in general not socially optimal. We explore the inefficiency from this contracting externality, and how this varies over the boom-and-bust cycle, in more detail in Section 7.

5 Estimation and identification

5.1 Overview and parameterization

Overall we estimate the model in two steps. After calibrating two parameters, in the first step we compute the evolution of the process for the demand shocks z_t 'offline'. Then we estimate the rest of the parameters via simulated method of moments. Table 4 provides an overview of the parameters and in which step they are estimated.

Calibrated parameters A period in the model is one month. Therefore we calibrate the discount factor, $\beta = 0.99$. We calibrate the per-period cost of operating a ship c =\$2,500

Table 4: Overview of parameters

Name	Object	Parameters	Stage
Discount factor		β	Calibrated
Ship operating cost		c	Calibrated
Switching cost	c_t^s	c^s	Calibrated
Demand process	z_t	γ_z	Stage 1
Charterer entry rate	e_t	γ_e	Stage 2
Charterer value	v_t	γ_v	Stage 2
Charterer survival process	η_t	$\gamma_{0,\eta}, \gamma_{1,\eta}$	Stage 2
Matching function	q_t	α	Stage 2
Duration-specific shock		σ	Stage 2
Bargaining parameter		δ	Stage 2
Extension parameter		$\mathbb{P}_{\text{extend}}$	Stage 2

per day, based on Stopford (2009); Appendix B.2 has details.

Switching costs c_t^s mainly come in the form of lost time when a ship is empty but traveling to its new match. Therefore we parameterize $c_t^s = c^s v_t$, where c^s is the time lost in transit. We assume that it takes half a month on average to move a ship from one match to another (consistent with typical ocean transit times), and therefore calibrate c^s to equal 0.5.[37]

Parametric forms We parameterize the match survival process as $\eta_t = 1 - \Phi(-\gamma_{0,\eta} - \gamma_{1,\eta} z_t)$ where $\Phi(.)$ is the standard normal CDF and $\gamma_{0,\eta}, \gamma_{1,\eta}$ are parameters.[38] We parameterize the match value as $v_t = \gamma_v z_t$ and the charterer entry process as $e_t = \gamma_e z_t$, where γ_v, γ_e are parameters. We assume the urn-ball matching function (Petrongolo and Pissarides, 2001), which implies the following match probabilities:

$$q_t^{ship}(\theta_t) = min\{1 - exp(-\alpha/\theta_t), 1, 1/\theta_t\} \tag{9}$$

[37] On average, a ship has to travel 3100 km when it is reallocated, which would take around 3-4 days at typical container-ship speeds of 18 - 25 knots. We assume the time lost is 15 days to account for potential port congestion (Brancaccio et al., 2024) and because a ship's capacity utilization tends to be low during the port calls just prior to beginning a new contract. We obtain similar results if the time lost in transit is calibrated to 7 days or set to zero.

[38] We have obtained similar results with other functional forms which ensure η is bounded between 0 and 1, such as an exponential function of z_t.

$$q_t^{charterer}(\theta_t) = min\{\theta_t(1 - exp(-\alpha/\theta_t)), 1, \theta_t\} \tag{10}$$

where α is a parameter capturing the efficiency of the matching process.

5.2 Identification

We discuss how we identify the stage 2 parameters in the simulated method of moments.

Charterer survival process η_t The parameters $\gamma_{0,\eta}, \gamma_{1,\eta}$, which underpin the charterer survival process η_t, are identified by matching the counter-cyclical patterns of dispersion in capacity utilization in the data. Concretely, we include two moments: the average dispersion in a high demand period (2006-2008) and the average dispersion in a low demand period (2009-2010). Similarly, the parameter γ_v, which underpins how the value of a match changes with the cycle, is identified through matching two moments corresponding to the average contracted prices of a ship in a boom and a bust.

Bargaining parameter δ Intuitively, the model allows for two reasons for lengthening contracts in booms. The first is that the match value may systematically change in the boom - this is already pinned down through the estimates for $\gamma_{0,\eta}, \gamma_{1,\eta}, \gamma_v$. The second is the bargaining parameter δ. A smaller bargaining parameter gives charterers a higher share of the match surplus, and this tends to amplify differences in the charterer's outside option in booms versus busts. This then tends to make differences in the match surplus of a long versus short contract more procyclical, and the corresponding equilibrium contract duration choice is also more procyclical.

Therefore, δ is identified by fitting the residual cyclicality of contract duration, once match-specific factors are controlled for. So, we identify δ by including two moments for the mean duration in a high-demand period (2006-2008) versus a low-demand period (2009-2010). We emphasize that this identification procedure could result in a potentially high δ, or a low δ, depending on the cyclicality of the match-specific factors.

Extension and duration-specific shock parameters The probability of extension $\mathbb{P}_{\text{extend}}$ is identified by matching how the average probability of an extension in the data differs between booms and busts (Appendix A.7 discusses how extensions are identified). The standard deviation of the logit shock σ in the contract duration choice is identified through a moment for within-period dispersion of contract duration.

Matching function parameters We identify the match efficiency α by directly matching the elasticity of contract duration with respect to market thickness after controlling for the demand state z_t (similar to the regression in Column (3) of Table 3). Intuitively, when α is larger, a given increase in market thickness will result in a larger increase in the probability that charterers can find a match (since matching frictions are smaller), making contract duration more responsive to changes in market thickness.

Charterer entry process With α identified, there is a one-to-one mapping between the proportion of ships under contract in each period and the (unobserved) number of search charterers. With the survival process η_t for charterers also identified, we can therefore identify the entry process e_t by matching how the share of ships under contract moves across the cycle. In practice we match a moment for the mean and standard deviation of the proportion of ships under contract in the bust, as well as the difference in the proportion of ships under contract in the boom vs the bust, to identify γ_e.[39]

5.3 Estimation

First stage: the demand state Since the demand state is not directly observed, we recover it "offline", similar to Jeon (2022). We first estimate demand for container-ship services, as given by the following equation:

$$q_t = \gamma_{0,z} + \gamma_{1,z} p_t + \gamma_{2,z} W_t + \varepsilon_t \tag{11}$$

[39] Appendix A.3 details how we calculate the proportion of ships under contract each period.

where q_t is the total amount of cargo carried by containerships during period t, and p_t is the price of hiring container-ships (which we measure using the time-charter rate index). W_t denotes observed demand shifters, while ε_t denotes idiosyncratic demand shocks. Following Jeon (2022), we instrument for the price p_t using the average size and age of ships, and the share of ships older than 20 years: these are all cost shifters since larger and newer ships are more cost-efficient. Appendix B.3 reports and discusses demand estimates.

Having estimated equation (11), we define the demand state z_t as:

$$z_t = \hat{\gamma}_{0,z} + \hat{\gamma}_{2,z} W_t + \hat{\varepsilon}_t \tag{12}$$

where $\hat{\gamma}_{0,z}$ and $\hat{\gamma}_{2,z}$ are estimated regression coefficients, and $\hat{\varepsilon}_t$ is the estimated demand residual. We then smooth z_t using a local polynomial regression, and estimate its AR(1) transition process, as described in Appendix B.3.

Second stage: simulated method of moments We use the moments described above in Section 5.2 using a diagonal weight matrix. We scale the weights on the diagonal so that the moments enter into the objective function with a common scale.

When constructing the moments, we exclude 3.9% of ships that are Post-Panamax or larger (i.e., above 5,000 twenty-foot equivalent unit (TEU) in capacity), so that the resulting sample is relatively homogeneous in terms of ship size. We also exclude 6% of the contracts that are by newly built ships, since we do not model the use of contracts as collateral to obtain cheaper financing terms. Likewise, we exclude a very small number of contracts longer than 8 years (about 0.2% of the sample), as some of them may be “capital leases" where the charterer acquires the ship upon contract expiry (Gavazza, 2010). The resulting sample has 2,826 contracts.

6 Results and model fit

Table 5 shows the estimated parameters. In the Appendix, Table 11 provides a comparison between the empirical moments and the simulated moments; the model fits the data well.

We find a matching efficiency parameter α of 0.14, indicating sizeable search frictions. As an illustration, if the market were perfectly balanced with an equal number of ships and charterers (so $\theta = 1$), an α of 0.14 implies a meeting probability of 13% per period for charterers. This corresponds with anecdotal evidence (see Section 2.1) that search frictions are a first-order feature of this market.

We next turn to the bargaining parameter δ. We again emphasize that the magnitude of δ is an empirical question and a high value of δ that favors ships is consistent with procyclical contract duration if the match-specific factors in the model (v_t, η_t) are also sufficiently procyclical. In our context, we find a $\delta = 0.12$, which implies that charterers capture most of the match surplus. However, it is hard to interpret δ in isolation. This is because there are two channels in the Nash bargaining solution that generate negotiation asymmetries: differences in the outside options and the δ. For example, an arguably more interpretable measure of the outcome of the negotiation process is the share of total *profit* rather than the surplus, which also includes the outside options, averaged over the sample. This value is 0.29, which implies that charterers still do well in the negotiation process but — because ships have a strong outside option — not as well as one might expect looking at the raw magnitude of δ.

Furthermore, the low δ for ships is consistent with institutional characteristics of the industry. For example, the market for leasing containerships is relatively new compared to our sample period: it did not exist before the 1990s (Stopford, 2009). However, the firms that charter ships have been around for decades offering liner services using their own ships. Therefore, these shipowners who specialize in leasing out ships are relatively new and potentially less sophisticated at negotiating.

Table 5: Parameter estimates

Parameter	Estimate	SE
Ship bargaining weight, δ	0.12	0.01
Matching function parameter, α	0.14	0.08
Std. dev. of logit shock, σ	0.41	0.06
Survival prob. parameter, $\gamma_{0,\eta}$	1.14	0.04
Survival prob. parameter, $\gamma_{1,\eta}$	0.18	0.02
Per-period value of leasing ship, γ_v	1.76	0.04
Entry rate, γ_e	8.14	0.25
Probability of contract extension, $\mathbb{P}_{\text{extend}}$	0.23	0.03

7 Counterfactuals

7.1 Quantifying misallocation and the effects of booms, busts, and the transition

We begin by documenting how the decentralized equilibrium contract length changes with the business cycle in our model and the implications for misallocation. We compare the results to a constrained social planner, who still faces search frictions and no information about the realization of future shocks, but is able to coordinate firms to set a contract length that optimally maximizes total welfare (measured in total joint profits) of the market.[40]

We present the results in Figure 6, with additional numerical results in Table 6. In Panel (a) of Figure 6 we plot the decentralized contract duration in the solid black line and the demand process in the dotted line. Consistent with the data, it is pro-cyclical.

We compute the constrained optimal duration in Panel (b). We find, strikingly, that the optimal contract length is actually counter-cyclical and decreases in the boom.

The difference in total profits of the firms (misallocation) moving from the decentralized contract to the optimal duration is pictured in Panel (c). The level of misallocation follows the level of the demand realization but with a lag. The lag occurs due because it takes time

[40]To compute the social planner's duration, we iterate over all possible contract durations, compute the equilibrium total welfare for each selection via forward simulation, and then choose the contract duration with the corresponding highest total welfare.

Table 6: Misallocation and utilization dispersion results for various counterfactuals

	(a) Duration (months)			(b) Dispersion in utilization		
	'06-'08	*'09-'10*	*Overall*	*'06-'08*	*'09-'10*	*Overall*
Baseline	9.7	6.2	7.5	0.157	0.174	0.164
Optimal contract	3.6	7.2	5.8	0.065	0.161	0.122
Intermediary - baseline duration	9.7	6.2	7.5	0.157	0.174	0.164
Intermediary	1.0	1.0	1.0	0.000	0.001	0.001

	(c) Total welfare (percent change from baseline)					
	(i) Transition			(ii) Overall		
	Quality	*Quantity*	*Total*	*Quality*	*Quantity*	*Total*
Optimal contract	11.7	-1.1	10.7	6.8	-1.2	5.6
Intermediary - baseline duration	0.0	0.0	0.0	0.3	2.6	2.9
Intermediary	35.8	0.0	35.9	37.8	3.7	41.5

Note: Panel (a) and (b): The period '06-'08 is the start of the data and corresponds to a boom; the period '09-'10 corresponds to the crash and the period where dispersion in utilization spikes in the data. Panel (c): Total welfare is measure in terms of the joint profits of the firms. 'Quality effect' determined by keeping the number of matches constant and measuring the change in total profit moving to the counterfactual. 'Quality effect' is then the remaining effect determined by changing the number of matches holding the total average profit of each match constant. The transition is the period in late 2008 - early 2009 between the boom and bust when the market crashes.

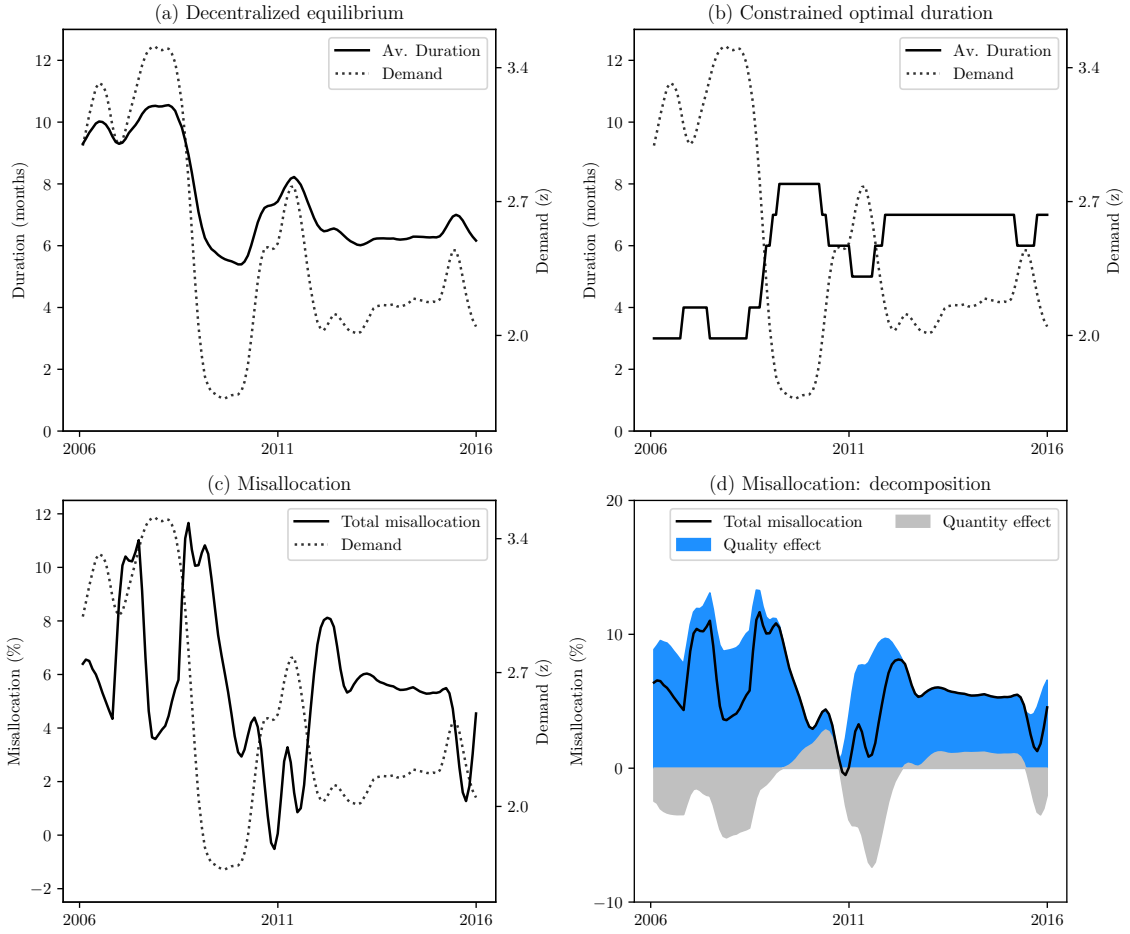
for the stock of existing contracts to accumulate (or deaccumulate in a crash). As a result, the corresponding effects on market thinness and misallocation can take time to appear.[41]

Therefore, misallocation is cyclical and increases in a boom. Moreover, the numbers in Table 6 — which simply split up the data into a boom vs bust based on the demand realization — understate the effects of the boom on misallocation because of the lagged effects. Indeed, in the transition in 2009, misallocation jumps to over 10% because of the contract overhang which was generated in the boom. We see a similar jump after a smaller demand shock in 2012.

These results illustrates a second implication of the efficient contract: it also allows the market to flexibly respond to aggregate shocks in booms. To put it another way, the supply

[41]This also highlights an advantage of our framework, which is that it is *not* simply a comparison between a steady-state boom and a steady-state bust. Rather, we accommodate that the market is constantly in transition.

Figure 6: Misallocation across the boom-bust cycle



Note: Panel (a): This shows the average duration in the baseline case (in the solid line) versus the demand shifter (in the dashed line). Panel (b): Shows the optimal duration for a social planner that is constrained to not have information about future shocks, and where search frictions are still present. There is a single monthly optimal duration at each time that is the optimal duration and so there are discrete jumps between time periods. Panel (c): Misallocation between the decentralized vs optimal duration, measured in terms of total joint profits.

chain rigidities that results from the long decentralized contract overhang in booms — which are a striking feature of the descriptive results — are inefficient.

Decomposition We develop intuition by decomposing total misallocation into a *quantity effect* and a *quality effect* in Panel (d) of Figure 6. To do so we begin with the decentralized contract duration. We then measure the change in total welfare keeping the number of matches fixed, but changing the profit per match to equal the level at the optimal contract duration. This will increase if the probability of an "alive" match is higher under the optimal duration. We then allow the number of matches to vary, and the resulting change in welfare is the quantity effect.

The decomposition in Panel (d) therefore illustrates the planner's tradeoff. On one hand, if contracts are too long then there may be lock-in of matches that have expired but are still under contract. These contracted ships could be reallocated to "alive" charterers who are unmatched. This is the quality effect.

On the other hand, if contracts are too short, ships who are matched to charterers who are "alive" at the end of the period will need to search again. But searching involves frictions where they risk being unmatched; ex-post they would rather have remained under contract. This is the quantity effect. Note that the planner's tradeoff is somewhat similar to the trade-off agents face when choosing the optimal contract length. The key difference is that the planner internalizes equilibrium effects since they are maximizing *total* welfare.

With the above trade-off in mind, consider the market in the boom when it is unbalanced with more searching charterers than ships. Therefore, the probability that an unmatched ship will match with a charterer in the search process increases. This implies that the contracting externality from long contracts also increases; it is better to thicken the search market and allow for ships to rematch with a high-probability with "alive" matches than to risk lock-in with a longer contract. Conversely, in a bust the number of searching agents on both sides of the market is more balanced. In this case, the probability a ship will

successfully match with a charterer if it chooses to search is relatively low; this provides incentives to the planner to sign a longer-contract.

Effects on dispersion in ship utilization Table 6(b) illustrates dispersion in ship utilization in booms, busts, and overall, across the counterfactuals. In the presence of search frictions, some dispersion is still efficient, as illustrated by the results for the optimal contract. Concretely, dispersion would be 0.164 under the decentralized contract baseline, but significantly lower at 0.122 under the optimal contract. Therefore, approximately 74.6 percent of dispersion is efficient and 25.4 percent is inefficient.

The intuition behind these findings is similar to the intuition behind the optimal contract. Dispersion rises due to contract lock-in which creates more "dead" matches. However, some risk of lock-in is optimal because a contract that is too short means that agents have to search more frequently, which (when there are frictions) can generate unmatched agents.

Role of search frictions/an intermediary The above findings all consider a "constrained optimal" contract length where a planner still faces search frictions. What is the role of search frictions in the results? To answer this question we implement an intermediary — intuitively, an Uber for containerships — who eliminates search frictions. We operationalize this implementing a 'frictionless' matching function of $\min\{a_t^{ships}, a_t^{charterer}\}$ and also setting $\mathbb{P}_{extend} = 1$. Agents then choose their privately-optimal contract. The results in Table 6(c) show that an intermediary would result in an extremely short contract in equilibrium (in part due to agents' ability to continuously extend the contract), and almost eliminate dispersion in utilization. Overall, an intermediary would increase welfare by 41.5%. Note that the benefits of an intermediary are amplified by the interaction between search frictions and endogenous contracting rigidities: if duration were held fixed at baseline levels, introducing the intermediary raises welfare by only 2.9%.

7.2 Policy implications: evaluating subsidy pass-through

Finally, we illustrate the implications of endogenous rigidities for evaluating maritime subsidies. These subsidies are large, totalling billions of euros per year in the EU alone, and directly target both sides of the market (see OECD/ITF, 2019 for a comprehensive list of subsidies). For example, on the charterer side of the market, there are tax exemptions for fuel and lower corporate tax through 'tonnage taxes'; these correspond to a constant increase in v_t in the model. On the shipowner side, subsidies directly reduce labor costs of the crews (which would correspond to a decrease in c in the model). Although the objective of these subsidies is primarily to support shipping to and from domestic markets, in practice most developed countries use subsidies and so we investigate global changes in these arrangements. Specifically, we ask: 1. what is the pass-through of these subsidies to the industry in the presence of endogenous rigidities 2. which side of the market should policymakers subsidise?

If rigidities are fixed then we would expect complete pass through: a dollar in subsidies increases the joint profits of the firms by one dollar, and this is independent of which side of the market is directly subsidised. However, when rigidities are endogenous, there is an additional effect that may reduce the efficacy of subsidies. If the subsidies also interact with the contract length (for example, they make a match more valuable thereby inducing the parties to sign longer contracts) then they will thin the market. This then creates an additional negative contracting externality, making it harder for other agents to find matches, which can reduce output (and also have potentially different effects in booms and busts).

We present the results in Table 7. Under fixed rigidities — keeping the contract length the same in the counterfactuals as in the baseline — the pass through is equal to 1.0. In the presence of endogenous rigidities the pass through is less than 1.0. The results favor a counter-cyclical subsidy, since contracting externalities are substantially lower in the bust.

Table 7: Policy counterfactuals: pass-through of subsidies to total industry profit

	Fixed rigidities	Endogenous rigidities		
		Boom	*Bust*	*Total*
Ship subsidy: c	1.00	0.13	0.47	0.36
Charterer subsidy: v_t	1.00	0.75	0.98	0.89

Note: These counterfactuals measure pass-through, defined as the dollar change in (joint) industry profits for a dollar change in the subsidy. We consider a subsidy that decreases the ship operating cost c and a subsidy that increases the charterer's net value of a match by a fixed amount at each period (v_t). 'Fixed rigidities' computes the pass-through holding the contract duration fixed at the baseline. 'Endogenous rigidities' computes pass-through allowing the contract duration to change. Pass through will be lower under endogenous rigidities if the subsidy causes a change in contract duration, leading to an equilibrium increase in inefficiency. 'Pct. Δ duration bust vs boom' documents the asymmetric increase in contract duration as a result of the subsidy in the bust versus the boom.

Subsidies have an asymmetric effect depending on which side of the market is directly affected, with pass-through relatively lower for a ship subsidy at 0.36. Overall, the policy recommendation that it does not matter which side of the market is subsidised — which corresponds to the simple fixed rigidities model — is qualitatively different in the endogenous rigidities case.

To understand the intuition behind the asymmetric effects of subsidies, recall that the individually-optimal contract duration is determined by a trade-off between the risk of lock-in with a bad match versus the option value of continued search. The reason for the asymmetric effect is due to the interactions of the subsidies with the outside options of continued search. When a ship subsidy (a c change) is implemented, since the bargaining parameter for ships is relatively low, the outside option of continued search does not increase much. As a result contract duration increases substantially and so do the corresponding externalities which cuts into the efficacy of the subsidy. Furthermore, as illustrated by the final column of Table 7, these effects on duration change with the boom-bust cycle. For a charterer subsidy (a v_t change), the effects on the outside options versus the match value are more balanced. Therefore, the equilibrium duration (and the resultant externalities) do not change as much.

8 Conclusion

This paper shows how endogenous rigidities — in the form of agents' choice of contract duration — affect physical capital reallocation in decentralized markets. To do so we exploit rich data on contracts and allocations in the market for leased containerships; a market that is also important by itself as a key part of the supply chain. Using these data and an empirical model, we argue that agents choose longer contracts in booms when the market becomes thin, and that this results in a contracting externality that implies contracts are too long in booms which leads to misallocation.

Using the model, we find that there is significant misallocation from endogenous rigidities, in particular in the market crash as it transitions from a boom to a bust. We also show that endogenous rigidities substantially reduce the efficacy of maritime subsidies, which is a common and large-scale form of industrial policy in the industry. Overall, given that contract duration rises across a range of other capital markets, our paper suggests that accounting for how endogeneous rigidities is central to understanding the process of capital reallocation in booms and busts.

References

Adland, Roar, Haakon Ameln, and Eirik A Børnes, "Hedging ship price risk using freight derivatives in the drybulk market," *Journal of Shipping and Trade*, 2019, *5* (1), 1.

_ , Haiying Jia, and Siri P. Strandenes, "The determinants of vessel capacity utilization: The case of Brazilian iron ore exports," *Transportation Research Part A: Policy and Practice*, April 2018, *110*, 191–201.

Ardelean, Adina and Volodymyr Lugovskyy, "It pays to be big: Price discrimination in maritime shipping," *European Economic Review*, 2023, *153*, 104403.

Asker, John, Allan Collard-Wexler, and Jan De Loecker, "Dynamic inputs and resource (mis) allocation," *Journal of Political Economy*, 2014, *122* (5), 1013–1063.

Balke, Neele and Thibaut Lamadon, "Productivity Shocks, Long-Term Contracts, and Earnings Dynamics," *American Economic Review*, 2022, *112* (7), 2139–2177.

BIMCO, "BOXTIME 2004," 2004. [Accessed 10-August-2024].

Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papageorgiou, "Geography, transportation, and endogenous trade costs," *Econometrica*, 2020, *88* (2), 657–691.

_ , _ , and _ , "Investment in Infrastructure and Trade: The Case of Ports," *National Bureau of Economic Research Working Paper*, 2024.

_ , _ , _ , **and Nicola Rosaia**, "Search Frictions and Efficiency in Decentralized Transport Markets," *The Quarterly Journal of Economics*, November 2023, *138*, 2451–2503.

Buchholz, Nicholas, "Spatial Equilibrium, Search Frictions, and Dynamic Efficiency in the Taxi Industry," *Review of Economic Studies*, March 2021, *89* (2), 556–591.

Castillo, Juan Camilo, "Who Benefits from Surge Pricing?," *Econometrica (conditionally accepted)*, November 2023.

Darmouni, Olivier, Simon Essig Aberg, and Juha Tolvanen, "Pulp Friction: The Value of Quantity Contracts in Decentralized Markets," 2023. SSRN working paper 3919592.

Eisfeldt, Andrea L. and Adriano A. Rampini, "Capital Reallocation and Liquidity," *Journal of Monetary Economics*, April 2006, *53*, 369=399.

Fréchette, Guillaume R., Alessandro Lizzeri, and Tobias Salz, "Frictions in a Competitive, Regulated Market: Evidence from Taxis," *American Economic Review*, August 2019, *109* (8), 2954–2992.

Gaineddenova, Renata, "Pricing and Efficiency in a Decentralized Ride-Hailing Platform," *Working Paper*, January 2022.

Ganapati, Sharat, Woan Foong Wong, and Oren Ziv, "Entrepot: Hubs, scale, and trade costs," *American Economic Journal: Macroeconomics (Forthcoming)*, 2023.

Gavazza, Alessandro, "Asset liquidity and financial contracts: Evidence from aircraft leases," *Journal of financial Economics*, 2010, *95* (1), 62–84.

_ , "Leasing and Secondary Markets: Theory and Evidence from Commercial Aircraft," *Journal of Political Economy*, 2011, *119* (2), 325–377.

_ , "The Role of Trading Frictions in Real Asset Markets," *American Economic Review*, June 2011, *101*, 1106–1143.

_ , "An Empirical Equilibrium Model of a Decentralized Asset Market," *Econometrica*, 2016, *84* (5), 1755–1798.

Greiner, Richard, "Ship operating costs: Current and future trends," 2017. Moore Stephens LLP. Accessed at https://www.moore-greece.gr/MediaLibsAndFiles/media/greeceweb.moorestephens.com/Documents/1-Richard-Greiner.pdf.

Haralambides, Hercules E., "Gigantism in container shipping, ports and global logistics: a time-lapse into the future," *Maritime Economics & Logistics*, 2019, *21* (1), 1–60.

Harris, Adam and Thi Mai Anh Nguyen, "Long-term Relationships and the Spot Market: Evidence from US Trucking," 2023. Working Paper.

Heiland, Inga, Andreas Moxnes, Karen Helen Ulltveit-Moe, and Yuan Zi, "Trade From Space: Shipping Networks and The Global Implications of Local Shocks," 2022.

Hubbard, Thomas N, "Contractual form and market thickness in trucking," *RAND Journal of Economics*, 2001, pp. 369–386.

Hummels, David, Volodymyr Lugovskyy, and Alexandre Skiba, "The trade reducing effects of market power in international shipping," *Journal of Development Economics*, 2009, *89* (1), 84–97.

Ifrach, Bar and Gabriel Y Weintraub, "A framework for dynamic oligopoly in

concentrated industries," *The Review of Economic Studies*, 2017, *84* (3), 1106–1150.

Jeon, Jihye, "Learning and investment under demand uncertainty in container shipping," *The RAND Journal of Economics*, 2022, *53* (1), 226–259.

Kalouptsidi, Myrto, "Time to Build and Fluctuations in Bulk Shipping," *American Economic Review*, 2014, *104* (2), 564–608.

_ , "Detection and Impact of Industrial Subsidies: The Case of Chinese Shipbuilding," *Review of Economic Studies*, 2018, *85* (2), 1111–1158.

Lanteri, Andrea and Adriano A. Rampini, "Constrained-Efficient Capital Reallocation," *American Economic Review*, February 2023, *113* (2), 354–395.

MacKay, Alexander, "Contract duration and the costs of market transactions," *American Economic Journal: Microeconomics (Forthcoming)*, 2022.

Notteboom, Theo, Thanos Pallis, and Jean-Paul Rodrigue, "Disruptions and resilience in global container shipping and ports: the COVID-19 pandemic versus the 2008–2009 financial crisis," *Maritime Economics & Logistics*, 2021, *23*, 179–210.

OECD/ITF, "Maritime Subsidies Do They Provide Value for Money?," 2019. Iternational Transport Forum, Policy Analysis Report.

Petrongolo, Barbara and Christopher A Pissarides, "Looking into the black box: A survey of the matching function," *Journal of Economic literature*, 2001, *39* (2), 390–431.

Rosaia, Nicola, "Competing Platforms and Transport Equilibrium," *Working Paper*, February 2023.

Stopford, Martin, *Maritime Economics*, New York: Routledge, 2009.

UNCTAD, *The Review of Maritime Transport 2010*, Geneva: UNCTAD secretariat, 2010.

Vreugdenhil, Nicholas, "Booms, Busts, and Mismatch in Capital Markets: Evidence from the Offshore Oil and Gas Industry," *Accepted, Journal of Political Economy*, 2023.

Wackett, Mike, "Yang Ming Set to Make Huge Profit from Sub-Letting Charter to Maersk," *The Loadstar*, November 2021. Accessed: 2024-10-06.

Yang, Ron, "(Don't) Take Me Home: Home Preference and the Effect of Self-Driving Trucks on Interstate Trade," *Working Paper*, May 2023.

Zahur, Nahim Bin, "Long-term Contracts and Efficiency in the Liquefied Natural Gas Industry," *Working Paper*, 2024.

Zhang, Abraham, Zhichao Zheng, and Chung-Piaw Teo, "Schedule reliability in liner shipping timetable design: A convex programming approach," *Transportation Research Part B: Methodological*, 2022, *155*, 499–525.

A Appendix

A.1 Data construction

Our full dataset of contracts includes over 16,000 time-charter contracts; for the period of 2005 - 2015, which is the main focus of our analysis, we observe around 4,109 contracts.

The contract data includes a number of key contract details, most notably the charterer name, charter period (or duration), and the contracted charter rate (in $/day). For some contracts, we also observe the “delivery location", which is the location where the ship is transferred from the shipowner to the charterer. The contract data also includes information on ship characteristics, such as vessel type, name, twenty-foot equivalent unit (TEU), deadweight tonnage (dwt tonnes), year of construction); we complement this ship-level information with a comprehensive dataset of containerships collected from Vessel Finder, which notably includes the International Maritime Organization (IMO) number for each ship.

Certain contract records retrieved from Clarkson’s database lacked precise chartering duration. However, these records typically included approximate chartering data ranges, such as “20-40 days" and “5-7 months". To facilitate the empirical analysis, we computed two measures of contract duration: mean values (e.g., 6 months for “5-7 months") and the maximum period (e.g., 7 months for “5-7 months" or the period until the next contract for the ship starts, whichever comes first). Our baseline analysis is carried out using the maximum period of each contract as the contract duration; all our results, however, are robust to using the other measure instead.

In addition to the contract data, Clarksons’ Shipping Intelligence Network provides aggregate indexes such as containership time-charter rate index, China Containerized Freight Index (CCFI), and Singapore bunker prices ($/Tonne).[42]

[42]The containership time-charter rate index is an index published by Clarksons based on daily rates of newly negotiated time-charter contracts. The China Containerized Freight Index (CCFI) is an index of container freight rates based on the price of containers leaving from all major ports in China. Bunker prices are

In order to merge port calls with fixtures data, we use the Vessel Finder database to match ship names with International Maritime Organization (IMO) numbers.

A.2 Measuring utilization

To measure utilization, we exploit the fact that for many port calls, we observe the draught of the ship, which is indicative of how much cargo the ship is carrying. We then define utilization as the percentage of the ship's capacity that is used for carrying cargo; if this number is low, it suggests the ship is not being fully utilized. While we do not directly observe this in the data, we can infer utilization from draught data, as we describe below.

A ship that is sailing without cargo is commonly known as sailing "in ballast". In practical terms, a ship is considered to be sailing in ballast if its draught is less than a specified threshold value known as the "ballast draught" (H_B). In maritime engineering literature, a weight of 0.55 is employed to establish the ballast draught (Heiland et al., 2022). Following this literature, we define the ballast draught (H_B) as 55% of the ship's scantling draught (H_S). The scantling draught, also referred to as the design draught, represents the ship's draught when fully loaded and is a constant value since the ship is constructed to operate at this specific draught. While we don't observe the scantling draught, we proxy for it by choosing the observed maximum draught for the specific ship in the data. We then compute utilization, defined as the percentage of the ship's capacity that is being utilized on a specific voyage, using the following formula:

$$\text{Utilization} = (H_A - H_B)/(H_S - H_B) \tag{13}$$

where H_A is the draught reported in the port call data. Note that in the analysis we always account for ship heterogeneity when using this measure (e.g., by including ship fixed effects), in order to ensure that the measure is comparable across vessels.

Finally, since ships only report draught information when they arrive at or depart from a

indicative of the fuel costs of operating containerships (paid for by charterers).

port, this measure of utilization only captures the intensive margin (i.e., how full the ship is, conditional on being non-idle). We separately measure the extensive margin by using the port call data to identify idle ships, as we discuss in the next section.

A.3 Measuring idleness

Identifying whether or not a ship is idle We use the port call data to identify whether or not a ship is idle. If a ship is not being utilized at all, it will typically stay moored at a single port for a longer period than is needed for the ship to unload and load.[43] We assume conservatively that it takes a maximum of 7 days for a ship to unload and load at a port (taking into account port congestion), and that any length of time it stays beyond that time is idle time. Based on this measure, we are able to calculate, for each ship, the number of days it is idle each month. Aggregating this across ships allows us to measure the overall share of idle ships at any point in time.

On average, across our sample from 2005 to 2015, ships are idle 6.8% of the time. As a validity check, we compare this data-driven measure of the share of idle ships with that published by Clarksons (but only from 2014 onwards).[44] Between 2014 and 2015, our method finds the share of idle ships to be 4.6%, which is reasonably close to the 5.0% share calculated from Clarksons data.

Measuring proportion of ships under contract The most direct way to measure the proportion of ships under contract at any point in time would be to simply count the number of ships in port call data with active contracts; however this is is likely to be an under-estimate since we do not see the universe of contracts.

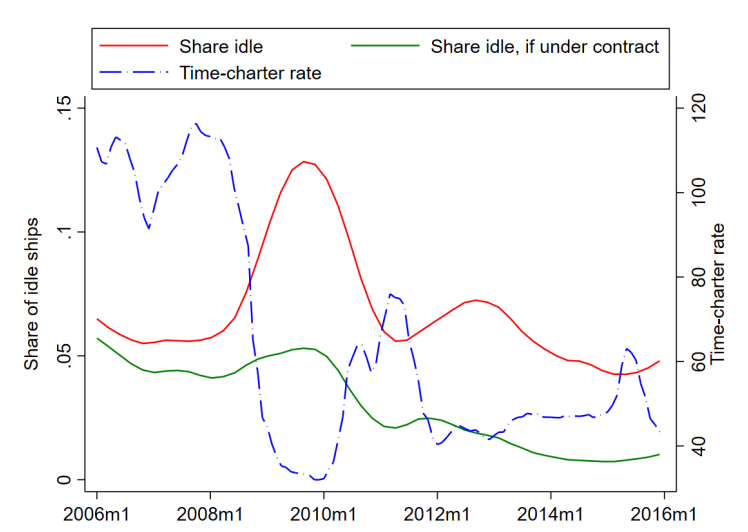
Instead, we leverage our measure of idleness in order to calculate the proportion of ships under contract. When we do this, we allow for the possibility that ships may be idle even when they are under contract: this could be because they are undergoing repairs, or because

[43] Note that a ship's draught is not reported except on days the ship arrives at a port or departs from it, which is why draught data itself is not enough to identify idleness.

[44] Clarksons utilizes AIS vessel tracking data to identify ships that have a very low average speed and therefore are idle.

the match has low value and the charterer is not able to find a productive use for the ship. As Figure 7 shows, while the share of ships with a contract that are idle is substantially smaller than the overall share of idle ships, it is still higher than zero.

Figure 7: Share of idle ships



Let u denote the share of ships that are not under contract (or unemployed). Let $i_{overall}$ denote the overall share of idle ships, and $i_{contract}$ the share of ships with a contract that are idle. Assuming that a ship without a contract is necessarily idle, we can decompose $i_{overall}$ as follows:

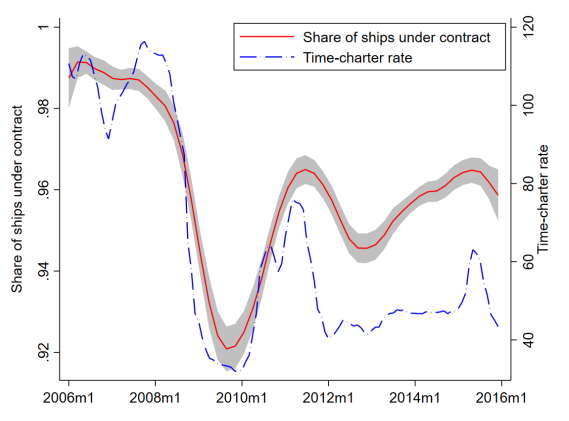
$$i_{overall} = u + (1-u)i_{contract}$$

Re-arranging this yields the following equation for u, the share of ships without a contract:

$$u = \frac{i_{overall} - i_{contract}}{1 - i_{contract}}$$

Figure 8 plots the share of ships under contract (i.e., $(1-u)$) over time, which (as expected) is pro-cyclical. Notice, however, that even during the bust, a very high share of ships (over 92%) is under contract, consistent with contract over-hang.

Figure 8: Share of ships under contract



A.4 Measuring reallocations

We use the port call data to measure ship reallocations. The key challenge is that since container-ships are operated like “buses", they may travel very long distances and stop at a large number of ports while remaining within a fixed schedule.[45] Thus, the mere fact that a ship physically travels from one port to another is not itself evidence the ship is being actively re-allocated from one use to another, since it may be simply fulfilling an itinerary that was decided many months ago.

We therefore develop an algorithm for identifying when a ship is reallocated. The idea behind the algorithm is that if a ship is reallocated, it is likely to stop at a new set of ports compared to those that were on its original itinerary. Thus, when we observe a ship visit a new port that it has not visited in recent months, we can infer that the ship has been spatially reallocated.

To be sure, sometimes a containership may visit a new port that involves a minimal deviation from its existing itinerary.[46] These are unlikely to be true reallocations of the

[45] For example, one containership in the data was observed to first stop at several ports in New Zealand, make its way up to North America (stopping at several Canadian and American ports), then travel to Western Europe, then return back to New Zealand (stopping in Colombia along the way). This sequence was repeated several times.

[46] An example of this would be a ship that is on an itinerary involving regular round-trips between Tokyo and Singapore, which at some point decides to make a stop at Manila along the way.

ship, and instead may simply represent extra voyages the carrier has decided to make while largely sticking to their original route. Thus, in order to not classify such minor deviations as reallocations, we also require that the new port that is visited be a sufficiently large distance away from any of the existing ports that the ship visited.

We now describe the algorithm we use to formalize this idea. For every port call, we calculate the *minimum* distance between coordinates of the current port call and all the other port calls in the last 6 months. The metric is assigned a value of 0 if the vessel has made a prior call at that port within the preceding 6 months. Conversely, when the port represents a new visit, the metric assumes a positive value, with its magnitude increasing proportionally as the port's distance from the current location grows, indicating a more significant alteration to the voyage schedule.

To rule out "false positives" caused by minor deviations from a set route, we classify the ship as having reallocated in a given time period only if the minimum distance metric exceeds a threshold value of 1,000 km. This threshold is large enough such that it would be very costly for a ship to temporarily deviate from an existing route by such a large distance; thus, we are more likely to pick up "true reallocations" where the ship's itinerary is substantially changed. We found this algorithm to work well in practice; the episodes it identifies as reallocations match well with what appear in the data to be true reallocations.[47]

A.5 Measuring dispersion in revenue

As we discussed in Section 2.2.1, we use cross-sectional dispersion in physical capacity utilization (residualized of ship heterogeneity) as our primary metric for inferring misallocation. However, dispersion in *physical* utilization may not necessarily imply an inefficient allocation of ships, if the prices of shipping services are heterogeneous across ships. For example, some charterers may operate on routes where the amount of cargo that needs to be moved each period is relatively limited (so that the charterer will have to

[47] We recognize that the choice of 1000 km as a threshold is somewhat arbitrary. Our results on reallocation are robust to other ways to measure reallocation, such as the average distance by which the ship was reallocated when it visited a new port.

operate at low physical utilization), but the value of moving that cargo is high.

As an alternative, therefore, we also construct a revenue-based dispersion measure. We multiply residualized capacity utilization by the freight rate of the itinerary the ship is operating on to obtain the revenue of the ship each month, and then calculate the standard deviation of the revenue each month. In this section, we discuss how we measure freight prices.

Our raw data does not include a direct measure of the freight price. In general, due to the confidentiality of the contractual agreements between charterers/shippers and downstream exporting firms, granular data on freight prices is impossible to obtain, except for a few specific importing countries.[48] Instead, our approach is to (i) classify all the itineraries/schedules[49] into one of a selected number of trade routes (ii) collect data on average freight rates in each of these trade routes, which we obtain from various industry sources (ii) match these freight rates to the specific route that each charterer is operating on at each time, which we identify by utilizing port call information.

We classify all the observed itineraries into the following major trade routes, using the criteria listed below:

- *Asia - North America (Trans-Pacific)*: port calls to Asian and North American ports account for at least 70% of all port calls on the itinerary, and each of the individual regions (Asia, North America) must separately account for at least 10% of the port calls.

- *Asia - Europe*: port calls to Asian and European ports account for at least 70% of all port calls on the itinerary, and each of the individual regions (Asia, Europe) must separately account for at least 10% of the port calls.

[48] For example, Ardelean and Lugovskyy (2023) collect transaction-level data for the Chilean import market in order to study price discrimination in liner shipping.
[49] See Section A.4 for the algorithm we use to classify port calls into itineraries.

- *Europe - North America (Trans-Atlantic)*: port calls to Asian and North American ports account for at least 70% of all port calls on the itinerary, and each of the individual regions (Europe, North America) must separately account for at least 10% of the port calls.

- *North - South*: any itinerary that is on one of the “North-South" trade routes, which includes Asia-South America, Asia-West Africa, Asia-Australia, Europe-South America, Europe -West Africa, Europe -Australia, North America-South America, North America - West Africa, and North America - Australia. For an itinerary to be classified into any one of these trade routes, the joint share of port calls to the two regions must be at least 70%, with each region having an individual share of at least 10%.

- *Intra - Asia*: at least 90% of the port calls on the itinerary are visits to Asian ports.

- *Others*: any itinerary that cannot be classified into one of the above trade routes.

We use the above classification because we have data on freight rates that can be used to construct a quarterly freight price index for each of the first five routes. For Asia - North America, we use the average of the CCFI (China Containerized Freight Index) for the China - US West Coast and China - East Coast trade routes. For Asia - Europe, we take the average of the CCFI for the China - Europe and China - Mediterranean routes.[50] For Europe - North America, we use the freight rate reported in UNCTAD up until 2009, and impute the freight rate from 2010 onwards.[51] The freight rate index for the North - South trade route is constructed as an average of the CCFI freight rate indices for the China - Australia/New Zealand, China - South America and China - West Africa routes. The

[50] An alternative approach for these two routes is to use freight rates reported in UNCTAD’s annual reports, similar to Jeon (2022); this yields a very similar measure of the Asia - North American freight rate.

[51] The imputation is needed as UNCTAD stopped reporting Europe - North America freight rates from 2010. We first regress the Europe - North America freight rate on CCFI freight rates in other routes (namely, China - Europe, China - US West Coast, China - West Africa, China - South America) in the pre-2009 period. Then from 2010 onwards, when we only observe the CCFI indices, we use the estimated regression equation to impiute the Europe - North America freight rate.

freight rate for the Intra - Asia trade route is the average of the CCFI indices for the China - Hong Kong, China - Japan, China - Korea, and China - South-east Asia routes. Finally, for all other trade routes, we assume their freight rate index equals CCFI's overall global freight rate index.

Note that our classification of trade routes follows classifications commonly used in industry reports. For example, UNCTAD's reports from 2010 - 2015 report "North-South" annual freight rates that are the average of the Shanghai-South America, Shanghai-Australia/New Zealand, Shanghai-West Africa, and Shanghai-South Africa freight rate indices; we follow the same regional classification when constructing our version of the freight rate index. Finally, we normalize all the freight rate indices so they equal 100 in the first quarter of 2003 (which is the first year for which we have complete freight rate data).

A.6 Measuring market thinness: details

In order to construct a market thinness measure we need to take a stand on how to define a "market". We define the relevant market for each contract as the number of unique ships that were within 5000 kilometers of the first port call on the contract, and were in this radius within 15 days on either side of the date the contract began.[52]

This measure proxies for market thinness by tracing out different regions across the earth, with geographically isolated markets (such as the west coast of Australia) producing lower measures than trading hubs (for example, ships located in a radius around Singapore). We construct two versions of the market thinness variable, one where we count all ships nearby in time and space, and one where we subset to only ships of the same type as the ship that was eventually under contract, reflecting that charterers may require a specific ship type.[53]

[52]The distance of 5000 kilometers is approximately the distance from the west coast to the east coast of the US. One limitation of this measure of market thickness is that we would like to also look at ships that were not just travelling in this market, but also were close to the end of their contract. Unfortunately our data, which only contain a subset of the total contracts, do not allow for this. Nevertheless, as we argue in the text, this measure is likely to still be a good *relative* proxy for thin and more geographically isolated markets, versus thicker markets.

[53]To construct ship type, we split ships into three bins based on their capacity (measured in twenty-feet equivalent).

A.7 Identifying contract extensions and subleases

Extensions We identify two types of contract extensions in the contract data.

First, contracts occasionally have a built-in option to renew, where the charterer reserves the right to extend the contract. These are sometimes recorded directly in the raw data.[54] In other cases, the raw variable recording the contract duration will include not just the duration of the first contract, but also the duration of the subsequent contract if the charterer were to exercise the option to renew.[55] We combine both types of information in order to measure the prevalence of such options: as mentioned in Section 3, these account for 3.3% of all contracts.

Second, even if a contract does not have a built-in option to renew (as is the case for the vast majority of contracts), the shipowner and charterer can and often do agree to extend the original contract. In that case, the extension will be recorded as a fresh contract in our contract micro-data.

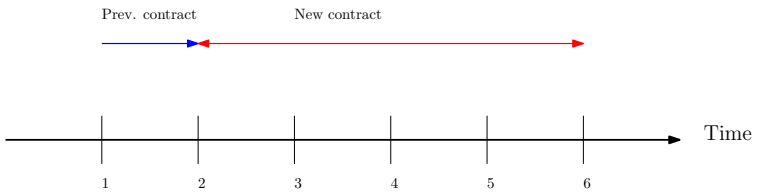
We identify extensions by looking for *consecutive* contracts agreed to by the same ship and the same charterer, with only a small gap between the end of the original contract and the start of the new one. Ideally, if our data had no measurement error, we should expect the extended contract to begin as soon as the original contract ends. An example of this is depicted in Figure 9, where the original contract ends in period 2 and the new contract begins right after: in such a case, as long as the charterer for the ship is the same, the new contract is very likely to be an agreed upon extension of the original contract.

But in practice, we often see consecutive contracts (for the ship with the same charterer),

[54] Normally, the delivery location is the location where the shipowner is expected to deliver the ship to the charterer. But if a contract is an extension of a previous contract with the same charterer that had an option to renew, there is of course no delivery location required (since the charterer already has possession of the ship). In such cases, instead of keeping the delivery location blank, Clarksons sometimes uses this variable to record whether or not the contract is based on an exercise of a previous renewal option.

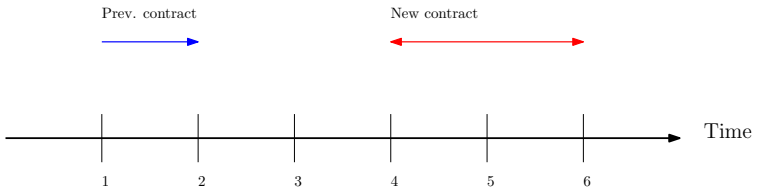
[55] For example, a contract with a 6-month duration and an option to renew for another 6 months might be recorded as having a duration of "6/6 months".

Figure 9: Consecutive contracts with no gap



but with a gap in between the two contracts: for example, the situation depicted in Figure 10. In such cases, provided the gap between the contract is fairly small, it is still quite likely that the original contract was extended, but due to measurement error, the start date of the extended contract is not recorded as beginning right after the end of the original contract. Of course if the gap is sufficiently long, then it is more likely that the ship and/or charterer were searching for interim matches in the period, before deciding on a brand-new contract, which is therefore not just a continuation or renewal of the previous contract.

Figure 10: Consecutive contracts with a gap



We therefore define a contract extension as any contract where the charterer is unchanged from that of the most recent contract by the same ship, with at most a gap of four months between them. Using this measure, we find that 42.6% of contracts in our data are extensions of previous contracts. The share of extended contracts we identify is quite robust to this choice, ranging from 40% - 43% as we vary the maximum gap from 1 to 6 months.

Subleases We identify subleases as contracts that satisfy the following criteria:

- The lease period is nested within the lease period of a previous contract for the same ship. We impose this criterion since a charterer cannot sublet beyond the period for which they have the ship.

- The charterer/lessee for the sublease must be different from that of the original lease (since one cannot sublet to oneself).

For example, suppose a ship were originally contracted from January 2006 to July 2007, and we subsequently see a contract for that same ship from January to July 2007, but to a different charterer: in this case, we consider the second contract a sublease by the charterer of the original contract.

A.8 Evidence on pro-cyclical contract duration from other leasing markets

Figure 1 in the introduction showed that contract duration is pro-cyclical not just for containerships, but also in the leasing markets for bulk carriers and drilling rigs. We describe here how we document this for other leasing markets.

Bulk carriers Bulk carriers are primarily leased via “trip-charters" (i.e., for one single voyage at a time), but about 10% of contracts for leasing bulk carriers are time-charter contracts where the owner and charterer decide in advance on the contract duration, similar to the contracts used for leasing container-ships (Brancaccio et al., 2020). We collected data on 10,629 bulk leasing contracts from past issues of the *Shipping Intelligence Weekly* published by Clarksons, covering the period from 2001 to 2016. We also collected monthly data on the Baltic Dry Index (the most commonly used shipping freight rate index for bulk shipping markets), from Clarksons’ Shipping Intellingence Network. In the second panel of Figure 1, we plot both the average contract duration (smoothed using a local polynomial regression), as well as the Baltic Dry Index.

Drilling rigs Data on leasing contracts for drilling rigs is obtained from Rigzone. We focus on jackup rigs used to drill wells in the Gulf of Mexico during the 2000 - 2010 period, prior to the Deepwater Horizon oil spill that triggered a drilling moratorium. Each contract specifies both a “dayrate" and a duration (typically 2 - 4 months). Vreugdenhil (2023) contains a further discussion of the industry and details of the dataset. Similar to Vreugdenhil (2023), we use movements in the natural gas price to capture booms and busts,

since wells drilled in the Gulf of Mexico contain more natural gas than oil. In the last panel of Figure 1, we plot both the average contract duration for rigs used in the Gulf of Mexico (again, smoothed using a local polynomial regression), as well as the Henry Hub natural gas price, finding that new contract duration is pro-cyclical.

Anecdotal evidence from other markets There is anecdotal evidence to suggest that lease duration is pro-cyclical in other markets as well, beyond the three markets for which we have detailed data. In the leasing market for aircrafts, for example, the use of shorter leases increased after the Covid-19 pandemic-induced collapse of demand for air travel.[56] Commercial office lease lengths also follow a pro-cyclical pattern.[57]

A.9 Pro-cyclical contract duration: robustness checks

Table 8 reports various regressions of contract duration on the logarithm of the time charter index (to capture pro-cyclicality of new contract duration) and the logarithm of the number of nearby ships (to account for the role of market thickness), as well as various controls. In Column (1), we repeat the same regression reported in Column (3) of Table 3 (with ship size fixed effects), but using our full sample of contracts between 2005 and 2015.[58] Column (1) shows that the contract duration is highly pro-cyclical, consistent with the pattern depicted in Figure 3, while contract duration decreases as the market becomes thicker. The coefficients on both contract duration and the market thickness proxy are very similar to those reported in Table 3.

In Column (2), we control for the contract dayrate (i.e., the daily price paid to the owner). Longer contracts have a higher dayrate, but even after controlling for the contract dayrate, duration is pro-cyclical, meaning that similarly priced contracts tend to have a longer duration when signed during a boom. The correlation between contract duration and the

[56] See https://www.iba.aero/insight/how-have-the-dynamics-of-the-leasing-market-changed-since-covid-19/.

[57] See https://www.avisonyoung.ca/w/average-term-on-office-leases-across-canada-is-on-the-rise.

[58] In Table 3, we had used our estimation sample which excludes very large ships, newly built ships and contracts longer than 8 years.

time-charter index remains very similar when we control for ship age (Column 3) and drop contracts signed by newly built ships (Column 4). Across all specifications, the effect of market thickness on contract duration is highly similar.

Table 8: Further evidence for pro-cyclical contract duration in containership leasing

	(1) Log(duration)	(2) Log(duration)	(3) Log(duration)	(4) Log(duration)
Log(time charter index)	0.85***	0.53***	0.56***	0.53***
	(0.033)	(0.066)	(0.066)	(0.068)
Log(# ships nearby, all types)	-0.14***	-0.14***	-0.14***	-0.16***
	(0.034)	(0.034)	(0.034)	(0.035)
Log(contract dayrate)		0.34***	0.31***	0.33***
		(0.061)	(0.061)	(0.063)
Age of ship			-0.01***	-0.01***
			(0.003)	(0.003)
1[mid-size ship]	0.08**	-0.03	0.00	-0.01
	(0.034)	(0.039)	(0.040)	(0.041)
1[large-size ship]	0.39***	0.19***	0.20***	0.17***
	(0.039)	(0.052)	(0.052)	(0.054)
Observations	3,303	3,298	3,298	3,096
Adjusted R^2	0.185	0.192	0.196	0.181

Note: Standard errors in parentheses. In Column (2), we control for the contract dayrate. In Column (3), we control for ship age. In Column (4), we additionally drop any leasing contracts by newly built ships.

B Additional proofs and results

B.1 Model details

Details on match payoff in Equation 5 Here we provide more detail that a previously unmatched, searching charterer's expected match payoff is $(1-\delta)W_t + \beta\eta_t \mathbb{E}_t U_{t+1}^{charterer}$. To see why, note that the charterer's payoff to a τ-duration contract, once this contract duration has been chosen and under Nash bargaining, is: $\Pi_t^{charterer} = (1-\delta)(S_{t,\tau} + \varepsilon_\tau) + \beta\eta_t \mathbb{E}_t U_{t+1}^{charterer}$. Denote $\tilde{S}_{t,\tau} = S_{t,\tau} + \varepsilon_\tau$. Then, the charterer's *expected* match payoff

(before matching has taken place and the ε_τ have been drawn) is:

$$\mathbb{E}_{\tilde{S}_{t,\tau}}[(1-\delta)\tilde{S}_{t,\tau}+\beta\eta_t\mathbb{E}_t U_{t+1}^{charterer}] = (1-\delta)\mathbb{E}_{\tilde{S}_{t,\tau}}\tilde{S}_{t,\tau}+\beta\eta_t\mathbb{E}_t U_{t+1}^{charterer} \quad (14)$$

$$= (1-\delta)W_t+\beta\eta_t\mathbb{E}_t U_{t+1}^{charterer} \quad (15)$$

which proves the result. A similar result can be derived for the ship's expected match payoff.

B.2 Calibration of ship operating costs

Industry measures of container-ship "operating costs" borne by the shipowner include both the crew cost, as well as maintenance and repair costs.[59] According to Stopford (2009), the total operating cost of a container-ship with capacity of 2,000 TEU (which is a median ship in our sample) is around $5,000/day. A very similar number is reported by Greiner (2017), who moreover finds that operating costs of containerships were fairly stable from year to year; therefore, when calibrating ship operating costs, we assume they does not change over time.[60]

When calibrating c based on these industry estimates, we have to account for the fact that only crew costs are true operating costs, in the sense that they are borne only when the ship is in use (for example, when it is under lease to a charterer). By contrast, a shipowner will likely have to incur costs of maintenance and repair even if the ship is idle, and therefore these are better thought of as fixed costs that are borne regardless of whether the ship is in operation. Since crew costs typically account for half of the total cost of operating a ship, we set c to be equal $2,500/day; this is consistent with typical industry estimates of the cost of crew ranging from $2,000-$3,000/day.

[59] These are the only variable costs incurred by shipowners; other voyage expenses (such as bunker costs or port charges) are borne by the charterer.

[60] Bunker costs will of course vary substantially over time as a function of fuel prices, but recall those are borne by the charterer, not by the shipowner, and as such are not part of c in our model.

B.3 Estimation of demand for shipping services

Here we provide more details on how we estimate demand (equation (11)) and construct the demand state z_t.

Variable construction We use our port call data to construct q_t. To do so, we first calculate how much cargo each ship carries each month, using information on idleness (i.e., the number of days the ship is carrying cargo) and utilization (i.e., the total proportion of the ship's capacity that is utilized on non-idle days).[61] We aggregate this across ships to calculate total cargo volume transported by container-ships every month, or q_t. Prices p_t are proxied using the time-charter rate index, which is the most granular price index available to us. Finally, the demand shifters W_t include a time trend, and an index for industrial production in the OECD.

Demand estimates Table 9 reports demand estimates. The first-stage regression indicates that as the average age of ships in the fleet increases, and as the share of older ships increases, prices increase (reflecting that it is costlier to deploy older ships that are slower and less fuel-efficient). Globally, demand for container-ship services is estimated to be somewhat price-inelastic, with an average elasticity around -1.36.[62] Demand increases as the industrial production index increases, but (controlling for industrial production) trends down over time.

Demand state We use the estimated demand coefficients from Table 9 together with equation (12) to construct the demand state z_t. We then smooth z_t by carrying out a local polynomial regression using an Epanechnikov kernel and a bandwidth of 3 months.[63] Figure 11 shows how z_t (both the raw measure, and the smoothed version we use in our structural model estimation) evolves over time. As expected, demand is high in the period prior to the financial crisis, but collapses during the Great Trade Collapse, before

[61] See Appendix A.2 for how we measure utilization, and Appendix A.3 for how we measure idleness.
[62] At the route level, demand is more elastic: Jeon (2022), who estimates demand at the trade route level, finds an elasticity of -3.89.
[63] Our results are very similar with other choices of bandwidth.

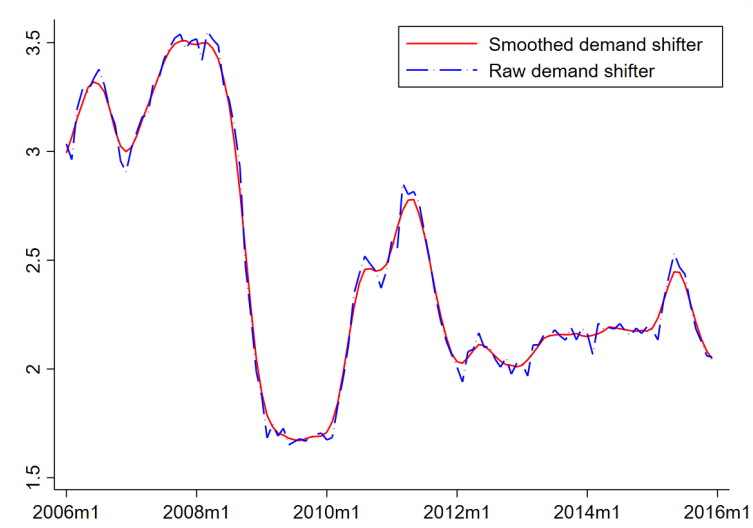
rebounding (but only partially) towards the end of 2010.

Table 9: Estimates of demand for container-ship services

	First-stage Price	2SLS Quantity
Price (timecharter rate)		-0.30
		(0.13)
Share of ships older than 20 years	1.07	
	(1.85)	
Average ship age	3.88	
	(3.01)	
Average ship size	-0.02	
	(0.04)	
Industrial Production Index (OECD)	4.02	1.34
	(0.21)	(0.54)
Time Trend	-0.80	-0.17
	(0.14)	(0.08)
N	120	120
Demand elasticity (mean)		-1.36

Standard errors in parentheses

Figure 11: Estimated demand shifter z_t, before and after smoothing



Transition process of demand state We model the evolution of the smoothed demand state, z_t, as an AR(1) process. We recover the parameters of the state transition process by

regressing z_t on z_{t-1} and a constant term. Table 10 shows the parameter estimates, as well as our estimate of the standard deviation of shocks to the AR(1) process.

Table 10: Estimates of AR(1) transition process for demand state

	Demand state
Lagged demand state	0.99
	(0.01)
Constant	0.02
	(0.03)
N	119
Standard deviation of demand shocks	0.07

Standard errors in parentheses

B.4 Algorithm for computing the equilibrium in the second stage

Overall, in the second stage of the estimation and for a fixed set of the parameters, we compute the equilibrium of the model (and the resulting moments) via an algorithm that involves an inner loop and an outer loop. In the outer loop we iterate over the 'perceived' state transitions and recompute these transitions to be consistent with equilibrium behavior in the inner loop. In the inner loop, given the current iteration of the state transitions, we compute the value functions and duration choice. Then we simulate the transitions over the sample period, and update the outer loop, continuing this process until the outer loop converges.

The detailed algorithm is as follows:

1. Initialize the algorithm at a guess of the outer loop 'perceived' transitions. Recall that we are taking the process for the demand state z_t as exogenous, and so we only need to iterate over the process for θ_t.

2. Given the current guess of the perceived transitions, compute six different value functions $(W_t, U_t^{charterer}, U_t^{ship}, M_t^{charterer}, M_t^{ship})$ via value function iteration:

   (a) Initialize these value functions as a set of nodes of the two aggregate states z_t, θ_t

(we compute intermediate values via linear interpolation).

(b) Using forward simulation, compute the (expected) match surplus at each node, for each contract length $\{1,2,...,\tau_{max}\}$.

(c) Compute the ex-ante surplus value at each node using the inclusive value formula:

$$W_t = \sigma \ln\Big(\sum_{\tau\in\{1,2,...,\tau_{max}\}} \exp(S_{t,\tau}/\sigma)\Big) + \sigma\gamma^{euler} \tag{16}$$

(d) Update the value functions and return to (a) until the value functions converge

3. Simulate the industry evolution over the sample period, using the empirical demand shock process.

4. Based on the simulated industry evolution, run a regression to update the perceived transition parameters: $\theta_t = \gamma_{\theta,0} + \gamma_{\theta,1}\theta_{t-1}$.

5. Return to step 1. until the transition parameters converge.

C Additional figures and tables

Figure 12: Average contract duration for container-ship time-charter contracts and time-charter index, 1999 - 2022

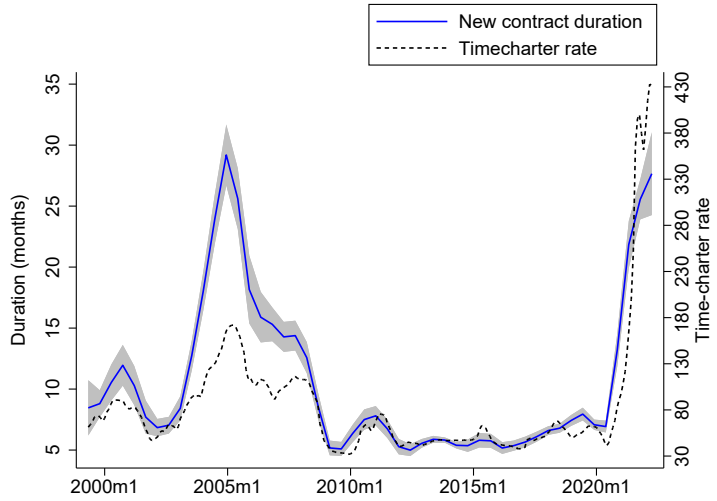
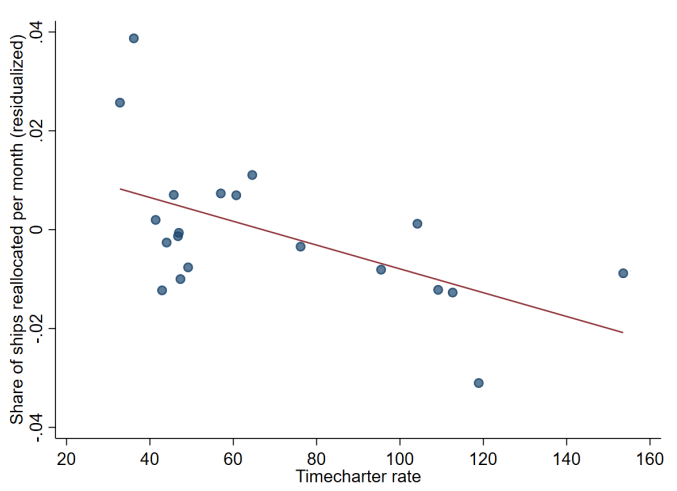


Figure 13: Reallocation (residualized) against the time-charter rate



Note: We regress our reallocation measure on ship fixed effects and an aggregate time trend, and obtained the residualized probability of reallocation for each ship, which we then aggregate across ships for each year-month. The figure shows a bin-scatter plot of the residualized share of containerships globally that were reallocated to a different route (vertical axis), against the containership time charter rate (horizontal axis).

Figure 14: Contract price, long vs. short contracts

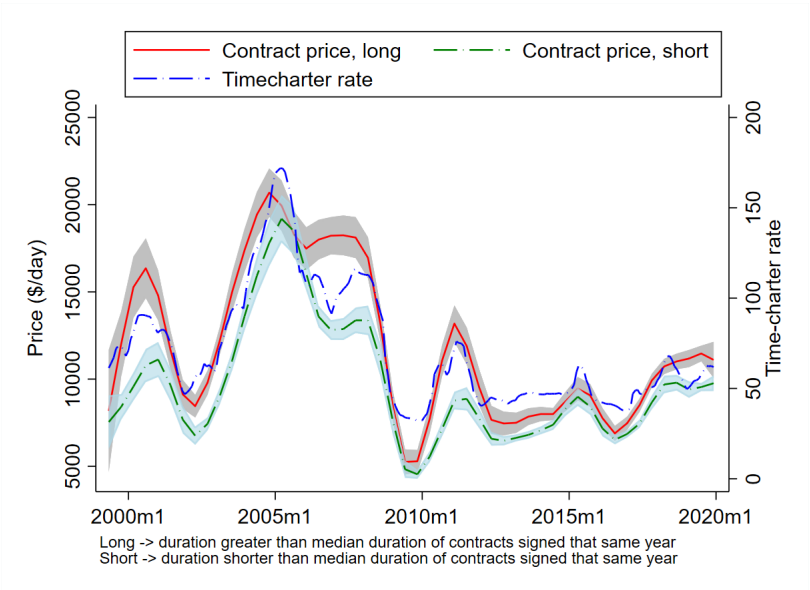


Table 11: Model fit: comparison of empirical and simulated moments

	Simulated	Empirical
Moments for share of ships under contract		
Mean share of ships under contract, bust	0.957	0.949
Mean share of ships under contract, (boom-bust)	0.033	0.031
Std. dev. in share of ships under contract	0.023	0.022
Dispersion moments		
Mean dispersion, 2006-08	0.157	0.156
Mean dispersion, 2009-10	0.174	0.178
Duration moments		
Mean duration, bust	6.346	6.171
Mean duration, 2006-08	9.754	9.871
Mean duration, 2009-10	6.2	6.167
Mean duration, boom	9.427	9.388
Std. dev. of contract duration	4.107	4.061
Elasticity of duration w.r.t. market thickness	-0.156	-0.15
Price moments		
Mean price, bust	0.72	0.753
Mean price, boom	1.185	1.169
Extension moments		
Prob. of extension, (boom - bust)	-0.061	-0.039