Alliance Capacity Sharing and Market Congestion in Container Shipping

Xiaoyan Zhuang* Anhua Chen†

October 2026

Abstract

This paper studies how alliance capacity sharing shapes long-term contracting, vessel availability, and carrier surplus in container shipping. Carriers lease vessels subject to uncertain demand, while alliance members can subsequently share capacity, improving capacity utilization. This encourages greater use of long-term contracts, which reduces vessel availability and creates congestion externalities for independent carriers. We develop and estimate a dynamic model of capacity contracting under demand uncertainty and search frictions using data on lease contracts and capacity utilization. The model captures the feedback between contracting and vessel availability, allowing us to quantify equilibrium spillovers. Counterfactual analysis shows that capacity sharing generates substantial utilization gains for alliance carriers but reduces independent carrier surplus by 1.73%. On net, these gains outweigh the losses, increasing aggregate surplus by 3.24%. The results reveal a channel through which an efficiency-enhancing alliance arrangement can create adverse spillovers for independent carriers by tightening access to capacity in a frictional input market.

*Department of Economics, The Pennsylvania State University. Email xzz5410@psu.edu.
†Cornerstone Research.

1 Introduction

Container shipping is central to global trade and is a key mode for moving manufactured goods across countries. The industry relies on lumpy vessel commitments usually made before future demand is known. Demand remains uncertain over the life of these commitments, exposing carriers to ex post demand-capacity mismatch. Strong demand can cause capacity shortages, delaying shipments and disrupting supply chains, while weak demand can leave committed vessels idle. The central problem is therefore how carriers manage capacity commitments under uncertain demand.

Leasing provides a key margin for managing capacity risk in the short run. In the upstream leasing market, carriers obtain vessels from shipowners through time-charter contracts lasting from several months to several years. Contract duration allows carriers to balance protection against future shortages with exposure to idle capacity. Each carrier negotiates its own leases and bears the vessel's freight-demand and utilization risk. In the downstream freight market, carriers sell transportation services to cargo owners. Some carriers participate in alliances whose members contribute vessels to joint services and share slots on member-operated ships (Ghorbani et al., 2022; International Transport Forum, 2018). Alliance members share vessel capacity use but do not lease vessels as a collective buyer (Federal Maritime Commission, 2026b).

Alliance capacity sharing allows members to use otherwise idle vessel slots to carry partners' cargo. Carrier-specific demand shocks can leave one member with idle capacity while another faces a shortage. For a given fleet and demand realization, sharing can therefore raise production and the return on committed capacity. This production benefit also reduces members' exposure to idle-capacity risk and changes their contracting incentives. Alliance carriers are therefore encouraged to commit to more capacity for longer periods, providing greater protection against future shortages. Both alliance and non-alliance carriers, however, procure vessels in the same leasing market, subject to significant search frictions (Brancaccio et al., 2023; Vreugdenhil, 2025). Longer commitments reduce the vessels available for matching and can create congestion externalities by making capacity harder for other carriers to obtain (Hosios, 1990). Capacity sharing can thus raise alliance members' production from committed capacity while imposing costs on non-alliance carriers that procure vessels in the same leasing market.

This paper thus asks two questions. First, how does capacity sharing change alliance carriers' vessel-contracting decisions and thereby affect vessel availability and search frictions for other carriers? Second, how do these changes affect surplus for alliance carriers, non-alliance carriers and the industry? The paper's main contribution is to quantify the congestion externalities generated by alliances' capacity sharing, which affects other carriers' ability to obtain capacity as demand fluctuates. This allows us to evaluate the effects of alliances' capacity sharing while accounting for spillovers through common frictional input markets.

We first examine contracting and capacity use by combining lease-contract data, service-level utilization, and carrier-level fleet information. The contract panel covers the universe of vessel lease contracts for 2020–2022, recording carrier identities, contract prices, and durations. Capacity utilization is constructed from actual capacity sailed relative to scheduled capacity, whose distribution provides evidence on short-run demand-capacity mismatch. The 2020–2022 Covid cycle provides substantial variation in aggregate demand and vessel availability, from the initial trade contraction to the subsequent rebound (UNCTAD, 2021). Three patterns emerge. First, contract duration varies with expected vessel availability. The scheduled-expiry specification suggests a more negative duration response to availability when more of a carrier's existing capacity is scheduled to expire. Second, alliance-joint services, which involve capacity sharing, have higher capacity utilization than comparable alliance-solo services. Third, alliance carriers use more long-term contracts in both high- and low-demand states. These patterns are consistent with longer contracts providing shortage insurance and capacity sharing improving utilization, thereby reducing the downside cost of committed capacity.

To quantify the mechanisms behind these patterns, we develop a dynamic model of carrier capacity contracting. Carriers differ by alliance status and by their portfolios of short- and long-term leased capacity. Production depends on previously committed capacity and realized demand, so carriers can face either capacity shortages or idle capacity. Carriers adjust future capacity through contracting in a leasing market with search frictions. Each month, they may meet a potential vessel supplier, with the meeting probability depending on aggregate vessel availability. Conditional on a meeting, a carrier chooses contract duration to match, trading off their capacity risk. Longer contracts secure future vessel access but extend exposure to idle-capacity risk when future demand falls. Their value thus depends on expected demand and vessel availability, as well as the carrier's existing capacity portfolio. Capacity sharing increases production from committed capacity by putting otherwise idle capacity into use, raising the value of persistent capacity holdings. Longer commitments keep vessels out of the leasing market and change future search conditions for all carriers. The model quantifies how these changes feed back into contracting decisions across carrier types in equilibrium. We then compare the baseline economy with a counterfactual without alliance capacity sharing to quantify changes in carriers' contracting choices and equilibrium vessel availability.

We estimate the model in two steps. First, we use an empirical transition approach to characterize the evolution of aggregate demand and vessel availability. Second, we estimate the parameters governing capacity sharing, search frictions, demand, and capacity costs by simulated method of moments. Contract duration responses to vessel availability inform search frictions. A key identification challenge is separating capacity-sharing gains from alliance carriers' larger scale. We use capacity utilization as a proxy for ex post demand-capacity mismatch, exploiting the timing that

scheduled capacity is committed before cargo demand is realized. To account for scale differences between alliance and non-alliance carriers, we allow type-specific demand scales and discipline them using each group's average capacity holdings. Accounting for scale within the model, we use alliance carriers' higher mean utilization, lower utilization dispersion, and greater use of long-term contracts relative to non-alliance carriers to identify capacity-sharing gains.

The estimates indicate that carriers face substantial demand uncertainty and search frictions to adjust capacity. Even with aggregate demand unchanged, a one-standard-deviation adverse carrier-specific shock reduces the carrier's own cargo demand by nearly half. This highlights the substantial uncertainty carriers face about their own future demand when committing capacity. The probability of meeting a suitable vessel supplier is 63.2% at low availability and 78.5% at high availability, limiting carriers' ability to search for new capacity. The estimated capacity sharing parameter implies that this arrangement absorbs 34.7% of otherwise idle capacity. Holding baseline capacity and demand fixed, this is equivalent to converting 8.2% of alliance carriers' total contracted capacity into production, which raises the production of committed capacity significantly.

The counterfactual results show that capacity sharing encourages alliance carriers to make larger and more persistent capacity commitments. Relative to the economy without capacity sharing, their total capacity increases by 5.9% and long-term capacity by 16.7%. With more capacity, alliance carriers experience fewer shortages. But these commitments also keep more vessels out of the leasing market, reducing aggregate vessel availability by 6.0%. The resulting decline in meeting probabilities creates congestion externalities for non-alliance carriers, which hold less capacity and face more frequent shortages.

The gains to alliance carriers exceed the losses imposed on non-alliance carriers. Alliance carriers' additional revenue outweighs their higher commitment costs, raising average net surplus by 12.56%. Non-alliance carriers' net surplus falls by 1.73% due to their frequent shortage and lost revenue. These losses offset approximately one-quarter of alliance gains, leaving a 3.24% increase in industry surplus. Capacity sharing therefore generates net gains for the industry, while passing on part of its costs to other carriers through the leasing market.

This paper contributes to four literatures. The first studies shipping industry dynamics, market structure, and the organization of capital. Existing work examines how time to build and demand uncertainty shape bulk shipping investment (Kalouptsidi, 2014), how uncertainty and competition affect container shipping investment (Jeon, 2022), and how capital turnover shapes fleet renewal (Peters, 2024). Closest to the alliance margin, Chen (2025a) studies the trade-off between vessel-reallocation efficiencies and market power in liner shipping alliances. Our paper complements this work by studying short-run capacity management through vessel leasing. It links stochastic demand and capacity commitments to quantify how alliance sharing affects capacity risk and surplus for both alliance and non-alliance carriers.

The second literature studies long-term contracting and contract design in markets with frictions. Contractual form and duration respond to market conditions and relationship-specific considerations (Joskow, 1987; Hubbard, 2001). Recent related work examines cancellation rights (Faltings, 2025), and long-term trucking relationships (Harris and Nguyen, 2025). Most closely related, Garcia-Osipenko et al. (2025) shows that pro-cyclical duration in the same vessel leasing creates contract overhang and ex post misallocation. We add a fleet-level analysis in which carriers choose contract duration considering shortage and idle-capacity risk across their existing portfolios, rather than evaluating each vessel individually. Alliance capacity sharing is also a technology that can reduce ex post mismatch between committed capacity and realized demand.

The third literature studies capacity constraints and the welfare effects of capacity sharing. Flexible or underused capacity can raise utilization and welfare in peer-to-peer accommodation (Farronato and Fradkin, 2022). Related work examines occupancy gains from hotel mergers under uncertain demand (Kalnins et al., 2017), the trade-off between market power and capacity-sharing efficiencies in hotels (McClure, 2025), and utilization gains from airline alliances (Chen and Chen, 2003). Our paper links ex post capacity sharing to ex ante capacity choice in a dynamic environment. Capacity sharing reduces idle-capacity risk but also changes how long firms commit to capacity and when that capacity returns to the market.

The fourth literature studies congestion externalities in decentralized markets. Hosios (1990) provides the classic efficiency benchmark for search markets. Empirical work studies congestion in housing search (Anenberg and Ringo, 2026), and the effects of search frictions on the dynamic allocation of taxis and vessels (Buchholz, 2022; Brancaccio et al., 2023). Our paper shows how congestion can arise from contract persistence. By solving for the equilibrium response, we also quantify the cost of congestion externalities created by alliance's capacity sharing.

The paper proceeds as follows. Section 2 describes the institutional setting and data. Section 3 presents motivating evidence for the model. Section 4 develops the dynamic contracting model. Section 5 discusses identification and estimation. Section 5.2 reports parameter estimates and model fit. Section 6 presents the counterfactual without alliance capacity sharing. Section 7 concludes.

2 Institutional Background and Data

2.1 Market Structure

The carrier side is heterogeneous. In the contract data described below, 186 carriers lease vessels. Ten are alliance carriers. The non-alliance group contains many smaller regional operators and a few larger independent carriers, but even the largest non-alliance carriers are much smaller than

Table 1: Carrier Size and Route Coverage by Alliance Status

Carrier type	Avg. routes	Avg. ships	Avg. fleet cap.
Alliance	126	290	2.01m TEU
Non-alliance	34	6	0.011m TEU

Notes: Fleet statistics are computed from the Clarkson carrier fleet in 2020.

alliance carriers[1]. This size contrast, together with differences in coordination form, motivates treating alliance and non-alliance carriers as two firm types. Table 1 reports these size and scope differences in the data. Because alliance networks require scale and scope to sustain coordinated services, size and alliance status are often correlated. We abstract from strategic capacity competition arising from alliance's potential downstream market power: carriers choose capacity to manage their own future capacity risk, taking their residual-demand process and the downstream revenue environment as given. The empirical analysis therefore distinguishes the capacity sharing channel from the higher baseline demand scale of alliance carriers.

Carriers' fleet capacity comes from owned vessels and leased vessels, but the two margins differ in adjustment horizon. Investment through newbuilding orders changes capacity only over long horizons because ships take time to build and new orders are lumpy and capital intensive (Kalouptsidi, 2014). Newbuilding is therefore a margin for long-run fleet renewal, expansion, and technology upgrading. Leasing reallocates the existing vessel stock and is the main short-run margin through which carriers adjust capacity when demand fluctuates. This margin is sizeable to the industry: leased vessels provided about 53% of operated capacity measured by TEU in 2019 (Global Ship Lease, Inc., 2019). Reliance on leasing also differs substantially across carriers. [2] This paper focuses on the short-run allocation and contracting of the existing vessel stock in the leasing market rather than the long-run determination of aggregate industry capacity. This scope matches the data: charter contracts are observed at high frequency, while newbuilding orders are infrequent and reflect expectations formed several years before delivery.

The upstream leasing market is not concentrated on the carrier side, so we do not treat strategic oligopsony in vessel procurement as the first-order mechanism.[3] On the owner side, the sample contains 355 owners, many of them small non-operating shipowners. These owners include independent shipowners, leasing companies, and financial investors that own vessels but do not run liner services. Appendix 8.1 reports additional statistics on carrier and owner concentration and on

[1] Large non-alliance carriers include ZIM, Wan Hai, PIL, SITC, X-Press Feeders, and KMTC, and ZIM and Wan Hai are 29.3% and 26.0% of the median alliance size, respectively.

[2] In 2026, carriers had a mean charter share of 40.2%, a median of 33.0%, and a standard deviation of 35.6 percentage points (Alphaliner, 2026).

[3] On the carrier side, the Herfindahl-Hirschman Index in leased capacity is about 1,100. We abstract from strategic oligopsony and treat the estimated charter-price schedule as the equilibrium cost faced by an individual carrier conditional on market conditions and contract characteristics.

Table 2: Major Container Shipping Alliances During the Sample Period

Alliance group	Member carriers
2M	Maersk; MSC
Ocean Alliance	CMA CGM; COSCO Shipping; Evergreen; OOCL
THE Alliance	Hapag-Lloyd; ONE; Yang Ming; HMM

carrier-owner contracting relationships.

2.2 Alliance Collaboration in Downstream Freight Market

Alliances are multi-year operational cooperation agreements among legally separate liner carriers.[4] During our sample period, the three main alliance groups were 2M, Ocean Alliance, and THE Alliance.

The industry moved from the old conference system toward the modern alliance organization after regulatory changes restricted explicit conference-based coordination over freight rates. In the modern system, strategic alliances emerged in the mid-1990s as operational collaborations built around vessel sharing, schedule coordination, and network coverage. There is no single global law or regulator that governs all alliance conduct. Instead, alliance agreements are reviewed or filed under jurisdiction-specific regimes[5]. These limits preserve commercial separation: alliances coordinate service operations and slot use, but member carriers remain separate sellers to cargo owners and procure vessel capacity individually through ownership and time-charter leases.[6]

The documented motives include increasing service frequency, exploiting economies of scale and scope, improving network coverage, and filling larger vessels more effectively (Ghorbani et al., 2022; International Transport Forum, 2018). These motives are closely related to capacity-risk management: by coordinating vessels and slots across member networks, alliance members can reduce the residual mismatch between realized cargo demand and committed vessel capacity.

The main form of alliance collaboration is the vessel sharing agreement. A vessel sharing agreement pools operating capacity on a service while leaving each carrier legally separate. On a given trade lane, alliance members jointly determine the service string, including port rotation,

[4] Alliance membership is not permanent. Alliance groups are formed, renewed, and reorganized over time, but the three major alliance groups were stable enough during the 2020–2022 sample period to provide a useful classification of carriers into alliance and non-alliance firms.

[5] In the United States, carrier agreements are filed with the Federal Maritime Commission. In the European Union, liner cooperation is governed by EU competition law.

[6] A shipper typically books with an individual carrier or with a freight intermediary, not with the alliance as a single seller. In U.S. regulation, ocean freight forwarders arrange or book space on behalf of shippers, while non-vessel-operating common carriers buy transportation from vessel-operating carriers and resell it to shippers. These intermediaries affect contracting and booking, but they do not turn the alliance into a joint seller (Federal Maritime Commission, 2026a,b).

sailing frequency, and the number and size of vessels assigned to the service. Members may contribute ships directly or exchange slots on partner-operated ships (Federal Maritime Commission, 2026b, 2024). Capacity sharing gives carriers access to space on partner-operated vessels, with alliance members receiving priority access to partners' unused slots.[7] Net slot positions can be settled through monthly payments, with net users compensating net providers according to agreed reference costs (International Transport Forum, 2018).

2.3 Upstream Leasing Market

Upstream procurement is decentralized. There is no centralized exchange that shows all available ships or posts a complete set of prices. Brokers help carriers and owners find each other by circulating inquiries, tracking open vessels, identifying potential counterparties, and negotiating the main terms before the parties finalize the charter party (Prasad, 2010; UNCTAD, 2022). The need for brokers is consistent with search frictions. Finding a vessel depends on ship characteristics, timing, location, owner identity, duration, delivery and redelivery terms, and optional clauses. Vessel availability is therefore a key market state. When fewer suitable ships are open, carriers face a lower probability of finding capacity on acceptable terms.

The main contract form is the time-charter lease. Under a time charter, the liner carrier hires a vessel for a fixed period and pays hire, usually quoted as a daily rate.[8] Contracts range from short-term contracts to multi-year leases and often include extension options (UNCTAD, 2022; Adland and Prochazka, 2021). Short-term contracts give carriers more flexibility when demand weakens, while long-term contracts secure future vessel access when vessel availability is low. In general, the owner pays for technical operation, including crewing, maintenance, insurance, and classification. The liner carrier pays charter hire and normally pays voyage expenses, including fuel, port and canal charges, and cargo-handling costs (HM Revenue & Customs, 2016; BIMCO, 2014). Once the ship is leased, the carrier must decide how to use it and thus bears the risk that freight demand is too weak to keep the ship fully utilized.

Contract duration creates a commitment because exit is costly and it is hard to reallocate a leased ship before the contract ends. A carrier cannot treat a time charter as a standardized asset that can be sold immediately in a secondary market. Reallocation usually requires a new negotiated contract, an agreed sub-charter or assignment right, or redelivery at maturity, all shaped by bilateral contract terms and brokered search (Prasad, 2010; UNCTAD, 2022). Long contracts therefore commit the carrier to capacity that may be valuable when ships are scarce but costly when demand turns out to be weak.

[7] Alliance agreement gives members first access to unused slots before these can be offered to outside carriers.

[8] A time charter gives the carrier commercial use of the vessel without transferring ownership or technical management. The lease rate is typically agreed at signing, fixing the owner's daily charter revenue over the contract term.

2.4 Data

The primary dataset for the empirical analysis is a contract-level panel of global container-ship time leases from Clarkson. The data cover contracts signed during 2020–2022 and include the contract date, delivery date, owner and carrier identities, ship name, ship characteristics such as age and TEU capacity, the daily contract rate, and contract duration. We exclude obvious renegotiations and contracts longer than seven years[9]. The resulting sample contains 3,983 contracts, with price observed for 3,594 of them, spanning 1,887 ships, 186 carriers, and 355 owners.

We consider contracting decisions that involve multiple ships. In the data, the same carrier frequently signs multiple contracts within a month, often for ships of similar size and with similar duration. In many cases, these contracts plausibly correspond to a common deployment plan, potentially tied to the same service or to closely related services. We therefore interpret such clusters as a bundling decision rather than as a sequence of independent ship-by-ship choices. In the empirical implementation, the relevant contracting unit is a bundle of contemporaneous capacity additions by a carrier.

Table 3 reports the main summary statistics for the contract sample. The average contract lasts roughly one year, but the standard deviation is large, reflecting the coexistence of short contracts and multi-year leases. The average ship in the sample is about 2,500 TEU and 14 years old, though the sample spans a wide range of vessel sizes.

Figure 1 shows the distribution of contract duration. Durations are highly bunched[10] and right-skewed, with many short contracts and a smaller number of long contracts. This pattern motivates the binary short-term versus long-term duration choice used below. The classification captures the main margin of duration choice[11] while preserving the economic contrast between flexible short contracts and persistent capacity commitments. We define short-term contracts as contracts with duration below three months and long-term contracts as contracts with duration of at least three months. This cutoff matches the operational horizon at which carriers typically adjust ship deployment on a quarterly basis.

[9] Charter terms of more than seven years are normally associated with the provision of off-balance-sheet finance supporting investment in newbuild vessels rather than operational leasing.

[10] Long-term contract durations cluster around common horizons, such as integer years.

[11] Because the model is monthly, the distinction between very long contracts, such as three-year and five-year leases, is less important than the distinction between short-term and long-term commitments.

Table 3: Summary Statistics for Lease Contracts, 2020–2022

Variable	N	Mean	SD	Min	Max
Contract duration (months)	3,983	12.06	13.84	0.13	60.00
Daily rate ($)	3,594	18,909.24	21,144.37	820	235,000
Daily rate / TEU ($/TEU)	3,594	9.71	9.07	0.79	68.97
Capacity (TEU)	3,983	2,507.53	1,982.44	326	14,036
Age (years)	3,983	13.62	4.89	1	32
Extended contracts			1,007		
Carriers			186		
Owners			355		
Ships			1,887		

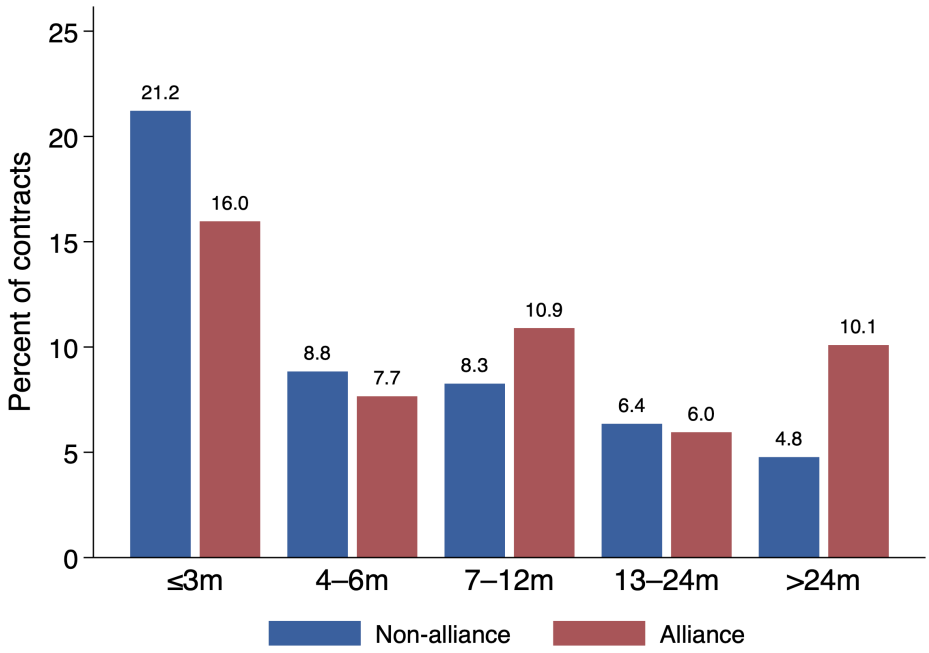


Figure 1: Distribution of Contract Duration by Alliance Status

We complement the contract data with carrier- and market-level information. A Clarkson fleet snapshot measures carrier size and fleet composition. Combining this snapshot with the contracting data, we construct a panel of firm-level fleet composition and contract exposure.[12] At the aggregate level, we use Clarkson's monthly containership idle rate, monthly OECD trade volume, and a Shanghai Shipping Exchange freight-rate index to construct the aggregate market states used in the estimation. The availability state, v_t, is the log of aggregate leased-capacity availability, measured

[12] We abstract from investment-driven changes in fleet composition because the sample period is short.

in units of 100,000 TEU. We measure availability by multiplying Clarkson's total idle-capacity percentage by a constant total leased-capacity base.[13]

Finally, we construct capacity utilization from Sea-Intelligence service deployment data in 2012–2020. Capacity utilization is the ratio of actual sailed capacity to committed scheduled capacity, both measured in TEU-weeks within a service-month. A service is a scheduled container-shipping loop on a trade lane, specified by port rotation, sailing frequency, and vessel assignment.[14] A trade lane can contain multiple services, and a service can be operated by one carrier or jointly by several alliance members. The numerator of capacity utilization is actual sailed capacity, observed weekly as the total capacity of vessels that sail on each service. The denominator is scheduled capacity, constructed from observed sailings and imputed for skipped sailings using nearby active weeks on the same service within the quarter. This construction assumes that schedules adjust quarterly, less frequently than transitory cargo demand.[15] By construction, utilization lies between zero and one, with one indicating that all constructed scheduled capacity was sailed. Because realized cargo is unobserved, we use capacity utilization as a proxy for cargo relative to committed capacity, assuming that carriers adjust sailing capacity to track cargo demand.[16]

Alliance carriers operate two kinds of services. Alliance-joint services operate under a joint vessel-sharing agreement and involve capacity sharing among alliance members. Alliance-solo services are run individually by one alliance carrier without capacity sharing with alliance partners. We therefore separate services into alliance-solo, alliance-joint, and non-alliance services. The purpose is to compare local capacity adjustment for a service across different service providers: alliance carriers operating individually, alliance carriers operating jointly, and non-alliance carriers.

3 Descriptive Evidence

We first examine how vessel availability relates to carriers' contract duration choices. We then ask whether alliance capacity sharing is associated with better use of committed capacity and different duration choices.

[13] Clarkson reports containership idle capacity as a share of the global fleet, measured in TEU. Industry descriptions of the Clarkson idle-capacity measure include vessels that are idle, laid up, or temporarily out of service. We estimate the leasing-market availability rate using this industry-wide idle rate, assuming that the same idle-rate movement applies to leased capacity.

[14] The data cover the seven major trade lanes.

[15] In industry terminology, a skipped scheduled sailing is often called a blank sailing. If a service is blanked for several months, we treat that period as inactive and exclude it from the analysis.

[16] Actual cargo could in principle be identified from bill-of-lading data, which record shipment-level carriers and the vessels carrying them. Following the capacity-utilization logic in Berndt and Fuss (1986) and Butters (2020), if k_{imt} is committed scheduled capacity and q_{imt} is actual sailing capacity, then $u_{imt} = q_{imt}/k_{imt}$ proxies for realized cargo relative to committed capacity. The proxy relies on vessel slots being a non-substitutable input for shipping cargo and on carriers adjusting q_{imt} to maintain target load factors.

3.1 Vessel Availability and Contract Duration

Contracting patterns provide suggestive evidence of search difficulties in low-availability markets. Table 4 summarizes these patterns across high- and low-availability states. In low-availability markets, fewer carriers successfully sign contracts, suggesting that the probability of finding a ship is lower.[17] The lead time between contract signing and delivery also more than doubles, consistent with precautionary contracting to secure future capacity expecting search difficulties.[18]

Table 4: Contracting Outcomes by Vessel Availability

	High-availability market	Low-availability market
Number of carriers with contracting per month	34.9	27.3
Lead time to contract start	13.7 days	34.6 days

Carriers use longer contracts when they expect lower future vessel availability. Lower availability makes searching for capacity harder, while a longer contract provides insurance against future shortages.

We estimate how contract duration varies with future availability at the firm-month level in equation (1).

$$\begin{aligned}\log D_{it} = \alpha_i + \delta_{m(t)} + \beta_v \widehat{v}_{t+1} + \beta_{\text{expiry}} \text{expiry}_{it} \\ + \beta_{\text{expiry},v} \big(\text{expiry}_{it} \widehat{v}_{t+1}\big) + X_{it}' \beta_X + \varepsilon_{it},\end{aligned} \tag{1}$$

The dependent variable is the log of the capacity-weighted average duration of the firm's contracting bundle each month, D_{it}. The main explanatory variable is the forward-looking availability measure, $\widehat{v}_{t+1}$. Because realized availability is partly shaped by current contracting decisions, we use two forward-looking measures based on information available at the time of contracting. The first is predicted availability, constructed from current market conditions.

$$\widehat{v}_{t+1|t} = \widehat{\mathbb{E}}[v_{t+1} \mid S_t], \tag{2}$$

where v_{t+1} is next-period log availability and S_t is information available at t.[19] The second is total scheduled future expiries, a predetermined supply shifter for vessel availability because expiry dates are set by contracts signed earlier.[20] When more capacity is scheduled to expire, more ships

[17] Contracting counts reflect both meeting opportunities and acceptance decisions.

[18] A longer signing-to-delivery time interval may also reflect longer delivery lags or delayed commencement.

[19] We use a linear predictor based on current availability, trade volume, freight-rate index, and scheduled expiring capacity.

[20] Scheduled expiries are observable to market participants because brokers track open positions, recent contracts, and upcoming redelivery windows.

are expected to return to the leasing market, improving availability.

The duration response should be stronger for carriers that are more exposed to capacity shortage risk. We therefore interact each availability measure with expiry_{it}, the carrier's expiring-capacity share, measured between zero and one. This share measures how much of its current capacity must be replaced shortly. The interaction coefficient $\beta_{\text{expiry},v}$ measures how the duration response to expected availability varies with the carrier's expiring-capacity share.

In Table 5, $\widehat{v}_{t+1}$ is measured by the forecast in equation (2) in column (1) and log aggregate scheduled future expiries in column (2). Both specifications include firm fixed effects α_i, log freight-rate index, and ship and contract controls X_{it}, and month-of-year fixed effects $\delta_{m(t)}$.

Table 5: Contract Duration and Expected Vessel Availability

	(1) log(duration)	(2) log(duration)
Measure of $\widehat{v}_{t+1}$	Forecast	Scheduled expiries
Expected availability $\widehat{v}_{t+1}$	-1.491***	0.174*
	(0.303)	(0.103)
$\text{expiry}_{it} \times \widehat{v}_{t+1}$	-0.364	-0.602**
	(0.471)	(0.286)
expiry_{it}	0.326	6.722**
	(0.544)	(3.227)
freight-rate controls	Yes	Yes
Ship and contract controls	Yes	Yes
Firm FE	Yes	Yes
Month-of-year FE	Yes	Yes
N	265	265
R^2	0.760	0.737

Notes: Standard errors are in parentheses. Ship and contract controls are the logged TEU-weighted averages of vessel capacity, daily charter rate, vessel age, and the contract-extension measure. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Column (1) of Table 5 associates lower predicted availability with longer contract duration. The interaction with the expiring-capacity share is negative but statistically insignificant. In column (2), the scheduled-expiry coefficient is positive for carriers with no expiring capacity, while the interaction is negative and significant.[21] The duration response to scheduled expiries therefore becomes more negative as carriers' replacement exposure increases. This pattern is consistent with improved future vessel access reducing the value of long-term shortage insurance, particularly for carriers with more capacity to replace. These findings motivate modeling duration choices jointly

[21] Each availability coefficient gives the duration response when the carrier's expiring-capacity share is zero. At other shares, the response is $\beta_v + \beta_{\text{expiry},v}\text{expiry}_{it}$.

with fleet-level shortage risk and expected search conditions.

3.2 Capacity Sharing and Capacity Utilization

Carriers set service schedules before cargo demand is realized. We use capacity utilization as a proxy for realized production relative to this committed capacity. Controlling for scheduled capacity, we examine whether services with capacity sharing have higher utilization. This comparison motivates the model's production mechanism, in which capacity sharing allows carriers to put otherwise idle capacity into use improving the utilization.

Joint services are services with capacity sharing among alliance members, whereas solo services are operated individually by an alliance carrier. We compare joint and solo services within the same alliance group.[22] This comparison controls for persistent heterogeneity across alliance groups, such as differences in fleet scale, brand and customer base. The service-level comparison focuses on local capacity adjustment within a service.

Let k_{imt} denote scheduled capacity for service i in market m and month t, let q_{imt} denote actual sailed capacity, and let $u_{imt} = q_{imt}/k_{imt}$ denote capacity utilization. The within-alliance service-month regression is

$$y_{imt} = \theta \mathbf{1}\{\text{Alliance joint}_i\} + \gamma_1 \log f_{imt} + \gamma_2 \log k_{imt} + \alpha_{mt} + \xi_{g(i)} + \varepsilon_{imt}. \tag{3}$$

The outcome y_{imt} is either utilization, u_{imt},[23] or log actual sailed capacity, $\log q_{imt}$. Scheduled capacity k_{imt} and scheduled frequency f_{imt} control for planned service design. The market-month fixed effect α_{mt} absorbs common demand conditions on each trade lane in each month. Here $\xi_{g(i)}$ is an alliance fixed effect for the alliance group associated with service i. Thus, θ summarizes the conditional difference between joint and solo services within the same alliance.

[22] Initial joint-service networks are planned as part of alliance formation. This advance planning supports treating joint-service status as plausibly predetermined relative to unobserved demand shocks. Selection based on expected demand or persistent service characteristics may nevertheless remain after conditioning on the controls and fixed effects. We therefore interpret the estimates as conditional differences.

[23] By construction, utilization is bounded above by one, with substantial mass at full utilization. We therefore use utilization in levels rather than logs.

Table 6: Within-Alliance Sailed Capacity Regressions

	Outcome	
	u_{imt}	$\log q_{imt}$
$\mathbf{1}\{\text{Alliance joint}_i\}$	0.070^{**}	0.098^{*}
	(0.035)	(0.052)
$\log f_{imt}$	-0.273	-0.263
	(0.395)	(0.470)
$\log k_{imt}$	-0.038	0.971^{***}
	(0.024)	(0.033)
Observations	5,016	5,016
Adjusted R^2	0.212	0.828

Notes: The sample is restricted to alliance-joint and alliance-solo services. The omitted group is alliance-solo. Standard errors clustered by service are in parentheses. All specifications include market-time fixed effects, alliance fixed effects, and scheduled-capacity weights. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 6 reports 7.0 percentage points higher utilization for joint services relative to solo services within the same alliance. The coefficient of the log sailed-capacity specification implies approximately 9.8% higher sailed capacity on joint services. Conditional on planned service design and common market-month conditions, jointly operated services sail more of their committed capacity. This pattern is consistent with capacity sharing improving the ex post allocation of capacity within a service. We also examine the probability of full utilization and utilization conditional on falling below full utilization. Appendix Table 19 shows that the utilization difference is mainly driven by the full-utilization margin.

Broader comparisons across alliance-joint, alliance-solo, and non-alliance services show the same ordering in both the unconditional utilization distribution and conditional utilization regressions. Alliance-joint services have the highest capacity utilization, followed by alliance-solo and non-alliance services. Appendix Figure 6 and Appendix Table 20 report these comparisons.

3.3 Alliance Carriers and Capacity Commitments

Alliance carriers use long-term contracts more frequently in both demand states, including when weak demand makes idleness a greater concern. We define these demand states using the Covid-era shipping cycle. The low-demand state corresponds to the contraction phase during the Covid disruption, while the high-demand state corresponds to the subsequent rebound. Figure 2 reports the share of accepted contracts with duration of at least three months. In high-demand states, the long-term-contract share is 0.85 for alliance carriers compared with 0.70 for non-alliance carriers. In low-demand states, the corresponding shares are 0.35 and 0.25. The alliance duration gap therefore persists even in low-demand states.

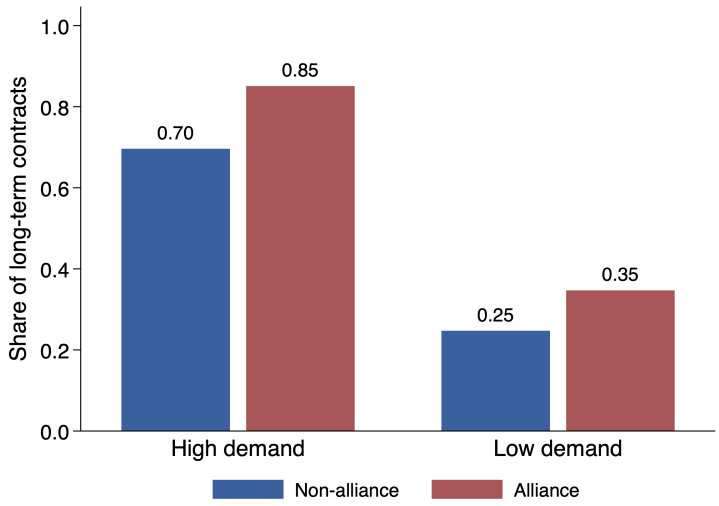


Figure 2: Share of Long-Term Contracts by Alliance Status and Demand State

Alliance carriers' greater use of long-term contracts even in low-demand states is consistent with capacity sharing reducing the downside of persistent capacity commitments. Together with the utilization evidence, this pattern motivates a model linking the gains from capacity sharing to carriers' contract duration choices.

4 Model

We develop a dynamic model of contract duration to quantify how alliance capacity sharing changes carriers' capacity commitments and vessel availability in the common leasing market. Carriers manage portfolios of short- and long-term leased capacity under uncertain shipping demand and limited opportunities to obtain replacement vessels. Capacity is chosen to manage future demand and procurement risk. Contract duration governs the persistence of these portfolios: longer commitments protect against future capacity shortages but incur holding costs even when demand leaves capacity idle. Alliance capacity sharing raises the revenue generated by otherwise idle capacity, changing the value of long-term commitments. These private choices also affect other carriers because capacity retained in existing contracts is unavailable for matching in the leasing market.

After defining the states and timing, we develop three components. The first specifies how demand, capacity constraints, and alliance capacity sharing determine operating payoffs. The second combines these payoffs with contract prices, search frictions, and capacity persistence to characterize carriers' optimal duration choices. The third specifies state transitions and equilibrium: carriers' contract portfolios determine vessel availability, which governs future meeting opportunities. We then use the model to explain the duration trade-off. Together, these components connect the pri-

vate gains from capacity sharing to its effects on capacity allocation across alliance and non-alliance carriers.

4.1 Environment, State Variables, and Timing

Time is discrete, and each period corresponds to one month. Carriers are indexed by i. Each carrier has a persistent type $\tau_i \in \{A, N\}$, where A denotes an alliance-affiliated carrier and N denotes a non-alliance carrier[24]. Let $\mathcal{I}_A$ and $\mathcal{I}_N$ denote the fixed sets of alliance and non-alliance carriers, respectively. The full carrier set is $\mathcal{I} = \mathcal{I}_A \cup \mathcal{I}_N$. We treat alliance status as exogenous and fixed over the sample period.[25]

Each firm's current contract portfolio is $(k_{l,it}, k_{s,it})$, and its type is τ_i. The contract states represent how much short-term and long-term capacity the firm carries into the period. [26] Long-term contracts decay slowly, while short-term contracts expire more quickly. The industry state has two components. The first is the aggregate downstream demand state, d_t, which is exogenous. The second is log aggregate vessel availability in the leasing market, v_t, measuring capacity not under existing contracts. Combining the firm state and the industry state, the complete set of state variables is

$$\mathbf{s}_{it} = (\tau_i, k_{l,it}, k_{s,it}, v_t, d_t),$$

The subscripts i and t index carriers and months, respectively. For ease of notation, $\mathbf{s}$ denotes the carrier's state vector and a its contract choice, with these indices suppressed.

At the time of contracting, carriers observe current aggregate demand d_t, while future demand remains uncertain. Within a period, timing is as follows:

1. State variables $(k_{l,it}, k_{s,it}, v_t, d_t)$ are realized, and carrier i observes its current capacity and current market conditions.

2. The carrier meets a potential vessel supplier with probability $\lambda(v_t)$.

3. Conditional on the opportunity, the carrier chooses an action

$$a_{it}(\mathbf{s}_{it}) \in \mathcal{D} \equiv \{\text{short-term}, \text{long-term}, \text{not match}\},$$

4. Carrier-specific demand is realized during the period. Realized demand and inherited capacity determine current production and operating profit.

[24] We do not distinguish among different alliance groups in the model; all alliance-affiliated carriers are treated as the same type, $\tau_i = A$.

[25] Alliance membership changes over longer horizons through formation, renewal, and reorganization. The three major alliance groups were stable during the sample period.

[26] Capacity is expressed in units of 100,000 TEU unless otherwise indicated.

5. Existing contract stocks evolve through contract expirations and rollover. Newly signed capacity bundles are delivered and added to the corresponding next-period stock, $k_{l,i,t+1}$ or $k_{s,i,t+1}$, according to the chosen duration.[27] The aggregate state also evolves.

4.2 Shipping Demand and Carrier Profits

Profit Function for Carriers. A carrier earns freight revenue from the capacity it puts into service and incurs holding costs on its entire committed fleet. Its realized operating profit is

$$\pi(\mathbf{s}) = m(d)Q_\tau(x,k) - \phi k - \Delta_\phi k_s, \tag{4}$$

Here x denotes realized firm-level latent shipping demand in capacity units before downstream alliance coordination, and $k = k_l + k_s$ is total committed capacity. The production function $Q_\tau(x,k)$ gives production for carrier type τ. Demand, capacity, and production are treated as continuous. Because individual freight rates are unobserved, we assume a common unit freight price $m(d)$ that depends on aggregate demand d. Each carrier takes this price as given.[28] The parameter ϕ is the baseline per-unit holding cost, and Δ_ϕ is the per-unit cost adjustment for short-term capacity k_s. The cost of carrying available capacity is incurred regardless of whether the capacity is used or idle. This captures the capacity commitment of leasing: once capacity is held, the carrier bears the cost of keeping it available. Thus, ϕ is broader than variable operating cost. It captures carrier-side costs required to provide scheduled service, including port charges, terminal handling fees, and other schedule-related costs that must be paid regardless of realized deployment. The term $\Delta_\phi k_s$ allows per-unit holding costs to differ flexibly between short- and long-term capacity.

Production depends on realized demand, committed capacity, and carrier type,

$$\begin{aligned} Q_\tau(x,k) &= \begin{cases} x + \rho \mathbf{1}\{\tau = A\}(k-x) & \text{if } x \le k, \\ k & \text{if } x > k, \end{cases} \\ &= \min\{x,k\} + \rho \mathbf{1}\{\tau = A\}(k-x)_+. \end{aligned} \tag{5}$$

Here $(k-x)_+ = \max\{k-x, 0\}$ denotes idle capacity before capacity sharing. The term $\min\{x,k\}$ captures capacity-constrained production under ex ante capacity commitment: realized service is limited by both realized demand and committed capacity, rather than chosen after demand is observed (Kreps and Scheinkman, 1983; Van Mieghem, 2003).

[27] This timing reflects the institutional fact that time-charter contracts usually have a lead time between signing and use.

[28] A carrier's capacity choice affects its production but does not affect freight prices or its latent shipping demand. The model abstracts from strategic competition in the downstream freight market. Carriers interact through the common vessel-leasing market, where capacity commitments affect future meeting opportunities.

The indicator $\mathbf{1}\{\tau = A\}$ equals one for alliance carriers and zero otherwise. The parameter $\rho \in [0, 1]$ is the share of idle capacity that an alliance carrier can put into production through capacity sharing. The second term thus captures the additional production generated by alliance access when demand falls below committed capacity.[29] We assume that the carrier providing the slot capacity captures all incremental rent from sharing. A carrier facing a shortage may use a partner's capacity, but receives no additional net payoff from the transfer in the model.[30] Under this rent-allocation assumption, ρ raises a carrier's payoff only when it supplies otherwise idle capacity, as captured by $(k - x)_+$.

Capacity sharing involves two stages: carriers first commit to capacity individually before demand is realized, after which alliance members can jointly allocate cargo across their vessels. The parameter ρ summarizes these ex post production gains in reduced form and incorporates their value into carriers' ex ante contracting decisions. We interpret ρ as carriers' belief about the expected production gain from capacity sharing, expressed as a fraction of otherwise idle capacity.[31] Appendix 8.5 illustrates how capacity sharing can raise production for a given fleet and demand realization.

We impose a hard capacity constraint because capacity must largely be secured in advance. Once demand is realized, acquiring and deploying additional vessels within the period is limited by contracting lead times and search frictions. Carriers therefore cannot purchase or lease vessels for immediate use after the demand shock is realized[32].

The nonlinearity in $Q_\tau(x, k)$ is central to the model. Additional capacity raises own-cargo revenue only when the firm expects to be capacity constrained. When demand is below capacity, extra committed capacity does not generate own-cargo revenue, which is left idle without alliance access. Formally, away from the kink at $x = k$, the marginal revenue value of capacity is

$$\frac{\partial\left[m(d)Q_\tau(x,k)\right]}{\partial k} = m(d)\begin{cases}\rho\mathbf{1}\{\tau = A\}, & x < k,\\ 1, & x > k.\end{cases} \tag{6}$$

Thus, the marginal revenue value of capacity is state dependent: high in shortage states, lower in excess-capacity states.[33]

[29] The common unit freight price assumes equal revenue per unit of own and partner cargo.

[30] Alliance members exchange slots and can settle imbalances through monetary payments monthly (International Transport Forum, 2018). Access to partner capacity can therefore help carriers fulfill existing shipping obligations. Thus, capacity sharing is compensated, while assigning all incremental rent to the capacity provider is a simplifying assumption.

[31] Holding this belief constant across aggregate demand states abstracts from variation in members' demand imbalances change over the state. Allowing $\rho(d)$ would partly capture such variation.

[32] Contracting lead times are approximately one month in low-availability markets and two weeks in high-availability markets (Table 4). Standard charter contracts also specify future delivery dates, and forward contracting is common.

[33] This nonlinearity in $Q_\tau(x, k)$ rules out a simple additive ship-by-ship valuation of capacity. The value of an

Shipping Demand. To specify the demand entering production, we restore the carrier and time indices and write x_{it}^{τ} for carrier i's latent demand in month t. We assume that alliance and non-alliance carriers share the same proportional sensitivity to aggregate demand, μ_1, and the same idiosyncratic log-demand dispersion, σ. Baseline demand scale remains type-specific. Demand for type $\tau \in \{A, N\}$ follows

$$x_{it}^{\tau} = \underbrace{\exp(\mu_{0,\tau})}_{\text{type-specific scale}} \cdot \underbrace{\exp\left(\mu_1 d_t + \sigma z_{it}\right)}_{\text{multiplicative demand variation}}, \qquad \tau \in \{A, N\}. \tag{7}$$

Here $\mu_{0,\tau}$ determines baseline demand scale, μ_1 measures sensitivity to aggregate demand d_t, and z_{it} is a standardized idiosyncratic demand shock. The parameter σ scales the shock and governs proportional demand risk. Define $\bar{x}_{\tau} \equiv \exp(\mu_{0,\tau})$ as the baseline demand scale. Then

$$x_{it}^{\tau} = \bar{x}_{\tau} \exp\left(\mu_1 d_t + \sigma z_{it}\right), \qquad \log x_{it}^{\tau} = \log \bar{x}_{\tau} + \mu_1 d_t + \sigma z_{it}. \tag{8}$$

The specification separates type-level scale, aggregate variation, and idiosyncratic variation. A larger σ implies greater proportional dispersion in demand conditional on the aggregate state. Differences in $\bar{x}_{\tau}$ capture persistent type-level scale differences, as alliance carriers are larger in size on average.

The common parameters impose equal proportional demand exposure and risk across types, conditional on baseline scale. They allow demand levels and fluctuations measured in TEU to differ because alliance carriers have a larger baseline scale. This restriction attributes differences in utilization to scale, capacity choices, and alliance sharing, rather than to lower underlying proportional demand risk for alliance carriers. In particular, ρ captures the effect of capacity sharing on production, while σ describes demand risk before sharing.

The idiosyncratic shock satisfies

$$z_{it} \overset{i.i.d.}{\sim} \mathcal{N}(0, 1), \qquad z_{it} \perp d_t. \tag{9}$$

The aggregate state d_t captures common trade shocks and broad shipping demand conditions. The model therefore allows carrier demands to be correlated and persistent through this common state. The idiosyncratic shock z_{it} captures carrier-specific exposure to routes, regions, or customer bases. Conditional on d_t, we treat z_{it} as independent across carriers and periods. Having both common and idiosyncratic components is a parsimonious way to summarize the unobserved covariance structure of carrier-level demand and keeps the state space tractable. At signing, future aggregate demand and

additional ship depends on the firm's existing fleet capacity and realized demand, so the model evaluates capacity at the portfolio level rather than treating each vessel as an independent separable unit.

carrier-specific shocks remain uncertain. Contract commitments expose carriers to both sources of demand risk.

Because carrier-specific demand shocks are independent across periods, we can integrating over these shocks to smooth the expected profits around the capacity constraint. Expected profits conditional on the state is

$$\bar{\pi}(\mathbf{s}) = m(d)\mathbb{E}_z[Q_\tau(x,k) \mid \mathbf{s}] - \phi k - \Delta_\phi k_s. \tag{10}$$

4.3 Dynamic Contracting Problem

Current production depends on previously committed capacity, while carriers adjust the capacity for future production through contracting. The contracting outcomes in Section 3.1 suggest that search frictions limit carriers' ability to adjust their fleets. We capture these frictions through random meetings with potential vessel suppliers. A meeting offers a capacity bundle and we assume the total size of contract bundle is fixed within carrier type.

Random Search. Each carrier participates in the leasing market in every period. The indicator $\iota \in \{0,1\}$ records whether the carrier meets a potential vessel supplier with suitable capacity, with

$$\iota \mid \mathbf{s} \sim \text{Bernoulli}(\lambda(v))\,. \tag{11}$$

The meeting probability $\lambda(v)$ depends on aggregate vessel availability v. Search is random, as in standard search models of decentralized markets, where frictions are summarized by meeting probabilities (Diamond, 1982; Pissarides, 2000; Gavazza, 2016).

We specify the meeting probability as

$$\lambda(v) = 1 - \exp\left(-\kappa_1 e^v - \kappa_2 e^{2v}\right), \qquad \kappa_1 \geq 0, \quad \kappa_2 \geq 0. \tag{12}$$

The nonnegative index $\kappa_1 e^v + \kappa_2 e^{2v}$ increases with available capacity and is mapped into a meeting probability between zero and one. A larger v represents a market with more available capacity, which makes it more likely that a searching carrier encounters a suitable vessel. The parameters κ_1 and κ_2 jointly determine how meeting probabilities respond to vessel availability, with κ_2 allowing additional curvature.

Search participation is taken as given, and all carriers face the same meeting probability $\lambda(v)$.[34] Conditional on a meeting, the carrier chooses whether to form a match and, if so, the contract duration.

[34] Brokers circulate vessel offers and connect carriers with potential suppliers (Prasad, 2010). This brokered structure motivates treating carriers as exposed to offers while leaving the acceptance decision to them.

Dynamic Contracting Problem. Let $V(\mathbf{s})$ denote the carrier's expected discounted value before the meeting and action-specific shocks are realized, and let $\beta \in (0,1)$ be the discount factor. The next-period state $\mathbf{s}'$ depends on the current state and action, with transitions specified in Section 4.4. The Bellman equation is[35]

$$V(\mathbf{s}) = \underbrace{\bar{\pi}(\mathbf{s})}_{\substack{\text{Current} \\ \text{operating profit}}} + \underbrace{\beta \mathbb{E}\left[V(\mathbf{s}' \mid \mathbf{s}, \text{not match})\right]}_{\substack{\text{Continuation value} \\ \text{without a match}}} + \underbrace{\lambda(v)}_{\substack{\text{Meeting} \\ \text{probability}}} \underbrace{\mathbb{E}_{\varepsilon}\left[\max_{a \in \mathcal{D}} \left\{ \begin{array}{l} \beta \mathbb{E}\left[V(\mathbf{s}' \mid \mathbf{s}, a)\right] \\ -\beta \mathbb{E}\left[V(\mathbf{s}' \mid \mathbf{s}, \text{not match})\right] - \mathcal{P}(\mathbf{s}, a) + \varepsilon_a \end{array} \right\}\right]}_{\substack{\text{Inclusive value of meeting surplus} \\ \text{(incremental to no matching)}}}. \tag{13}$$

Current operating profit depends on existing capacity. Conditional on a meeting, the carrier therefore chooses the action that maximizes discounted continuation value, net of the contract payment, as well as the action-specific shock,

$$a^*(\mathbf{s}, \varepsilon) \in \arg\max_{a \in \mathcal{D}} \left\{\beta \mathbb{E}\left[V(\mathbf{s}' \mid \mathbf{s}, a)\right] - \mathcal{P}(\mathbf{s}, a) + \varepsilon_a\right\}. \tag{14}$$

The continuation value $\mathbb{E}\left[V(\mathbf{s}') \mid \mathbf{s}, a\right]$ depends on the fleet composition implied by duration choice a, expected demand, and future search conditions over the life of the resulting commitments. Upon meeting a supplier, the carrier observes action-specific shocks ε_a before choosing whether to reject the opportunity or accept a short- or long-term contract. These shocks are i.i.d. Type-I extreme value draws with scale ω_ε. They capture unobserved match-specific considerations, such as how well the offered vessel fits the carrier's service plans.[36]

Accepting a contract requires the discounted payment $\mathcal{P}(\mathbf{s}, a)$, which depends on carrier type, aggregate states, and contract duration. Contract payments are usually made periodically, but the model charges their discounted value at signing. Carriers take this schedule as given. Appendix 8.4 rationalizes the underlying price schedule using an auxiliary shipowner model. Rejecting the match has no payment, so $\mathcal{P}(\mathbf{s}, \text{not match}) = 0$. Without a meeting, the carrier must choose $a = \text{not match}$.

[35] Appendix 8.3 defines the admissible action processes and presents the corresponding expected discounted-payoff formula.

[36] We assume the carrier learns these details upon meeting a supplier and assessing the offer, consistent with search preceding negotiation (Prasad, 2010).

4.4 State Transitions and Equilibrium

The contract stocks evolve according to

$$\begin{bmatrix} k_s'(a) \\ k_l'(a) \end{bmatrix} = \begin{bmatrix} (1-\delta_S)k_s + \zeta k_l + \Delta k_s^\tau(a) \\ (1-\delta_L)k_l + \Delta k_l^\tau(a) \end{bmatrix}, \tag{15}$$

where the action-specific capacity increments satisfy

$$\Delta k^\tau(a) \equiv \begin{pmatrix} \Delta k_s^\tau(a) \\ \Delta k_l^\tau(a) \end{pmatrix} = \begin{cases} (0,0)' & \text{if } a = \text{not match}, \\ \big(\Delta k_s^\tau(\text{short-term}), 0\big)' & \text{if } a = \text{short-term}, \\ \big(0, \Delta k_l^\tau(\text{long-term})\big)' & \text{if } a = \text{long-term}. \end{cases} \tag{16}$$

Thus, a short-term contract adds capacity only to the short-term stock, a long-term contract adds capacity only to the long-term stock, and no new contract adds no capacity. The parameters δ_S and δ_L are the monthly expiry rates from the short- and long-term contract stocks, respectively. The coefficient ζ captures the rollover of long-term capacity into the short-term stock.[37] The capacity increments $\Delta k_s^\tau(a)$ represent the bundles available at a meeting. We use the same type-specific representative capacity bundle for the short- and long-term actions. The action thus determines only whether the bundle is assigned to the short-term or long-term stock.

Aggregate market availability is endogenously determined by the total amount of capacity not already tied up in existing contracts. Let $\bar{S}$ denote the total stock of vessels relevant for the leasing market and assume to be constant. Then

$$v = \bar{S} - \sum_{i \in \mathcal{I}} \left(k_{l,i} + k_{s,i}\right). \tag{17}$$

When carriers hold more capacity through long-duration contracts, more capacity remains tied to incumbent carriers and less capacity remains available for future matching.

The aggregate demand state follows a general exogenous transition with shock ε_d,

$$d' = g(d, \varepsilon_d). \tag{18}$$

Aggregate-State Equilibrium. We adopt an aggregate-state equilibrium concept in which carriers condition on a low-dimensional aggregation state with consistent beliefs. In principle, current contracting activity and future market availability depend on the full cross-sectional distribution of carriers' states. Carriers' contract stocks shape both current duration choices and the evolution of

[37] We interpret ζ as a partial rollover coefficient. Only part of the capacity leaving the long-term stock enters the short-term stock, so ζ need not equal δ_L.

market availability: they determine which capacity expires, which capacity remains tied to incumbent carriers, and thus how much vessel capacity is available for future matching. Tracking the full distribution would require each carrier to forecast a high-dimensional industry state, which is computationally heavy and unrealistic given the large number of carriers. The key idea is to approximate the full distribution with a low-dimensional summary of aggregate availability, in the spirit of Krusell and Smith (1998) and the moment-based dynamic oligopoly framework of Ifrach and Weintraub (2017). In this application, v summarizes the distribution of carriers' contract-stock states through the amount of vessel capacity available for future matching, as defined in equation (17). Related empirical applications of reduced aggregate-state approaches include Vreugdenhil (2025) and Chen (2025b).

Assumption 1 (Atomic Firms). Each carrier is small relative to the aggregate leasing market. A single carrier's duration choice has a negligible direct effect on market availability. Carriers take the forecasting rule for vessel availability, $\Gamma(v, d; \gamma)$, as given when choosing contract duration.

Assumption 2 (Beliefs and Forecasting). At the time of contracting, carriers observe the current public aggregate state (v, d), where v measures leasing-market availability and d measures downstream demand. They do not track the full vector of all carriers' contract positions. Instead, they forecast future market availability through

$$v' = \Gamma(v, d; \gamma). \tag{19}$$

Together with the demand transition in equation (18), this rule defines carriers' perceived law of motion for the aggregate state. The assumption treats (v, d) as a sufficient summary for forecasting future availability. This is consistent with a brokered leasing market in which carriers use public information on market availability and demand conditions to form expectations about future availability.

Equilibrium Definition. Given the discounted contract bundle-payment mapping $\mathcal{P}(\mathbf{s}, a)$ and the search technology $\lambda(v)$, the equilibrium consists of:

- a perceived law of motion for aggregate availability, $\Gamma(v, d; \gamma)$, defined in equation (19), together with the demand transition in equation (18);
- a Markov duration policy $a^*(\mathbf{s}, \varepsilon)$ that solves the choice problem in equation (14), with contracting opportunities governed by the meeting process in equation (11). Its associated value function $V(\mathbf{s})$ satisfies the Bellman equation (13), given the perceived law of motion for aggregate states;

- a belief-consistency condition that carriers' belief about the law of motion for v is consistent with the law of motion induced by their Markov duration policies, contract-stock transitions in equation (15), and the aggregation rule of market availability in equation (17).

4.5 The Duration Trade-Off

Carriers evaluate each contracting opportunity against their existing capacity portfolios. Additional capacity is more valuable when future shortages are likely, but can increase future idleness. Contract duration determines how long these benefits and costs persist, linking each contracting decision to the carrier's future capacity portfolio.

Conditional on a meeting, short- and long-term contracts add the same capacity bundle next period. The persistence difference emerges in subsequent periods, as long-term capacity leaves its contract stock more slowly. The value of choosing a long-term rather than a short-term contract is

$$\begin{aligned} &\beta\left(\mathbb{E}[V(\mathbf{s}' \mid \mathbf{s}, a = \text{long-term})] - \mathbb{E}[V(\mathbf{s}' \mid \mathbf{s}, a = \text{short-term})]\right) \\ &- \left(\mathcal{P}(\mathbf{s}, \text{long-term}) - \mathcal{P}(\mathbf{s}, \text{short-term})\right) \\ &+ \left(\varepsilon_{\text{long-term}} - \varepsilon_{\text{short-term}}\right). \end{aligned} \tag{20}$$

The three terms capture the difference in continuation values, the additional discounted payment for a long-term contract, and the difference in duration-specific shocks, respectively.

Mechanism 1. Shortage Insurance in Low-Availability Markets. High-demand states often coincide with low vessel availability. When future demand is high and vessel availability is low, carriers need more capacity but find it difficult to obtain new ships. A long-term contract has value in this state because it locks in capacity before searching becomes difficult. In this sense, holding committed capacity has an option-like value: it pays off especially in future states where demand is high and search frictions are large, similar to the real-options logic of investment under uncertainty (Dixit and Pindyck, 1994).

The value of long-term contracting can increase the continuation-value advantage in equation (20), particularly for carriers likely to face future shortages. Long-term contracts also require a larger discounted payment due to their longer duration. Thus, carriers tradeoff between the continuation-value difference and the payment difference.

Mechanism 2. Capacity Sharing Reduces the Downside of Capacity Commitments. Long-term contracts preserve access to capacity but also keep carriers committed when their own demand falls below available capacity. In these idle states, capacity produces little revenue while continuing to incur holding costs.

Capacity sharing allows alliance carriers to put idle capacity into production, recovering part of this revenue loss. Both short- and long-term capacity benefit from sharing, but long-term contracts expose carriers to these downside states for longer. By offsetting this additional downside, capacity sharing can increase the continuation-value advantage of long-term contracting.

5 Estimation and Empirical Results

We estimate the model in two steps. The first step determines representative capacity bundles, demand and availability transitions, discounted contract payments, contract expiry and rollover rates, and the unit revenue.

The second step estimates the remaining structural parameters by simulated method of moments, taking the empirical law of motion for (v, d) as carriers' belief about future aggregate states. This follows the two-step empirical dynamic approach in Pakes et al. (2007): observed state transitions discipline continuation behavior. The estimated parameters govern search frictions, type-specific demand scale and dispersion, capacity holding costs, and alliance capacity sharing.

5.1 First-Step Estimation

Table 7: First-Step Estimation Results

Term	Parameter	Estimate	S.E.
Panel A: Demand State Process d_t			
Intercept	c_d	0.048	(0.054)
Lagged demand d_t	ρ_d	0.947	(0.055)
Innovation SD	σ_d	0.330	–
Panel B: Availability Process v_t			
Intercept	c_v	0.169	(0.157)
Lagged availability v_t	ρ_v	0.883	(0.104)
Lagged demand d_t	$\rho_{v,d}$	-0.010	(0.045)
Innovation SD	σ_v	0.189	–
Panel C: Contract-Stock Rollover			
Short-term retention ($k_s \to k'_s$)	$1-\delta_S$	0.587	(0.024)
Long-to-short rollover ($k_l \to k'_s$)	ζ	0.072	(0.010)
Long-term retention ($k_l \to k'_l$)	$1-\delta_L$	0.789	(0.021)

Aggregate State Transition. The demand state, d_t, is the standardized log of monthly merchandise trade volume for OECD countries. We subtract the sample mean of log trade and divide by its sample SD over 2020–2022. We use this series because container-shipping demand mainly

comes from traded-goods demand, and OECD monthly trade statistics provide a monthly demand measure (OECD, 2026). The availability state v_t is measured as described in Section 2. Figure 3 plots the availability and demand states. Higher demand coincides with lower market availability. The demand recovery from the early-2020 contraction accompanies a decline in vessel availability through early 2022. Availability subsequently recovers while demand remains elevated.

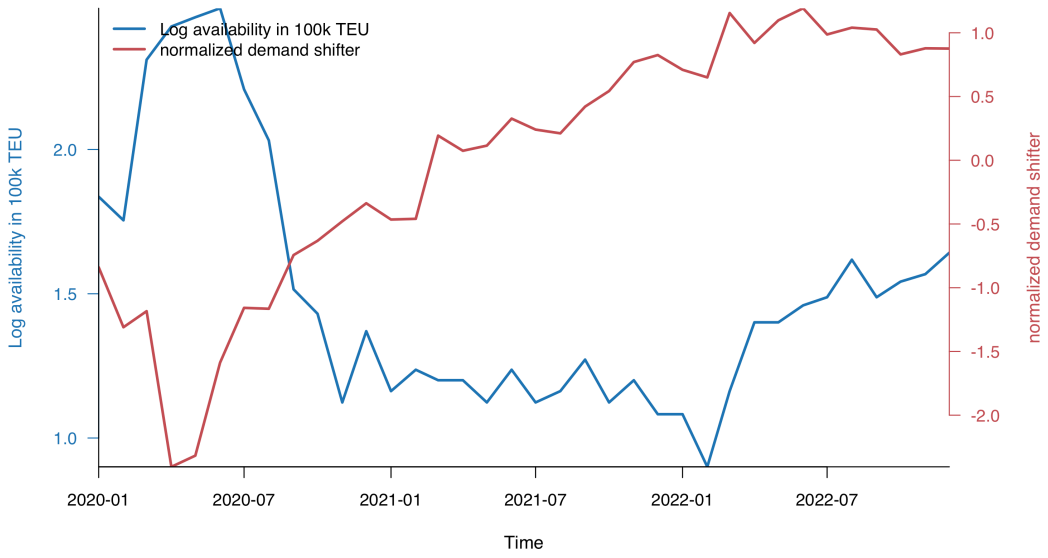


Figure 3: Aggregate Availability and OECD Trade Demand States

We assume an AR(1) process for demand,

$$d_{t+1} = c_d + \rho_d d_t + \varepsilon_{d,t+1}, \qquad \varepsilon_{d,t+1} \sim \mathcal{N}(0, \sigma_d^2). \tag{21}$$

For availability, we impose the conditional linear transition

$$v_{t+1} = c_v + \rho_v v_t + \rho_{v,d} d_t + \xi_{v,t+1}, \qquad \xi_{v,t+1} \sim \mathcal{N}(0, \sigma_v^2). \tag{22}$$

Table 7 reports the first-step estimates. The demand process is highly persistent, with a lag coefficient of 0.947, while its innovation standard deviation is 0.330. Log availability is also persistent, with a lag coefficient of 0.883. The negative conditional demand coefficient, -0.010, indicates that higher demand decreases future availability, although the estimated effect is small after conditioning on current availability.

Price Schedule. We regress log daily per-ship contract rates on aggregate demand, vessel availability, a long-term contract indicator, and its interaction with availability. Controls include log

vessel capacity, log vessel age, a contract extension indicator, and month-of-year fixed effects. The estimates indicate that higher demand and lower vessel availability raise contract prices. Long-term contracts carry a lower average unit price than short-term contracts. This discount is larger when fewer vessels are available in the leasing market.

Let $\widehat{p}_B^{\tau}(v, d, a)$ denote the predicted monthly payment for the representative bundle of carrier type τ. The total discounted contract payment is

$$\mathcal{P}(\mathbf{s}, a) = \sum_{j=0}^{h(a)-1} \beta^j \, \widehat{p}_B^{\tau}(v, d, a), \tag{23}$$

where $h(a)$ is the contract duration in months.

Table 8 summarizes the representative bundles and their predicted monthly prices per TEU. For a given bundle, the longer duration of long-term contracts outweighs the unit-price discount, resulting in a larger total discounted payment. The representative bundles account for a size discount in alliance carriers' contract payments, since larger vessels have lower unit prices per TEU. The estimated price schedule determines contract payments in the carrier's dynamic problem.

Table 8: Representative Contract Bundles and Predicted Prices

	Alliance	Non-alliance
Representative ship size (TEU)	10,000	3,000
Ships per payment bundle	3	1
Total payment bundle size (TEU)	30,000	3,000
Average Price per TEU (USD per Month)		
Short-term (3 months)	163.95	285.89
Long-term (12 months)	156.91	273.61

Notes. Monthly bundle payments $\widehat{p}_B^{\tau}$ equal predicted monthly per-TEU prices times bundle size. Reported prices average predictions equally over months in 2020–2022, using 30-day months. Representative vessels are 14 years old, with extension and seasonal effects set to zero.

Contract Stock Transition. We estimate the contract state law-of-motion parameters in the rollover model given by equation (15), using observed monthly stock transitions. We first subtract capacity added through newly observed contracts from the next month's observed stock. The remaining stock evolution identifies short- and long-term contract persistence and the rollover of expiring long-term capacity into the short-term stock. Panel C of Table 7 reports short-stock retention of $1 - \widehat{\delta}_S = 0.587$, implying a monthly expiry rate of $\widehat{\delta}_S = 0.413$. Long-stock retention is $1 - \widehat{\delta}_L = 0.789$, implying a monthly exit rate from the long-term stock of $\widehat{\delta}_L = 0.211$. The long-to-short rollover coefficient is $\widehat{\zeta} = 0.072$. This reflects partial rollover, as only part of the ca-

pacity leaving the long-term stock enters the short-term stock.[38] Table 8 reports the representative contract bundle sizes for each carrier type.

Revenue Shifter. We calibrate unit revenue as a function of the aggregate demand state d_t using the freight-rate index,

$$m(d_t) = \exp(m_0 + m_1 d_t), \tag{24}$$

with calibrated coefficients $m_0 = 7.780$ and $m_1 = 0.489$. To convert revenue per TEU per sailing into monthly unit revenue, we assume one sailing per vessel per month.[39] Under this assumption, $m(d_t)$ is measured in USD per TEU per month. At the mean demand state, the implied unit revenue is approximately \$2,392 per TEU per month.

5.2 Structural Estimation

5.2.1 Empirical Moments and Identification

Empirical Moments. We target three groups of empirical moments, reported in Table 9. Mean leased capacity measures the scale of each carrier type's commitments. Long-term contract shares and their differences across availability states and carrier types describe duration choices. Utilization means and standard deviations in low-demand months capture the level and dispersion of capacity use.

[38] Grouping contracts around the three-month cutoff introduces approximation error because contracts within each category have different expiry dates. The constant monthly rates summarize average stock transitions rather than reproduce each contract's fixed maturity.

[39] Each vessel completes a rotation approximately once per month for long-haul routes.

Table 9: Empirical Moments Used in SMM

Moment	Empirical definition
Panel A: Capacity	
$\overline{k}_A$	Mean total leased capacity, alliance carriers
$\overline{k}_N$	Mean total leased capacity, non-alliance carriers
Panel B: Contract Duration	
$\ell_{N,\text{high } v}$	Non-alliance long-term contract share, high availability
$\ell_{N,\text{low } d}$	Non-alliance long-term contract share, low demand
$\ell_{N,\text{low } v} - \ell_{N,\text{high } v}$	Non-alliance low-minus-high availability contrast
$\ell_{A,\text{high } v}$	Alliance long-term contract share, high availability
$\ell_{A,\text{low } v} - \ell_{A,\text{high } v}$	Alliance low-minus-high availability contrast
$\ell_{A,\text{low } d} - \ell_{N,\text{low } d}$	Alliance–non-alliance contrast, low demand
Panel C: Utilization	
$\overline{u}_{A,\text{low FR}}$	Mean alliance utilization, low freight-rate months
$\overline{u}_{N,\text{low FR}}$	Mean non-alliance utilization, low freight-rate months
$\text{std}_{u,A,\text{low FR}}$	Std. dev. of alliance utilization, low freight-rate months
$\text{std}_{u,N,\text{low FR}}$	Std. dev. of non-alliance utilization, low freight-rate months

Table 10: Capacity Utilization in Low-Demand Periods by Carrier Type

Carrier type	Mean	Std. Dev.
Alliance joint	0.895	0.101
Non-alliance	0.808	0.194

Notes: Low-demand periods are identified using the freight-rate index. Statistics are weighted by total scheduled capacity. Alliance-solo services are excluded.

We classify high and low availability using cutoffs from the empirical distribution of v_t. For contract-duration moments, low d denotes the Covid low-demand state. Because the service data cover a different period, we use the freight-rate index within that sample to define low-demand months for utilization moments. Using these utilization moments alongside the contract data requires a stationarity assumption. Conditional on the aggregate freight-rate state and month fixed effects, the underlying demand-capacity mismatch process is assumed comparable across the two sample periods.

We construct total capacity utilization at the carrier-group level in low-demand periods. For each group-month, we divide total sailed capacity by total scheduled capacity, aggregating services within the same alliance or non-alliance carrier group. Table 10 reports capacity-weighted means and standard deviations for these low-demand periods.

Only accepted contracts are observed, so long-term contract shares are calculated relative to all accepted contracts in each carrier-type and state group. The data do not directly measure meeting

probabilities or rejected opportunities.

We estimate the parameters jointly. The following discussion explains the identification for each parameter.

Demand Scale and Risk. The demand scales $\mu_{0,N}$ and $\mu_{0,A}$ in equation (7) determine baseline demand $\bar{x}_\tau = \exp(\mu_{0,\tau})$. Mean capacity holdings and utilization rates for each carrier type jointly discipline these demand scales. The parameter σ governs proportional private-demand risk. A larger σ makes realized demand harder to predict and can increase utilization dispersion. Non-alliance utilization dispersion is therefore informative about σ. Because sharing changes realized alliance utilization, we use the common σ assumption to infer demand risk for both types from the non-alliance moments.

Capacity Costs. The holding-cost parameter ϕ in equation (4) governs the cost of committed capacity regardless of its use. Capacity holdings discipline this common component of capacity cost. The short-term cost wedge Δ_ϕ governs the relative cost of short-term and long-term capacity. Non-alliance long-term contract shares in high-availability are informative about this wedge, when shortage risk are lower.

Meeting Probabilities. The parameters (κ_1, κ_2) govern how the meeting probability $\lambda(v)$ changes with vessel availability in equation (12). The sensitivity of duration choices to vessel availability helps identify the intensity of search frictions. The parameter κ_2 governs curvature in the relationship between availability and meeting probabilities.

Capacity Sharing. The parameter ρ in equation (5) allows alliance carriers to generate production from otherwise idle capacity. To distinguish this gain from scale effects, we use utilization as a proxy for ex post demand-capacity mismatch. Scheduled capacity is committed before cargo demand is realized, so utilization measures how much of that capacity is subsequently used. The comparison between alliance and non-alliance carriers accounts for their different scales through type-specific demand parameters. After accounting for scale differences, alliance carriers' higher mean utilization and lower dispersion in low-demand periods inform how much otherwise unused capacity is absorbed through capacity sharing. The alliance–non-alliance duration gap provides complementary information about how these gains affect contracting incentives.

5.2.2 Estimation Procedure

Our datasets provide complementary information on contracting, capacity holdings, and utilization, but do not form a complete panel of carrier states and utilization outcomes. We therefore

use a moment-based estimator. The model moments depend on carriers' capacity states, which evolve endogenously through contracting decisions and are incompletely observed in the data. We therefore simulate capacity portfolios under the model's contracting policies to compute the corresponding moments conditional on the observed aggregate-state path.

For each candidate parameter vector, we solve the carrier's dynamic contracting problem in equation (13), given the estimated price schedule and the empirical transition rule $\Gamma(v, d; \gamma)$ in equation (19). We approximate the value function with basis functions in the state variables defined in Section 4.1, rather than solve the model on a full tensor grid. Conditional on an initial distribution of carrier states, we hold the aggregate state path $\{v_t, d_t\}_{t=1}^{T}$ fixed to the observed path and forward simulate meeting opportunities from equation (12), duration choices from equation (14), private demand shocks from equation (7), and contract-stock transitions from equation (15). We then estimate the parameters by minimizing the distance between simulated and empirical moments in equation (25). Appendix 8.6 gives the computation details.

$$\widehat{\theta} \in \arg\min_{\theta} \left[m^{\text{sim}}(\theta) - m^{\text{data}}\right]' \Omega \left[m^{\text{sim}}(\theta) - m^{\text{data}}\right], \tag{25}$$

where Ω is a weighting matrix.[40]

5.2.3 Parameter Estimates and Moment Fit

Table 11 shows that the model captures capacity differences across carrier types and contract duration responses to availability. The fit is weaker for non-alliance mean utilization and alliance utilization SD, both of which the model overstates. The common demand-risk restriction may contribute to the overstated alliance utilization SD if alliance carriers face lower underlying proportional demand risk than non-alliance carriers.

[40] We use an identity weighting matrix after normalizing each capacity moment by its corresponding empirical mean.

Table 11: Moment Fit

Moment	Data	Model
Mean capacity, A	1.500	1.479
Mean capacity, NA	0.150	0.140
Low- minus high-v long-term share, NA	0.606	0.559
Low-d long-term share, NA	0.257	0.321
High-v long-term share, NA	0.210	0.275
Low-d A–NA long-term-share gap	0.103	0.186
Low- minus high-v long-term share, A	0.641	0.637
High-v long-term share, A	0.287	0.333
Low freight-rate firm-level utilization SD, A	0.101	0.165
Low freight-rate firm-level utilization SD, NA	0.194	0.177
Low freight-rate firm-level mean utilization, A	0.895	0.869
Low freight-rate firm-level mean utilization, NA	0.808	0.911

Notes: A and NA denote alliance and non-alliance carriers. Mean capacity is in units of 100,000 TEU per carrier. The table reports the empirical targets and simulated model moments evaluated at the SMM estimate.

Table 12 reports the parameter estimates.

Table 12: Parameter Estimates

Parameter	Value	Interpretation
ρ	0.347	Share of idle capacity absorbed through capacity sharing within alliances
$\mu_{0,A}$	0.329	Log baseline demand scale for alliance carriers
$\mu_{0,N}$	-1.669	Log baseline demand scale for non-alliance carriers
σ	0.627	Dispersion of private demand shocks
ϕ_L	432.901	Cost per TEU-period of long-term capacity
Δ_ϕ	-420.417	Additional cost per TEU-period for short-term capacity
κ_1	0.245	Linear coefficient in the contracting-opportunity arrival process
κ_2	0.017	Quadratic coefficient in the contracting-opportunity arrival process

Notes: The monthly discount factor $\beta = 0.99$, aggregate demand sensitivity $\mu_1 = 0.126$, and idiosyncratic duration-choice shock scale $\omega_\varepsilon = 0.050$ are calibrated and held fixed. Standard errors are not reported.

The estimated capacity sharing parameter is $\rho = 0.347$, so alliance access absorbs about 34.7% of otherwise unused capacity raising production from committed capacity. At each alliance carrier-state, the direct production gain due to capacity sharing is $\rho \max\{k - x, 0\}$. Holding baseline capacity and demand fixed, this additional production is equivalent to 8.23% of alliance carriers' contracted capacity and 3.18% of total industry contracted capacity. Relative to baseline production, the same gain accounts for 9.74% of alliance production and 3.64% of industry production.

Capacity sharing therefore allows a given fleet to produce more by absorbing demand-capacity mismatch across members.

At average aggregate demand conditions, estimated baseline demand scales imply monthly cargo demand is approximately 169,000 TEU per alliance carrier and 23,000 TEU per non-alliance carrier.[41] Alliance carriers have greater demand scale than non-alliance carriers, consistent with their larger fleet sizes and broader service networks.

Private log-demand dispersion is $\sigma = 0.627$. Holding aggregate demand fixed, a one-standard-deviation adverse carrier-specific shock reduces the carrier's own cargo demand by 46.6%. For comparison, a one-standard-deviation adverse aggregate innovation reduces a carrier's own demand by 4.1% on impact. The large private shocks highlight the uncertainty carriers face about their own cargo demand when choosing capacity. Committing capacity under this demand uncertainty thus exposes them to both capacity shortages and idleness. Aggregate shocks remain important because they persist over time.

The estimated monthly holding costs are approximately $433 per capacity TEU for long-term capacity and $12.5 for short-term capacity. The higher long-term holding cost can be interpreted as the cost of maintaining service, including port charges and terminal fees. Short-term additions may meet temporary capacity needs within existing services, requiring fewer additional commitments. While short-term capacity costs arise mainly from contract payments (Table 8). For comparison with industry figures, Maersk's operating costs imply approximately $535 and $595 per TEU per month in 2021 and 2022 (A.P. Moller–Maersk, 2023). These industry figures include some contract-payment expenses, whereas the model accounts for contract payments separately.

The meeting coefficients are $\kappa_1 = 0.245$ and $\kappa_2 = 0.017$. Through equation (12), these estimates imply meeting probabilities of 63.2% at low availability and 78.5% at high availability. A carrier therefore fails to meet a suitable vessel supplier in roughly one-third of months at low availability, compared with one-fifth at high availability, indicating significant search frictions. Lower meeting probabilities in low-availability markets make new capacity harder to find, raising the value of securing future vessel access through longer contracts. This explains the higher long-term contract share in these markets.

Figure 4 compares empirical availability with the path simulated under observed aggregate demand, using the rule for availability in equation (17), expressed in logs. This checks whether simulated contracting reproduces market availability dynamics, which govern search frictions in the counterfactuals. The model captures the broad decline but overstates variation and misses the 2022 recovery. As contracted capacity approaches total supply, the logarithmic rule amplifies small changes in commitments, potentially contributing to excess variation. The rule also holds total supply fixed, so omitted supply-side shocks in the leasing market may contribute to the weaker fit.

[41] At $d = 0$, mean demand is $\exp(\mu_{0,\tau} + \sigma^2/2)$.

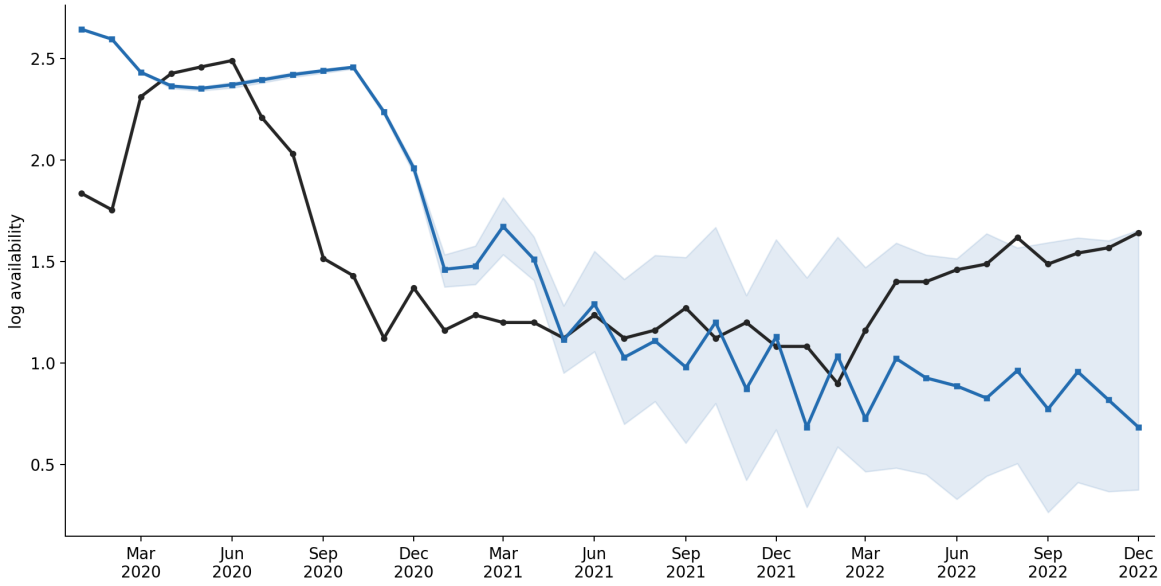


Figure 4: Fit of the Availability Path

Notes. Black line shows empirical log availability; blue line shows the simulation mean. The shaded band spans the 25th to 75th percentiles across simulations using the observed demand path. Total supply $\bar{S}$ is calibrated to match the mean of empirical availability.

6 Counterfactual Analysis

Alliance capacity sharing improves members' use of committed vessel capacity, but it also changes their contracting incentives and other carriers' access to vessels. Assessing alliance efficiencies therefore requires quantifying both members' production gains and congestion externalities in the leasing market. We compare the baseline economy with a counterfactual in which existing alliance members cannot share capacity.

In the baseline economy, capacity sharing absorbs a share $\rho = 0.347$ of otherwise idle alliance capacity. The counterfactual sets $\rho = 0$, representing a stylized restriction on capacity sharing. Alliance membership and all other primitives remain fixed, including both carrier types' demand processes and the estimated contract-price coefficients. The exercise thus isolates the effects of capacity sharing within existing alliances.

Removing capacity sharing lowers the return on committed alliance capacity and may weaken incentives to sign long-term contracts. Shorter commitments could release vessels into the leasing market, improving access for non-alliance carriers. We first examine alliance carriers' contract duration, capacity composition, shortages, and idleness. We then examine equilibrium changes in vessel availability and meeting probabilities, together with their effects on non-alliance carriers' capacity holdings and shortages. Finally, we report changes in net surplus for each carrier group

and the industry.

6.1 Effects of Alliance Capacity Sharing

For each economy, we solve the belief-consistent equilibrium described in Appendix 8.6. Carriers solve their dynamic contracting problems, their duration choices update contract stocks, and aggregate availability is recomputed. We update carriers' beliefs about the law of motion for vessel availability until they are consistent with the one generated by their contracting behavior.

6.1.1 Alliance Capacity Commitments and Capacity Risk

Table 13 shows that capacity sharing increases alliance carriers' total capacity by 5.9% and long-term capacity by 16.7% relative to the counterfactual. Short-term capacity falls by 17.4%. Conditional on a meeting, the probability of choosing a long-term contract rises from 70.4% to 82.7%. Capacity sharing therefore strengthens alliance carriers' incentives to commit to long-term capacity, shifting their portfolios toward longer commitments and increasing their total capacity stock.

Shortage probability measures how often demand exceeds capacity. Capacity sharing reduces this probability from 47.0% to 44.0%. The larger, more persistent capacity portfolio provides some insurance against future shortages, making alliance carriers less likely to lack capacity when demand is high.

Table 13: Alliance Carriers' Outcomes

Variable	Counterfactual	Baseline	Capacity sharing effect
Total capacity, k (kTEU)	153.34	162.36	5.9%
Long-term capacity, k_l (kTEU)	104.65	122.15	16.7%
Short-term capacity, k_s (kTEU)	48.69	40.20	-17.4%
Pr(LT | meeting) (%)	70.44	82.67	12.2 pp
Shortage probability (%)	46.99	44.02	-3.0 pp
Mean utilization (%)	79.89	85.58	5.7 pp
SD of utilization (%)	24.60	16.64	-8.0 pp
Idle rate (%)	21.36	15.31	-6.0 pp

Notes. Capacity is in kTEU (1,000 TEU) per carrier. Effects are relative to the counterfactual, in percent for capacity and percentage points otherwise. Idle rates are capacity weighted.

Capacity sharing raises alliance mean utilization from 79.9% to 85.6%. The idle rate measures the share of capacity that remains unused; it falls from 21.4% to 15.3%. By putting idle slots into production, capacity sharing allows alliance carriers to hold more capacity while leaving a smaller share unused. Also, the standard deviation of utilization falls from 24.6 to 16.6 percentage points, indicating less variability in capacity utilization.

6.1.2 Search Congestion and Spillovers to Non-Alliance Carriers

Alliance carriers' increased use of long-term contracts keeps more vessel capacity committed and reduces its availability in the leasing market. Table 14 reports the aggregate response. Capacity sharing raises total contracted capacity by 0.7% and reduces mean vessel availability by 6.0% relative to the counterfactual. These more persistent commitments leave less capacity available for other carriers to obtain.

The reduction in vessel availability creates a congestion externality that lowers the meeting probability for all carriers. Meeting probability falls by 2.23 percentage points, from 69.93% to 67.70%, making it harder for carriers to find vessels.

Table 14: Market Availability

Outcome	Capacity sharing effect
Total contracted capacity	$+0.7\%$
Long-term contracted capacity	$+5.6\%$
Short-term contracted capacity	-8.8%
Mean vessel availability	-6.0%
Meeting probability	-2.23 pp

Notes. Effects are relative to the counterfactual, in percent except for meeting probability (percentage points).

Table 15: Non-Alliance Carriers' Outcomes

Variable	Counterfactual	Baseline	Capacity sharing effect
Total capacity, k (kTEU)	14.65	14.31	-2.3%
Long-term capacity, k_l (kTEU)	9.48	9.36	-1.2%
Short-term capacity, k_s (kTEU)	5.18	4.95	-4.3%
Pr(LT | meeting) (%)	65.55	66.31	0.8 pp
Shortage probability (%)	67.89	69.21	1.3 pp
Mean utilization (%)	89.85	90.35	0.5 pp
SD of utilization (%)	18.63	18.23	-0.4 pp
Idle rate (%)	11.03	10.59	-0.4 pp

Notes. Capacity is in kTEU (1,000 TEU) per carrier. Effects are relative to the counterfactual, in percent for capacity and percentage points otherwise. Idle rates are capacity weighted.

Non-alliance carriers receive no direct production benefit from capacity sharing but face the resulting congestion in the leasing market. Table 15 shows that their probability of choosing a long-term contract conditional on a meeting rises from 65.55% to 66.31%. This modest shift is consistent with protecting against increase in search frictions. Despite this response, non-alliance carriers' total capacity falls by 2.3%, with long-term capacity falling by 1.2% and short-term capacity by

4.3%. Choosing long-term contracts more often conditional on a meeting does not prevent their capacity stocks from falling as search conditions worsen in the market.

The smaller capacity stock leaves non-alliance carriers more exposed to shortages. Their shortage probability rises from 67.89% to 69.21%, even as mean utilization increases from 89.85% to 90.35% and the idle rate falls from 11.03% to 10.59%. Higher utilization here accompanies less capacity and more frequent shortages, rather than a direct production gain from capacity sharing. Non-alliance carriers adjust their contract duration choices but still face greater capacity risk.

6.1.3 Carriers' Surplus

Table 16 reports average monthly net carrier surplus for alliance and non-alliance carriers. Net surplus equals revenue from production minus capacity holding costs and contract payments.[42]

Table 16: Carriers' Monthly Surplus Decomposition

Component	Counterfactual	Baseline	Difference	% of counterfactual
Alliance				
Revenue	3.83	4.30	0.47	15.62
Total cost	-0.82	-0.91	-0.09	-3.06
Net surplus	3.01	3.39	0.38	12.56
Non-alliance				
Revenue	7.49	7.37	-0.12	-2.11
Total cost	-1.85	-1.83	0.02	0.38
Net surplus	5.63	5.54	-0.10	-1.73

Notes. Group totals are in billion USD per month; costs enter negatively. Differences are baseline minus counterfactual, and percentages are relative to each group's counterfactual net surplus.

For alliance carriers, capacity sharing increases production by putting otherwise idle slots into use and encouraging larger capacity commitments. The resulting revenue gain comes with higher total costs. A larger long-term capacity stock raises holding costs, even though long-term contracts carry a lower monthly contract price. The monthly revenue gain of $0.47 billion exceeds the additional costs of $0.09 billion, raising alliance carriers' net surplus by 12.56% relative to the counterfactual.

Non-alliance carriers receive no direct production benefit from sharing but face greater search frictions through the leasing market. Congestion reduces their capacity stock and leaves them short of capacity more often, lowering production and revenue. Their smaller fleets also reduce holding costs and contract payments, but these savings only partly offset the revenue loss. Non-alliance

[42] Total costs include discounted contract payments at signing, without amortization over the contract term, consistent with the model.

net surplus therefore falls by 1.73%. This loss reflects the negative congestion externalities from alliance carriers' commitments.

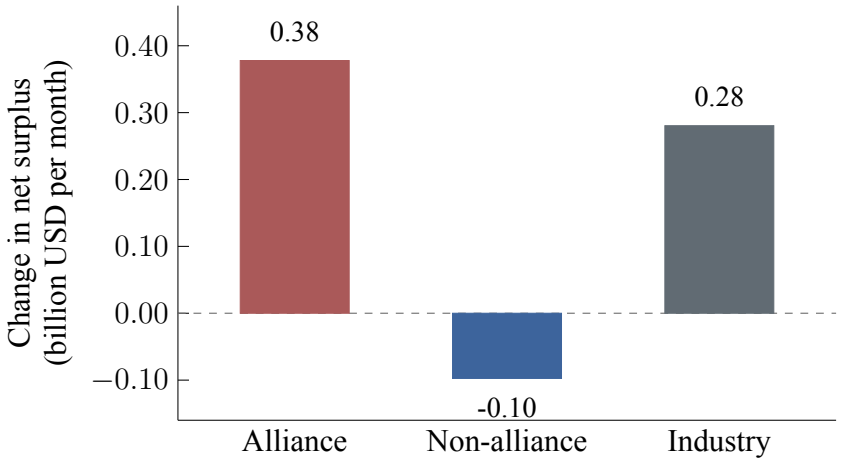


Figure 5: Net Surplus Gains and Losses from Capacity Sharing

Notes. Baseline minus counterfactual, aggregated by carrier group. Industry is the sum of alliance and non-alliance effects.

Figure 5 compares the gains and losses across carrier groups. Alliance carriers' gains exceed non-alliance carriers' losses. The non-alliance loss offsets 25.8% of the alliance gain, leaving a net increase of $0.28 billion per month, or 3.24% of counterfactual industry surplus. Capacity sharing therefore raises industry carrier surplus by 3.24% relative to the economy without capacity sharing, even after accounting for the congestion externalities it imposes on carriers outside the alliance.

This measure covers industry surplus rather than total welfare, which also depend on shipowners, cargo owners' preference for timely and reliable service, and alliance market power (Hummels and Schaur, 2013). Quantifying those channels is left for future research.

7 Conclusion

This paper studies how downstream alliance coordination changes capacity contracting in the vessel leasing market. Carriers commit to vessels before future demand is known, so longer contracts reduce shortage risk but raise idle-capacity risk when demand is weak. Capacity sharing within alliances allows members to absorb some ex post demand-capacity mismatch across services. We estimate this capacity sharing channel and quantify how it changes contract-duration choices and market availability.

Capacity sharing improves alliance carriers' use of committed capacity but encourages longer commitments that reduce vessel availability for other carriers. In the estimated model, alliance

carriers' gains exceed non-alliance carriers' losses, increasing industry carrier surplus. Evaluating alliance efficiencies therefore requires accounting for both members' production gains and congestion externalities in the common leasing market.

References

Roar Adland and Vit Prochazka. The value of timecharter optionality in the drybulk market. *Transportation Research Part E: Logistics and Transportation Review*, 145:102185, 2021. doi: 10.1016/j.tre.2020.102185.

Alphaliner. Alphaliner top 100: Operated fleets. https://alphaliner.axsmarine.com/PublicTop100/, 2026.

Elliot Anenberg and Daniel Ringo. Housing market congestion: The effects of new for-sale supply on home prices and sale hazards. Working paper, dated March 29, 2026; SSRN 5369943, 2026. URL https://ssrn.com/abstract=5369943.

A.P. Moller–Maersk. Annual report 2022. Annual report, 2023. Ocean operating results and operating-cost disclosures, pp. 45 and 91.

Ernst R. Berndt and Melvyn A. Fuss. Productivity measurement with adjustments for variations in capacity utilization and other forms of temporary equilibrium. *Journal of Econometrics*, 33 (1–2):7–29, 1986. doi: 10.1016/0304-4076(86)90025-4.

BIMCO. Bunker terms 2014, 2014. URL https://www.bimco.org/contracts-and-clauses/bimco-clauses/current/bunker_terms_2014.

Giulia Brancaccio, Myrto Kalouptsidi, Theodore Papageorgiou, and Nicola Rosaia. Search frictions and efficiency in decentralized transport markets. *Quarterly Journal of Economics*, 138(4):2451–2503, 2023. doi: 10.1093/qje/qjad023.

Nicholas Buchholz. Spatial equilibrium, search frictions, and dynamic efficiency in the taxi industry. *Review of Economic Studies*, 89(2):556–591, 2022. doi: 10.1093/restud/rdab050.

R. Andrew Butters. Demand volatility, adjustment costs, and productivity: An examination of capacity utilization in hotels and airlines. *American Economic Journal: Microeconomics*, 12(4): 1–44, 2020. doi: 10.1257/mic.20170056.

Anhua Chen. Co-opetition on the seven seas: Global strategic alliance in the container shipping industry. Working paper, SSRN 5355071, 2025a. URL https://ssrn.com/abstract=5355071.

Fisher Chia-Yu Chen and Chialin Chen. The effects of strategic alliances and risk pooling on the load factors of international airline operations. *Transportation Research Part E: Logistics and Transportation Review*, 39(1):19–34, 2003. doi: 10.1016/S1366-5545(02)00025-X.

Luming Chen. The dynamic efficiency of policy uncertainty: Evidence from the wind industry. Working paper, 2025b. URL https://lc929.github.io/workingpaper/Chen_Luming_JMP.pdf.

Peter A. Diamond. Aggregate demand management in search equilibrium. *Journal of Political Economy*, 90(5):881–894, 1982. doi: 10.1086/261099.

Avinash K. Dixit and Robert S. Pindyck. *Investment under Uncertainty*. Princeton University Press, Princeton, NJ, 1994.

Richard Faltings. Pricing the right to renege in search markets: Evidence from trucking. Working paper, arXiv 2506.01650, revised March 2026, 2025. URL https://arxiv.org/abs/2506.01650.

Chiara Farronato and Andrey Fradkin. The welfare effects of peer entry: The case of airbnb and the accommodation industry. *American Economic Review*, 112(6):1782–1817, 2022. doi: 10.1257/aer.20180260.

Federal Maritime Commission. The ocean alliance agreement, fmc agreement no. 012426. https://www2.fmc.gov/FMC.Agreements.Web/Public/AgreementHistory/1214, 2024. Synopsis and agreement history.

Federal Maritime Commission. Ocean transportation intermediaries, 2026a. URL https://www.fmc.gov/about/bureaus-offices/bureau-of-enforcement-investigations-and-compliance-beic/office-of-compliance/ocean-transportation-intermediaries/.

Federal Maritime Commission. How to file an agreement and agreement types. https://www.fmc.gov/about/bureaus-offices/office-of-the-general-counsel-ogc/how-to-file-carrier-mto-agreements/, 2026b.

Maria Garcia-Osipenko, Nicholas Vreugdenhil, and Nahim Bin Zahur. Endogenous rigidities and capital misallocation: Evidence from containerships. Working paper, SSRN 5217587; Queen's Economics Department Working Paper No. 1536, 2025. URL https://ssrn.com/abstract=5217587.

Alessandro Gavazza. An empirical equilibrium model of a decentralized asset market. *Econometrica*, 84(5):1755–1798, 2016.

Mohammad Ghorbani, Michele Acciaro, Sandra Transchel, and Pierre Cariou. Strategic alliances in container shipping: A review of the literature and future research agenda. *Maritime Economics & Logistics*, 24(2):439–465, 2022. doi: 10.1057/s41278-021-00205-7.

Global Ship Lease, Inc. Form 424b5 prospectus supplement, 2019. URL https://www.sec.gov/Archives/edgar/data/1430725/000119312519294860/d759148d424b5.htm. Filed with the U.S. Securities and Exchange Commission on November 15, 2019.

Adam Harris and Thi Mai Anh Nguyen. Long-term relationships and the spot market: Evidence from U.S. trucking. Working paper, 2025.

HM Revenue & Customs. Tonnage tax manual: The meaning of ship operation: Time charters, 2016. URL https://www.gov.uk/hmrc-internal-manuals/tonnage-tax-manual/ttm15120. Updated March 17, 2016.

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

Thomas N. Hubbard. Contractual form and market thickness in trucking. *RAND Journal of Economics*, 32(2):369–386, 2001.

David L. Hummels and Georg Schaur. Time as a trade barrier. *American Economic Review*, 103 (7):2935–2959, 2013. doi: 10.1257/aer.103.7.2935.

Bar Ifrach and Gabriel Y. Weintraub. A framework for dynamic oligopoly in concentrated industries. *Review of Economic Studies*, 84(3):1106–1150, 2017. doi: 10.1093/restud/rdw047.

International Transport Forum. The impact of alliances in container shipping. Technical Report 62, OECD Publishing, Paris, 2018.

Jihye Jeon. Learning and investment under demand uncertainty in container shipping. *RAND Journal of Economics*, 53(1):226–259, 2022. doi: 10.1111/1756-2171.12406.

Paul L. Joskow. Contract duration and relationship-specific investments: Empirical evidence from coal markets. *American Economic Review*, 77(1):168–185, 1987.

Arturs Kalnins, Luke Froeb, and Steven Tschantz. Can mergers increase output? evidence from the lodging industry. *RAND Journal of Economics*, 48(1):45–62, 2017. doi: 10.1111/1756-2171.12172.

Myrto Kalouptsidi. Time to build and fluctuations in bulk shipping. *American Economic Review*, 104(2):564–608, 2014. doi: 10.1257/aer.104.2.564.

David M. Kreps and Jose A. Scheinkman. Quantity precommitment and Bertrand competition yield Cournot outcomes. *The Bell Journal of Economics*, 14(2):326–337, 1983. doi: 10.2307/3003636.

Per Krusell and Jr. Smith, Anthony A. Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896, 1998. doi: 10.1086/250034.

Jonathon McClure. Markups and costs under capacity constraints: The welfare effects of hotel mergers. Job market paper, Purdue University, 2025. URL https://users.ssc.wisc.edu/~jcmcclure2/jmp_jcmcclure_nov23.pdf.

OECD. International merchandise trade statistics. OECD Data Explorer, 2026. URL https://data-explorer.oecd.org/vis?df[ds]=dsDisseminateFinalDMZ&df[id]=DSD_IMTS@DF_IMTS.

Ariel Pakes, Michael Ostrovsky, and Steven Berry. Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *The RAND Journal of Economics*, 38(2):373–399, 2007. doi: 10.1111/j.1756-2171.2007.tb00073.x.

Allen Peters. Beached assets? capital turnover and emissions in shipping. Job market paper, 2024. URL https://alpeters.github.io/assets/pdf/Peters_JMP.pdf.

Christopher A. Pissarides. *Equilibrium Unemployment Theory*. MIT Press, Cambridge, MA, 2 edition, 2000.

K. P. Krishna Prasad. Changing role of ship-brokers: A study of the impact of modern communication in practical ship-broking. The Nautical Institute, 2010. URL https://www.nautinst.org/static/f958c794-0ccc-45e8-899ca30ef1a4e16f/HE00915-CHANGING-ROLE-OF-SHIP-BROKERS-A-study-of-the-impact-of-modern-communication-in-practical-ship-broking.pdf.

UNCTAD. Container shipping in times of covid-19: Why freight rates have surged and implications for policy makers. Policy Brief No. 84, 2021. URL https://unctad.org/publication/container-shipping-times-covid-19-why-freight-rates-have-surged-and-implications-policy.

UNCTAD. Time charterparties. Training course synopsis, 2022. URL https://unctad.org/system/files/information-document/Day_2_Session_3_%26_4_Time_Charterparties_Synopsis.pdf.

Jan A. Van Mieghem. Capacity management, investment, and hedging: Review and recent developments. *Manufacturing & Service Operations Management*, 5(4):269–302, 2003. doi: 10.1287/msom.5.4.269.24882.

Nicholas Vreugdenhil. Booms, busts, and mismatch in capital markets: Evidence from the offshore oil and gas industry. *Journal of Political Economy*, 2025. doi: 10.1086/739827.

8 Appendix

8.1 Leasing Market Structure and Carrier-Owner Relationships

Table 17 reports the number of charterers and owners observed in the 2020–2022 contract sample and the share accounted for by the top 10 firms on each side of the market. Table 18 describes the density of carrier-owner relationships in the same sample.

Table 17: Summary Statistics for Contracts (2020–2022)

	Charterer	Owner
No. of firms	186	355
Top firms	top 10 (52%)	top 10 (23%)

Table 18: Charterer-Owner Relationships (2020–2022)

	Mean	p10	p50	p90
No. of owners per charterer	63.6	6	49	155
No. of charterers per owner	14.1	3	11	31

8.2 Service-Level Utilization Distribution

Figure 6 plots the empirical cumulative distribution of service-month utilization by service type. The figure complements the service-month regressions in Section 3.2 by showing the unconditional distribution of realized deployment relative to constructed scheduled capacity.

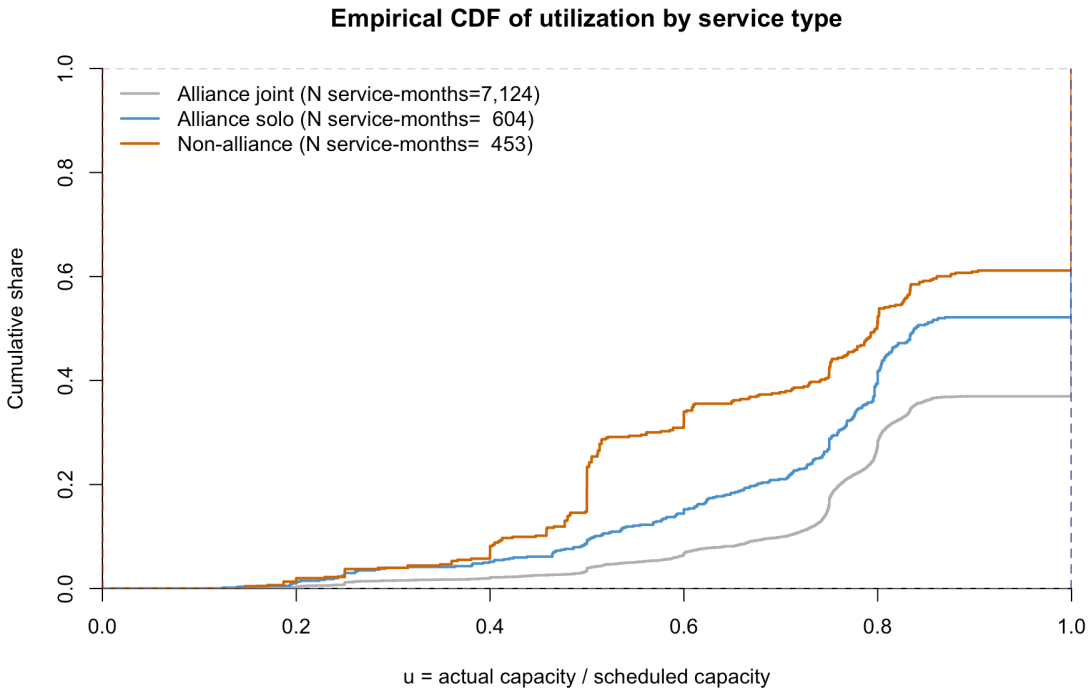


Figure 6: Empirical CDF of Utilization by Service Type

Notes: Utilization is actual capacity divided by constructed scheduled capacity at the service-month level. The dashed vertical line at one marks full realization of the precommitted schedule. Alliance-joint services place less mass below one than alliance-solo and non-alliance services.

Table 19: Within-Alliance Deployment Margins

	Outcome		
	$\log u_{imt}$	$\mathbf{1}\{u_{imt} = 1\}$	$u_{imt} \mid u_{imt} < 1$
$\mathbf{1}\{\text{Alliance joint}_i\}$	0.098^{*}	0.206^{**}	0.010
	(0.052)	(0.092)	(0.021)
$\log f_{imt}$	-0.263	-1.681	1.557^{***}
	(0.470)	(1.790)	(0.416)
$\log k_{imt}$	-0.029	-0.261^{***}	0.068^{***}
	(0.033)	(0.073)	(0.016)
Observations	5,016	5,016	1,923
Adjusted R^2	0.208	0.185	0.208

Notes: The sample is restricted to alliance-joint and alliance-solo services. The omitted group is alliance-solo. Standard errors clustered by service are in parentheses. All specifications include route-month fixed effects, alliance fixed effects, and scheduled-capacity weights. Indicator outcomes are estimated as linear probability models. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 20: Deployment Regressions by Service Type

	Outcome			
	$\log q_{imt}$	$\mathbf{1}\{u_{imt} = 1\}$	$u_{imt} \mid u_{imt} < 1$	u_{imt}
$\mathbf{1}\{\text{Alliance solo}_i\}$	-0.074^{**}	-0.159^{***}	-0.016	-0.053^{**}
	(0.033)	(0.058)	(0.018)	(0.022)
$\mathbf{1}\{\text{Non-alliance}_i\}$	-0.113	-0.261^{***}	-0.009	-0.084^{*}
	(0.069)	(0.088)	(0.039)	(0.046)
$\log f_{imt}$	0.338	0.802	0.186^{***}	0.268
	(0.447)	(1.983)	(0.040)	(0.396)
$\log k_{imt}$	1.019^{***}	-0.148^{***}	0.077^{***}	-0.002
	(0.024)	(0.048)	(0.013)	(0.017)
Observations	8,181	8,181	3,225	8,181
Adjusted R^2	0.848	0.153	0.200	0.187

Notes: The omitted service group is alliance-joint. Standard errors clustered by service are in parentheses. All specifications include route-month fixed effects and scheduled-capacity weights. Indicator outcomes are estimated as linear probability models. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

8.3 Admissible Actions and Lifetime Payoffs

Using the notation in the main text, the set of admissible action processes starting from state $\mathbf{s}$ is

$$\mathcal{A}(\mathbf{s}) = \left\{ \{a_t\}_{t\geq 0} : \begin{array}{l} a_t \text{ is chosen after } (\mathbf{s}_t, \iota_t, \varepsilon_t) \text{ is realized, and} \\ a_t \in \begin{cases} \mathcal{D}, & \iota_t = 1, \\ \{\text{not match}\}, & \iota_t = 0 \end{cases} \end{array} \right\}. \tag{26}$$

The corresponding expected lifetime payoff is

$$V(\mathbf{s}) = \max_{\{a_t\}_{t\geq 0} \in \mathcal{A}(\mathbf{s})} \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \left\{ \bar{\pi}(\mathbf{s}_t) - \iota_t \mathcal{P}(\mathbf{s}_t, a_t) + \iota_t \varepsilon_{t,a_t} \right\} \,\middle|\, \mathbf{s}_0 = \mathbf{s} \right]. \tag{27}$$

This formulation yields the Bellman equation (13) in the main text.

8.4 Auxiliary Owner Model for the Contract Price Schedule

This appendix gives one way to rationalize the reduced-form price schedule $p(v_t, d_t, a)$ used in the carrier problem. The auxiliary model shows why, after conditioning on vessel characteristics, aggregate availability, demand, and duration can summarize the relevant price variation.

Consider a shipowner with one vessel of characteristics χ. The owner can accept a time-charter contract of duration choice a with representative length $h(a)$, or keep the vessel available for fu-

ture opportunities. Let $V^{\text{wait,owner}}(v_t, d_t; \chi)$ denote the owner's value of waiting with an available vessel at public state (v_t, d_t). This waiting value depends on (v_t, d_t) because these public states summarize the owner's future outside option. Aggregate availability affects the probability that the owner meets a prospective charterer for the open vessel, while aggregate freight demand affects the expected contracting price conditional on a meeting. Low vessel availability raises the scarcity value of an open ship, and high freight demand raises the expected value of future contracts. Under this Markov-state assumption, the owner's outside option can be written as a function of (v_t, d_t), conditional on vessel characteristics.

If the owner accepts duration a at per-period hire p, the discounted value is

$$V^{\text{accept,owner}}(v_t, d_t, a; p, \chi) = \left[p - c^O(\chi)\right] \frac{1-\beta^{h(a)}}{1-\beta} + \beta^{h(a)} \mathbb{E}\left[V^{\text{wait,owner}}(v_{t+h(a)}, d_{t+h(a)}; \chi) \mid v_t, d_t\right], \tag{28}$$

where $c^O(\chi)$ is the owner's per-period technical operating cost. The first term is the net payment flow during the charter. The second term is the continuation value after the vessel is redelivered and becomes available again.

The owner's reservation offer for duration a is the contract rate that makes the owner indifferent between accepting and waiting:

$$\begin{aligned} \bar{p}^O(v_t, d_t, a; \chi) = c^O(\chi) &+ \frac{1-\beta}{1-\beta^{h(a)}} V^{\text{wait,owner}}(v_t, d_t; \chi) \\ &- \frac{(1-\beta)\beta^{h(a)}}{1-\beta^{h(a)}} \mathbb{E}\left[V^{\text{wait,owner}}(v_{t+h(a)}, d_{t+h(a)}; \chi) \mid v_t, d_t\right]. \end{aligned} \tag{29}$$

This expression makes the opportunity-cost logic explicit. Accepting a longer contract gives the owner a longer stream of payments, but it also locks the vessel out of future opportunities for more periods. When future market conditions are expected to be strong, the waiting value is high and the reservation hire rises. When aggregate vessel availability is high or freight demand is weak, the waiting value is lower and the reservation hire falls.

We impose two assumptions to connect this owner-side object to the price schedule in the main model. First, conditional on vessel characteristics χ, aggregate availability v_t and aggregate freight demand d_t summarize the owner's future charter opportunities. Under this assumption, the reservation hire depends on the public state and duration through (v_t, d_t, a), with duration entering through $h(a)$.

Second, the auxiliary model treats owners as price takers on the supply side. This assumption is supported by the market structure in Appendix 8.1: the sample contains 355 owners, and the top 10 owners account for only 23% of observed contracts. In this competitive benchmark, the observed

market price equals the marginal owner's reservation price:

$$p(v_t, d_t, a; \chi) = \bar{p}^{O}(v_t, d_t, a; \chi). \tag{30}$$

In the empirical implementation, vessel characteristics are controlled for when estimating contract prices. The carrier model then uses the reduced-form schedule $p(v_t, d_t, a)$, suppressing χ in the notation.

8.5 A Simple Capacity-Constrained Illustration of Internal Transfers

This appendix gives a simple static example for the second-stage internal transfer described in Section 4. Consider two alliance members, A and B. Carrier $j \in \{A, B\}$ enters the period with realized cargo demand x_j and predetermined vessel-slot capacity k_j. Without an internal slot transfer, each carrier can serve only its own cargo on its own capacity, so realized output is capped by a hard capacity constraint:

$$Q^{\text{no capacity sharing}} = \min\{x_A, k_A\} + \min\{x_B, k_B\}. \tag{31}$$

The hard constraint is carrier specific: carrier A cannot use carrier B's idle slots, and carrier B cannot use carrier A's idle slots.

With capacity sharing, the alliance can reassign cargo or slots after demand is realized. The same physical capacity is now constrained at the alliance level rather than at each carrier's boundary. The resulting output under capacity sharing is

$$Q^{\text{capacity sharing}} = \min\{x_A + x_B, k_A + k_B\}. \tag{32}$$

This equation captures internal transfer because it is the value of the second-stage allocation problem after member boundaries are relaxed.

The capacity sharing gain is weakly positive. Since $\min\{x_A, k_A\} \leq x_A$ and $\min\{x_B, k_B\} \leq x_B$,

$$Q^{\text{no capacity sharing}} \leq x_A + x_B. \tag{33}$$

Since $\min\{x_A, k_A\} \leq k_A$ and $\min\{x_B, k_B\} \leq k_B$,

$$Q^{\text{no capacity sharing}} \leq k_A + k_B. \tag{34}$$

Equations (33) and (34) imply

$$Q^{\text{no capacity sharing}} \leq \min\{x_A + x_B, k_A + k_B\} = Q^{\text{capacity sharing}}. \tag{35}$$

The inequality is strict when demand-capacity mismatch is offsetting within the alliance. Let $e_j = (x_j - k_j)_+$ denote excess demand and $s_j = (k_j - x_j)_+$ denote slack capacity. For two carriers,

$$Q^{\text{capacity sharing}} - Q^{\text{no capacity sharing}} = \min\{e_A + e_B, s_A + s_B\}. \tag{36}$$

Thus capacity sharing creates additional realized output only when one member has excess demand and another member has slack capacity.

The same logic has direct utilization implications. Without capacity sharing, carrier-level utilization is

$$u_A^{\text{no capacity sharing}} = \frac{\min\{x_A, k_A\}}{k_A}, \qquad u_B^{\text{no capacity sharing}} = \frac{\min\{x_B, k_B\}}{k_B}. \tag{37}$$

Now consider the case above in which carrier A has slack and carrier B has excess demand. Let

$$t_{BA}^* = \min\{k_A - x_A, x_B - k_B\} \tag{38}$$

be the maximum internal transfer from carrier B's excess cargo to carrier A's slack slots. After this transfer, utilization of the capacity-owning carriers is

$$\begin{aligned} u_A^{\text{capacity sharing}} &= \frac{x_A + t_{BA}^*}{k_A} = u_A^{\text{no capacity sharing}} + \frac{t_{BA}^*}{k_A} \geq u_A^{\text{no capacity sharing}}, \\ u_B^{\text{capacity sharing}} &= u_B^{\text{no capacity sharing}} = 1. \end{aligned} \tag{39}$$

The borrowing carrier is already fully utilizing its own capacity, so the physical utilization gain appears on the slack carrier's slots. If carrier B has slack and carrier A has excess demand, the same expression holds with the carrier labels reversed.

Aggregating across the two carriers gives the alliance-level implication. Let capacity-weighted mean utilization be

$$\bar{u}^r = \frac{k_A u_A^r + k_B u_B^r}{k_A + k_B}, \qquad r \in \{\text{no capacity sharing}, \text{capacity sharing}\}. \tag{40}$$

Then

$$\bar{u}^{\text{capacity sharing}} - \bar{u}^{\text{no capacity sharing}} = \frac{Q^{\text{capacity sharing}} - Q^{\text{no capacity sharing}}}{k_A + k_B} = \frac{\min\{e_A + e_B, s_A + s_B\}}{k_A + k_B} \geq 0. \tag{41}$$

Capacity sharing therefore raises utilization by moving the low-utilization carrier toward fuller deployment.

8.6 Estimation and Computation Details

This appendix summarizes the computation behind the second-step SMM estimator and the belief-consistent simulation used for counterfactuals.

Approximate Value Functions. Let θ collect the structural parameters. For a given θ and a candidate public forecasting rule Γ, the carrier's dynamic problem is solved over the state

$$\mathbf{s} = (\tau, k_l, k_s, v, d). \tag{42}$$

Because the state includes both contract stocks and two market conditions, exact solution on a dense grid is computationally expensive. We therefore approximate continuation values with a basis expansion

$$\widehat{V}(\mathbf{s};\, b) = B(\mathbf{s})'b, \tag{43}$$

where $B(\mathbf{s})$ is a vector of basis functions in (k_l, k_s, v, d, τ) and b is the corresponding coefficient vector. This approximation is used to evaluate expected continuation values and to interpolate the value function away from simulated support points. In the implementation, $B(\mathbf{s})$ uses flexible quadratic spline terms in k_s and k_l, linear spline terms in v and d, and interactions for $k_s \times v$, $k_l \times v$, $k_l \times d$, and $v \times d$.

SMM Estimator. The current second-step estimator is SMM conditional on the empirical aggregate-state transition. The value-function solution uses the estimated forecasting rule $\Gamma(v, d; \gamma)$ to integrate over future states, while the panel simulation evaluates the resulting policies along the realized historical path $\{v_t, d_t\}_{t=1}^{T}$. We initialize carriers from the empirical distribution of capacity states over the sample data. Later capacity states are generated by simulated opportunities, duration choices, and contract-stock transitions.

The moment vector uses contract stocks, contract-duration choices, and low freight-rate utilization, as listed in Table 9. Figure 7 gives the computation algorithm.

Algorithm: Conditional SMM Estimation

1. **Set inputs and targets.** Take as given the historical path $\{v_t, d_t\}_{t=1}^{T}$, the empirical initial distribution of (τ, k_l, k_s), the price schedule $p(v, d, a)$, the forecasting rule $\Gamma(v, d; \gamma)$.

2. **Fix simulation draws.** Draw R independent sets of opportunity-arrival, duration-choice, and private-demand shocks. Generate these draws once and hold them fixed across all candidate parameter vectors.

3. **Solve and simulate.** For each candidate θ, solve the dynamic contracting problem and simulate R carrier panels conditional on the observed aggregate path $\{v_t, d_t\}_{t=1}^{T}$.

4. **Construct simulated moments.** Compute each target moment within each simulated panel and average across the R replications:

$$m^{\text{sim}}(\theta) = R^{-1} \sum_{r=1}^{R} m^{\text{sim},r}(\theta).$$

   Evaluate the SMM criterion in equation (25).

5. **Optimize.** Search over θ subject to parameter bounds.

Figure 7: Computation algorithm for the conditional SMM estimator

Belief-Consistent Inner Loop for Counterfactuals. For counterfactuals, we solve for a belief-consistent pair of policy functions and belief parameters. This routine closes the aggregate feedback from carriers' simulated duration choices to aggregate availability. Figure 8 summarizes the inner algorithm.

Algorithm: Belief-Consistent Fixed Point

1. **Input:** estimated parameters, calibrated supply, common initial states, and demand paths and firm draws fixed across cases and belief iterations.
2. **Initialize:** guess coefficients and innovation dispersion $\gamma^{(0)}$ of the perceived availability law.
3. **For** $j = 0, 1, 2, \ldots$:
   (a) Construct quadrature transitions from $\gamma^{(j)}$ and re-solve the Bellman equation and duration policies.
   (b) Simulate each replication using those policies, realized meetings and utilization shocks, and smoothed choices conditional on a meeting. Update capacity after contract expirations and rollover.
   (c) At each simulated state compute v from residual supply using equation (17); do not generate v from a realized innovation to the perceived law.
   (d) After burn-in, pool replication-period observations and regress next-period v on current v and d. Estimate the innovation dispersion to obtain $\widehat{\gamma}^{(j)}$.
   (e) Stop if $\max_\ell |\widehat{\gamma}_\ell^{(j)} - \gamma_\ell^{(j)}| < 0.01$. Otherwise damp the coefficient update and repeat. No separate availability-path convergence condition is imposed.

Figure 8: Computation algorithm for the belief-consistent fixed-point algorithm