Booms, Busts, and Endogeneous Rigidities: Evidence from Containerships

Link to most recent version here.

Maria Osipenko Nicholas Vreugdenhil Nahim Bin Zahur*

September 2023

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 counter-cyclical 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. 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). In contrast, more recent papers emphasize the role of dynamic inputs which are costly to adjust in inhibiting reallocation (e.g. Asker et al., 2014): what appears to be misallocation from a static perspective may actually be efficient after accounting for adjustment costs.

In this paper we argue that these capital adjustment rigidities may be endogeneous and that this can lead to inefficiency and misallocation. We focus on decentralized leasing markets for physical capital assets (like ships, aircraft, and oil and gas drilling rigs). 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.

Our empirical setting for this paper is the global market for leased containerships.[1] 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 exact location and amount of cargo that each ship is car-

[1] We also have similar results about pro-cyclical contract duration in other markets where contract data are available, including bulk shipping and shallow-water offshore oil and gas rigs. We will include this evidence in a future draft.

rying. 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. In this industry, inefficiently long contracts and a failure to adjust to shocks may be a crucial component of supply chain fragility.

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 contractual rigidities inhibit reallocation, defined as ships moving to alternative routes across the globe. Next, we show that dispersion in fleet efficiency (defined by dispersion in ship utilization after controlling for ship heterogeneity) is counter-cyclical, which suggests more reallocation may be beneficial during economic downturns. Consistent with this, we observe an increase in spatial reallocation following the bust, though such reallocation is limited by contractual rigidities. Finally, we 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 matches and larger ships.

Based on these facts, we develop an illustrative model of the market. The model shows that in the presence of search frictions and market thinness, agents sign longer contracts during booms, further amplifying market thinness. The increase in contract duration is socially inefficient as it leads to lock-in, consistent with the stylized facts observed in the data (where ships are rarely reallocated during the latter stages of medium-term contracts). In a future draft we will estimate a version of this model using the data in order to quantify the degree to which endogenous contract

duration affects inefficiency and also to consider how commonly-proposed policies (like counter-cyclical subsidies) might improve efficiency.

2 Industry and Data

2.1 Container Shipping Industry

Our analysis focuses on the global container shipping industry. Container ships engage in the transportation of containerized goods. The majority of world trade in goods is carried out by container ships. Container ships generally operate on fixed schedules where they pass through a designated set of ports, typically carrying containerized cargo in each one of these trips. In this respect, container ships can be thought of as similar to buses. 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).

Container ships are operated by liner companies that specialize in the transport of containerized goods across the world (for example, Maersk, CMA CGM or 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. These “charter-owned" ships are leased out to liner companies using time-charter contracts that range from 12 - 36 months in length, and specify a fixed day-rate that the liner company (or the “charterer") pays the shipowner. The focus of our empirical analysis is on these leasing contracts, and the dataset we have compiled primarily focuses on charter-owned ships that are regularly leased, though we have also collected data for a limited set of owner-operated ships.

2.2 Data

We use several data sources for the empirical analysis. Clarkson's Shipping Intelligence Network provide the data on latest containership timecharter fixtures (contract records) each month starting from May 1999. 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).[2] Table 1 provides descriptive statistics on the duration of the contracts, rates, age of the ships contracted, and aggregate indexes. Our full dataset of contracts includes over 16,000 timecharter 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 3,500 contracts.

[2]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.

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 detailed information on worldwide containerships movements and every port call made by each container ship, including the dates of arrival and sailing, and the exact locations of the ports visited on each port 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).

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).

Table 1: Summary statistics for leasing contracts

Variable	Obs	Mean	Std. dev.	Min	Max
Duration (months)	16,177	10.83	13.64	0.17	216.0
Rate ($/day)	15,348	11,725.7	10,889.0	0	235,000.0
Time-charter rate index	16,180	79.8	68.2	32	433.0
Ship Age (years)	16,063	10.65	5.66	0	36
Capacity (nominal TEU)	16,180	2,362.9	1,823.8	170	14,952.0
Bunker price ($/ton)	16,180	362.0	135.0	85	733.3
Volume of Contracting (TEU)	16,180	123,794.3	173,250.4	0	1,109,916.0
CCFI (Freight Index)	14,948	1,093.2	547.8	641.5	3,510.8

2.3 Measuring fleet efficiency and reallocation

We use the port call data both to measure fleet efficiency, as well identify if a ship is reallocated. Both fleet efficiency and reallocation are critical for understanding how container ships are utilized and how contracting affects.

Fleet efficiency To capture fleet efficiency, 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. A natural measure of fleet efficiency is 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 utilized efficiently. While we do not directly observe this in the data, we can infer fleet efficiency 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. Since we don't observe design draught, we construct it by choosing the observed maximum draught for the specific ship in the data. We compute fleet efficiency, defined as the percentage of the ship's capacity that is being utilized on a specific voyage using the following formula:

$$\text{Fleet Efficiency} = (H_A - H_B)/(H_S - H_B) \tag{1}$$

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.

Reallocation Next, 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 sequences of port calls. 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.[3]

[3]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.

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.[4] 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 to work well in practice; the episodes it identifies as reallocations match well with what appear in the data to be true reallocations.[5]

[4] 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.

[5] 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.

3 Stylized Facts

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

3.1 Contract duration over the business cycle

We begin by providing descriptive evidence on how contracting behavior changes with the business cycle, and how in turn contracting affects re-allocations of ships over space. We highlight two key stylized facts. First, firms sign substantially longer contracts during booms. Thus, in the aftermath of the boom, a greater share of ships are tied up under long-term contracts, and this overhang of contracts persists even after the boom evaporates. Second, ships are significantly more likely to reallocate across space at the time when they sign a new contract, but there is little evidence of reallocation happening afterwards while the ship is reallocated. Taken together, these two empirical facts suggest that the lengthening of contracts caused by boom-bust cycles may inhibit subsequent reallocation of container ships which are locked into longer contracts.

Contract duration and the business cycle Figure 1 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).[6] 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.

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 2,

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

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.

Figure 1: Average contract duration and time-charter index, 1999 - 2019

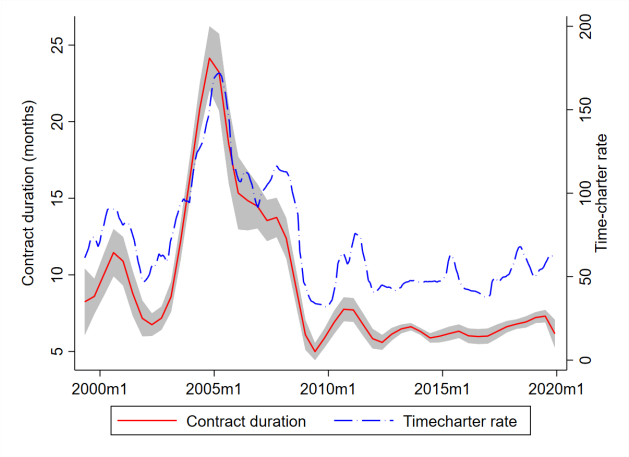
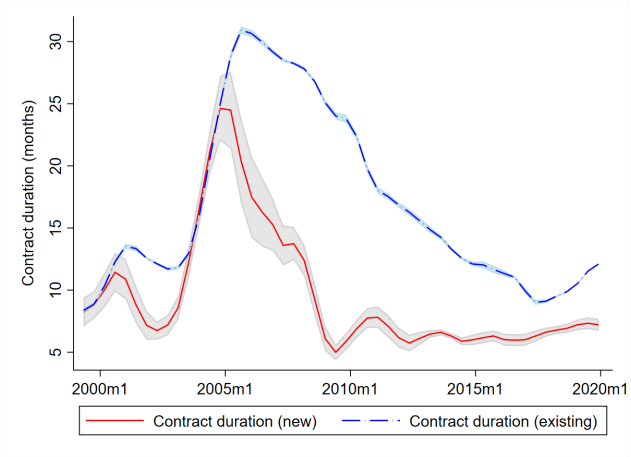


Figure 2: Average duration: new vs. existing contracts, 1999 - 2019



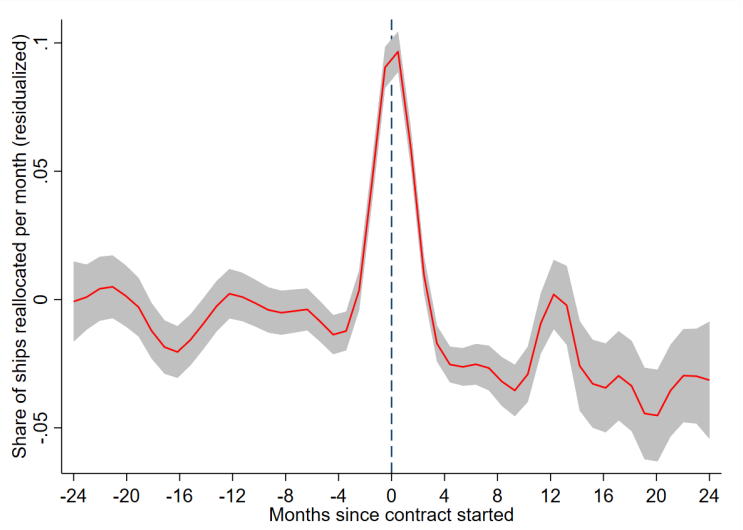
Contracting and reallocation The overhang of contracts, caused by the lengthening of contracts during booms, may have negative consequences for efficiency, in the presence of contractual rigidities. In the containership leasing market, contracts typically specify a fixed per-day rate received by the shipowner. Thus shipowners have very 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 probability 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 3 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.

Figure 3: Contract rigidities affect reallocation: share of ships reallocated/month (residualized)



Note: 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.

3.2 Reallocation and dispersion in fleet efficiency

A key feature of the container-ship industry that we document is that during busts, there is evidence of increased dispersion in productivity; thus, although demand for shipping services is lower during busts, the efficiency gains from reallocating are potentially larger. Consistent with this, we find that reallocations of ships across space are strongly counter-cylical, with container ships more likely to move to a different trading route during busts. Reallocation, however, is costly, as evidenced by the fact that the efficiency with which ships operate is disrupted every time they are reallocated. In the remainder of this section, we present descriptive evidence illustrating each of these points.

Dispersion in fleet efficiency We measure fleet efficiency by the share of the ship's capacity that is used to carry cargo. We consider this a reasonable though imperfect proxy for how efficiently the ship is being utilized. Fleet efficiency is naturally influenced by changes in shipping demand over time; however, cross-sectional differences in fleet efficiency across different ships are unlikely to reflect aggregate fluctuations in demand, and instead are more likely to capture underlying differences in efficiency across ships.

As we might expect, fleet efficiency is pro-cylical (Appendix Figure 11): while fleet efficiency is on average around 60% prior to the financial crisis, it drops to below 55% in the aftermath of the 2008 - 09 crisis.[7] The decline in demand for shipping during this time period (as shown by the drop in the timecharter rate index) is a natural explanation for why container ships were carrying less cargo during this period.

What is more striking, though, is that during economic downturns, there is greater dispersion in fleet efficiency across ships. Indeed, as Figure 4 illustrates, dispersion in fleet efficiency appears to be significantly counter-cyclical: the coefficient of variation in fleet efficiency (measuring variability across ships in their fleet efficiency) rises significantly after the financial crisis. Container ships operate on varying routes and different parts of the world; the counter-cyclical pattern of fleet efficiency suggests that during economic downturns, there is greater dispersion in the profitability of operating on different routes. Thus paradoxically, the potential gains from reallocating are higher during an economic downturn, even as average demand (and therefore average profits) is lower.

Reallocation of ships Next, we examine how ship reallocation evolves over time and in response to demand fluctuations. Figure 5 shows how the share of container ships reallocated each month changes over time. We find that capital reallocation follows a striking 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. We show in Appendix Figure 9 that this finding is robust to the inclusion of ship fixed effects and an aggregate time trend. Appendix Figure 10 further corroborates this finding by illustrating that there is a negative correlation between the probability of reallocation and the time-charter rate index (our

[7] This pattern is robust to the inclusion of ship fixed effects and a time trend (Appendix Figure 12).

Figure 4: Dispersion in fleet efficiency over the business cycle

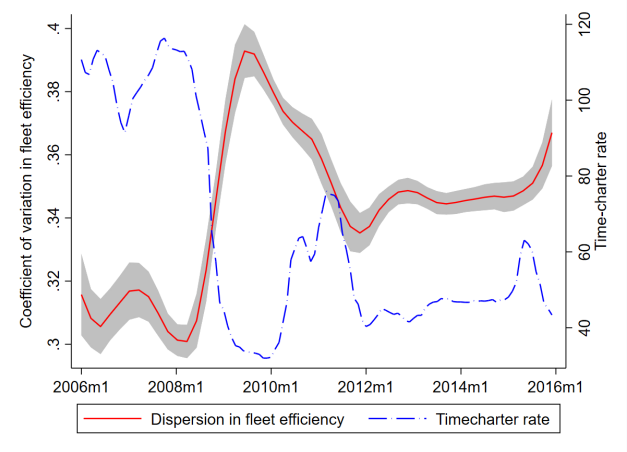
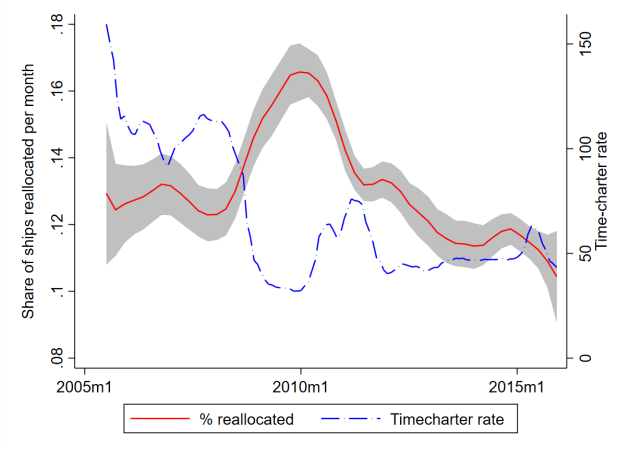


Figure 5: Reallocation over the business cycle



proxy for aggregate demand in this market).

These empirical patterns are consistent with the fact that during booms, there is reduced dispersion in profitability across different regions. Thus, the incentive for a container ship to reallocate from one geographic route to another is lessened during booms. By contrast, during a bust, there is an increase in the dispersion in profitability between different trade routes, and thus more opportunities for profitable reallocation. Our descriptive analysis suggests that container ships do, on the whole, respond to these changing incentives by reallocating more often during economic downturns.

Cost of reallocation Reallocating ships across space, however, is costly. Appendix Figure 14 shows how fleet efficiency changes upon reallocation, controlling for ship fixed effects, year fixed effects and time-varying aggregate determinants of demand and costs (i.e., the time-charter rate and bunker rate). As we can see, fleet efficiency declines significantly during the two months just prior to reallocation. Part of this decline may be due to the fact that when a container ship is being transferred from one part of the world to another, it may have to make that trip carrying less cargo than usual. But we also find that fleet efficiency is low even *after* the ship has already been reallocated to its new trade route, and stabilizes only around two months after the reallocation takes place. This suggests that reallocation can potentially lead to a short-run disruption in how efficiently the ship is being utilized: integrating the ship fully into existing liner schedules may

take time even after the ship has already been physically moved to a new region.

In addition to these (potentially unavoidable) costs of reallocation, contracts may also endogenously increase the cost of reallocating. As illustrated before, an overhang of longer contracts from an earlier boom may inhibit reallocation of ships during a subsequent bust.

3.3 Determinants of long and short contracts

How do firms choose their contract length in the micro-data? We present descriptive evidence for several key considerations: longer contracts are associated with thinner markets, longer contracts are associated with high match quality, and longer contracts are associated with larger ships (ship heterogeneity). 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 2. 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 contractor 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.[8]

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

[8]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.

Table 2: Longer contracts are associated with thinner markets and larger ships

	(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

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

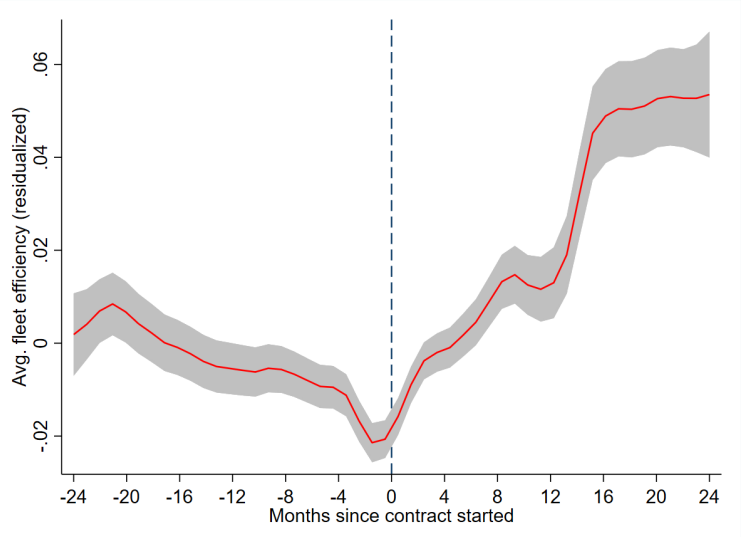
than trading hubs (for example, ships located in a radius around Singapore). We also consider two versions of the market thinness variable, one where we count all ships nearby in time and space, and one where we subset to only ships of the same type as the ship that was eventually under contract, reflecting that charterers may require a specific ship type and so direct their search.

Overall, the results across the four specification in Table 2 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.

Heterogeneity Table 2 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.[9]

Match quality Figure 6 shows that match quality tends to improve with time on average, which would favor signing longer contracts if aggregate conditions remained constant. There is a dip in efficiency as the contract starts, but over time the efficiency of the ship ramps up.[10] In the containership context, a primary reason for this improvement is due to better integration over time into the charterer's schedule. This is reminiscent of the labor literature where accumulation of match-specific experience improves the productivity of the match over time.

Figure 6: Evolution of efficiency over the course of contract



Note: The figure plots the residualized fleet efficiency for ships, for each month in a 4 year window around the month when a new contract is signed. The residualized fleet efficiency is obtained for each ship-month observation by regressing fleet efficiency 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.

[9]These results, though not included in the present draft, are available upon request.

[10]The initial drip in efficiency can be attributed to the fact that ships are more likely to be reallocated at the start of the contract, temporarily disrupting their efficiency. For ships that do not reallocate at the onset of the contract, no initial decline in efficiency is observed, as we can see from Appendix Figure 8.

4 Illustrative Model [in progress]

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

Setup

The firms in the industry consist of “charterers", which are shipping companies that lease container ships and use them to transport container goods across global shipping markets; and “shipowners" (which for simplicity we refer to as ships) who own ships and lease them out to charterers. There are C charterers, denoted by $j = 1,...,J$.. There are S ships, $i = 1,....,i$, each owned by a distinct shipowner.

Time is denoted by t, ranging from $t = 1,....,T$. For simplicity, we assume that $T = 2$, so that there are only two time periods. When charterers negotiate leasing contracts with ships, therefore, they choose between one-period and two-period contracts.

We make the following simplifying assumptions:

- We assume away the possibility that either charterers, or shipowners, exercise market power.[11]
- We also assume all ships are identical. (In practice, ships differ by capacity, age etc., and may differ in operating costs due to differences in fuel efficiency).

Costs and demand Operating costs equal c for each ship i, and are borne by the shipowner. The value of a ship to the charterer is v_{jt} for each charterer j in period t. We assume for simplicity that $v_{j1} = v > c$: there is no uncertainty or heterogeneity in charterer values for the ship the first

[11] During the main period of our study, from 2005 - 2015, global container ship leasing markets were fairly unconcentrated. On the shipowner side, the largest 40% of owners only owned about 50% of all ships (and 75% of all ship capacity). On the charterer side, the Herfindahl-Hirschman Index (HHI) is only around 500. Market power has emerged as a significant issue in more recent years with the consolidation of alliances in the global liner market.

period.[12]

However, we allow for the possibility that over time, the value of a ship to the charterer may fluctuate. Thus, we allow for uncertainty in v_{j2}, and assume that it follows an exponential distribution. There are two types of charterers (who exist in equal probability), type A and B: type A have higher valuations on average.There are two types of industry states, booms H and busts L.

The distribution of v_{j2} is then given by:

$$v_{j2} \sim Exp(1/\bar{v}_{AH}) \text{ High type A, boom}$$
$$v_{j2} \sim Exp(1/\bar{v}_{AL}) \text{ High type A, bust}$$
$$v_{j2} \sim Exp(1/\bar{v}_{BH}) \text{ Low type B, boom}$$
$$v_{j2} \sim Exp(1/\bar{v}_{BL}) \text{ Low type B, bust}$$

The expected valuation of a high type A exceeds that of a low type B: $\bar{v}_{AH} > \bar{v}_{BH}, \bar{v}_{AL} > \bar{v}_{BL}$. The expected valuation of any type during a boom exceeds their expected valuation during a bust: $\bar{v}_{AH} > \bar{v}_{AL}, \bar{v}_{BH} > \bar{v}_{BL}$.

We can summarize all of these assumptions by using $\bar{v}_j$ to capture the expected valuation of any given charterer in period 2 (which will take one of the above 4 values). Let $\bar{v}$ be the average of the valuation of the high types (A) and low types (B).

Timing The timing of the game, in each of the two periods, is:

1. **Decision to search:** Each charterer without a prior contract observes their v_{jt} and decides whether to search for ships. Each ship without a prior contract searches for charterers.

2. **Matching:** The number of searching charterers is C_t. The number of searching ships is S_t. The number of successful matches is $n(C_t, S_t)$, where $n(C_t, S_t) \leq min(C_t, S_t)$.

   The probability that a charterer finds a ship equals $r_t^c = n(C_t, S_t)/C_t$. The probability that a ship finds a charterer equals $r_t^s = n(C_t, S_t)/S_t$.

[12] Allowing for heterogeneity in v_{j1} does not fundamentally change the results.

3. **Contracting:** Each matched charterer and ship negotiate (i) contract duration (1 or 2 periods) (ii) a contract price p.[13] The price is determined by Nash bargaining with the ship's bargaining power equal to δ, and disagreement results in neither party having anyone else to trade with in that period.[14] Let L_t denote the number of pairs who decide to sign two period contracts.

Solving the model

We solve the model backwards, starting in period 2.

Period 2: L_1 charterers/ships have already signed 2-period contracts in period 1. There remain $S_2 = S - L_1$ ships and $C_2 = C - L_1$ charterers.

Decision to search: A charterer will only search if their idiosyncratic value, v_{jt}, exceeds c. The number of charterers who search is then given by $\tilde{C_2} = C_2 Pr(v_{j2} > c)$. The share of charterers who search, $\tilde{C_2}/C_2$, is going to be larger in a boom, and smaller in a bust. Let $\mu_{2j} = Pr(v_{j2} > c)$ be the probability that any given charterer searches.

Matching: The number of matches is $n(\tilde{C_2}, S_2)$.

Contracting: Each matched charterer and ship sign a 1-period contract, since there are positive gains from trade.

The total value of contracting is $v_{jt} - c$ (with $t = 2$). Under Nash bargaining, the parties will choose p_{ij} such that:

$$(v_{jt} - p_{ij})^{1-\delta}(p_{ij} - c)^{\delta}$$

which yields:

$$p_{ij} = \delta v_{ij} + (1-\delta)c$$

[13] In period 2, they can only negotiate one-period contracts.
[14] In period 1, of course, they can wait till period 2 and hope to be matched then.

Notice that if $\delta = 1$ (ship has all bargaining power), $p_{ij} = v_{ij}$; whereas if $\delta = 0$ (charterer has all bargaining power), $p_{ij} = c$.

Ship's payoff equals: $p_{ij} - c = \delta(v_{ij} - c)$.

Charterer's payoff equals: $v_{ij} - p_{ij} = (1-\delta)(v_{ij} - c)$.

Expected payoff at start of period 2, before observing v_{jt}:

The expected value to charterer type j conditional on deciding to search equals $E(v_{j2}|v_{j2} > c) = \bar{v}_j + c$. The average of this expected value across all charterer types is $\bar{v} + c$.Thus, charterer j's expected payoff, conditional on deciding to search and being matched, equals $E_j((1-\delta)(v_{ij} - c)) = (1-\delta)(\bar{v}_j + c - c) = (1-\delta)\bar{v}_j$.

The ship does not know which type of charterer they will be matched to. So the ship's expected payoff, conditional on being matched, equals: $E(\delta(v_{ij} - c)) = \delta(\bar{v} + c - c) = \delta\bar{v}$.

Ship's expected payoff = Prob. of matching × Expected payoff conditional on being matched

$$\pi_2^s = r_2^s \times \delta\bar{v}$$

Charterer j's expected payoff = Prob. of searching × Prob. of match × Expected payoff cond. on match

$$\pi_{2j}^c = \mu_{2j} \times r_2^c \times (1-\delta)\bar{v}_j$$

Period 1: We start with $S_1 = S$ ships and $C_1 = C$ charterers (half of which are high type and low type).

Decision to search: Since there is no period 1 heterogeneity in values, and as $v > c$, all charterers will search and all ships will search.

Matching: The number of matches is $n(C_1, S_1) = n(C, S)$.

Contracting: The key decision parties face is whether to sign 1 period contracts or 2 period

contracts.

The joint value of a contract is:

$$\begin{aligned} V(\text{1-period contract}) &= (v-c)+\pi_2^s+\pi_{2j}^c \\ &= (v-c)+r_2^s\delta\bar{v}+\mu_{2j}r_2^c(1-\delta)\bar{v}_j \\ V(\text{2-period contract}) &= (v-c)+(E(v_{j2})-c) \\ &= (v-c)+(\bar{v}_j-c) \end{aligned}$$

The parties will sign 2-period contracts if:

$$\bar{v}_j - c > r_2^s\delta\bar{v}+\mu_{2j}r_2^c(1-\delta)\bar{v}_j \tag{2}$$

Types of equilibria

Equilibrium contracting in this model (as summarized by equation 2) will depend on (i) whether the industry is in a boom or bust (ii) difference between high and low types (iii) matching technology (and any frictions) (iv) market thinness (v) bargaining frictions. We assume for simplicity that no matter what, there are always more charterers looking for ships, than ships looking to be chartered.[15] Then we consider a number of different cases, starting with the simplest one:

Case 1. No matching frictions, no heterogeneity across firms If there is no heterogeneity across charterers in their expected valuation, $\bar{v}_j = \bar{v}$. This further implies that $\mu_{2j} = \mu_2$ (every charterer has the same probability of searching for ships).

If there are no matching frictions, and given that there are fewer ships than charterers, the number of matches will equal simply the number of ships: that is, $n(C_t, S_t) = S_t$. In the absence of matching frictions, all ships will get matched.

[15]This requires that $C_t < S_t$. For this to be true, we need both $C > S$, and that $\tilde{C}_2 > S_2$ for any L_1.

- The probability that a ship is matched in period 2, r_2^s, equals 1.

- The probability that a charterer is matched in period 2, r_2^c, equals:

$$\begin{aligned} r_2^c &= \frac{S_2}{\tilde{C}_2} \\ &= \frac{S - L_1}{Pr(v_{jt} > c)(C - L_1)} \\ &= \frac{S - L_1}{\mu_2(C - L_1)} \end{aligned}$$

  where recall that L_1 is the number of 2-period contracts signed in period 1.

Using equation (2), any given pair will sign 2-period contracts as long as:

$$\begin{aligned} \bar{v} - c &> r_2^s \delta \bar{v} + \mu_2 r_2^c (1 - \delta)\bar{v} \\ \bar{v} - c &> \delta \bar{v} + \mu_2 \frac{S - L_1}{\mu_2(C - L_1)}(1 - \delta)\bar{v} \\ \bar{v} - c &> \delta \bar{v} + (1 - \delta)\frac{S - L_1}{C - L_1}\bar{v} \end{aligned} \tag{3}$$

Notice that the right-hand side (the value of signing a short-term contract) is *decreasing* in L_1, since $\frac{S-L_1}{C-L_1}$ is decreasing in L_1.[16]. That is, if other pairs sign long-term contracts, the joint value of a short-run contract to the charterer and ship decreases, because it becomes harder for the charterer to find a ship next period. So long-term contracts exert negative externalities on other charterers; and long-term contracts are strategic complements in this model.[17]

We now consider how the equilibrium choices of contracting depend on demand fluctuations, by looking at a number of different cases:

(A) *Low demand:*

$$\bar{v} - c < \delta \bar{v} \Rightarrow \bar{v} < \frac{c}{1 - \delta}$$

[16] To see why, notice that if we set $L_1 = S$ (the maximum value it can take), then the term equals 0, but if $L_1 = 0$, the term equals S/C, a positive number

[17] The strategic complementarity of contracts is also true in other general equilibrium models of vertical integration such as McLaren (2000) and Grossman and Helpman (2005).

In this case, the right-hand side of equation (3) is strictly larger than the left-hand side. Even if all other pairs sign long-term contracts, any given charterer-ship pair will still find it optimal to sign short-term contracts. Thus, when demand (as captured by $\bar{v}$) is sufficiently low, the equilibrium will involve everyone signing short-term (1-period) contracts.

(B) *High demand:* Suppose that

$$\bar{v} - c > \delta\bar{v} + \frac{S}{C}(1-\delta)\bar{v} \Rightarrow \bar{v} > \frac{c}{(1-\delta)(1-\frac{S}{C})}$$

This ensures that even with beliefs that are most pessimistic scenario for signing long-term contracts (where no one else signs long-term contracts, $L_1 = 0$), it is still preferable to sign long-term contracts over signing short-term contracts. That is, the LHS of equation (3) is guaranteed to be strictly larger than the RHS (no matter what L is). In this case, the only equilibrium will involve everyone signing long-term contracts.[18] Intuitively, when demand is sufficiently high, charterers have a strong incentive to sign longer contracts with ships, since the market becomes so thin that the probability of finding a ship in the future is low.

(C) *Intermediate demand:* Suppose that we are in the intermediate case where

$$\frac{c}{(1-\delta)(1-\frac{S}{C})} \geq \bar{v} \geq \frac{c}{(1-\delta)}$$

Now, a multiplicity of equilibria is possible. If no one else signs long-term contracts, then any given pair will also not sign long-term contracts. If everyone else signs long-term contracts, then any given pair will also sign long-term contracts.

Thus, the model predicts that agents are more likely to utilize longer contracts during booms, just as we see in practice for container leasing contracts. Long-term contracts are socially inefficient here, because they lead to the possibility of inefficient lock-in (where the charterer keeps using a ship even if their $v_{j2} < c$). Long-term contracts are still signed in equilibrium because the charterer

[18] Because $1 - S/C < 1$, this threshold is higher than the previous threshold.

fears they will not be matched next period and some other charterer will be matched instead; but from a social welfare standpoint, welfare is maximized by short-term contracts. The thinness of the market for ships (S/C) plays an important role in the lengthening of contracts during booms. If the market for ships is very thin, then the threshold beyond which we see only long-term contracting will decrease. That is, the thinner the market for ships, the stronger the incentive to sign long-term contracts (similar to what we observe in the data).

Case 2. Exogenous matching frictions (CRS), no heterogeneity across firms Suppose now there are matching frictions, so that even agents on the short side of the market (i.e., ships) do not always find a match. Let ρ denote the probability that a ship gets matched in the second period. If $\rho = 1$, we are back to the previous case.

The no. of matches is

$$n(C_t, S_t) = \rho S_t$$

(with $S_t < C_t$).

Using equation (2), any given pair will sign 2-period contracts as long as:

$$\begin{aligned} \bar{v} - c &> r_2^s \delta \bar{v} + \mu_2 r_2^c (1-\delta)\bar{v} \\ \bar{v} - c &> \rho(\delta \bar{v} + (1-\delta)\frac{S - L_1}{C - L_1}\bar{v}) \end{aligned} \tag{4}$$

This is similar to the previous case (equation (3)) except that the value of not signing 2-period contracts (the right-hand side of (4)) is scaled by ρ: the lower ρ is, the greater the attractiveness of contracting.

But now, long-term contracting need not always be inefficient. This is because 2-period contracts help address matching frictions.At the same time, the negative consequence of contracting (which is that charterers may have low value ex-post) is still present.

For long-term contracting to be socially efficient, we need to compare the value of 2-period contracts with the *social* value of a 1-period contract. When doing this, we have to recognize that

as long as the ship gets matched to *some* charterer (does not matter who) in period 2, the social value of that match is $\bar{v}$. But this will only happen with probability ρ. So signing a long-term contract is socially efficient only if:

$$\begin{aligned} \bar{v} - c &> \rho \bar{v} \\ \bar{v} &> \frac{c}{1-\rho} \end{aligned} \tag{5}$$

This implies that long-term contracting is only efficient when demand is very high (so that avoiding any possibility that ships are unmatched is socially valuable). As matching frictions decrease (and ρ increases), the set of possible values of $\bar{v}$ for which long-term contracting is efficient also shrinks.

Case 3. Endogenous matching frictions (IRS), no heterogeneity across firms

As before, let ρ denote the probability that a ship gets matched in the second period. But now, instead of being an exogenous parameter, ρ is an endogenous function of market thickness. To make it very simple, suppose that ρ is strictly increasing in S and in C: the thicker the market is, the more likely any ship will be matched.[19]

Using equation (2), any given pair will sign 2-period contracts as long as:

$$\begin{aligned} \bar{v} - c &> r_2^s \delta \bar{v} + \mu_2 r_2^c (1-\delta) \bar{v} \\ \bar{v} - c &> \rho(C - L_1, S - L_1)(\delta \bar{v} + (1-\delta)\frac{S-L_1}{C-L_1}\bar{v}) \end{aligned} \tag{6}$$

The only thing that has changed from Case 2 is that ρ depends on how many pairs sign long-term contracts, or L_1. The negative externality from signing 2-period contracts is now more severe than before, because now we have a new effect: every 2-period contract signed worsens matching frictions faced by everyone else.

[19] An example of such a matching function is

$$n(C_t, S_t) = \rho(C_t, S_t) S_t$$

.

References

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

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

Grossman, Gene M and Elhanan Helpman, "Outsourcing in a global economy," *The Review of Economic Studies*, 2005, *72* (1), 135–159.

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

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.

McLaren, John, ""Globalization" and vertical structure," *American Economic Review*, 2000, *90* (5), 1239–1254.

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

A Appendix

Figure 7: Average contract duration and time-charter index, 1999 - 2022

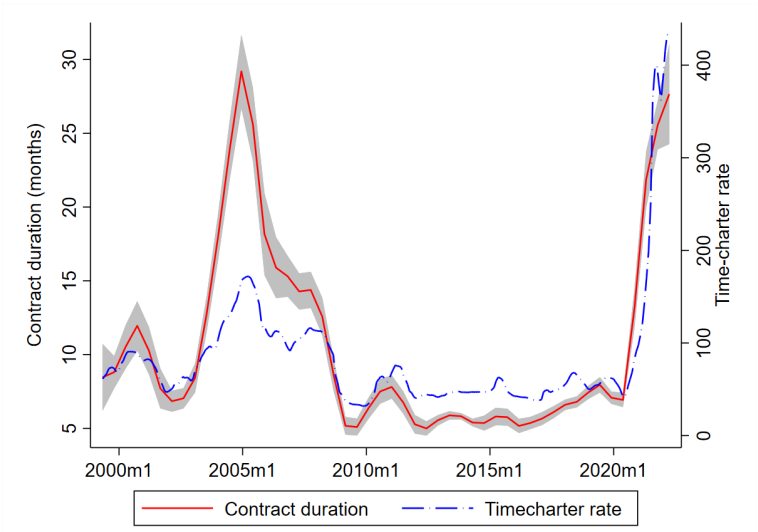
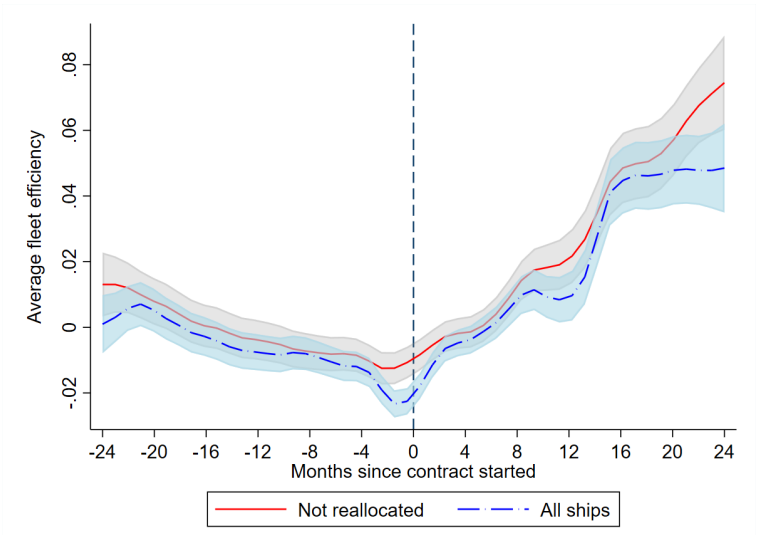
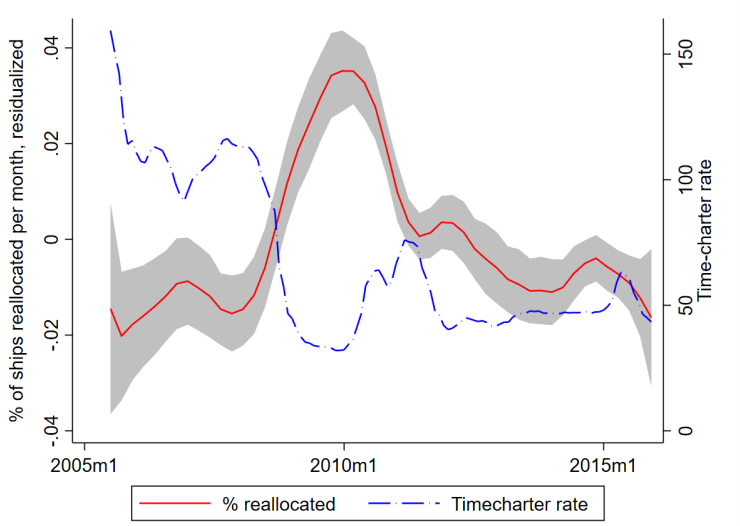


Figure 8: Evolution of efficiency over the course of contract: effect of reallocation



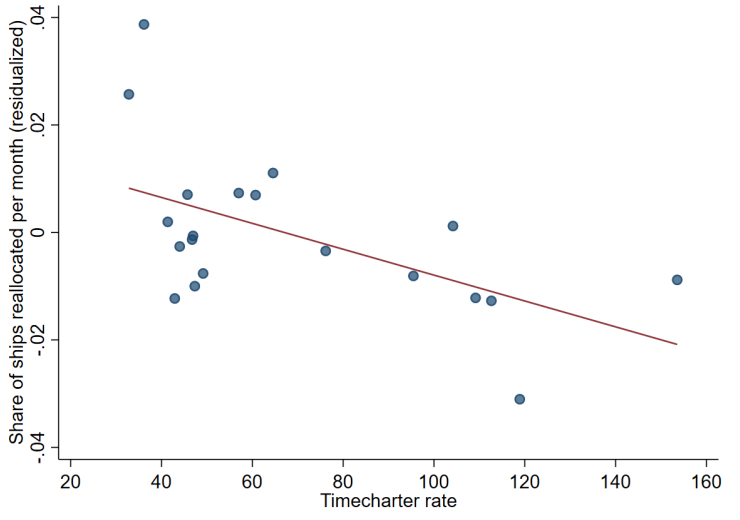
Note: The figure plots the residualized fleet efficiency for ships, for each month in a 4 year window around the month when a new contract is signed. We plot this separately for all ships (dashed blue line), as well as ships that were not reallocated within +/- 4 months of the contract start date (solid red line). The residualized fleet efficiency is obtained after absorbing fleet efficiency on ship fixed effects, year fixed effects, and the time-charter rate and bunker rate.

Figure 9: Reallocation over the business cycle, residualized



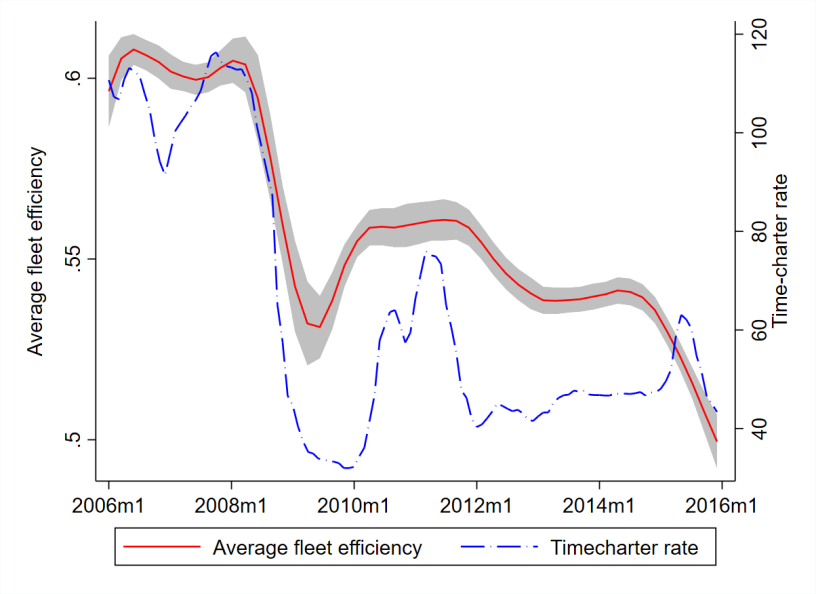
Note: We regress our reallocation measure on ship fixed effects and an aggregate time trend, and obtained the residualized probability of reallocation for each ship, which we then aggregate across ships for each year-month. The left vertical axis plots the residualized share of container ships globally that were reallocated to a different route. The right vertical axis plots the containership time charter rate published by Clarkson's.

Figure 10: Reallocation (residualized) against the time-charter rate



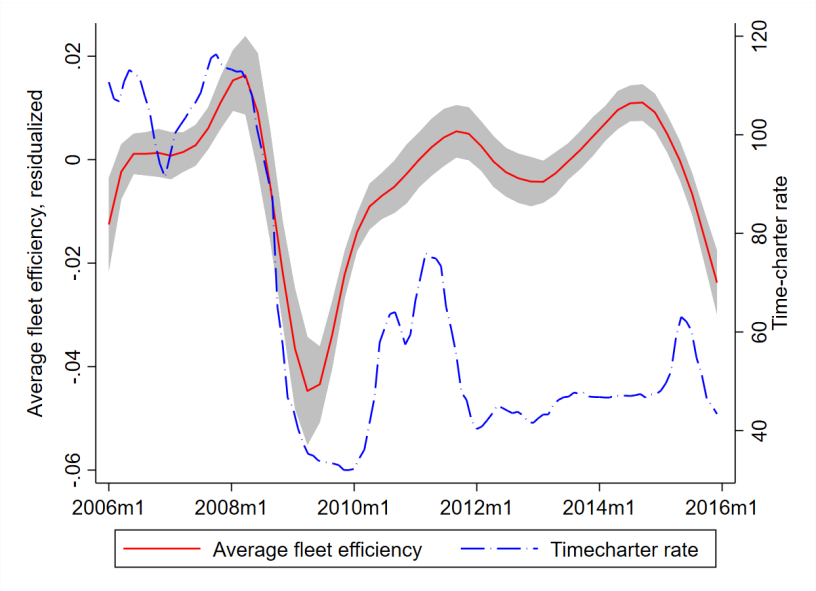
Note: We regress our reallocation measure on ship fixed effects and an aggregate time trend, and obtained the residualized probability of reallocation for each ship, which we then aggregate across ships for each year-month. The figure shows a bin-scatter plot of the residualized share of container ships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

Figure 11: Average fleet efficiency over the business cycle



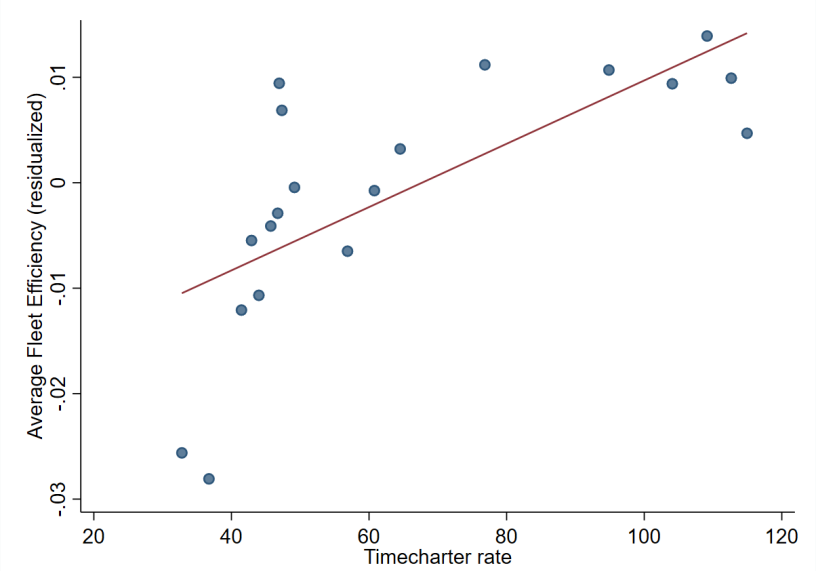
Note: The left vertical axis plots the average fleet efficiency at any given year-month. The right vertical axis plots the containership time charter rate index.

Figure 12: Average fleet efficiency over the business cycle, residualized



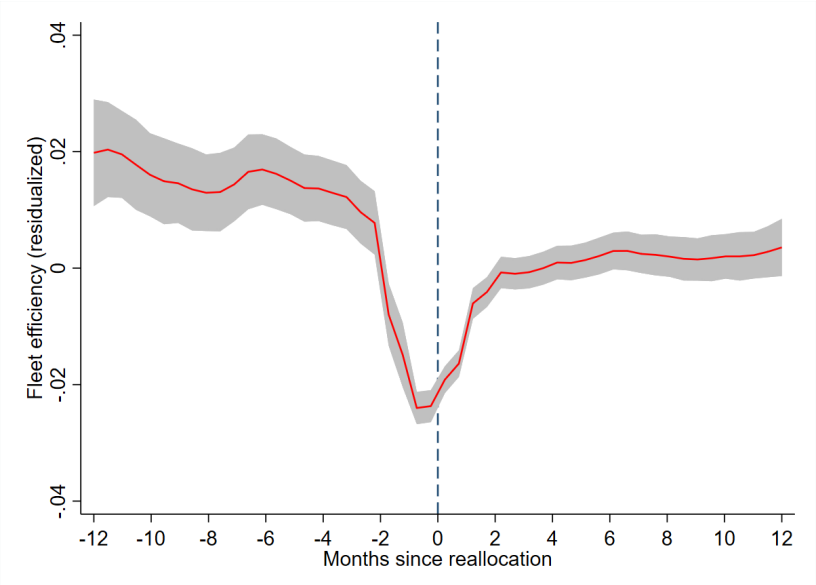
Note: The left vertical axis plots the average of the residualized fleet efficiency (across ships) at any given year-month, with the residuals obtained from a regression of each ship-month's fleet efficiency on ship fixed effects and a time trend. The right vertical axis plots the containership time charter rate index.

Figure 13: Correlation between average fleet efficiency (residualized) and timecharter rate



Note: We plot the average of the residualized fleet efficiency (across ships) at any given year-month, with the residuals obtained from a regression of each ship-month's fleet efficiency on ship fixed effects and a time trend, against the containership time charter rate index.

Figure 14: Efficiency before and after reallocation



Note: The figure plots the residualized fleet efficiency for ships, for each month in a 2 year window around the month when the ship was reallocated. The residualized fleet efficiency is obtained for each ship-month observation by regressing fleet efficiency 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.