Booms, Busts, and Endogeneous Rigidities: Evidence from Containerships

Link to most recent version here.

[Preliminary and incomplete]

Maria Osipenko Nicholas Vreugdenhil Nahim Bin Zahur*

February 2024

Abstract

This paper investigates 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 inhibits the adjustment of these markets to productivity shocks. We show evidence for these mechanisms in the market for containership leasing contracts. We provide a framework that captures the details of the market and illustrates the tradeoffs conceptually. Overall, the results have implications for policies in this industry like shipping subsidies, as well as the fragility of the supply-chain to shocks.

*Vreugdenhil: Department of Economics, Arizona State University (email: nvreugde@asu.edu); Osipenko: Department of Economics, Arizona State University (email: mosipenk@asu.edu); Zahur: Department of Economics, Queen's University (email: nahim.zahur@queensu.ca). The authors do not currently have any financial ties to firms operating in transportation markets, but for full disclosure Nicholas was an employee of Uber from July 2019-June 2020.

1 Introduction

Capital (mis)allocation plays a crucial role in determining aggregate productivity across countries, industries, and in booms and busts, as illustrated by an expansive literature. This literature emphasizes the role of dynamic inputs which are difficult to adjust in inhibiting reallocation. A primary implication is that what appears to be misallocation may actually be efficient after accounting for exogenous adjustment costs and rigidities.[1]

In this paper we argue that these capital adjustment rigidities may be endogenous and that this can lead to inefficiency and misallocation. We focus on decentralized leasing markets for physical capital assets (like ships and aircraft). These markets are often highly cyclical. Adjustment rigidities in the form of term contracts prevent immediate reallocation of capital after a productivity shock.

In these industries contract duration may be an endogenous choice and can increase with market thinness as it becomes harder to find trading partners. Therefore, in a boom when the market for capital assets becomes thin, firms secure capacity by signing longer contracts. This produces an externality where even fewer assets are available. In equilibrium, this behavior generates an inefficiency where in booms firms sign contracts that are too long and the market is too thin, leading to an inefficiently low level of reallocation and inhibiting the adjustment of these markets to aggregate shocks.

Overall, we answer the questions: what are the effects of endogeneous rigidities in booms and busts, and what are the implications for capital misallocation as well as policy? Our empirical setting is the global market for leased containerships.[2] Here charterers (companies like COSCO who then transact with downstream exporters) lease ships from shipowners. 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

[1] See e.g. Asker et al. (2014). Earlier papers argue that capital is often inefficiently allocated, motivated by dispersion in the cross-section of the marginal revenue product of capital (Restuccia and Rogerson, 2008; Hsieh and Klenow, 2009).

[2] In the Appendix we show similar patterns for bulk shipping, another market where detailed contract data are available.

exact location and amount of cargo that each ship is carrying. As well as being an excellent example of a decentralized leasing market with fluctuations, this industry is also important in its own right and has taken on recent prominence as a central part of the supply chain. For example, although the focus of our analysis is on an earlier boom-bust period, during the Covid shock contract durations increased substantially.

We begin by documenting several facts about this industry using the raw data. We focus on the period 2005-2015: 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. Furthermore, in this period we observe fluctuations in market conditions.

First, we show that contract duration is cyclical and lengthens dramatically during booms, consistent with market participants locking in extra capacity when the market is thin. Matches formed under these contracts stretch long after the shipping industry bust in 2009. Second, we show that dispersion in capacity utilization — a proxy for productivity dispersion — is counter-cyclical. In particular, it is highest shortly after the market crashes when there are many long-term contracts signed in the boom that persist into the bust. Over time, reallocation — defined as moving a ship to a new route and schedule (where it might be better utilized) — also increases counter-cyclically and the market slowly adjusts.

Third, we perform a contract-level analysis to show in more detail how adjustment rigidities in the form of contract duration connect with these aggregate patterns. We show that contracts inhibit quick reallocation to shocks: the start of a contract is associated with a large jump in the probability a ship is reallocated to a new route. We also provide cross-sectional evidence about the determinants of contract duration and show that longer contracts are associated with thinner regional markets within a time period, as well as better ex-ante matches. Finally, we explore how misallocation in the leasing market can spill over to other parts of the supply chain through 'blank sailings', instances when a ship misses a scheduled stop, which can cause disruptions to downstream exporters and importers.

Based on these facts, we develop a model of the market. The model is dynamic and charterers (demand) need to search and match with ship-owners (supply) in order to lease ships which they then allocate to scheduled routes for their downstream exporters. Charterers enter the market and are heterogeneous in their valuation of a match, as well as a probability that their value for a match will expire in each period. Ships are homogeneous. In boom times there is higher entry rate of charterers. Upon meeting a particular ship, 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 match.

In our framework, it is an empirical question about whether contracts will be too long compared to the socially optimal level, and also whether contracts will get longer in a boom. Inefficiency may arise because agents do not fully internalize the contracting externality they create on the rest of a market when they sign a long contract.

In this draft we have not yet completed the estimation of the model, but we highlight its qualitative implications using numerical simulations. First, we explore how the decentralized contract length may differ from the optimal contract length (retaining search frictions) and how this varies in booms and busts. Second, we simulate a market crash (like we see in the data in 2009) and show how the industry would evolve moving from endogenous rigidities, to fixed rigidities (where the contract length is exogenously fixed), to the optimal contract. We argue that endogenous rigidities are essential for understanding the patterns of productivity dispersion in the data, amongst other things.

Finally, we evaluate the benefits of large-scale industrial policies in this industry that aim to construct more ships. Failing to account for endogenous rigidities may under-value the benefits of these subsidies. This is because adding ship capacity has an extra effect: it reduces the contract duration of existing ships endogenously towards the social optimum, leading to even more capacity available on the market.

Contributions We make three main contributions in this paper. The first is new descriptive evidence of how endogenous rigidities affect reallocation and other outcomes in a decentralized capital market. The market for leased containerships is a type of firm-to-firm market, but these kinds of markets have typically been difficult to study. This is because key objects like contracts between firms, as well as allocations of capital within and across firms at the level of granularity of a particular piece of physical capital, are typically secret. The containership industry, as well as other similar markets like bulk shipping, provide an exception to this rule: we observe a relatively large sample of contracts between firms. Furthermore, due to the spatial nature of the 'capital allocation' — adding a particular ship to a route — we can leverage port call data to observe the universe of allocations in this industry.

The second contribution is a new empirical framework that incorporates endogenous rigidities into a search and matching model with booms and busts. A key difference from previous work is that we embed a choice of *how long* to match for into these models. Standard applications of search and matching models are in labor economics where at most one-sided commitment is possible: firms can commit to a contract but employees cannot be forced to work and can quit employment at any time. As a result, although the expected duration of a match enters indirectly into agents' decisions, it cannot be directly contracted over.[3] Physical capital markets involve agreements between two firms, and so can operate quite differently with two-sided commitment and explicit agreements over duration.

The third contribution is a new set of results that showcase the consequences of endogenous rigidities in capital markets. One group of results relates to understanding patterns in the macro-productivity literature that are present in more aggregated data. Specifically, we observe counter-cyclical productivity dispersion in our data, consistent with broader patterns over a wide range of industries (Kehrig, 2015). From the point of view of a 'Schumpeterian' theory of creative destruction this fact is puzzling: one would expect low productivity arrangements to be destroyed in recessions leading to pro-cyclical productivity dispersion. We show that endogenous rigidities provide a novel explanation for observed counter-cyclical productivity dispersion. Simpler models

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

involving fixed rigidities or an optimal contract either fail to match the transitional dynamics of productivity dispersion, or the level of counter-cyclical dispersion.

The second group of results relates to policy implications for the industry, such as the billions of dollars spent on industrial policies subsidising shipbuilding and shipping. If extra ships are added to the market, we show there can be an extra benefit to increased capacity, in the form of reduced equilibrium rigidities, which improves allocations of the existing capital.

1.1 Related literature

This paper is related to four strands of literature. The first is a set of papers about the inner workings of decentralized asset markets. As well as those already discussed, some of these papers highlight the role of search frictions, transaction costs, and market thinness in determining efficient allocations e.g. Gavazza (2011a), Gavazza (2011b), Gavazza (2016). Other papers show the central role of capital reallocation in determining efficient allocations e.g. Lanteri and Rampini (2023). In this paper we combine elements of these papers while adding the additional consideration of contract duration, and our results hinge on how all the elements interact in equilibrium with booms and busts.

The second strand is the literature that investigates how contract duration, and more broadly relationships, are affected by market thinness and transaction costs. Hubbard (2001) shows that trucks who operate in thinner markets sign longer contracts. MacKay (2022) estimates an empirical auction framework where agents choose contract duration in the presence of transaction costs. Darmouni et al. (2023) quantifies the insurance value of quantity contracts in the presence of trading frictions in the context of the pulp and paper industry. We build on the idea that contract duration is a choice and influenced by market conditions and other factors, with a focus on how these forces interact in equilibrium. For instance, in our framework, contract length is influenced by the market, but equilibrium market thinness is also determined by aggregating these individual contract duration choices.

In terms of the equilibrium effects of organizational form, Harris and Nguyen (2023) explore the effects of relationships in the trucking industry. They show that relationships can increase equilibrium market thinness to a level that is not necessarily socially optimal. Our results are complementary to this work, but our focus is instead on the effects of a different organizational form: term contracts where match duration is a choice that is made when the contracts are signed. The key difference is that such term contracts involve two-sided commitment over time; relationships are effectively zero-sided commitments. As a consequence, amongst other things, contracts can lead to lock-in ex-post, and long contracts formed in a boom can persist well into a bust inhibiting productive reallocation.

The third strand is the literature that studies shipping markets. These papers include Kalouptsidi (2014), Kalouptsidi (2018), Jeon (2018), Ganapati et al. (2023), amongst others. We build on many of the institutional details described in these papers in our analysis.

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 container ships operating on regular schedules involving many trips. As a result, 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, which account for the majority of world trade in goods.[4] Container ships can be thought of as the 'buses of the ocean', typically

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

Figure 1: Container Shipping Industry Structure

(a) Container Ship's Course Visualization **(b)** Value Chain

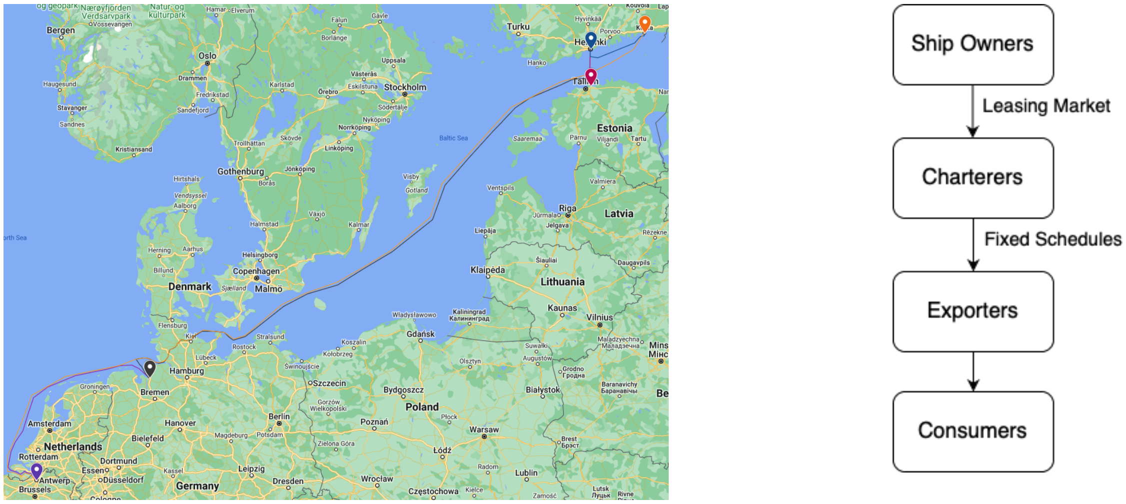

Note: Panel (a): This shows a geographical map of a ship in our data performing its scheduled stops. Here, the ship has the following schedule: Antwerp - Bremerhaven - Muuga - Helsinki - Kotka - Antwerp. Note that the ship uses the Kiel canal through Denmark. Panel (b): One component not pictured (but 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.

operating on fixed schedules where they pass through a designated set of ports.[5] At each port a ship drops off and picks up a portion of its containerized cargo. Figure 1a demonstrates a map with the route for one of the container ships in our sample.

The value chain of the maritime shipping industry is visualized in Figure 1b. Cargo owners (or exporters) are the enterprises that use maritime transport providers and other suppliers of services to import/export their cargo. Container ships are operated by liner companies that specialize in the transport of containerized goods across the world (for example, COSCO). About half of the container ships in the world are owned by these liner companies; these "owner-operated ships" are rarely leased out to other companies, and form the core of the fleet of liner companies. The remaining container 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 liner companies.

[5]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).

These ships are leased out to liner companies using time-charter contracts that range from a few months to over five years in length, and specify a fixed day-rate that the liner company (or the “charterer") pays the shipowner. It is this leasing market that we focus on in this paper. Unlike other chartering methods, time charters do not specify the ship’s routes or ports of call, providing the charterer with the flexibility to choose these based on cargo requirements. Ships are available for lease globally, and the transit time for relocating a ship to various international regions typically ranges from 20 to 45 days. For example, the standard transit time for ocean freight from the USA to Singapore is approximately 15 to 25 days, whereas shipping from China to the UK generally takes between 35 and 45 days.[6]

Search frictions The market for leasing containerships on time charters is decentralized, and no widely used central platform for connecting charterers (carriers) and shipowners currently exists. Both sides of the market employ specialized brokers, such as Bertling, Clarksons and Maersk Broker, to assist with the chartering process, with commission fees typically ranging from 1.25%-1.5%.[7] The fact that firms commonly hire brokers is itself indicative of the presence of search frictions.

The decentralized nature of the market makes it challenging for any one firm to keep track of all of the information on potential matches. Much of the relevant information is shared via emails, which means that shipping companies have to collate and process information from a large number of emails.[8] As one manager in the industry observed: “Many shipping companies face e-mail 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."[9]

In response to these challenges, some firms have now developed software which parses through

[6]See https://www.maersk.com/logistics-explained/transportation-and-freight/2023/09/27/sea-freight-guide

[7]For example, in 2013, Seaspan (a shipowner) reported paying a commission fee of 1.25% to brokers on its time-charters to MSC. See https://www.seaspancorp.com/wp-content/uploads/2014/10/SSW_2013_Annual_Report_20-F.pdf.

[8]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.

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

these emails and provides more structured information to shipping companies and brokers in exchange for subscription fees. There have also been recent attempts to develop centralized platforms to connect owners and charterers, but none that has as yet succeeded in commanding wide usage. Moreover, such efforts to develop platforms have largely focused on dry and wet bulk chartering, with only a handful of startups that have targeted container ship chartering.[10]

One reason why it may be challenging to find a match is that the market is quite unconcentrated and fragmented, with a large number of agents searching on both sides of the market. For example, in our study period 2005-2015, the HHI for charterers is 415, and for ship-owners who lease out their ships the HHI is 136.[11] 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.[12]

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 of contracts includes over 16,000 time-charter contracts; for the period of 2005 - 2015, which is the main focus of our analysis (since that is the period for which we have detailed ship movement data), we observe around 4,109 contracts. Table 1 provides descriptive statistics on the duration of the contracts, rates, age of the ships contracted, capacity utilization, probability of reallocation, and aggregate indexes for observations in 2005-2015.

The other key dataset we utilize in our analysis is port call data for 2005-2015 time period 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 port

[10]See https://www.pwc.com/gr/en/industries/Chartering_marketplaces.pdf.

[11]Here, we compute market shares by the total ship capacity in dead-weight-tonnes in the leasing market.

[12]Note that there did exist alliances in the 2005-2015 period, but these were relatively rare. 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.

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

Variable	Obs	Mean	Std. dev.	Min	Max
Duration (months)	4,109	10.74	14.3	0.17	180.0
Rate ($/day)	4,103	10,397	6,620	2,400	59,950
Ship age (years)	4,057	8.9	6.09	0	36
Capacity utilization	906,857	0.53	0.24	0	1
Number of port calls	1,655,140				

call. For a subset of port calls, we also observe information on the ship's "draught", or the the vertical distance between the waterline and the bottom of the hull. Draught data is valuable as it allows us to infer how much cargo the ship is carrying, since a ship that is carrying more cargo will sink deeper into the water (causing its draught to increase). We explain exactly how we construct utilization from draught data in Appendix A.2.

In order to merge port calls with fixtures data, we use information from Vessel Finder website that allows us to match ship names with IMO numbers. A total of 1,655,140 port call observations are used in this paper, with 299,903 of them matched with contracts data (3,546 unique contracts are matched to the port call data).

2.3 Measuring reallocations and productivity

Reallocations We use the port call data to identify when a ship reallocates from one itinerary to another. This is in general a challenging problem, since we do not directly observe shipping itineraries in our raw data. Instead, what we observe are repeated sequences of port calls. We therefore identify reallocations as large deviations across space from old port call sequences to a new set of port call sequences. We describe the algorithm for constructing reallocations in Appendix A.3.

Productivity The production function in this market is relatively simple. Containerships are the sole input, and this maps to an output equal to a total amount of cargo transported per time unit. Charterers are differentiated by heterogeneous productivity, which is exogenous and can change over time. Note that — as is typically the case when measuring productivity — we do not observe productivity directly in the data. However, more productive charterers will choose to transport more cargo per trip, after controlling for ship heterogeneity; therefore, capacity utilization moves one-to-one with productivity. As a result, in the descriptive evidence, we proxy for dispersion in productivity by considering the corresponding observed dispersion in residualized capacity utilization. Later when we estimate the model we also recover the productivity process directly.

The above method of measuring productivity dispersion is consistent with the way that practitioners view productivity in the industry e.g., Adland et al. (2018). It is also consistent with prior economics literature on measuring productivity in other transport markets e.g. Gavazza (2011b) who uses a similar measure and conceptual 'production function' in the market for aircraft.

3 Descriptive evidence

In this section, we describe some of the key empirical patterns that underlie our analysis.

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

Figure 2a shows how the average duration of newly signed leasing contracts changes over time. In the same figure, we also illustrate the time series evolution of the containership timecharter rate index (which reflects changes in aggregate demand).[13] As we can see, firms sign significantly longer contracts during booms. This effect was especially pronounced during the boom in container ship markets during the mid-2000s, when the average duration of newly signed contracts increased from less than 10 months to more than 24 months at the peak of the boom. More recently, during the upsurge in container shipping demand after the Covid-19 pandemic, there was a similar increase in the duration of newly signed contracts mirroring the rise in container freight rates, as illustrated in Appendix Figure 7.

[13]The containership timecharter rate index is a measure of industry container freight rates published by Clarksons.

A consequence of this is that there may be an overhang of contracts left over from the boom, long after the boom has evaporated. This contract overhang is strikingly illustrated by Figure 2b, which compares the average duration of newly signed contracts (at each point in time) against the average duration of existing contracts (i.e., contracts that are still active). Following the economic crisis between 2008 - 2009, which led to a sharp decline in containership freight rates, new contracts signed were much shorter in length (less than 10 months on average after 2009). However, the average duration of existing contracts remained high, as the longer contracts signed during the boom continued to persist long after the boom ended.

Overall, the patterns in Figure 2a and Figure 2b suggest that the lengthening of contracts caused by boom-bust cycles may inhibit subsequent reallocation of container ships which are locked into longer contracts. As a result, we next discuss patterns around reallocation and productivity dispersion.

Observation 2: Reallocation, and dispersion in capacity utilization, are counter-cyclical.

Counter-cyclical dispersion in capacity utilization As previously mentioned, we proxy for productivity dispersion using dispersion in ship capacity utilization, residualized by ship type. We then consider how this measure changes across the business cycle in Figure 2c.

Before discussing the results, note that it is primarily *dispersion* in capacity utilization that we are interested in, and not *average* capacity utilization. This is because changes in average utilization over time confound changes in productivity with aggregate demand shifts. As a result, there are at least two possibilities that could explain differences in average capacity utilization across periods: (i) there are alternative matches that exist where ships could produce more output, but they are not allocated to them (which is suggestive of misallocation) (ii) in that period all matches have equally low capacity utilization due to a low demand shock. By contrast, dispersion within a time period in capacity utilization indicates that there are more productive matches available, but ships are not matched to them, which is a necessary condition for misallocation.

Figure 2: Industry evolution over the business cycle

(a) Average contract duration and time-charter index

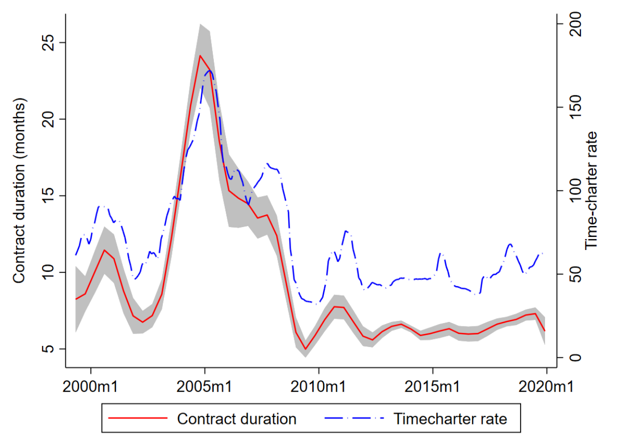

(b) Average duration: new vs. existing contracts

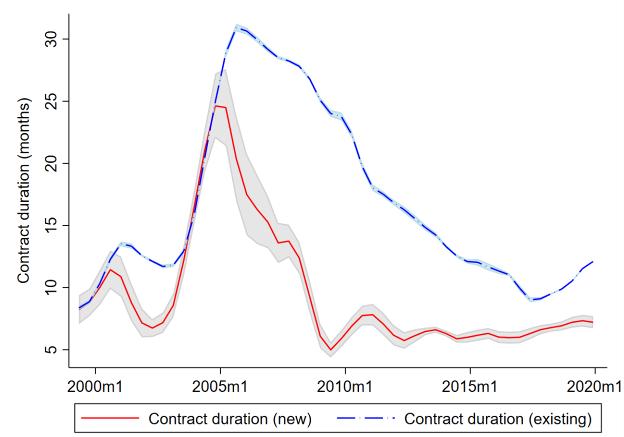

(c) Dispersion in residualized capacity utilization (a proxy for productivity dispersion)

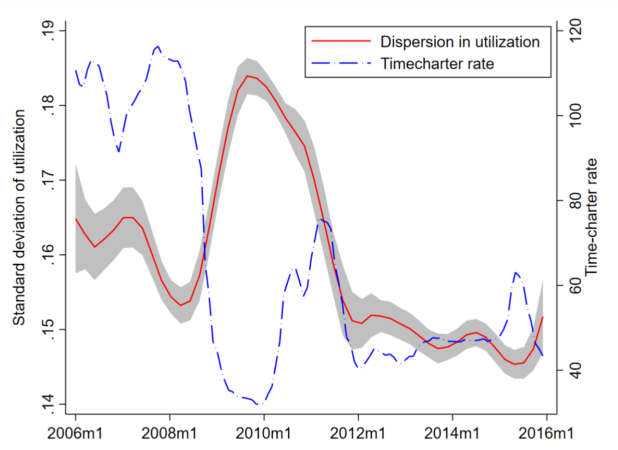

(d) Probability of reallocation per month

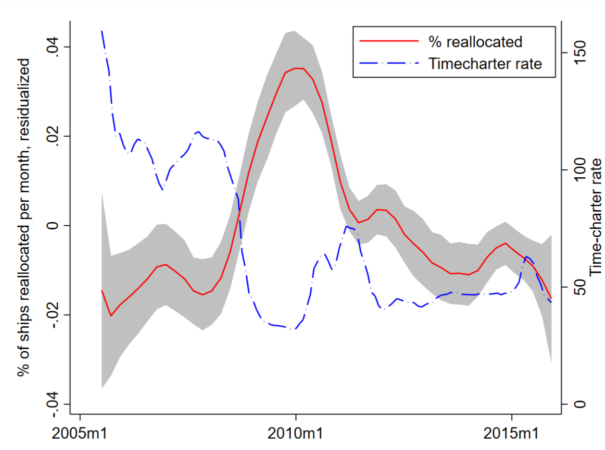

Note: Panel (c): Plots dispersion in residualized utilization across ships every month, where residualized utilization is obtained by regressing utilization on ship fixed effects and a time trend. Panel (d): plots the residualized probability of reallocation each month over time. We obtain this by regressing our reallocation measure on ship fixed effects and an aggregate time trend, and then calculate the residualized probability of reallocation for each ship, which we then aggregate across ships for each year-month.

Figure 2c illustrates that dispersion in capacity utilization (measured by the the standard deviation in utilization across ships) appears to be significantly counter-cyclical, and rises significantly after the financial crisis.[14]

Counter-cyclical capital reallocation Figure 2d shows how the share of container ships reallocated each month changes over time. We find that capital reallocation follows a counter-cyclical pattern. In particular, after the financial crisis in 2009, and the associated downturn in global shipping markets, the share of ships that reallocate to more distant trade routes increased by almost one-third. Appendix Figure 9 corroborates this finding by illustrating that there is a negative correlation between the probability of reallocation and the time-charter rate index (our proxy for aggregate demand in this market).

These empirical patterns are consistent with the previously discussed patterns around counter-cyclical dispersion in capacity utilization. In a bust, there is an increase in the dispersion in productivity (and profitability) between different trade routes, and thus more opportunities for profitable reallocation. Overall our descriptive analysis suggests that container ships do, on the whole, respond to these changing incentives by reallocating more often during economic downturns.

Observation 3: Contract rigidities inhibit reallocation.

The overhang of contracts in Figure 2b may have negative consequences for efficiency in the presence of contractual rigidities. In the containership leasing market, contracts typically specify a fixed per-day rate received by the shipowner. Since these contracts feature two-sided commitment, shipowners have little incentive to re-negotiate lengthy contracts signed at the peak of the boom (typically at high prices). As such, contract overhang may lead to ships being locked into specific relationships, inhibiting reallocation.

To examine the relationship between contracting and reallocation, we look at how the prob-

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

Figure 3: Contract-level analysis

(a) Contract rigidities inhibit reallocation
(share of ships reallocated/month; residualized)

(b) Capacity utilization (residualized) over the contract

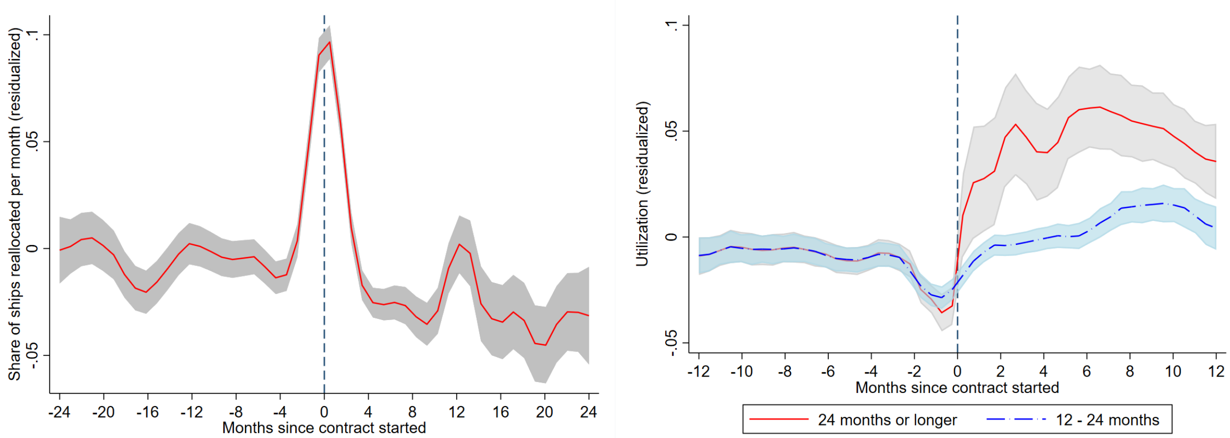

(c) Longer contracts are associated with thinner markets

	(1) Log(duration)	(2) Log(duration)	(3) Log(duration)	(4) Log(duration)
Log(# ships nearby, same type)		-0.12***		-0.17***
		(0.031)		(0.030)
Log(# ships nearby, all types)	-0.15***		-0.21***	
	(0.033)		(0.032)	
Log(time charter index)			0.98***	0.98***
			(0.031)	(0.031)
1[mid-size ship]	0.10***	0.10***	0.11***	0.10***
	(0.033)	(0.034)	(0.033)	(0.033)
1[large-size ship]	0.44***	0.47***	0.44***	0.48***
	(0.039)	(0.039)	(0.038)	(0.038)
Year fixed effects	Yes	Yes	No	No
Observations	3,531	3,531	3,531	3,531
Adjusted R^2	0.243	0.242	0.250	0.248

Note: Panel(a): The figure plots the share of container ships 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. Panel (b): The figure plots the residualized utilization for ships, for each month in a 2 year window around the month when a new contract is signed. We do so separately for contracts that are 12 - 24 months in duration (dashed blue line), and contracts that are longer than 24 months (solid red line). The residualized utilization is obtained for each ship-month observation by regressing utilization on ship fixed effects, year fixed effects, and the time-charter rate and bunker rate (to capture aggregate changes in demand and costs), and obtaining the residuals. Panel (c): Standard errors in parentheses, * (p<0.10), ** (p<0.05), *** (p<0.01).

ability that a ship is reallocated to a different geographical route changes over the lifetime of a contract. If contracts did not inhibit reallocations, one would expect reallocations to happen more or less independently of which stage of the contract the ship happens to be in. Figure 3a shows, however, that the probability of reallocation jumps towards the start of the contract (controlling for ship fixed effects and a time trend). Once the ship is a few months into the contract, the probability of reallocation decreases significantly.

Ships are therefore mostly reallocated at the start of a contract, and rarely afterwards (until they begin another new contract). Thus the lengthening of contracts during the boom limits the extent to which ships reallocate in response to changing profitability of operating in different regions, which may lead to inefficiency in how ships are allocated. This contributes to the counter-cyclical dispersion in capacity utilization that we observe in the industry: indeed we find that during busts, there is increased dispersion in capacity utilization among ships under longer contracts, compared to ships under shorter contracts (see Appendix Figure 12).

Observation 4: Firms sign longer contracts when (i) the market is thinner, and (ii) there is higher ex-ante match quality.

How do firms choose their contract length in the micro-data? We present descriptive evidence for two key considerations: longer contracts are associated with thinner markets, and longer contracts are associated with high ex-ante match quality. These considerations need 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.

Market thinness We present results for how market thinness is associated with contract duration in Table 3c. We focus on isolating this relationship in the cross-section, controlling for time-varying factors that might also affect contract duration (for example, the match value might be systematically higher in booms, causing agents to sign longer contracts for reasons other than thin markets). In order to construct a market thinness measure we need to take a stand on how to define a market. To do so we exploit that it is costly to move a ship long distances, and that charterer needs are relatively time-specific. We therefore 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.[15]

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 also consider 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 and so direct their search. In order to construct ship type we split ships into three bins based on their capacity (measured in twenty-feet equivalent or TEU).

Overall, the results across the four specification in Table 3c 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 35% 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 and the modeling exercise. This is because the thinness of the regional markets is highly correlated across booms and busts.[16]

[15]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.

[16]Table 3c also illustrates that larger ships tend to operate under longer contracts. This is in part because larger ships often operate on longer shipping routes across more distant ports. In addition, the market for larger markets is thinner, creating an incentive for charterers to sign longer contracts for such ships (for fear of not being able to match with them in the future). Consistent with this, we observe that increases in contract durations induced by booms (which we observe for all types of ships) are most pronounced for larger ships. These results, though not included in the present draft, are available upon request.

Match quality Figure 3b shows that match quality tends to be higher for longer contracts: average utilization over the first 12 months is higher for ships that are signed under contracts that are 24 months or longer in length, compared to those with contracts shorter than 24 months. This suggests that higher-value charterers select into longer contracts.

Observation 5: Periods with high dispersion in capacity utilization are associated with more blank sailings, a key cause of supply chain fragility.

So far the descriptive evidence we have presented is about the effects of endogenous rigidities on the containership leasing market. However, as previously illustrated in Figure 1b, the containership leasing market is connected to downstream exporters and so inefficiencies in this market may have follow-on effects to other parts of the supply chain.

Concretely, in this setting, we investigate how matches with low capacity utilization affect blank sailings. Blank sailings are occasions when a ship skips a port on its fixed schedule. 'Blanking' a port — as it is known in the industry — can be highly disruptive to downstream exporters. One particular example is that many importers use lean manufacturing processes, where inventory costs are minimized through careful scheduling of inputs that arrive 'just-in-time' for production. Here, shipment reliability is critical; even infrequent blank sailings can have high-magnitude consequences.

We identify blank sailings using our port call data using an algorithm detailed in Appendix A.4. We plot the probability of a blank sailing in Appendix Figure 10. Blank sailings spiked in periods like 2010-2011, where there is a high degree of dispersion in capacity utilization in the market. Appendix Figure 11 shows that periods with higher capacity utilization dispersion feature a higher probability of blank sailings. These patterns are consistent with anecdotal evidence from market participants, who state that blank sailings are used by charterers to reduce costs on low-demand routes.[17]

[17] See, for example: https://www.gcllog.com/post/what-are-blank-sailings-and-why-are-they-becoming-more-common

4 Model

In this section, we describe a simple dynamic model of the container-ship leasing industry that illustrates the relationship between boom-and-bust cycles, market thinness, and contract duration.

4.1 Setup

Time is discrete and denoted by t. There are two types of agents in our model, shipowners and charterers. All agents are forward-looking and have discount factor β. We assume that each shipowner owns a single ship, and treat ships and shipowners interchangeably from here onward.[18] There are n homogeneous container ships. The per-period cost of operating a ship is c: this includes bunker fuel costs, labor costs and port charges.

Shipowners lease out ships to charterers, who use them to transport container goods across global shipping markets. Each charterer is distinguished by their type $x = \{v, \eta\} \in X$. v is the per-period value to the charterer of leasing the ship, while η is the probability that the value will survive to the next period (equalling 0 otherwise). Let v_x and η_x refer to the v and η of a charterer of type x. The set of possible types X is discrete. Without loss of generality, we assume that the lowest possible v_x exceeds c (since any charterer with a valuation lower than the cost would have no reason to enter).

Timing In each period, the timing is as follows:

1. *Entry of charterers:* $e_{x,t}$ charterers of type x enter the market. Each charterer's type x is realized when they enter the market, cannot change over time, and is publicly observable.
2. *Search and matching:* Available charterers search for ships. Searching charterers and ships are randomly matched.

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

3. *Contract duration and price:* If agents are matched, they choose both how long to match for (the contract length), and a fixed price paid by the charterer to the ship.[19] 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.

4. *Contract expiry and charterer exit:* At the end of the period, any past contracts that were due to end in period t expire, and the ships and charterers in those contracts become available to match in the next period. The valuations of all charterers are updated: among all existing charterers of type x, the valuation remains at v_x with probability η_x ("alive" charterers). Any charterer whose valuation turns to 0 and who is not currently under contract will exit the market.

Search and matching Matching takes place in a single global market.[20] The mass of available charterers of each type x equals $a_{x,t}^{charterer}$, so the total mass of available charterers equals $a_t^{charterer} = \sum_x a_{x,t}^{charterer}$. 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.

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. Let these probabilities equal $q_t^{charterer}$ (for charterers) and q_t^{ship} (for ships).

[19]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.

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

4.2 Contract choice

When a charterer is matched to a ship, they choose the duration of the contract, τ, to maximize their joint surplus, and then Nash bargain over the price. In the event of disagreement, both the charterer and ship wait until the next period, when they may enter the pool of searching agents. 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, and also the contract duration choice does not affect the outside options of the agents. Therefore, maximizing total match surplus leads both sides of the match to be maximally well-off once this surplus is split via Nash bargaining.

Choice of contract duration Let $S_{x,t,\tau}$ denote the total lifetime surplus (or the *match surplus*) from a τ-period contract agreed between a ship and type-x charterer in period t. We describe this object in more detail below. At the time when agents match, they also receive 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 information, but the ε_τ are idiosyncratic to each particular match, as well as to each contract length.The contract duration is chosen to maximize the realized total surplus. Denote $W_{x,t}$ as the ex-ante surplus (the expected value of the surplus before the ε_τ are drawn):

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

where τ_{max} is the maximum possible contract duration. Let $P_{x,t,\tau}$ denote the probability that a matched type-x charterer chooses a contract of length τ:

$$P_{x,t,\tau} = \frac{exp(S_{x,t,\tau}/\sigma)}{\sum_{\tau'} exp(S_{x,t,\tau'}/\sigma)} \tag{2}$$

Match surplus We now define the match surplus in terms of the value functions of searching agents. Later we describe how these value functions can be derived. Let U_t^{ship} denote the value of a ship that will search for a charterer in period t. Likewise, let $U_{x,t}^{charterer}$ denote the value of a charterer of type x that will search for a ship in period t. The match surplus $S_{x,t,\tau}$ is the joint value to the charterer and ship from a τ-period contract, minus the sum of their outside options (i.e., the sum of their respective values from being unmatched):

$$S_{x,t,\tau} = \underbrace{\sum_{k=0}^{\tau-1} \beta^k(\eta_x^k v_x - c) + \beta^\tau \mathbb{E}_t U_{t+\tau}^{ship} + \beta^\tau \eta_x^\tau \mathbb{E}_t U_{x,t+\tau}^{charterer}}_{\text{Value of match}} - \underbrace{\left(\beta \mathbb{E}_t U_{t+1}^{ship} + \beta \eta_x \mathbb{E}_t U_{x,t+1}^{charterer}\right)}_{\text{Outside options}} \quad (3)$$

To explain how we arrive at the above equation, note that the outside options of the ship and the charterer are to wait till the next period, and attempt to match again. The present value of that outside option equals $\beta \mathbb{E}_t U_{t+1}^{ship}$ for ships and $\beta \eta_x \mathbb{E}_t U_{x,t+1}^{charterer}$ for type-x charterers. (A charterer will only search next period if they retain a positive valuation, which happens with probability η_x, which is why their effective discount factor is $\beta\eta_x$.)

The value of the match consists of three components. The first term, $\sum_{k=0}^{\tau-1} \beta^k(\eta_x^k v_x - c)$, is the discounted sum of the expected per-period surplus generated by the match over the length of the contract (τ). Recall that the charterer's probability of retaining a positive valuation next period is η_x, so the expected benefit to the charterer from operating the ship k periods from today is $\eta_x^k v_x$. Once the contract expires, the ship and charterer become unmatched again; the expected value of being unmatched τ periods from now, discounted back to today, equals $\beta^\tau \mathbb{E}_t U_{t+\tau}^{ship}$ for ships and $\beta^\tau \eta_x^\tau \mathbb{E}_t U_{x,t+\tau}^{charterer}$ for type-x charterers.

Value of searching agents We now describe the value of searching ships and charterers. A searching charterer of type x 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. Their value can therefore be written as:[21]

$$U_{x,t}^{charterer} = q_t^{charterer} \underbrace{\left((1-\delta)W_{x,t} + \beta\eta \mathbb{E}_t U_{x,t+1}^{charterer}\right)}_{\text{(Expected) payoff if matched}} + (1-q_t^{charterer}) \underbrace{\beta\eta \mathbb{E}_t U_{x,t+1}^{charterer}}_{\text{Payoff if not matched}} \quad (4)$$

[21] We prove that the expected match payoff can be written in this form in Appendix B.

A searching ship finds a match with probability q_t^{ship}. Conditional on finding a match, it matches with a type-x charterer with probability equal to the share of type x charterers as a fraction of total available charterers: $h_{x,t} = a_{x,t}^{charterer}/a_t^{charterer}$. Thus, the value of a searching ship can be written as:

$$U_t^{ship} = q_t^{ship} \sum_x h_{x,t} \underbrace{\left(\delta W_{x,t} + \beta \mathbb{E}_t U_{t+1}^{ship}\right)}_{\text{(Expected) payoff if match type-x charterer}} + (1 - q_t^{ship}) \underbrace{\beta \mathbb{E}_t U_{t+1}^{ship}}_{\text{Payoff if not matched}} \tag{5}$$

4.3 States and transitions

States The detailed industry state in period t, s_t, consists of both the distribution of searching agents, and the distribution of current matches. Each agent takes the industry state as given.

As discussed above, the mass of charterers available to match in period t is $a_{x,t}^{charterer}$: this includes charterers who just entered, as well as existing charterers whose contracts have expired and who still have positive value v_x. The mass of ships available to match in period t equals $a_{x,t}^{ship}$.

Each current match can be described by (i) the charterer's type x (ii) the number of periods left on the contract τ (iii) whether or not the charterer retains a positive valuation. Let $m_{x,t,\tau}^{contract}$ denote the mass of agents matches of type x with τ periods remaining, regardless of whether or not the charterer's valuation is positive or not ("contract" matches). Let $m_{x,t,\tau}^{alive}$ denote the mass of contracted matches of type x with τ periods remaining where the charterer's valuation remains positive ("alive" matches).

The distribution of current matches can therefore by described fully by $\mathbf{m}_{x,t}^{contract}$ and $\mathbf{m}_{x,t}^{alive}$ which correspond to vectors of the current "contract" matches and "alive" matches respectively:

$$\mathbf{m}_{x,t}^{contract} = (m_{x,t,1}^{contract}, m_{x,t,2}^{contract}, \ldots, m_{x,t,\tau}^{contract}, \ldots, m_{x,t,\tau_{max}}^{contract}) \tag{6}$$

$$\mathbf{m}_{x,t}^{alive} = (m_{x,t,1}^{alive}, m_{x,t,2}^{alive}, \ldots, m_{x,t,\tau}^{alive}, \ldots, m_{x,t,\tau_{max}}^{alive}) \tag{7}$$

Transitions At the start of each period, $e_{x,t}$ charterers of type x enter the market, and join the pool of "alive" charterers from previous periods who are available to match (either because their contracts just expired, or because they were unable to find a match last period). The mass of searching charterers of type x therefore equals:

$$a_{x,t}^{charterer} = \underbrace{\eta_x(1-q_t^{charterer})a_{x,t-1}^{charterer} + \eta_x m_{x,t-1,1}^{alive}}_{\text{Alive unmatched charterers}} + \underbrace{e_{x,t}}_{\text{Entering charterers}} \tag{8}$$

where $m_{x,t-1,1}^{alive}$ is the mass of charterers of type x who had 1 period left on their contract in period $t-1$ (and whose contracts therefore just expired), while $(1-q_t^{charterer})a_{x,t-1}^{charterer}$ is the mass of searching charterers last period who were unable to find a match.

All ships that are not currently under a contract are available to match, and therefore the mass of searching ships equals:

$$a_t^{ship} = n - \left(\sum_x \sum_{k=1}^{\tau_{max}} m_{x,t,k}^{contract}\right) \tag{9}$$

Finally, to update the distribution of current matches ($\mathbf{m}_{x,t}^{contract}$ and $\mathbf{m}_{x,t}^{alive}$), we add any new matches and count down all existing contracts by 1 period at the end of each period.

4.4 Equilibrium

For this draft of the paper, we consider a steady-state equilibrium, where all agents view the industry state s_t as fixed when choosing contract duration and price.[22] A steady state equilibrium is characterized by a mass of searching agents ($a_x^{charterer}$, a_x^{ship}), a distribution of contract and alive matches, ($\mathbf{m}_x^{contract}$, $\mathbf{m}_x^{alive}$), contract duration choice probabilities, and prices, such that the following conditions are satisfied:

1. Agents optimally choose contract duration, according to equations (1)-(5).

[22] In a future version of the paper, we will incorporate state transitions as well as agents' beliefs about future states.

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 4.3.

4.5 Discussion

We now discuss some of the key properties of the model, and comment on modelling assumptions.

Our model builds on standard search and matching frameworks that have been widely used to study the role of search frictions (Hosios, 1990), but allows agents to endogenously choose the duration of their contracts. The model highlights how agents' choices of how long to match for respond to search frictions, and these choices in turn dynamically affect market thickness and result in endogenous rigidities. During a boom, when the charterer entry rate is higher (represented by an increase in $e_{x,t}$), there are many more searching charterers than ships. Charterers and ships that are matched during a boom have a unilateral incentive to match for longer, to guard against the possibility that the charterer may be left without a match in future periods. However, these longer contracts endogenously result in a thinner market and lead to lock-in, where some relationships persist due to contractual rigidities, despite having low match value. We further explore these channels and the misallocation caused by endogenous contractual rigidities in counter-factual simulations described in Section 6.

We currently abstract away from contract extensions: once a contract between a given charterer and ship expires, the charterer and ship both enter the pool of searching agents on both sides of the market, and have to find a match again. The majority of leasing contracts during the study period are not extended, suggesting this is not a first-order concern. Nevertheless, in a future version of the paper, we plan to incorporate contract extensions into the model, by allowing the ship and charterer to decide (at the time the contract is due to expire) whether to extend the contract for another period.

In practice, once a charterer has leased a ship under a time-charter contract, they integrate this ship into their existing fleet of owned and leased ships, and determine which trading route to allocate the ship to, as well as the schedule of stops the ship will make. They then provide shipping services to downstream exporters, competing with other charterers (carriers). We do not explicitly model these choices or the interaction between charterers and downstream exporters, since the focus of our model is on explaining the contract duration choices in the upstream leasing market. Instead, we capture these in a reduced-form manner in our model, by allowing heterogeneity across charterers in their per-period value (v), the probability of survival (η), and the entry rate e, and allowing these to depend on market conditions. For instance, an increase in downstream shipping demand in a particular geographic region can be translated within our model into an increase in v and e for charterers that specialize in those regions.

5 Estimation [in progress]

Estimation for this paper is still in progress. For the counterfactuals we therefore present numerical exercises using the model to illustrate the main economic forces at work in the paper. In the numerical simulations below, we assume the urn-ball matching function (Petrongolo and Pissarides, 2001), which implies the following match probabilities:

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

$$q_t^{charterer}(\theta_t) = min\{\theta_t(1 - exp(-\alpha/\theta_t)), 1, \theta_t\} \tag{11}$$

where α is a parameter capturing the efficiency of the matching process: the higher α is, the lower matching frictions are.

6 Counterfactuals [in progress]

6.1 Decentralized equilibrium versus a constrained social planner

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 but is able to coordinate firms to set a contract length that optimally maximizes the total output of the market. In this market total output is defined as: [# alive matches] $\times v -$ [# contracts] $\times c$.

We present the results in Figure 4. Across these subplots, we show results for different values of the bargaining parameter δ as well as different values of the charterer entry rate e. The charterer entry rate e indexes whether the market is in a boom or a bust, with higher values corresponding to a boom where there are many searching charterers for the fixed number of ships. We chose the minimum of the charter entry rate on the graphs to align with the probability that a ship will be unemployed in a bust in the data. In this industry, a bust still features a high number of ships under contract (on the order of 90%). All the results in this figure are constructed at a steady state equilibrium, so we are abstracting away from transitions between booms and busts, amongst other things.[23]

Decentralized contract duration In the left panel of Figure 4 we plot the decentralized contract duration for three different bargaining parameters. We choose the value $\delta = 0.2$, where the charterers have more bargaining power, a value of $\delta = 0.8$ where ship-owners have more bargaining power, and an in-between case where $\delta = 0.5$. In all three cases the decentralized contract length increases in the boom. The lengthening of contracts is most pronounced for smaller values of δ (when charterers have more bargaining power). The intuition behind these results is as follows: during a boom, the charterer is less likely to be able to find another match in the near future, and so will prefer a longer contract over a shorter one. Provided that the charterer has some bargaining

[23]Note that we intend to incorporate transitions into the model — as well as agents' expectations over these transitions — in a future draft of the paper both in these counterfactuals and in the estimated model.

Figure 4: Comparing the decentralized equilibrium versus the constrained optimum

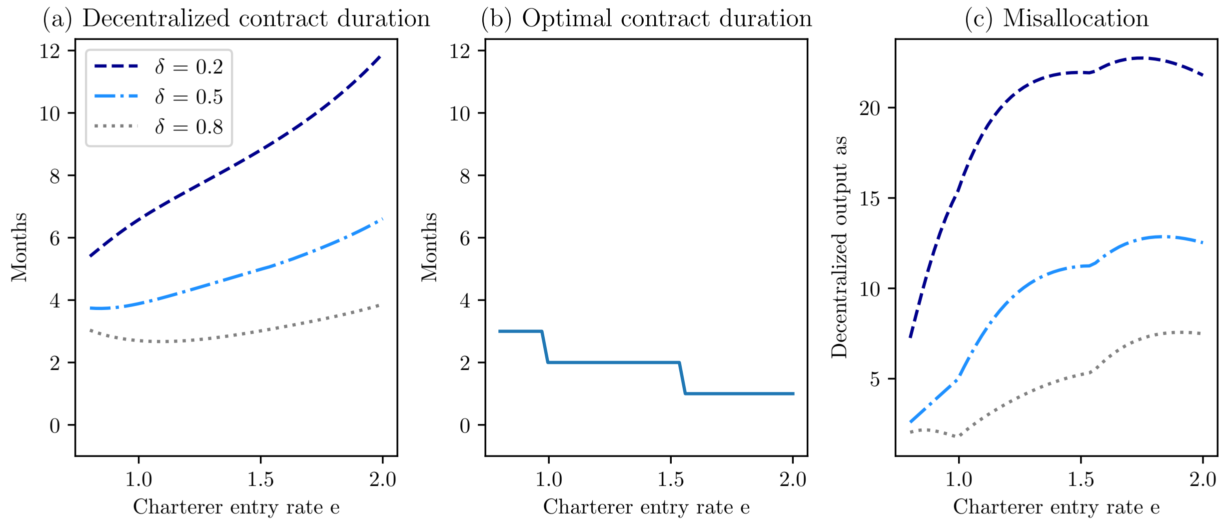

Note: To compute these graphs we assume that the market is in a steady state that could correspond to a 'boom' steady state, where there are many charterers who need a ship and e is high, or a 'bust' steady state, where there are fewer charterers. The optimal duration counterfactuals retains search frictions but solves for the contract length that maximizes total output of the market as a whole.

power, they will be able to induce the shipowner to agree to sign a longer contract by providing them a higher share of the surplus from matching, in exchange for giving up the option of re-matching in future periods.[24] Although it is not pictured nor supported by the empirical data, note that the model can also accommodate contract duration decreasing in booms by giving ship-owners all the bargaining power at $\delta = 1$. In other words, whether contract duration is pro-cyclical or counter-cyclical could go either way in our theoretical model and is an empirical question with the data pointing to a pro-cyclical relationship.

Optimal contract duration In the center panel of Figure 4 we plot the optimal contract duration from the point of view of a constrained social planner. That is, this is a social planner who still faces search frictions but can coordinate an industry-wide contract length. To construct this figure we iterate over all possible contract durations, compute the equilibrium total output for each selection,

[24]This is why the lengthening of contracts during booms is most pronounced if δ is high: the shipowner knows that if they refuse the contract, they are highly likely to find a match next period: but because they only have limited bargaining power, they will only get part of the surplus from any future short contracts, whereas they can lock in a bigger share of the surplus from signing a longer contract today.

and then choose the contract duration with the corresponding highest total output.

Here the optimal contract length is actually counter-cyclical and decreases in the boom. The intuition is as follows. In choosing the optimal contract length the planner faces a 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. 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 tradeoff is somewhat similar to the individual decision that agents will make when choosing the optimal contract length. The key difference is that the planner internalizes equilibrium effects since they are maximizing *total* output.

With the above trade-off in mind, consider the market in "bust" where the number of searching agents on both sides of the model is more balanced (i.e. lower values of e).[25] 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 and the optimal contract is around 3-4 months in the Figure.

Why does the optimal contract get shorter in the boom (as e increases)? As the market moves to a boom, it becomes 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.

Misallocation In the right panel of Figure 4 we plot misallocation, defined as the percentage change in total output moving from the decentralized solution to the optimal contract length. Con-

[25]We have solved this model for values of e lower than the minimum on the graphs, which corresponds to a more unbalanced market with so few charterers searching that the unemployment rate for ships falls far below its empirical minimum. In this case, consistent with the logic in this section, we see a symmetric contract length decrease to the middle panel.

sistent with the degree of the contracting externality increasing in booms, misallocation rises initially and contracts get shorter. Eventually, the optimal contract reaches its minimum length (1 period), where all ships are available to match in every period. At this bound, the percentage misallocation decreases between the two markets as e continues to increase.

6.2 Responsiveness to aggregate shocks

We next consider the model's predictions for how the market will transition after an aggregate shock in Figure 5. We focus on an immediate and long-term aggregate shock to demand that causes the market to switch from a boom to a bust. Here we decrease both the entry rate of charterers e as well as the probability that a match remains alive η. We compare three scenarios, moving from endogenous contract duration, to a fixed duration, to the optimal contract length, and evaluate the effects on productivity dispersion and misallocation.

Transitions with endogenous rigidities Our benchmark model has an endogenous choice of contract length and we plot the results in the dark blue line on Figure 5. In panel (a) we can see that this generates a 'contract overhang' similar to what we see in the actual market following the 2009 bust in Figure 2b. That is, as soon as the aggregate shock occurs, new contracts immediately fall, but the average duration contract remaining in the market (which both includes these new contracts as well as existing contracts that were signed before the shock) falls more slowly. In hindsight, from the perspective of the bust, these contracts are too long.

In panel (b) we can see the corresponding effects on productivity dispersion, which is computed as the standard deviation of the distribution of "alive" matches vs "dead" matches under contract. The adjustment of the market to this shock is slow-moving under endogenous rigidities, driven by the contract overhand pictured in panel (a). Productivity dispersion rises quickly due to lock-in of bad matches, caused by the persistence of long-duration contracts. This corresponds directly to the observed empirical patterns (Figure 2c), where dispersion in the marginal product of capital rises sharply after the 2009 crash. Over time, as these longer contracts expire and the ships can re-match on shorter contracts, productivity dispersion falls.

Figure 5: Effect of endogenous rigidities on transitions to aggregate shocks

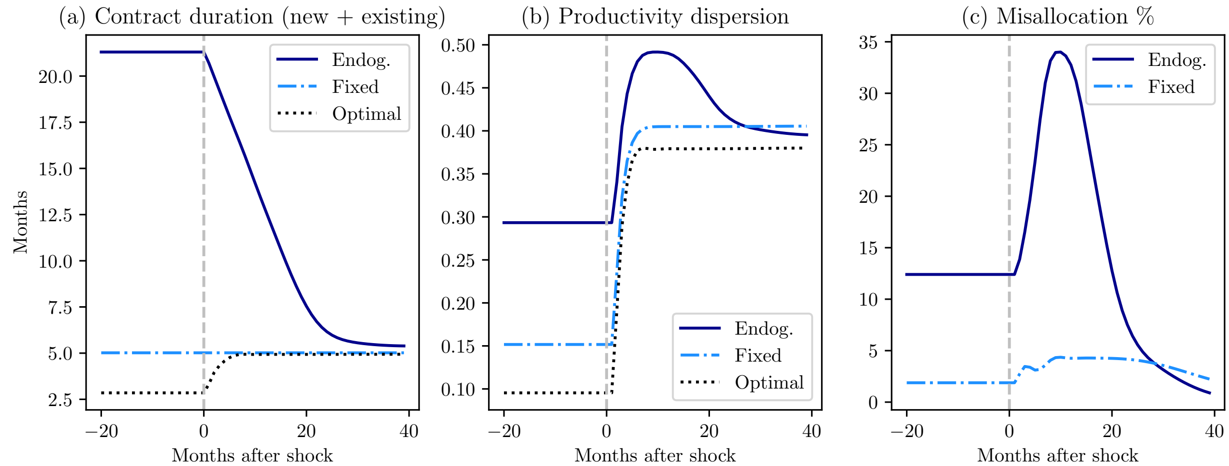

Note:

In panel (c) we plot the corresponding misallocation. We define misallocation as the percentage change in output compared to the scenario with the optimal contract length that maximizes market-level output. The transition period, under endogenous rigidities, is a major contributor to misallocation. Over time, the market self-corrects and reduces this misallocation by reducing the contract duration. In other words, although agents' may not be entirely aligned with the objectives of a social planner, they do sometimes face similar incentives — lock-in of a bad match is both socially and privately costly.

Transitions with fixed rigidities We next consider the case with fixed rigidities. For the fixed contract length we need to take a stand on what duration to choose here; we choose a contract length of 5 months since in this numerical example that is the optimal contract length in a bust, and close to the endogenous choice in a bust. Since the fixed contract length is closer to the socially optimum length both in booms — which is due to our choice of which 'fixed' rigidity contract to analyze — the productivity dispersion in booms is much lower in panel (b) than under endogenous rigidities. One element that is not affected by the specific choice of fixed length is that there is no 'contract overhang'. Therefore, the market does not exhibit a temporary spike in productivity dispersion after the shock; rather, this is a permanent change. In terms of misallocation, note that there is still some which occurs shortly after the shock; this is because the optimal contract length is not a fixed length. Rather, the optimal contract is countercyclical and shorter in the booms.

Transitions with the optimal contract The optimal contract duration is relatively short and countercyclical. Therefore, productivity dispersion is lower both before, during, and after, the aggregate shock.

6.3 Policy implications: evaluating shipping subsidies

Finally, we illustrate the implications of endogenous rigidities for evaluating shipping subsidies. The scale of these shipping subsidies is enormous, with China alone estimated to subsidise its domestic shipbuilding industry by $11 billion per year between 2006-2013 (Jia Barwick et al., 2021); many countries also subsidise shipowners and shipping companies (International Transport Forum, 2019). Therefore, an important question is the extent to which this spending translates into market benefits once the ships are built. If rigidities are fixed then an extra ship may in principle not generate much change in total output once it is added to the market. This is due to cannibalization effects: some of the observed cargo that this new ship is carrying may just be reallocated from existing ships. However, when rigidities are endogenous, there is an additional effect that may improve the efficacy of subsidies. Here, an additional ship thickens the market, causing agents to endogenously respond in equilibrium by signing shorter contracts. This then creates an additional positive contracting externality, and to the extent that booms are associated with particularly thin markets for ships, this effect may be the largest in the booms.

We show a numerical example of these effects in Figure 6. Here we consider the benefits of increasing the number of ships by 10% at different points in the business cycle (indexed by the charterer entry rate e). Due to time to build, which is a key feature of the shipbuilding industry (e.g. (Kalouptsidi, 2014)), an immediate 10% increase in the number of ships is infeasible. Instead, the interpretation of this counterfactual is that ships are long-lived assets in a market shaped by frequent cycles. As a result, the lifetime value of a ship is determined in part by these booms and busts whenever it arrives in the cycle. Furthermore, these results also connect directly with subsidies that prevent exit of ships, since these kinds of policies are not subject to time-to-build.

Figure 6: Effects of subsidies that increase the number of ships in the market

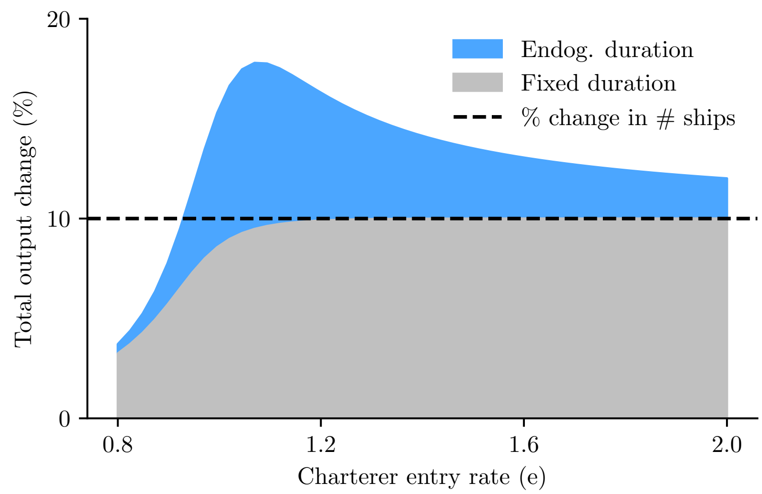

Note: In this figure we increase the number of ships by 10% at different values of e, which index whether the market is in a bust or a boom, with higher rates corresponding to a boom. The gray area corresponds to the increase in total output in the market at equilibrium contract duration before the 10% increase in ships. The blue area corresponds to the extra increase in the total output of the market allowing for contracts to vary endogenously with market thinness.

Finally, for simplicity, in this particular numerical example we are setting a bargaining parameter $\delta = 0$, the baseline number of ships as $n = 10$, and the probability that a match does not expire $\eta = 0.9$. We emphasize that these numbers are chosen for illustrative purposes only and may change when we estimate the model.

In Figure 6, a low charterer entry rate corresponds to a bust. Here, we see that the increase in total output from increasing the number of ships by 10% is less than 5%. The reason is due to cannibalization - not every match for this new ship is an extra match for the market; some of these matches are just reallocated from existing ships. These qualitative effects are similar in a bust both keeping the contract length fixed at the level before the 10% increase, and allowing the contract length to change endogenously after the extra ships are added.

As the market moves towards a boom and e increases, the corresponding contract length increases and may move further from the efficient benchmark, as previously discussed in Figure 4. Therefore, the positive externalities from extra ships that arise due to thickening the market may also increase. In this numerical simulation, for some values of e, these externalities are relatively large. In the most extreme case, we observe increases in total output that are *greater* than 10%, which is the increase in total output from adding an extra ship absent any cannibalization effects. For example, at $e = 1.05$, adding a single extra ship generates an increase in total output equivalent to the average output of around 1.75 ships when contract length is endogenous.

References

Adland, Roar, 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.

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.

Barwick, Panle Jia, Myrto Kalouptsidi, and Nahim Bin Zahur, "Industrial Policy Implementation: Empirical Evidence from Chinaâs Shipbuilding Industry," *Working Paper*, 2021.

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

_ , _ , _ , 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.

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," *Work-*

ing 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, "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.

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.

Hosios, Arthur J, "On the efficiency of matching and related models of search and unemployment," *The Review of Economic Studies*, 1990, *57* (2), 279–298.

Hsieh, Chang-Tai and Peter J Klenow, "Misallocation and manufacturing TFP in China and India," *The Quarterly journal of economics*, 2009, *124* (4), 1403–1448.

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

International Transport Forum, *Maritime Subsidies: Do They Provide Value for Money?*, OECD Publishing, 2019.

Jeon, Jihye, "Learning and Investment under Demand Uncertainty in Container Shipping," 2018. Department of Economics, The Boston University.

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.

Kehrig, Matthias, "The Cyclical Nature of the Productivity Distribution," 2015. Working Paper.

Lanteri, Andrea and Adriano A. Rampini, "Constrained-Efficient Capital Reallocation," *Ameri-*

can Economic Review, February 2023, *113* (2), 354–395.

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

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.

Restuccia, Diego and Richard Rogerson, "Policy distortions and aggregate productivity with heterogeneous establishments," *Review of Economic dynamics*, 2008, *11* (4), 707–720.

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

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

A Appendix

A.1 Data construction

The 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 container ships 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 new 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, Clarkson’s Shipping Intelligence Network provides aggregate indexes such as containership time-charter rate index, China Containerized Freight Index (CCFI), Singapore bunker prices ($/Tonne), and the volume of containership contracting (TEU).[26] Table 1 provides descriptive statistics on the duration of the contracts, rates, age of the ships contracted, and aggregate indexes.

[26]The containership time-charter rate index and China Containerized Freight Index (CCFI) are two separate (albeit highly correlated) measures of container freight rates. Bunker prices are useful as they are indicative of the fuel costs incurred by container ships.

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 design 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:

$$\text{Utilization} = (H_A - H_B)/(H_S - H_B) \tag{12}$$

where H_A is the draught reported in the port call data. Note that later in the analysis we always control for ship heterogeneity when using this measure, in order to ensure that the measure is comparable across vessels. Finally, 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 of utilization by using the port call data to identify idle ships.

A.3 Measuring reallocations

Depending on the particular itinerary a container ship is operating on, they may travel very long distances and stop at a large number of ports while remaining within a fixed schedule.[27]

We therefore develop an algorithm for identifying when a ship is reallocated. The idea behind the algorithm is that if a container 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 container ship may visit a new port that involves a minimal deviation from its existing itinerary.[28] These are unlikely to be true reallocations of the ship, and instead may simply represent extra voyages the operator of the ship 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

[27] For example, one container ship 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.

[28] 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.

to work well in practice; the episodes it identifies as reallocations match well with what appear in the data to be true reallocations.[29]

A.4 Measuring blank sailings

Blank sailings, instances where a ship omits a scheduled port call, are critical in understanding maritime shipping patterns and efficiencies. We develop an algorithm to identify these blank sailings, providing insights into operational practices.

We begin with defining a ship's routine, identified as a sequence of regularly visited ports. For instance, a ship might routinely travel from port A to B, and then from B to C. We classify a sequence as a routine if it has been consistently repeated at least six times in both the past and future, establishing a clear pattern of regularity. In this process, we also introduce a variable to capture deviations from the routine, such as when a ship travels directly from A to C, skipping the usual stop at B.

After establishing the routines, the algorithm flags a blank sailing when it detects a skipped port. To ensure that this skip is not a mere anomaly but part of a broader pattern, the algorithm scrutinizes the ship's itinerary. It checks whether, in the three months before or after the skipped port, the ship maintained a routine corresponding to this route. Further, to distinguish blank sailings from regular variations in the route, the algorithm limits these occurrences to a maximum of two times in either the past or future. This restriction ensures that identified blank sailings are not simply routine alterations but deviations from the schedule warranting attention.

B Additional proofs and results

Details on match payoff in Equation 4 Here we provide more detail that the charterer's expected match payoff is $(1-\delta)W_{x,t} + \beta\eta\mathbb{E}_t U^{charterer}_{x,t+1}$. To see why, note that the charterer's payoff to a

[29] 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. In our future draft, we plan to carry out more extensive robustness analysis related to reallocation measures.

τ-duration contract, once this contract duration has been chosen and under Nash bargaining, is: $\Pi_t^{charterer} = (1-\delta)(S_{x,t,\tau} + \varepsilon_\tau) + \beta\eta\mathbb{E}_t U_{x,t+1}^{charterer}$. Denote $z_{x,t,\tau} = S_{x,t,\tau} + \varepsilon_\tau$. Then, the charterer's *expected* match payoff (before matching has taken place and the ε_τ have been drawn) is:

$$\mathbb{E}_{z_{x,t,\tau}}[(1-\delta)z_{x,t,\tau} + \beta\eta\mathbb{E}_t U_{x,t+1}^{charterer}] = (1-\delta)\mathbb{E}_{z_{x,t,\tau}} z_{x,t,\tau} + \beta\eta\mathbb{E}_t U_{x,t+1}^{charterer} \tag{13}$$

$$= (1-\delta)W_{x,t} + \beta\eta\mathbb{E}_t U_{x,t+1}^{charterer} \tag{14}$$

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

C Additional figures

Figure 7: Average contract duration and time-charter index (in logs), 1999 - 2022

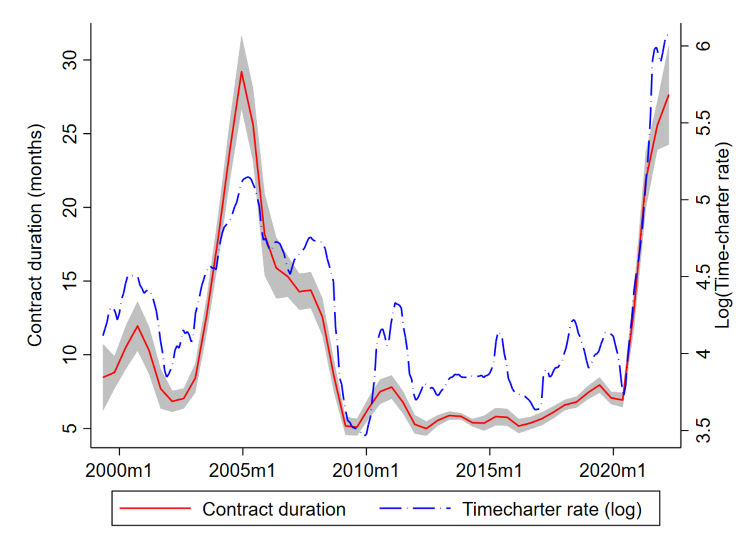

Figure 8: Average contract duration for bulk time-charters and Baltic Dry Index, 2001 - 2016

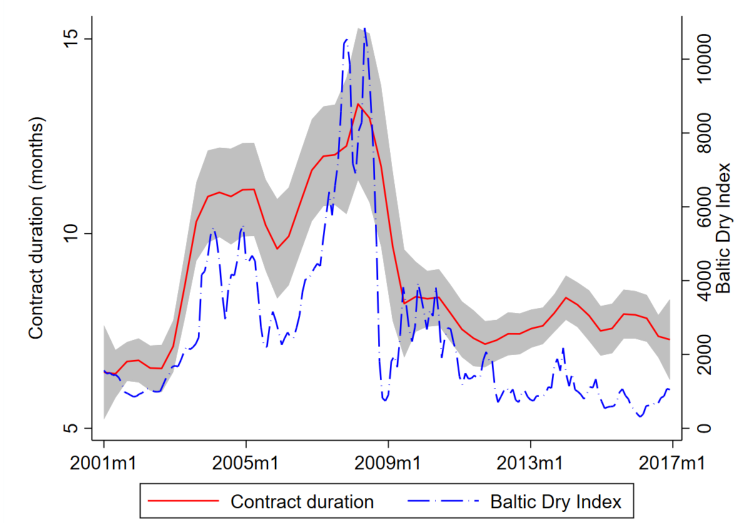

Figure 9: Reallocation (residualized) against the time-charter rate

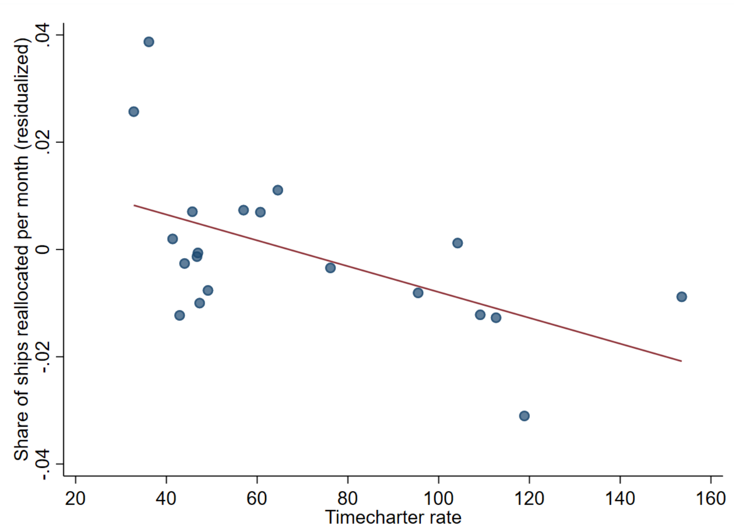

Note: We regress our reallocation measure on ship fixed effects and an aggregate time trend, and obtained the residualized probability of reallocation for each ship, which we then aggregate across ships for each year-month. The figure shows a bin-scatter plot of the residualized share of container ships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

Figure 10: Probability of blank sailings over the business cycle, residualized

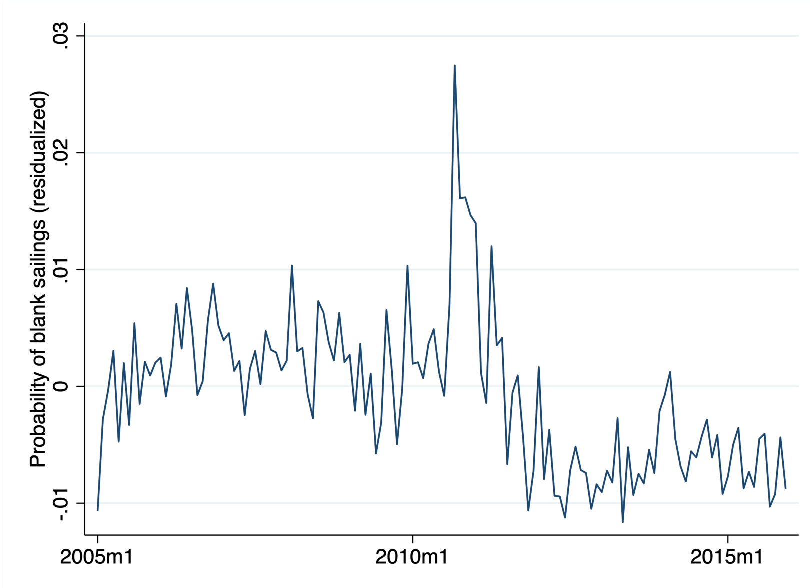

Note: We plot the residualized probability of blank sailing at any given year-month, with the residuals obtained from a regression of each port call's indicator for blank sailing on ship's age, TEU and the area of operation.

Figure 11: Probability of blank sailings (residualized) against dispersion in residualized capacity utilization (a proxy for productivity dispersion)

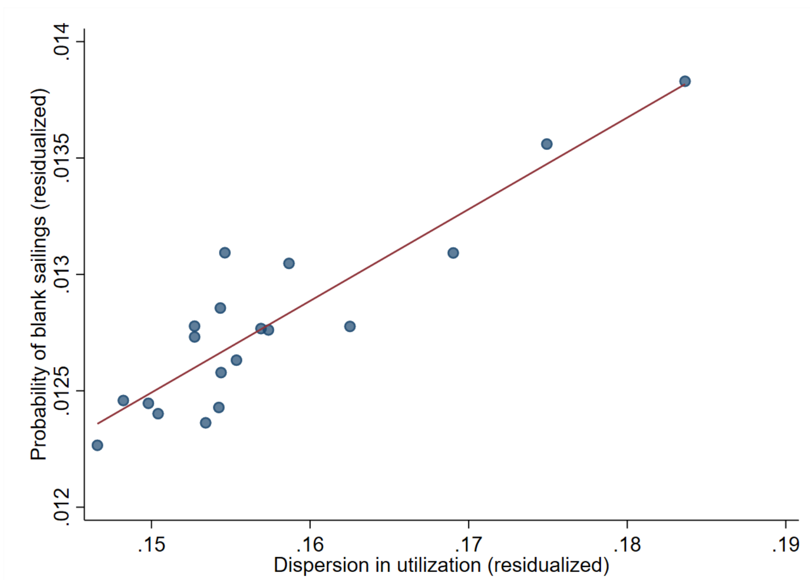

Note: This shows a bin-scatter plot of the residualized probability of blank sailing against standard deviation of residualized utilization across ships every month. The residualized probability of blank sailing is obtained from the residuals of a regression of each port call's indicator for blank sailing on ship's age, TEU and the area of operation. Residualized utilization is obtained by regressing utilization on ship fixed effects and a time trend.

Figure 12: Difference in dispersion in residualized capacity utilization (a proxy for productivity dispersion), between ships under longer and shorter contracts

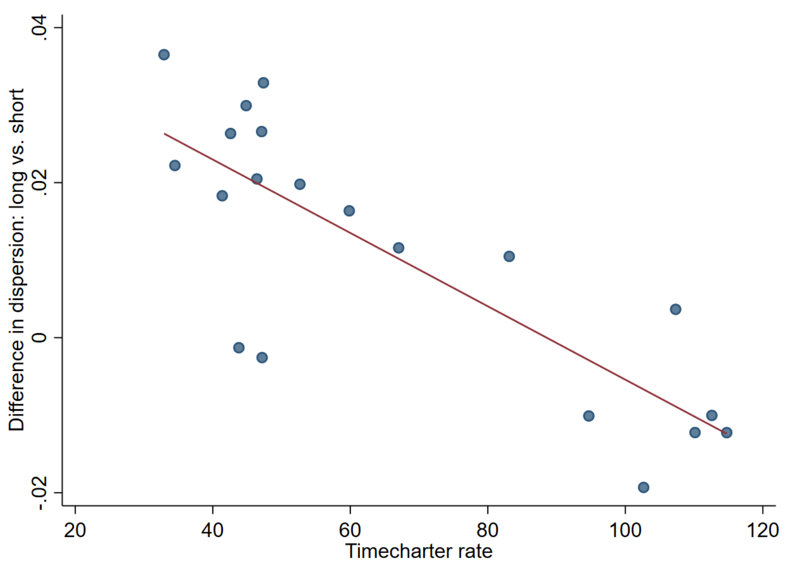

Note: This shows a bin-scatter plot of the difference in dispersion in residualized capacity utilization, between ships under longer and shorter contracts. Residualized utilization is obtained by regressing utilization on ship fixed effects and a time trend. We separately calculate the standard deviation in residualized capacity utilization across ships under longer contracts (where the contract was signed at least 12 months prior) and those under shorter contracts (where the contract was signed within the last 12 months). We then calculate the difference between the two: a positive value means there is greater dispersion in capacity utilization across ships under longer contracts than ships under shorter contracts.