How Big Is Too Big?

The Competitive Effects of Shipping Technology Innovation

Anhua Chen

anhua_chen@g.harvard.edu

March 5, 2025

Abstract

Technological innovation drives economic progress but can also lead to unintended consequences, such as market over-consolidation. This paper examines the competitive effects of technological advancement in the container shipping industry, focusing on the relationship between vessel size and market structure, and its impact on welfare. Using a dynamic oligopoly model and proprietary data, we quantify the economies of scale associated with vessel size and explore how it interact with market competition and investment behavior. We find that a 10% increase in vessel size reduces operational costs by 3.4% but contributes to market concentration, potentially offsetting consumer benefits. Counterfactual analysis suggests that the welfare-optimal vessel size is around 20,000 TEU under current demand level, as larger vessels risk over-consolidating the market. Additionally, smaller innovation steps promote competition, while larger steps drive consolidation. These findings highlight the need for policymakers to balance the efficiency gains of technological advancements with their competitive impacts, particularly when designing infrastructure investments in industries with strong economies of scale.

Keywords: Oligopoly, Innovation, Market Power, Transportation, International Trade

JEL Codes: L13, O31, L13, R40, F13

1 Introduction

Technological innovation has long been a cornerstone of productivity growth and economic progress, fundamentally reshaping industries and global markets over the past century. However, as technology continues to advance, it raises an important question: can innovation go too far? While technological advancements often promise efficiency gains and lower costs, they can also generate unintended consequences, such as market over-consolidation and reduced competition. These effects may undermine the very welfare benefits that innovation is intended to deliver. Understanding and accurately accounting for the competitive effects of technological innovation is therefore essential for policymakers seeking to balance the trade-offs between efficiency, competition, and long-term economic welfare.

This paper examines the competitive effects of technological innovation in the container shipping industry, with a focus on assessing whether the advancement of shipping technology—particularly the trend toward larger vessels—has surpassed its welfare-optimal level. One of the most significant technological advancements in container shipping over the past few decades is the steady increase in vessel size, aimed at reducing transportation costs through economies of scale.[1] However, the growing prevalence of ultra-large container ships has drawn criticism for its potential risks to maritime infrastructure and its role in exacerbating market concentration, as the industry has experienced substantial consolidation in recent years.[2]

Specifically, this paper investigates the relationship between increasing vessel size and

[1] See Merk (2018), Imai et al. (2006), Haralambides (2019), Murray (2016), Yang et al. (2011), Veldman (2011)

[2] The Dali incident and the Ever Given blockage highlight the growing risks that ultra-large container vessels pose to critical maritime infrastructure. In March 2021, the Ever Given, a 400-meter-long, 20,000+ TEU container ship, ran aground in the Suez Canal, blocking one of the world's most vital trade routes for nearly a week. The disruption halted global trade, stranding over 400 vessels and causing billions in economic losses. More recently, in March 2024, the Dali, a 10,000+ TEU vessel, lost power and collided with the Francis Scott Key Bridge in Baltimore, Maryland, leading to the catastrophic collapse of the bridge and severe disruptions to regional port operations. Both incidents underscore the infrastructure challenges posed by the increasing size of container ships—larger vessels require more precise navigation, place greater strain on port and canal infrastructure, and magnify the consequences of accidents, raising concerns about the long-term sustainability of ever-growing ship sizes. There are also discussion of how large vessel has facilitated the market consolidation in the container shipping industry. See also link.

the consolidation of the container shipping market. It seeks to answer the central question: how large is too large when the potential competitive effects of vessel size are considered? While the pursuit of larger vessels may drive down freight costs through efficiency gains, a more concentrated market structure could offset these benefits by pushing freight prices back up. Furthermore, this study explores how industry equilibrium outcomes vary under different technological frontiers (maximum vessel sizes) and innovation step sizes (the pace at which carriers can upgrade their fleets). These questions are critical for understanding the broader implications of technological innovation on competition and welfare in the container shipping industry and beyond.

Our paper also contributes to the ongoing discussion on the returns to investment in maritime infrastructure (see Brancaccio et al. (2024), Ganapati et al. (2024)). By incorporating the effects of market structure and competition into the analysis of technological advancements, we provide a more precise estimation of the potential welfare benefits of infrastructure investments in a highly concentrated industry such as container shipping. This framework enables a more accurate cost-benefit analysis, which is crucial for evaluating billion-dollar infrastructure investment projects.[3]

We developed a structural, model-based empirical framework to estimate the demand and cost structure of the container shipping industry. Specifically, we construct and estimate a non-stationary dynamic oligopoly model with stochastically alternating moves, following the approach of Igami and Uetake (2020). In the stage game, container carriers compete in a Cournot fashion, with operating costs primarily determined by the average vessel size of their fleet—larger vessels incur lower costs due to economies of scale.[4]

To estimate demand, we utilize proprietary monthly container trade volume and price data from Container Trade Statistics, covering 22 major trade lanes to estimate a gravity-

[3]The Biden administration has enacted significant investments in transportation infrastructure, primarily through the Bipartisan Infrastructure Law (BIL), also known as the Infrastructure Investment and Jobs Act, signed into law in November 2021. This historic $1.2 trillion legislation represents the most substantial federal investment in infrastructure in decades. $42 Billion is directed at modernizing airports and maritime ports.

[4]In this context, carriers refer to shipping companies, while shippers are customers of shipping services.

style demand model. To address price endogeneity, we construct an instrumental variable exploiting the re-routing of container ships from the Suez Canal to the Cape of Good Hope following the Red Sea attacks by the Houthi group. Only some trade lanes are directly affected by these attacks, resulting in an exogenous, supply-driven increase in shipping prices for those routes. For cost estimation, we leverage the prediction that carriers with lower costs tend to supply larger quantities, particularly during demand downturns. Using proprietary capacity utilization data from Sea Intelligence, we estimate operational costs within the stage game framework and infer the economies of scale associated with vessel size.

The dynamic game of investment, entry, and exit is modeled using a stochastic alternating-move framework to represent a non-stationary dynamic oligopoly. In each period, Nature selects a single carrier (the "mover") to make decisions regarding whether to upgrade their average vessel size, remain idle, or exit the market. As mentioned in Igami and Uetake (2020), this framework addresses the issue of multiple equilibria by transforming the problem into a single-agent dynamic discrete choice model for each period. Additionally, the stochastic process of selecting movers mitigates concerns about first-mover advantages inherent in predetermined move orders. By employing a non-stationary equilibrium concept, we substantially reduce computational complexity, as the model can be solved using backward induction.

Due to the infrequent consolidation events in our dataset, we employ the method of simulated moments, focusing on the transition paths of key industry moments—specifically, the number of carriers or alliances and the industry's average vessel size—to estimate investment, maintenance, and entry costs in the dynamic game. Finally, we conduct counterfactual analyses by altering the maximum vessel size and innovation step size to evaluate the welfare implications of the technology frontier and the speed of innovation.

We found significant economies of scale associated with vessel size. On average, a 10% increase in vessel size reduces the average cost of operation by 3.4% on the Asia–Northern Europe trade lane. This finding is consistent with the existing maritime economics literature.[5] The strong economies of scale create a substantial cost advantage for carriers with larger average vessel sizes in their fleets, providing a strong business-stealing incentive for carriers to invest in upgrading vessel size.

Our dynamic estimation reveals a nuanced and non-linear relationship between competition and innovation. The effect of competition on innovation differs for market leaders and followers. Increased competition incentivizes leaders to innovate more due to the preemptive business-stealing effect, as they aim to force laggard competitors out of the market. In contrast, a more concentrated market motivates followers to innovate, as the potential profit gains are larger. Additionally, our results show that the incentive to innovate is higher when firms are more neck-and-neck, consistent with the central prediction of Aghion et al. (2005). Furthermore, we find that the average incentive to innovate is higher for followers than for leaders, reflecting the higher cost of pushing the technological frontier compared to catching up or imitating.

Our counterfactual analysis addresses the welfare implications of the technology frontier and innovation step size. Under current aggregate demand levels, we find that the optimal shipping technology frontier lies at an average vessel size of approximately 20,000 TEU. The current 24,000+ TEU maximum ship size risks over-consolidating the shipping market. Increasing the average vessel size beyond 18,000 TEU contributes minimally to consumer surplus, primarily shifting surplus from consumers/shippers to carriers. The optimal technology frontier is demand-dependent: the pro-competitive effects of technology dominate its anti-competitive effects when demand is high, and the opposite holds true when demand is low. This underscores the importance of incorporating demand forecasts into technology and infrastructure investment policies when accounting for the competitive effects of technological advancements.

Our simulation also sheds light on how innovation step size influences industry equilibrium. Even with the same potential technology frontier, different innovation step sizes can

[5] See Merk (2018), Imai et al. (2006), Haralambides (2019), Murray (2016), Yang et al. (2011), and Veldman (2011).

lead the industry to distinct long-run equilibria. Smaller step sizes reduce the incentive to upgrade, resulting in a more competitive market with smaller vessels, whereas larger step sizes drive greater consolidation, with remaining firms achieving the technology frontier.

To summarize, our counterfactual analysis provides three key policy lessons:

1. The design of an optimal technology frontier must balance the pro-competitive effects of technological advancement with the anti-competitive effects of market consolidation, particularly when technology exhibits strong economies of scale.

2. The optimal technology frontier is contingent on aggregate demand, making demand forecasting a critical component of technology policy design.

3. Innovation step size play a crucial role in determining whether the potential technology frontier can be achieved and at what cost.

These findings have significant implications for the container shipping industry. A key motivation for port and canal infrastructure investments is to accommodate larger vessels, thereby reducing transportation costs. However, our analysis indicates that such investments could lead to market over-consolidation, offsetting the cost-reduction benefits of larger vessels, particularly in the presence of significant economies of scale and high entry barriers. This suggests that the benefits of transportation infrastructure investments may be overestimated if the competitive effects of technology are not considered.

Moreover, even if port and other transportation infrastructure are upgraded to accommodate larger vessels, whether carriers choose to invest in fleets to achieve the maximum size depends on the speed at which they can upgrade their fleets. Therefore, policies and regulations related to shipbuilding, investment, and financing must be coordinated with infrastructure investment authorities to maximize welfare outcomes, especially when innovation step size has a significant impact on industry dynamics.

Related Literature

Our paper contributes to the literature on the relationship between competition and innovation. The seminal work of Schumpeter (1942) argues that under certain conditions, increased competition may discourage innovation, a phenomenon known as the Schumpeterian effect. By incorporating both the Schumpeterian effect and the "escaping competition" effect, Aghion et al. (2005) demonstrate an inverted-U relationship between innovation and competition, where competition discourages laggard firms from innovating but encourages neck-and-neck firms to do so.[6]

Several studies explore the relationship between competition and innovation in dynamic settings, including Benkard (2004), Goettler and Gordon (2011), Kim (2013), Igami (2017), Igami (2018) and Igami and Uetake (2020). Our paper is closest to Igami and Uetake (2020) in terms of model structure and setting, as we adopt their stochastically alternating-move game framework to construct our non-stationary dynamic oligopoly model.[7] However, while Igami and Uetake (2020) focus on optimal merger policies when both innovation and mergers are endogenous, our paper examines policies influencing the trajectory of technological development (e.g., the technology frontier and innovation step size) and their implications for competition and welfare.[8]

Our work also closely relates to Marshall and Parra (2019), who highlight a complex and potentially non-monotonic relationship between competition and innovation. They emphasize that the relationship between the profit gap (between leaders and followers) and the number of firms plays a critical role in determining the effect of competition on innovation.

[6]Our findings align with this framework, providing more detailed results from our structural empirical model of the container shipping industry. Specifically, we find that market followers are significantly less motivated to innovate in highly competitive markets, while the incentive to invest increases when firms are closer to being neck-and-neck. However, we also show that market leaders have reduced incentives to innovate in more concentrated markets, as their business-stealing incentives diminish when the probability of driving smaller competitors out is low.

[7]Recent studies that also utilize this framework include Otani (2024) and Garg and Saxena (2023).

[8]Unlike a focus solely on competition authorities, we emphasize the broader policy implications for agencies such as transportation authorities responsible for infrastructure investment or subsidies. This underscores the need for coordinated policy efforts across multiple agencies, including but not limited to transportation and competition authorities.

Consistent with their findings, we show empirically that the incentive to innovate is directly linked to the business-stealing effect, which is influenced by the profit gap.[9] In summary, while much of the existing literature emphasizes the impact of competition on innovation, our primary focus is on the competitive and welfare effects of different technological innovation characteristics.

Our paper also contributes to the growing literature on the shipping industry and infrastructure investment in the transportation sector. Brancaccio et al. (2020) and Asturias (2020), among others, endogenize transportation costs by carefully modeling the sector and examining the implications of transportation markets on global trade. Our contribution to this strand of literature lies in a detailed exploration of the economies of scale associated with vessel size and the role of market structure. Jeon (2022) investigates the role of beliefs in explaining investment patterns in the container shipping industry, while Kalouptsidi (2014) analyzes how build times contribute to boom-and-bust cycles in shipbuilding. Our paper adds to this body of work by providing deeper insights into how investments in larger vessels (i.e., better shipping technology) interact with market structure.

Additionally, recent research has focused on maritime infrastructure investments.[10] Ganapati et al. (2024) studies the role of hub-and-spoke structures in global trade, emphasizing how large-scale hubs reduce trade costs through economies of scale and improved connectivity. The paper highlights the dual effects of hubs, where they improve efficiency but also concentrate market power, raising important questions for infrastructure investment policy. While their focus is on port-level dynamics, our paper investigates a similar efficiency-concentration trade-off but at the carrier level. Brancaccio et al. (2024) examines how port infrastructure investments impact global trade flows, shipping costs, congestion, and eco-

[9] Marshall and Parra (2019) find that if the profit gap is (weakly) increasing with the number of firms, competition will promote innovation. Conversely, competition reduces innovation when the profit gap decreases as the number of firms increases. Our results further indicate that a significant portion of the product-market business-stealing effect arises from the ex-ante dynamic incentive to force competitors out of the market. For seminal work on dynamic incentives, see Berry and Pakes (1993). A more detailed discussion of the relationship between their findings and ours is provided in the estimation section.

[10] Brancaccio et al. (2024), Ganapati et al. (2024), and Fuchs and Wong (2024)

nomic welfare under demand volatility and network effects. Their findings suggest that the returns on port investments are positive for only a subset of U.S. ports, underscoring the importance of targeted investment strategies. Our paper complements this discussion by demonstrating how competition concerns can influence welfare estimates of maritime infrastructure investments.

2 Data, Industry and Stylized Facts

2.1 Industry

Approximately 80% of global merchandise trade by volume is transported via sea routes, with containerized shipping accounting for about 35% of this volume and over 60% of its commercial value.[11] Since its introduction in the late 1950s, container shipping has significantly improved the efficiency of maritime transportation. Bernhofen et al. (2016) provides evidence that containerization has been a major driver of 20th-century economic globalization.[12] Subsequent innovations in shipping technology have largely focused on increasing vessel size. Over the past two decades, container ship sizes have grown substantially, driven by the pursuit of economies of scale to further reduce transportation costs (see Section 2.3.1).

Unlike bulk (or tramp) shipping, which operates more like an "ocean taxi," container shipping (also referred to as liner shipping) functions more like an "ocean bus." Container shipping services are provided by carriers to shippers on *fixed routes* with *regular schedules*, typically on a weekly basis with predetermined port calls. The number of service lines offered by carriers across specific regional origin-destination pairs (referred to as tradelanes) has remained relatively stable over time.[13] As noted by Wong (2022), container shipping exhibits

[11] See this link.

[12] For additional literature on the economic impact of containerization, see Coşar and Demir (2018) and Brooks et al. (2018).

[13] In industry terminology, broad regional origin-destination pairs are called tradelanes. For example, Transpacific East Bound (TPEB) refers to the tradelane from Asia to North America, while Far East West Bound (FEWB) covers tradelanes from Asia to Europe. These represent the coarsest level of organization for carriers' shipping networks.

a “round-trip” effect, where the head-haul leg carries the majority of the merchandise, while the back-haul leg typically operates with lower quantities. Our demand estimation utilizes data from both head-haul and back-haul legs, but the main empirical analysis focuses on the head-haul segment.

The bus-like nature of the container shipping industry means carriers primarily rely on capacity management to address demand volatility. Among the various tools, blank sailing—where carriers cancel scheduled sailings on certain routes during specific weeks—is the most frequently used for short-term capacity adjustments. Blank sailing during demand downturns helps carriers improve per-vessel capacity utilization without incurring excessive operational overhead.[14] However, even during blank sailings, carriers still incur fixed costs such as amortization, maintenance, or charter fees for idle vessels. This necessitates efficient medium- and long-term fleet capacity management. In our model, the static game captures short-term capacity management decisions, while the dynamic investment game models long-term capacity planning.

The container shipping industry is highly concentrated. Unlike the more competitive bulk shipping market (see Kalouptsidi (2014) and Brancaccio et al. (2020)), the high-sea container shipping market is dominated by approximately eight major carriers, which are organized into three strategic alliances.[15] Historically, container shipping operated under global shipping conferences, which acted as cartels to suppress freight rate competition and limit entry by maintaining excess capacity. Although shipping conferences were dismantled due to antitrust regulations (e.g., the Shipping Act of 1984 and the Ocean Shipping Reform Act of 1998 in the U.S., and Regulation 1419/2006 in the EU), they were replaced by global strategic alliances. Similar to airline alliances, these alliances feature vessel sharing agreements (VSA), allowing members to share vessel capacity and slots. This arrangement

[14] Industry interviews suggest that most carriers aim to maintain a per-vessel capacity utilization rate between 65% and 85%.

[15] High-sea container shipping refers to intercontinental long-haul shipping. As discussed in Ganapati et al. (2024), the shipping network follows a hub-and-spoke structure, with high-sea shipping encompassing hub-to-hub routes.

enables members to deploy larger vessels, leveraging economies of scale and scope to reduce operational costs. Over the past two decades, the market share of global alliances has risen to 90%, and the number of alliances has consolidated to three, with almost no new independent carriers entering the high-sea shipping market.

2.2 Data

Our empirical analysis is based on two primary datasets. First, we use monthly container trade flow volume and price index data from CTS to estimate the demand for container shipping services. This dataset covers 22 origin-destination subregional trade flows, aggregated from more granular monthly port-to-port shipping manifest information. Additionally, we obtained origin-destination container price indices at the subregional level. A key advantage of this price index is that it reflects the *actual* freight rates paid by shippers to carriers, incorporating both spot freight rates and contractual rates.[16] Because the price index is derived from the *universe* of shipping manifests, it provides a comprehensive and realistic representation of transportation costs by combining spot and contractual rates.[17] Furthermore, the CTS data includes information on the *actual* number of container boxes transported, which is essential for estimating the marginal cost function for carriers, as discussed in the model and estimation sections. For years not covered by CTS, we supplement the analysis with monthly shipping indices from Drewry.

Our second key data source is the weekly capacity and vessel deployment dataset from Sea Intelligence.[18] This dataset provides information on the average vessel size of the fleets operated by different firms and alliances, as well as a unique measure of capacity utilization.

[16] In the container shipping industry, spot freight rates are commonly referred to as “Freight of All Kinds” (FAK), while contractual rates are referred to as “Named Account” (NAC).

[17] Due to legal and regulatory concerns, container shipping freight rate data is rare and difficult to obtain. Excessive disclosure of freight rate information could facilitate tacit collusion or coordinated pricing among carriers. As a result, we rely on an aggregated measure of freight rates, as carrier-specific prices are strictly prohibited from being disclosed by CTS. Wong (2022) also highlights the significant challenges in accessing detailed container shipping freight rate data. However, further insights into carriers’ pricing behavior using more detailed proprietary data are provided in the author’s other work with logistics technology firms.

[18] Sea Intelligence uses this dataset in its flagship TCO reports.

Specifically, it records all canceled sailings (known as "blank sailings") for major trade routes from 2012 to 2020.[19] This measure allows us to observe variations in capacity utilization across carriers during periods of fluctuating aggregate demand. The data reveals a negative correlation between capacity utilization and average vessel size during demand downturns, indicating the cost advantage of operating larger vessels. This unique variation is crucial for identifying the economies of scale in our analysis.

2.3 Stylized Facts

We present four sets of stylized facts in this section that will motivate our modeling, and empirical framework.

2.3.1 Increasing Vessel Size

One of the most notable trends in the container shipping industry over the past decade is the rapid increase in vessel size. As illustrated in Figure 1, the average vessel size for most alliances grew from 10,000 TEUs in 2013 to 18,000 TEUs in 2020.[20] This trend reflects the industry's pursuit of economies of scale (Merk (2018), Imai et al. (2006), Haralambides (2019), Murray (2016), Yang et al. (2011), Veldman (2011)). Larger vessels generally achieve lower per-unit transportation costs because the increases in building, manning, fuel, and other operational expenses are proportionaly smaller than the increase in capacity. As noted by Veldman (2011), while costs at sea decline with vessel size, port-related costs tend to rise, suggesting that economies of scale are more significant on routes with longer sea distances and port infrastructure optimized for mega ships. The extent of these economies of scale

[19] In the shipping industry, blank sailings are primarily used as a capacity-optimization tool, particularly during periods of low demand. However, during the COVID-19 pandemic, blank sailings became more common due to network disruptions caused by port congestion during periods of high demand. For this reason, we avoid using the COVID-19 period to estimate our model.

[20] TEU, or Twenty-foot Equivalent Unit, is the standard unit of measurement in container shipping used to describe the capacity of container ships and terminals. One TEU represents the volume of a standard 20-foot-long container, which is 8 feet wide and approximately 8.5 feet high. For reference, a 40-foot container, commonly used in shipping, is equivalent to 2 TEUs. TEU is widely used in the industry to quantify cargo capacity and facilitate comparisons across vessels and operations.

Figure 1: Increasing vessel size

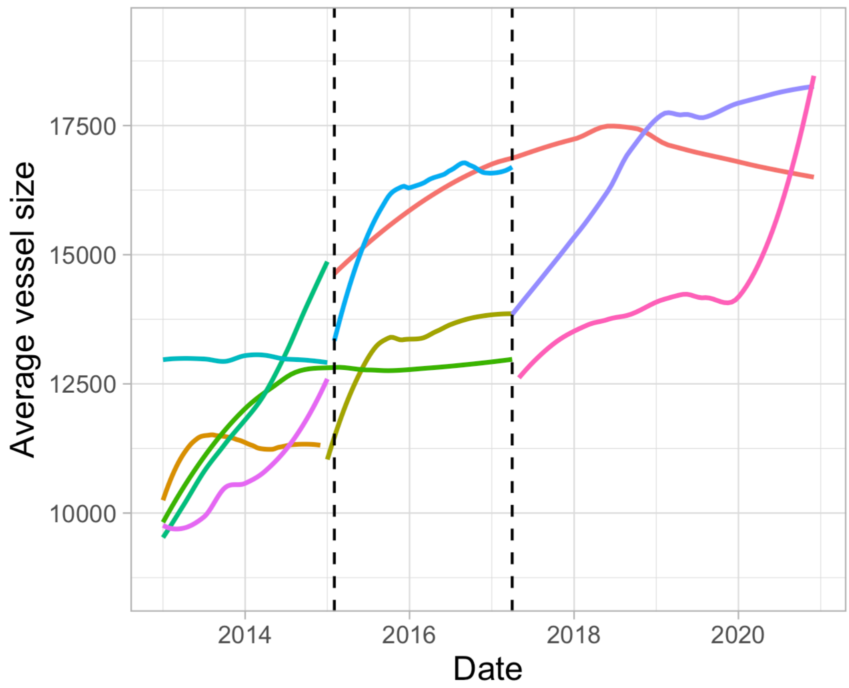

will be quantified later in the paper.

2.3.2 Increasing Consolidation

Another key trend in the container shipping industry is rapid market consolidation. As shown in Figure 2, the number of shipping alliances and independent carriers has declined sharply over the past decade, dropping from eight in 2012 to just three major alliances by 2017. This consolidation is closely tied to, and arguably driven by, developments in shipping technology, particularly the increase in vessel size (UNCT (2022), Merk (2018)). As vessel size and capacity have grown faster than shipping demand, carriers' profit margins have been squeezed. Additionally, the economies of scale associated with larger vessels have altered the cost structure, requiring high fixed costs to maintain global service networks. These shifts have necessitated significant consolidation within the industry to sustain profitability.

Industry consolidation has occurred through exits, mergers and acquisitions, and the formation of cooperative strategic alliances. For instance, Hanjin Shipping filed for bankruptcy

Figure 2: Increasing consolidation

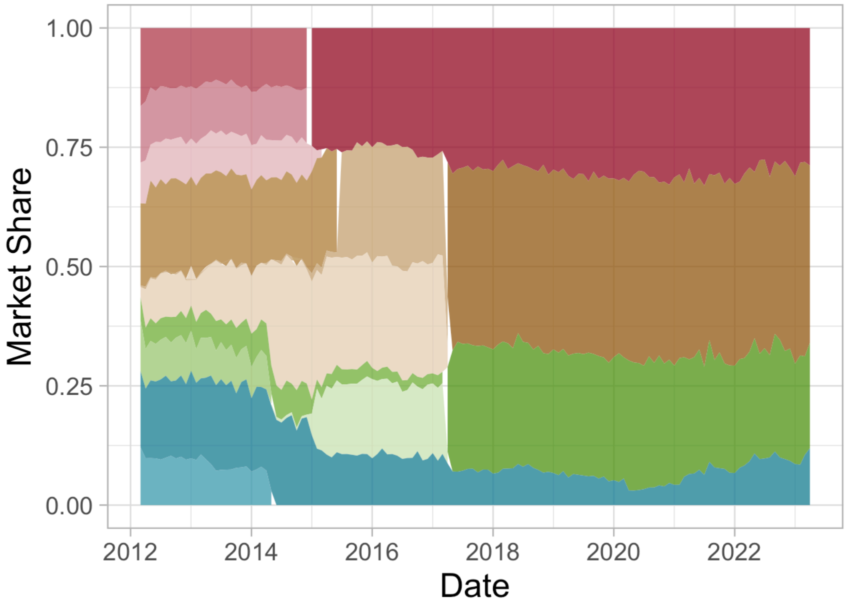

in 2016, and a wave of mergers and acquisitions reshaped the market over the past decade. Larger vessels also prompted carriers to form alliances where members share slot space to optimize network utilization. For a more detailed discussion of market structure and market power, refer to [[]] in the Appendix. However, this paper does not focus on the specific mechanisms of consolidation. Instead, it examines the number of market players the industry can sustain under different technology transition scenarios.

2.3.3 Lower Capacity Utilization For Carriers with Smaller Vessels

Our data offers a unique measure of fleet capacity utilization ratios in the container shipping industry, which operates similarly to the airline or bus transportation systems. In each period (typically quarterly), carriers decide on the number and size of ships to deploy for a given service. However, due to demand volatility, carriers often deploy only a fraction of their total fleet capacity once demand is realized, aiming to stabilize freight rates and maintain profitability. This allows us to identify the fleet and vessel characteristics associated with capacity utilization under fluctuating demand conditions.

Figure 3: Lower capacity utilization for carriers with smaller vessels

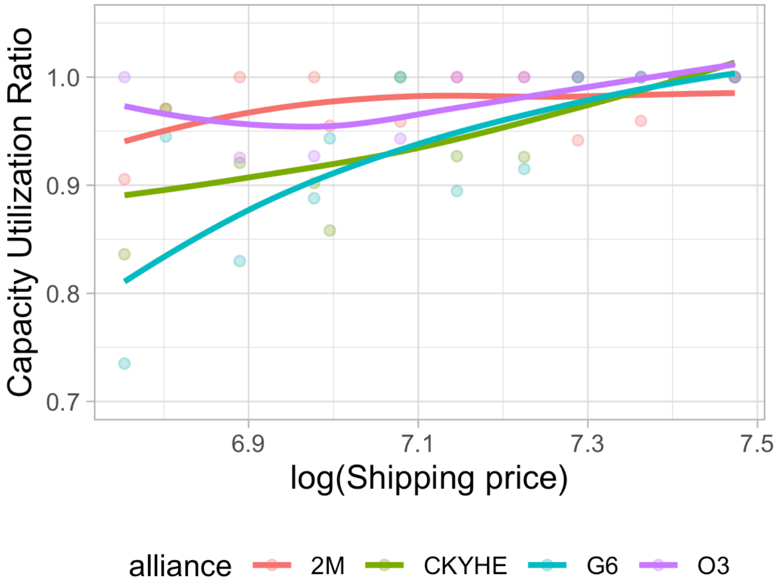

Note: This figure plots the capacity utilization ratio of fleet against the log freight price for 2015. Each dot represents a month-alliance observation. Lines are the LOESS fit.

In 2015, freight rates on Asia–Northern Europe routes faced significant downward pressure due to excess capacity from the introduction of new 18,000 TEU-class vessels and weak demand. During this downturn, we observed that alliances with smaller vessels reduced their capacity utilization significantly more than those with larger vessels. In Figure 3, we plot the monthly capacity utilization ratio against the log freight price for four alliances during the 2015 demand downturn. The data show that CKYHE and G6 alliances (blue and green lines) reduced their capacity utilization ratios more drastically compared to 2M and O3 alliances. As illustrated in Figure **??**, 2M and O3 operated larger vessels on average than CKYHE and G6. This negative correlation between average vessel size and capacity utilization during the demand downturn suggests that carriers with larger vessels reduced capacity utilization less, leveraging the cost advantages of larger ships through economies of scale. This variation in capacity utilization serves as a key source of identification for estimating the economies of scale associated with ship size, as detailed in Section 4.2.

2.3.4 Panama Canal expansion leads to larger vessel and less players

To further illustrate the positive relationship between market concentration and the increase in vessel size, we analyze the Panama Canal expansion in 2016 as an event study. The expansion project significantly enhanced the canal's capacity by allowing larger vessels to transit. Prior to the expansion, the canal accommodated only Panamax ships with a maximum capacity of approximately 5,000 TEUs. The project, completed in mid-2016, added a third set of locks, widened and deepened navigation channels, and upgraded infrastructure, enabling the passage of New Panamax (Neo-Panamax) ships with capacities of up to 14,000 TEUs. This more than doubled the canal's cargo capacity, facilitating the use of larger and more efficient vessels on this critical maritime route.

As shown in the top panel of Figure 4, vessel sizes operated by alliance and non-alliance carriers were comparable before the expansion, with both groups largely capped at 5,000 TEUs. Following the commencement of the third lock in mid-2016, alliance carriers began deploying significantly larger vessels on this tradelane compared to non-alliance carriers, which primarily consisted of smaller independent operators. Alliance carriers had the advantage of larger fleets, enabling them to quickly relocate bigger vessels from other routes (e.g., Asia-Europe) to the Asia–North America East Coast route, capitalizing on the cost efficiencies of larger vessels made possible by the expansion. For further details on vessel reallocation across markets, refer to the appendix [[]]. The lower panel in Figure 4 illustrates the impact on market structure: non-alliance carriers rapidly lost market share after the expansion (marked by the dotted black line), and by the second half of 2018, the tradelane was entirely dominated by alliance carriers. This case demonstrates how an exogenous increase in vessel size can lead to market consolidation.

The Panama Canal expansion illustrates the core policy question of this paper: how changes in the frontiers of shipping technology influence both technology adoption and market concentration. This case highlights that the frontier of shipping technology depends not only on the potential vessel size but also the maritime infrastructure. Large vessels achieve their cost advantages only when supported by modernized ports and canals. However, as demonstrated by the Panama expansion, these changes in technology frontier also drive increases in vessel size and market concentration. In our counterfactual analysis, we explore scenarios with varying economies of scale to identify the "optimal" level of scale for shipping technology. In the context of the Panama expansion, this translates to determining the "optimal" expansion size.

3 Model

We developed a non-stationary dynamic oligopoly model of the container shipping industry. In the stage game, carriers compete in a quantity-setting framework, incorporating economies of scale. Additionally, carriers make dynamic investment decisions to increase their vessel sizes, capturing the strategic interactions involved in such investments.

3.1 Environment and Timeline

The model is non-stationary and finite horizon as we are interested in both the long-run equilibrium and transition paths under different technology innovation trajectory. Similar to the quality ladder literature, I discretize the average vessel size of the fleet into finite number of levels: $\kappa_1, \kappa_2, ..., \kappa_K$, and state variables are

$$\mathbf{s}_t = \{N_t(\kappa_1), ..., N_t(\kappa_K); D_t\}$$

where

- $N_t(\kappa_1)$ is the number of carriers in the vessel size class κ_1[21]

[21] We treat strategic alliance as a single carrier/firm. Even though the arrangement within strategic alliances are mostly on the supply side, and they are not allowed for joint marketing or pricing. However, the VSA allows the alliance members to perfectly coordinate on their capacities. Within the Cournot game framework in our stage game, this indicates carriers within the same alliance could be modeled as a single entity. Asturias (2020) also made the same assumption in modeling the oligopoly in the container shipping industry.

Figure 4: Panama Expansion: Vessel Size and Market Share

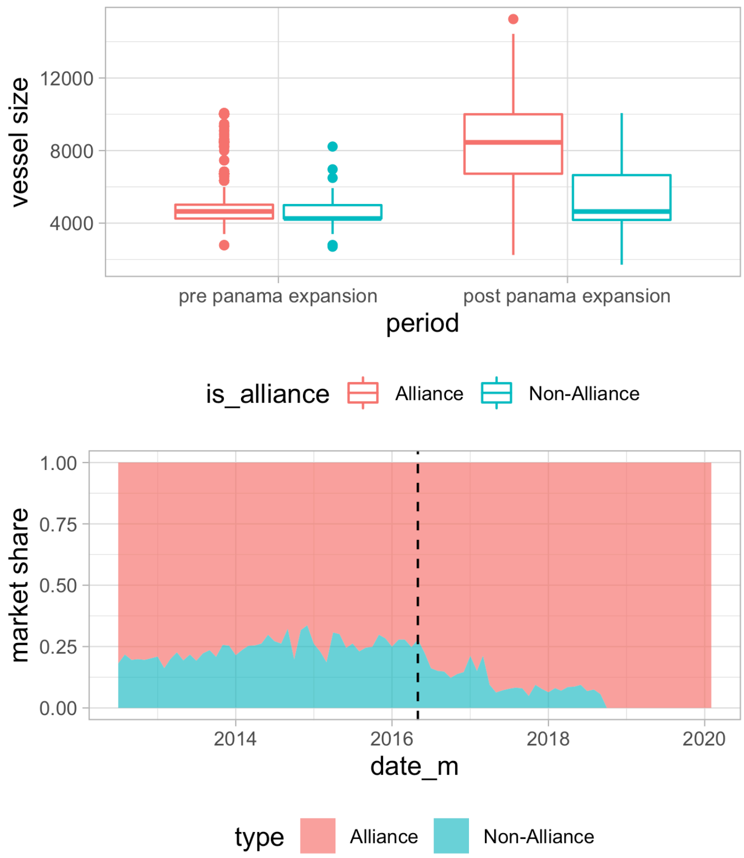

Note: This figure presents the vessel size distribution and market share of alliance and non-alliance carriers. Alliance carriers generally operate significantly larger fleets composed of bigger vessels globally, which explains their ability to deploy New-Panamax class ships from other markets. The data reflect vessel size and market share for all carriers on the Asia–North America East Coast tradelane, including certain routes transiting westbound through the Malacca Strait and Europe. This accounts for the outliers in the vessel size distribution chart prior to the Panama Canal expansion, as no vessel exceeding 5,000 TEUs should have been able to traverse the canal at that time. The dotted line in the bottom panel marks April 2016, when the third lock of the Panama Canal became operational.

- D_t is the aggregate demand shifter[22]

Industry start at the state $\mathbf{s}_t$. Nature will randomly pick a mover in this period. The mover will make a dynamic decision between $\{invest, idle, exit\}$. The mover is restricted to increase their average vessel size by one level up. Then all incumbents engage in quantity-setting stage game. We assume there exists one potential entrant in each period, and it makes entry decision after the stage game. The mover makes its dynamic action and the state evolves to the next period.

3.2 Stage Game

We assume carriers face a constant elasticity demand over aggregate quantity, and the demand function take a log-linear form:

$$\log(Q_{odt}) = \log(D_{odt}) - \sigma \log(P_{odt}) \tag{1}$$

where Q_{odt} is the total container volume on route od at time t, and P_{odt} is the per-TEU freight price. D_{odt} is the route-time specific demand shifter and $-\sigma$ is the price elasticity of demand which will be estimated later.

The cost function of each carrier has a variable cost and fixed cost component:

$$C(\rho_t^i; \kappa_t^i). = \underbrace{c(q_t^i, \tilde{\kappa}(\rho_t^i, \kappa_t^i))}_{\text{variable cost}} + \underbrace{f \cdot \rho_t^i \cdot n_t^i \cdot (\kappa_t^i)^\alpha}_{\text{fixed cost}} \tag{2}$$

In each period, we assume carrier i cannot change its average vessel size κ^i. However, due to the volatility in demand, carrier would be able to change their sailing frequency $\rho^i \in [0, 1]$.

[22]We assume the demand shifter follows a distribution: $D_t \sim \Psi(\mu_D, \sigma_D)$. To simplify computation burden, we also assume the moving carrier will only calculate the value function based on the mean of the demand shifter: μ_D, implying the long-run capacity management is based on the average demand. However, the capacity utilization or blank sailing will depend on the realization of D_t, implying the variance of the demand shifter will influence carriers' short term capacity management strategy.

The effective capacity supplied by carrier i in period t is

$$\tilde{\kappa}(\rho_t^i, \kappa_t^i) = \rho_t^i \kappa_t^i n_t^i \tag{3}$$

where n_t^i is the number of scheduled sailing in period t[23]. Following Ryan (2012), we construct a linearly increasing marginal cost curve for carrier i

$$mc\left(q_t^i, \kappa_t^i\right) = \begin{cases} c + b\frac{q_t^i,}{\tilde{\kappa}(\rho_t^i,\kappa_t^i)} & \text{if } q_t^i \leq \tilde{\kappa}(\rho_t^i, \kappa_t^i) \\ \infty & \text{otherwise.} \end{cases} \tag{4}$$

This variable cost function specification assumes the marginal cost increases as the firm gets closer to its effective capacity. Another convenient property of this marginal cost assumption is the aggregate supply function in the market could be represented as

$$mc(Q_t, \tilde{K}_t) = c + b\frac{Q_t}{\tilde{K}_t} \tag{5}$$

where $\tilde{K}_t = \sum_i \tilde{\kappa}(\rho_t^i, \kappa_t^i)$. The equilibrium price is determined by the intersection of demand curve in Equation 1 aggregate supply curve in Equation 5:

$$P_t = c + b\frac{Q_t}{\tilde{K}_t} \tag{6}$$

This equilibrium condition states that the price is determined by the total utilization rate of the effective capacity: $\frac{Q_t}{\tilde{K}_t}$. Since we have data for P_t, Q_t and $\tilde{K}_t$, we would be able to estimate the key marginal cost function parameters b and c using a linear regression.

The fixed cost component in the carrier's total function (Equation 2) is $f \cdot \rho_t^i \cdot n_t^i \cdot (\kappa_t^i)^\alpha$. The fixed cost will increase as the average vessel size (κ_t^i) increases, but not in proportion as the size increase. The economies of scale of larger average vessel size comes in mostly

[23]We treat n_t^i as exogenous in the stage game because the number of service and scheduled weekly sailing are hardly changed even as the market consolidates and the vessel size gets bigger.

through this fixed cost component because the marginal cost will be the same for all carriers in equilibrium as determined in Equation 5.[24] This is a very reasonable depiction of the cost structure of container shipping in reality, as the economies of scale of larger ships comes through the fact that the fixed cost component of operating a ship will not increase in proportion with the vessel size.

In stage game, carrier i tries to maximize its per-period profit by choosing their sailing frequency:

$$\pi_t^i(\kappa_t^i, \mathbf{s}_t) = \max_{\rho_t^i} P_t q_t^i - C(\rho_t^i; \kappa_t^i) \tag{7}$$

where $\mathbf{s}_t = (\{\kappa^i\}, D_t)$ is the state variable which include the average vessel size of all carriers and the demand state. Note that the fixed cost of operation exhibits economies of scale for parameters $\alpha < 1$ which means the per-unit cost of fixed cost of operation is lower for bigger ships. We use a fixed point algorithm to solve for the equilibrium of stage game.

3.3 Investment Game

In each period, Nature will randomly pick a mover from incumbents. For example, if we have N_t carriers in period t, the probability of a certain carrier getting picked is $\frac{1}{N_t}$. Only the mover would be able to make a dynamic decision between $\{idle, upgrade, exit\}$ by solving a discrete choice problem as in Equation 4.3. It will draw a private idiosyncratic shock for each potential action: $\{\varepsilon^{upgrade}, \varepsilon^{idle}, \varepsilon^{exit}\}$ from a Type-I extreme value distribution. If the mover chooses to upgrade, it can only increase its average vessel size by one level. And the

[24]This is a similar modeling strategy as in Jeon (2022). However, even the marginal cost is the same for all carriers, carrier with a larger average vessel size will supply more quantities as the *individual* slope of the marginal cost function will be flatter for larger carriers, leading them to sell more quantities in the equilibrium.

upgrade/investment cost it needs to pay is $\left(\iota \cdot (\kappa_{t+1}) + \varepsilon^{\text{upgrade}}\right)$.[25]

$$V_{i,t}^{I} = \max\left\{\underbrace{-\left(\iota \cdot (\kappa_{t+1}) + \varepsilon^{\text{upgrade}}\right) + \Lambda_{i,t+1}\left(\boldsymbol{s}\left(t_{t+1}\right) \mid \boldsymbol{s}\left(t_{t}\right), \text{upgrade}\right)}_{\text{upgrade}}, \right.$$
$$\underbrace{-\varepsilon^{idle} + \Lambda_{i,t+1}\left(\boldsymbol{s}_{t+1} \mid \boldsymbol{s}_{t}, idle\right)}_{\text{stay idle}},$$
$$\left. \underbrace{\phi_{i,t}^{exit} - \varepsilon^{\text{exit}}}_{\text{exit}} \right\}.$$

where $\phi^{exit}(\kappa)$ is the scrap value.[26] $\Lambda_{t+1}(\boldsymbol{s}_{t+1})$ is the *expected* value function in the next period before the Nature picks a mover. It's formulated as the average of expected value functions weight by the probability each carrier will be chosen as the mover:

$$\Lambda_{i,t+1}(\boldsymbol{s}_{t+1}) = Pr(\text{i is chosen}) \cdot \mathbb{E}V_{t+1}^{I}(\boldsymbol{s}_{t+1}) + \sum_{j \neq i} Pr(\text{j is chosen}) \cdot W_{i,t+1}(\boldsymbol{s}_{t+1}; \text{j is the mover})$$

where $W_{t+1}(\boldsymbol{s}_{t+1}; \text{j is the mover})$ is the expected value function of firm i where another firm

[25] We assume the deterministic part of the investment to be proportional to the average vessel size. This is because it will cost more to upgrade a fleet from the average size of 16,000 TEU to 18,000 TEU than from the average size of 14,000 TEU to 16,000 TEU.

[26] Different from our model specification here, we model the exit process differently when it comes to estimation. Rather than directly estimate the scrap value, we assume a per-period maintenance cost $m(\kappa) = m \cdot \kappa$ which is proportional to the average vessel size. Carriers need to incur this maintenance cost as long as they are in operation. This is similar to the idea of the adjustment cost in the macroeconomic literature. Therefore the actual dynamic equation we are estimating is:

$$V_{i,t}^{I} = \max\left\{\underbrace{-\left(\iota \cdot (\kappa_{t+1}) + \varepsilon^{\text{upgrade}}\right) - m \cdot \kappa_{t} + \Lambda_{i,t+1}\left(\boldsymbol{s}\left(t_{t+1}\right) \mid \boldsymbol{s}\left(t_{t}\right), \text{upgrade}\right)}_{\text{upgrade}}, \right.$$
$$\underbrace{-\varepsilon^{idle} - m \cdot \kappa_{t} + \Lambda_{i,t+1}\left(\boldsymbol{s}_{t+1} \mid \boldsymbol{s}_{t}, idle\right)}_{\text{stay idle}},$$
$$\left. \underbrace{0 - \varepsilon^{\text{exit}}}_{\text{exit}} \right\}.$$

we make this adjustment as this makes the estimation easier to converge. On the other hand, the scrap value is just the NPV of the future maintenance cost.

j is chosen to be the mover:

$$\begin{aligned} W_{i,t+1}(\boldsymbol{s}_{t+1}; \text{j is the mover}) = & Pr(\text{j expands}) \cdot \Lambda_{i,t+1}(\boldsymbol{s}_{t+1}; \text{j expands}) + \\ & Pr(\text{j idles}) \cdot \Lambda_{i,t+1}(\boldsymbol{s}_{t+1}; \text{j idles}) + \\ & Pr(\text{j exits}) \cdot \Lambda_{i,t+1}(\boldsymbol{s}_{t+1}; \text{j exits}) \end{aligned}$$

3.4 Entrants' Problem

We assume there exists one potential entrants in each period. The potential will draw a private idiosyncratic cost shocks for both of its potential action of entering and staying out. After observing the private cost shocks, the potential entrant decides whether to enter by solving the following problem :

$$V^E = \max \left\{ - \left(\zeta^E + \varepsilon^{\text{enter}} \right) + \Lambda \left(s \left(t' \right) \mid s \left(t \right); \text{ enter } \right), -\varepsilon^{\text{stay out}} \right\}$$

The potential entrant can only start from the lowest state if it chooses to enter.

3.5 Equilibrium

To close the model, we assume the state stops evolving after T, and all carriers continue to receive the stage game payoff in perpetuity for t beyond T. The terminal value is then the net present value of the stage game payoff:

$$\Lambda_i \left(\boldsymbol{s}_t \right) = \frac{\pi_i \left(\boldsymbol{s}_T \right)}{1 - \beta}$$

The strategies played by carrier in Equilibrium are type-symmetric and we can solve this finite horizon, random mover model by backward induction. Given our setup of a random mover picked by Nature in each period, this simplify our model solution to a single agent discrete choice problem in each period.

4 Estimation and Empirical Results

4.1 Estimating Demand

Given our log-linear demand system, the specification we use is the following:

$$\ln Q_{odt} = -\sigma \ln P_{odt} + \gamma_{ot} + \gamma_{dt} + \gamma_{od} + \varepsilon_{odt} \tag{8}$$

where o, d, t represent origin, destination, and month, respectively. Q_{odt} denotes the TEU volume on tradelane od at time t, and P_{odt} represents the per-TEU container freight price. We include fixed effects for origin-time (γ_{ot}), destination-time (γ_{dt}), and origin-destination pairs (γ_{od}). However, there exists an issue of price endogeneity, as certain origin-destination-time-specific factors in ε_{odt} may simultaneously influence both freight price P_{odt} and container volume Q_{odt}. To address this, we employ an instrumental variable framework.

In late December 2023, the Houthi group began attacking ships passing through the Suez Canal, prompting carriers to gradually divert their vessels to the Cape of Good Hope (see Figure 5). This diversion significantly increased shipping prices on origin-destination routes previously reliant on the Suez Canal (see Figure 6 for a comparison of price dynamics between routes affected and unaffected by the attack). Leveraging proprietary monthly data on container trade volumes and prices across 22 origin-destination pairs from CTS, we observe substantial increases in freight prices on the affected routes. To address the price endogeneity issue, we use the recent Red Sea Crisis as a supply-side instrument for price:

$$\ln P_{odt} = \mathbb{1}\left[\text{o-d route is affected at t}\right] + \gamma_{ot} + \gamma_{dt} + \gamma_{od} + \epsilon_{odt} \tag{9}$$

The identification of our demand estimation relies on the time-series and cross-sectional variations across the 22 main trading routes in our data, under the assumption of constant price elasticity of demand.

We present our demand estimation results in Table 4. Our analysis estimates the price

Figure 5: Red Sea Crisis

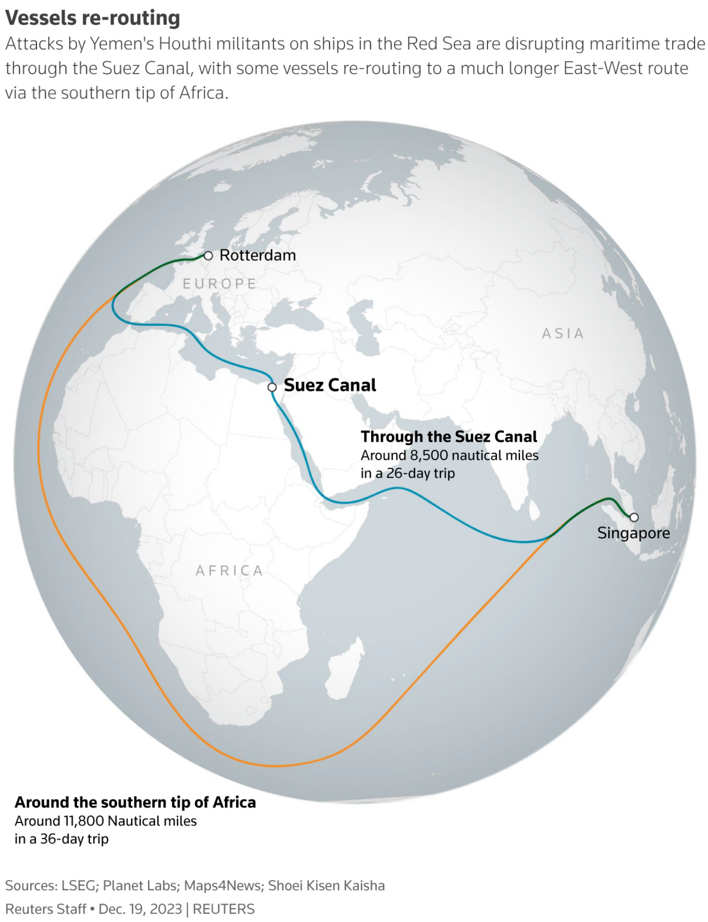

Figure 6: Shipping price for routes affected and not affected by the Houthi's attack

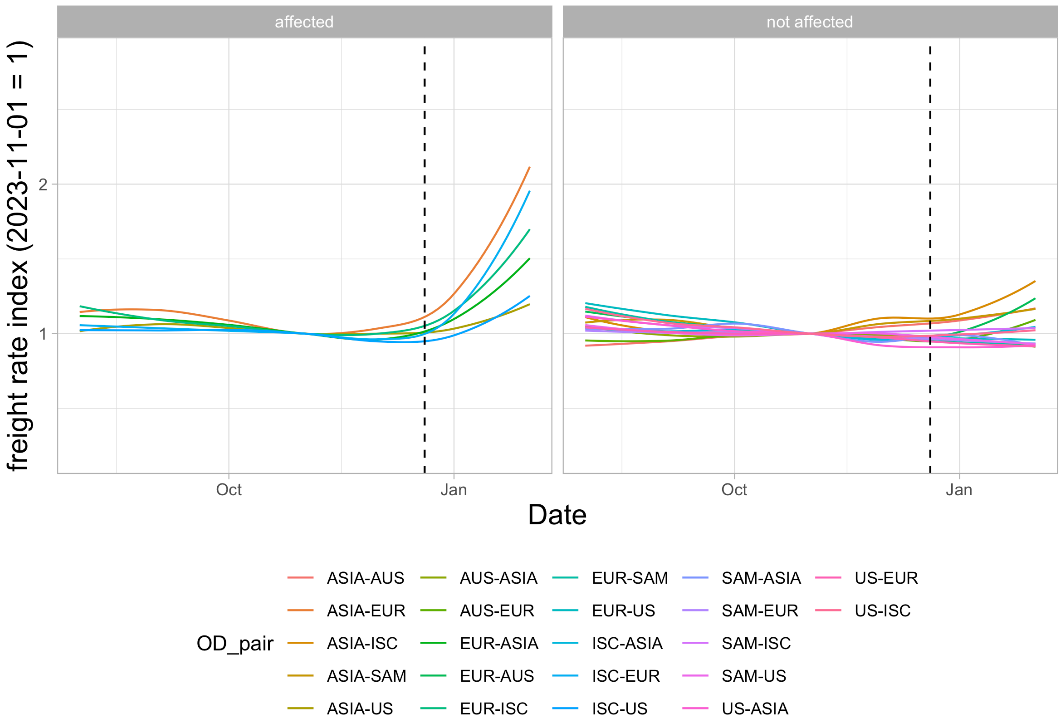

elasticity of container shipping demand to be -1.22, meaning a 10% increase in freight prices would result in a 12.2% decline in total volume, on average.[27] This elasticity is relatively low compared to estimates in the existing literature. For instance, Kalouptsidi (2014) used ship size and age as instruments for price and estimated the demand elasticity for bulk shipping at -6.17. Similarly, Jeon (2022) applied a comparable strategy and estimated container shipping demand elasticity at -3.89. Wong (2022) employed the round-trip effect as an instrument and found an elasticity of -3. Asturias (2020) used population as an instrument, estimating an elasticity of -5. More recently, Otani (2024) employed a similar approach to Jeon (2022) and estimated an elasticity of -0.89 for the 1966–1990 period, attributing this low value to the limited availability of alternative transportation methods during that era.

Our elasticity estimate is also on the lower end compared to these studies. One pos-

[27]For comparison, we also show the results of a hedonic price regression in column (3) in Table 4. As we can see from the results, the price-elasticity will not be significantly from 0 if we do not employ our instrumental variable technics.

Table 1: Demand Estimation

	IV	IV	OLS
	Stage 1	Stage 2	
	$\ln P_{odt}$	$\ln Q_{odt}$	$\ln Q_{odt}$
disruption dummy	0.120^{*}		
	(0.059)		
P_{odt}		-1.183^{**}	-0.140
		(0.558)	(0.128)
Obs	88	88	88
origin × month	X	X	X
origin × destination	X	X	X
destination × month	X	X	X
Adjusted R^2	0.849	0.998	0.996
F-stat	8.74	-	-
Note:	*p<0.1; **p<0.05; ***p<0.01		

sible explanation is that our estimate reflects short-run demand elasticity, where the lack of alternative transportation modes constrains substitution. When combined with additive shipping costs, our results align with a trade elasticity in the range of 7–8. With estimted demand elasticity σ, we can then estimate the demand shifter D_{odt}. For further details on the demand estimation process, please refer to Appendix A.

4.2 Estimating Economies of Scale

We estimate the parameters in per-period profit function in two steps. In the first step, we use OLS regression to recover the marginal cost function parameter (b, c) in Equation 6 and results are summarized in Figure 18 and Table 5 in Appendix B.

$$P_t = c + b \times \frac{Q_t}{\tilde{K}_t} + \epsilon_t \tag{10}$$

Then in order to recover the parameters in fixed cost, we search over the grid for parameter (f, α) to minimize the sum of squared difference between the sailing frequency observed in data and the sailing frequency simulated in our stage game equilibrium:

$$\min_{f,\alpha} \sum_{i,t} (\rho_t^i - \hat{\rho}_t^i)^2 \tag{11}$$

where $\hat{\rho}_t^i$ is obtained by solving the system of non-linear equation in Equation 7 given any guesses of (f, α). Our estimation results are in Table 2. We also showcase our in-sample and out-of-sample model fit in Appendix C.

Table 2: Cost Function Estimation

b	c	f	α
2.20	-0.8	1.36	0.66

4.3 Estimating Dynamic Parameters

There are four dynamic parameters we need to estimate: investment cost, maintenance cost, entry cost and the scale of the private idiosyncratic errors. The investment cost to upgrade the average vessel size to the next level is proportional to the average vessel size of the fleet: $\iota(\kappa) = \iota \cdot \kappa$. We also estimated the maintenance cost $m(\kappa)$ which all active carriers will incur as long as they are in operation. Exist and consolidation will help to identify the maintenance cost. Since there isn't entry in our data sample, we cannot point identify the entry cost. However, by revealed preference, we can use the estimated expected value to provide a lower bound for the entry cost. Finally we assume the private idiosyncratic cost shock for investment, maintenance and entry cost are all drawn from the same Type-I Extreme Value distribution: $\varepsilon^{\text{idle, upgrade, entry, exit}} \sim T1EV(\sigma \cdot \kappa)$. Note that we scale the variance of the distribution by the average vessel size.

Thanks to the assumption that the private idiosyncratic cost shocks are drawn from a T1EV distribution, we can right the probability of various dynamic action for a mover as an analytical expression of the expected value functions. For example, if carrier i with the current average vessel size of κ is chosen as the mover, then its probability to invest to upgrade the vessel size is

$$Pr(upgrade|\mathbf{s}_t) = \frac{\exp(V_{\kappa,t}^{I,upgrade}/\sigma^{\kappa})}{\exp(V_{\kappa,t}^{I,idle}/\sigma^{\kappa}) + \exp(V_{\kappa,t}^{I,upgrade}/\sigma^{\kappa}) + \exp(V_{\kappa,t}^{I,exit}/\sigma^{\kappa})} \tag{12}$$

where $V_{\kappa,t}^{I,upgrade}$ is the value function if carrier i chooses to upgrade its fleet vessel size as in Equation 4.3.

$$\begin{aligned}
V_{\kappa,t}^{I,upgrade} =& -\left(\iota \cdot (\kappa_{t+1}) + \varepsilon^{\text{upgrade}}\right) - m \cdot \kappa_t + \Lambda_{i,t+1}\left(\boldsymbol{s}\left(t_{t+1}\right) \mid \boldsymbol{s}\left(t_t\right), \text{upgrade}\right) \\
V_{\kappa,t}^{I,idle} =& -\varepsilon^{idle} - m \cdot \kappa_t + \Lambda_{i,t+1}\left(\boldsymbol{s}_{t+1} \mid \boldsymbol{s}_t, idle\right) \\
V_{\kappa,t}^{I,idle} =& 0 - \varepsilon^{\text{exit}}
\end{aligned}$$

And similarly, we could calculate the probability for potential entrant to enter

$$Pr(enter|\mathbf{s}_t) = \frac{\exp(V_{\kappa_0,t}^{I,enter}/\sigma^{\kappa_0})}{\exp(V_{\kappa_0,t}^{I,enter}/\sigma^{\kappa_0}) + 1} \tag{13}$$

where the payoff function for entering is defined as

$$V_{\kappa_0,t}^{I,enter} = \zeta + \varepsilon^{\text{enter}} + \Lambda_{i,t+1}(\boldsymbol{s}(t_{t+1}) \mid \boldsymbol{s}(t_t), \text{enter})$$

We use a full solution approach through method of simulated moments to estimate the dynamic parameters by minimizing the transition path of average vessel size in data and that generated by our model. With a guess of the dynamic parameters, we start the estimation using a backward induction approach. We assume the state stops evolving after $T = 100$ months, and all carriers continue to receive the stage game payoff in perpetuity beyond T. This provides us with a terminal payoff estimates $\Lambda_i(\mathbf{s}_T)$. Then we could solve the dynamic discrete choice problem for period $T-1$ all the way to the initial state. This solution will provide us with a time-specific estimates of the upgrading/idling/exit/entry probability for

Table 3: Dynamic parameter estimates

ι	m	ζ[31]	σ
9.25	0.675	280	5.6

each potential state. Then we start simulate the evolution of the industry state using these estimated transition probability. We perform a grid search over the parameter space for our dynamic parameters to find the one minimizing the distance between simulated industry state transition path and those observed in the data. More specifically, we focus on two moments: i) the industry mean of average vessel size, and ii) number of firms.[28]

We summarized the estimated dynamic parameters in Table 3. Just to have a reality check of our estimates, our estimates of the investment/upgrade cost $\iota = 9.25$ implies the price of a 20,000 TEU container ship will be price at 277.5 million USD.[29] It falls within the range of 150-300 million USD of the same ultra-large container ship price in the ship building industry.[30] For more details on the dynamic model fit, please refer to the Appendix D.

4.3.1 Market Structure and Investment Incentive

The estimated policy function provides insights on how market structure influences investment and innovation incentives. Figure 7 illustrates the probability of upgrading for a moving carrier[32] under various market structures and industry states.

First, when the moving carrier leads the industry in average vessel size (depicted in the left region of the figure), the incentive to upgrade increases with the number of firms (as indicated

[28]Most literature(for example, Igami (2017); Igami and Uetake (2020)) estimates the dynamic parameter using a maximum-likelihood estimation approach to maximize the likelihood of the observed state transition. However, a unique challenge of my setting is the scarcity of the consolidation, which makes the MLE approach not feasible and finalize on the method of simulated moments.

[29]This is calculated as 9.25 per TEU*20,000TEU*1500 USD = 277.5 million USD. 1500 USD is the average freight rate on the Asia-Northern Europe route.

[30]Source: https://gegcalculators.com/whats-the-cost-of-a-normal-cargo-ship/?utm_source=chatgpt.com

[32]A moving carrier is the firm chosen by 'Nature' to make a dynamic decision in each period. Here, we assume the mover currently operates vessels with an average size of 18,000 TEU.

by the blue and green lines being above the red line). Conversely, this incentive diminishes as the mover transitions from a leader to a follower. This suggests that more competitive market structures amplify the *dynamic* incentive to invest or innovate for market leaders but not for followers. The rationale is that in a competitive market, leaders face a stronger business-stealing incentive: by investing or innovating, they can pressure weaker competitors to exit. This creates a dynamic business-stealing incentive for leaders. However, in more concentrated markets, innovation provides insufficient cost advantages to drive competitors out, significantly dampening the incentive to invest (as seen in the near-zero probability for leaders in a 3-player market, represented by red dots in Figure 7). For followers, the opposite holds true. In concentrated markets, the larger share of the market for each player intensifies the business-stealing effect, enhancing the incentive to invest.[33]

Second, the incentive to invest or innovate exhibits an inverted U-shape with respect to the industry's average state. Market leaders and followers have a larger incentive to invest when they are closer to the industry average. For market leaders, the incentive diminishes as their vessel size diverges further from the industry average, reflecting reduced benefits of pulling even further ahead. Conversely, for followers, the incentive to invest is strongest when they are closer to the industry average or leaders, as the potential to catch up is more achievable.

Third, market followers generally exhibit a greater incentive to invest than leaders. As shown in Figure 7, the probability of investment is higher on average in the right region of the figure, where the firm is a follower. This aligns with the intuition that investing to catch up is often less costly than investing to stay ahead. Consequently, this dynamic limits asymmetries among firms, as carriers tend to grow their vessel sizes at relatively similar rates. This pattern is corroborated by the data, where carriers have been observed expanding their vessel sizes at comparable paces.

[33] Further explanation is needed to reconcile this finding with existing literature... [[]]

Figure 7: Probability to Invest: Market Structure

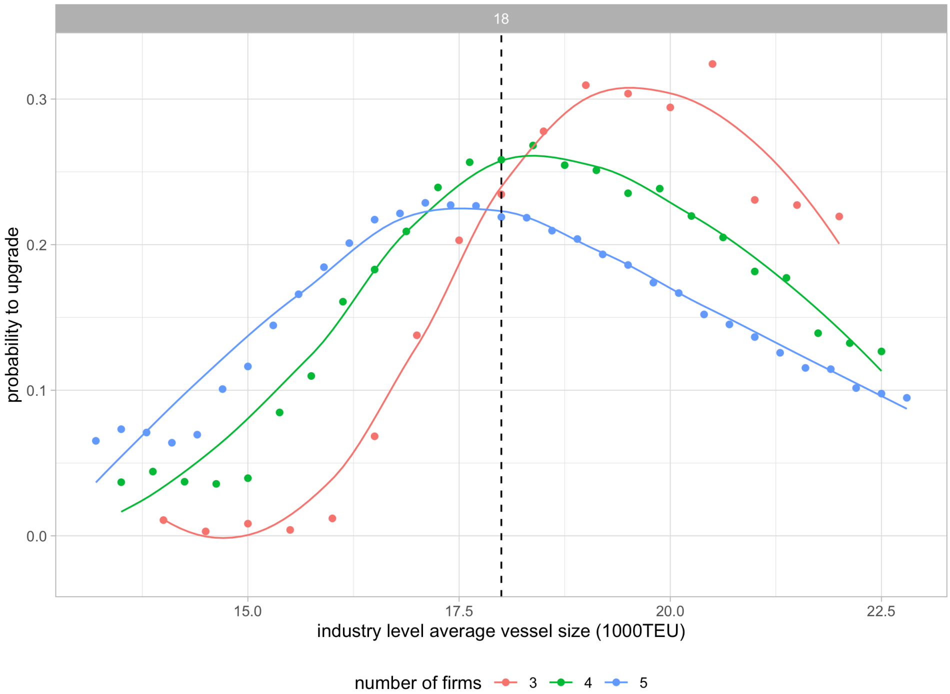

Note: This figure plots the probability to invest or to upgrade to the next level of vessel size for carrier chosen to move. Each dot represents the probability for a specific state (we reduced the dimension of the industry state by calculating the *average* vessel size in the industry). Since the current state of the mover is 18,000 TEU, represented by the vertical dotted black line. The left region of the figure represents the state where the mover is a market leader where its average vessel size is larger than the industry average. The right half represents the case where the mover is a market follower. Different color represents different market structures (#firm = 3, 4, 5). We use the LOESS with a span of 0.1 to plot the smoothed policy function.

4.3.2 Aggregate Demand and Investment Incentive

In Figure 8, we present the probability of investment across various industry states under different expected aggregate demand scenarios, as derived from our dynamic estimation.[34] As expected, the probability of upgrading average vessel size is higher when aggregate demand is greater (Figure 8, with red and blue lines above the green line in the right half). However, higher demand predominantly increases the investment incentives for market followers, while the incentives for market leaders are comparatively lower.

This asymmetry can be attributed to the dynamics of business-stealing potential. Under higher aggregate demand, it becomes significantly more challenging for market leaders to force competitors out of the market, thereby limiting their dynamic incentive to invest. Conversely, for market followers, the opportunity to gain market share through investment is amplified under higher demand, resulting in a stronger incentive to invest.

To summarize the findings from the dynamic estimation, we find that, on average, the incentive to invest is higher for market followers and increases with aggregate demand. However, this aggregate trend masks significant differences in how market structure and demand influence investment incentives for market leaders versus followers. In more competitive and challenging environments—characterized by more competitive market structures and lower demand—market leaders exhibit stronger incentives to invest, driven by a heightened business-stealing motive to push weaker competitors out. In contrast, more concentrated markets and higher aggregate demand levels enhance the investment incentives for followers, as the potential profit gains from catching up become more substantial. These dynamics suggest that industry-wide technology adoption accelerates in the early stages when market is more competitive, led by leaders striving to maintain their advantage, but followers eventually close the gap more quickly as the market consolidates.

[34] The middle-level demand corresponds to the average demand shifter estimated from the data, while high-level demand assumes a 25% increase in the demand shifter, and low-level demand assumes a 25% decrease.

Figure 8: Probability to Invest: Aggregated Demand

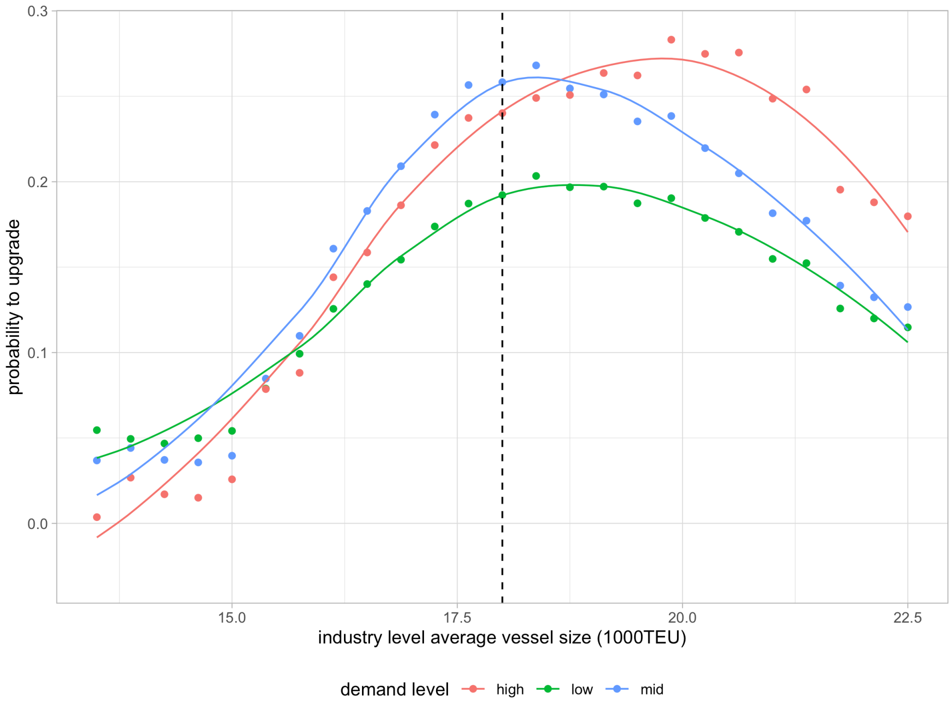

Note: This figure plots the probability to invest or to upgrade to the next level of vessel size for carrier chosen to move. Each dot represents the probability for a specific state (we reduced the dimension of the industry state by calculating the *average* vessel size in the industry). Since the current state of the mover is 18,000 TEU, represented by the vertical dotted black line. The left region of the figure represents the state where the mover is a market leader where its average vessel size is larger than the industry average. The right half represents the case where the mover is a market follower. Different color represents different market structures (#firm = 3, 4, 5). We use the LOESS with a span of 0.1 to plot the smoothed policy function.

5 Counterfactual Analysis

In this counterfactual analysis, we investigate how the progression of shipping technology (measured by average vessel size) and market structure (represented by the number of market players) would vary under different scenarios of i) economies of scale and ii) expected aggregate demand over a 10-year period.

5.1 Maximum Vessel Size

In this counterfactual exercise, we examine scenarios where the maximum possible average vessel size varies. Specifically, we consider cases where the maximum average vessel size is capped at 18,000 TEU, 20,000 TEU, and 24,000 TEU. Our analysis focuses on both the long-run equilibrium and the transition path under each scenario. For each case, we simulate the evolution of the number of players and average vessel size (Figure 9), as well as the median freight price (Figure 10) over a 15-year period (180 months), while keeping expected demand and innovation step size constant.[35]

In the long-run equilibrium, surviving carriers upgrade their fleets to the maximum vessel size across all scenarios. However, when the maximum vessel size reaches 24,000 TEU, the market consolidates into a duopoly.[36] In contrast, the scenarios with maximum vessel sizes of 18,000 TEU and 20,000 TEU stabilize with three firms in the long run (Figure 9, top panel). Regarding freight prices, the scenario with a 20,000 TEU cap results in the lowest median freight price in the long run (Figure 10, green line), as this scenario achieves the highest total market capacity (Figure 11, green line). When the maximum vessel size is too large (24,000 TEU), market consolidation into a duopoly reduces total capacity (Figure 11, blue line). Conversely, when the maximum vessel size is too small (18,000 TEU), the market stabilizes with three firms, but the limited vessel capacity prevents the market from reaching

[35] The step size is assumed to be 2,000 TEU.

[36] Currently, the largest container vessel, MSC Irina, has a capacity of 24,346 TEU. While the average vessel size on the Asia-Northern Europe tradelane is below this level, our setting interprets maximum vessel size as the upper bound for a fleet's average vessel size.

its full potential (Figure 11, red line). These results suggest an optimal maximum vessel size that balances the cost-reducing benefits of larger vessels with the anti-competitive effects of market concentration.

We also find that the transition path of vessel size growth is steeper when the maximum vessel size is larger. Larger vessels offer forward-looking firms greater cost advantages, increasing their business-stealing incentives. Interestingly, the scenario with the largest maximum vessel size (24,000 TEU) results in the lowest freight prices during the early transition phase, as vessel size growth outpaces market consolidation. However, the anti-competitive effects of market concentration dominate in the long run, leading to higher prices once the market evolves into a duopoly.

Finally, we present the welfare calculations from our simulation in Table 4. As shown in the top panel, the equilibrium with a 20,000 TEU maximum vessel size generates the highest consumer surplus over the 15-year simulation period, exceeding the case with an 18,000 TEU cap by 43 billion USD and the 24,000 TEU case by 113 billion USD. However, producer surplus is significantly higher in the scenario with more market concentration and larger vessels of 24,000 TEU. While the total surplus is highest in the 24,000 TEU scenario, this reflects a substantial transfer of surplus from consumers to producers, exceeding 100 billion USD.

To further examine the welfare dynamics, we break down the calculations into the transition period (first six years) and the stabilizing period (years seven to fifteen). During the transition period, consumer surplus may experience a temporary reduction, ranging from 22 to 28 billion USD compared to other equilibria, before recovering to generate significantly higher consumer welfare in the long run.

These welfare calculations offer a basis for evaluating the cost-benefit implications of maritime infrastructure investments to accommodate larger vessels. If the costs of infrastructure modernization exceed the improvements in consumer surplus estimated here, the net returns on such investments may be negative, raising concerns about their economic viability.

Figure 9: Different Maximum Vessel Size: Transition Path

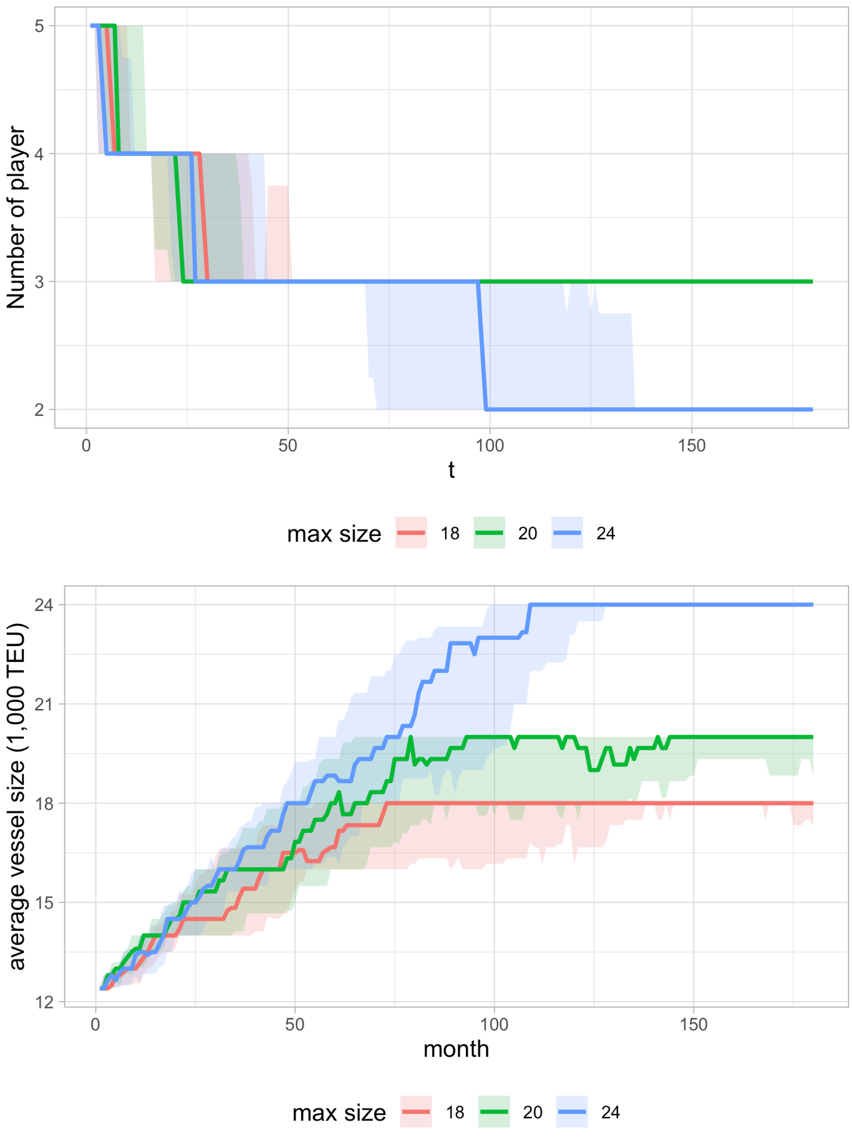

Note:

Figure 10: Different Maximum Vessel Size: Price

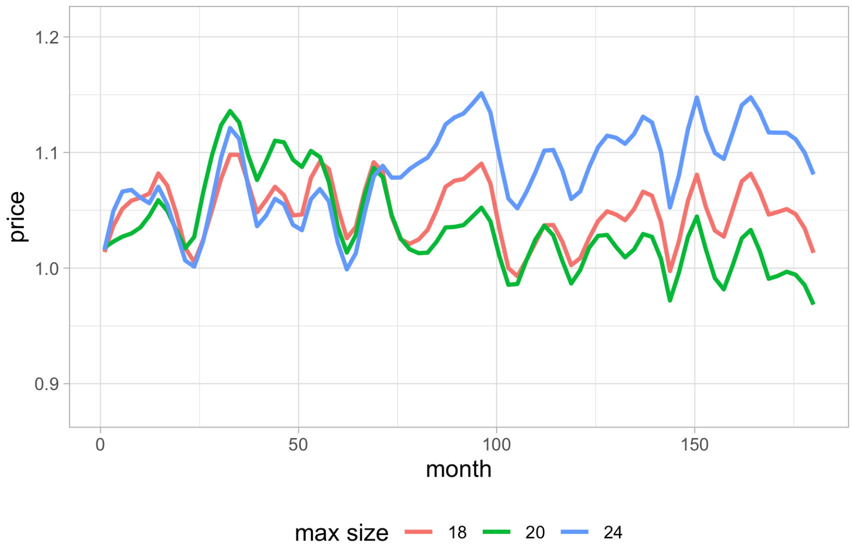

Note:

Figure 11: Different Maximum Vessel Size: Total Capacity

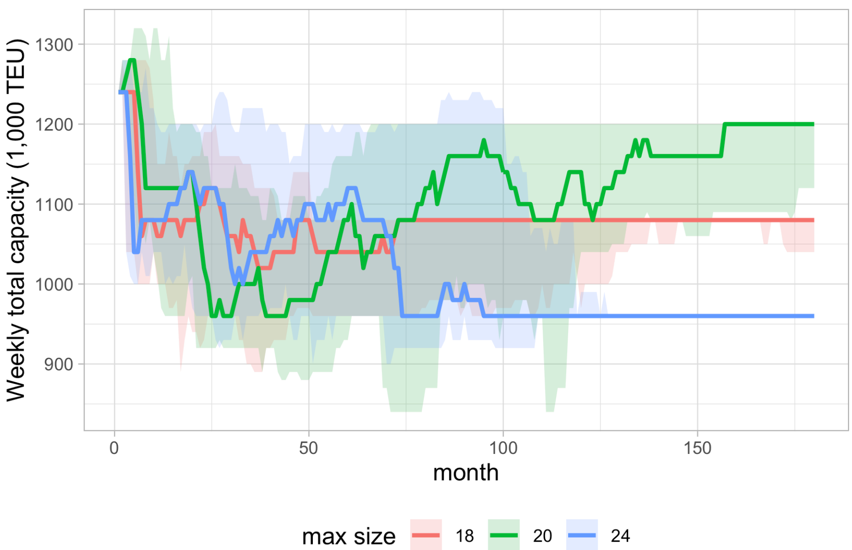

Note:

Table 4: Welfare Calculation Across Different Maximum Vessel Sizes (in Billion USD)

Breakdown	Maximum Vessel Size	Consumer Surplus	Producer Surplus	Total Surplus
Total	18,000	0.00	0.00	0.00
Total	20,000	43.31	−12.60	30.71
Total	24,000	−70.35	127.45	57.10
Transition[a]	18,000	0.00	0.00	0.00
Transition	20,000	−21.74	6.43	−15.31
Transition	24,000	6.14	−17.74	−11.60
Long-run[b]	18,000	0.00	0.00	0.00
Long-run	20,000	50.12	−25.83	24.29
Long-run	24,000	−120.61	150.24	29.63

[a] Transition refers to the simulated dynamics during the first 72 months.
[b] Long-run refers to the steady-state equilibrium values after the transition period (month 73-180).
Note:

This counterfactual exercise highlights the importance of considering the competitive effects of technological advancements, particularly when they involve significant economies of scale. While a rapid "arms race" in technology adoption may reduce consumer prices during the transition phase, it risks leading to excessive market concentration in the long run. In the context of container shipping, this exercise underscores the need for competition and maritime authorities to carefully evaluate the impact of vessel size limits. Allowing excessively large vessel sizes could result in socially suboptimal market concentration, while overly restrictive size limits could prevent the industry from fully realizing the cost-saving benefits of economies of scale. This calls for incorporating the competitive implications into policy decisions regarding maritime infrastructure improvements, such as port and canal expansions.

5.2 Aggregate Demand and Maximum Vessel Size

In our previous analysis, we assumed constant aggregate demand. Here, we simulate various aggregate demand scenarios to examine how the industry's transition dynamics and long-run equilibrium respond under differing demand conditions. Specifically, we simulate the transition paths for maximum vessel sizes of 18,000 TEU, 20,000 TEU, and 24,000 TEU across low, middle, and high demand levels.[37]

Our results indicate that the optimal maximum vessel size depends on the level of aggregate demand. As shown in Figure 12, a maximum vessel size of 20,000 TEU generates the lowest long-run freight prices under low and middle demand scenarios. However, this is not optimal when aggregate demand is 25% higher. This is because the trade-off between the anti-competitive effects of market concentration and the pro-competitive benefits of technological upgrades shifts with demand levels.

Under low demand, the market supports only two players in the long run, regardless of the maximum vessel size (Figure 13, top left panel). Carriers lose their incentive to invest further once their average vessel size reaches approximately 20,000 TEU, making 20,000 TEU the optimal maximum vessel size (Figure 13, bottom left panel). In contrast, under high demand, the market can sustain either a four-player scenario with an 18,000 TEU average vessel size or a three-player scenario with a 24,000 TEU average vessel size. In this high-demand environment, the downside of market concentration is less pronounced. A maximum vessel size of 18,000 TEU leverages the competitive benefits of having four players, while a maximum vessel size of 24,000 TEU maximizes the pro-competitive effects of technology upgrades.[38]

5.3 Innovation Step Size

We previously examined the competitive effects of changing the technology frontier in our counterfactual analysis. In this section, we shift our focus to the competitive implications of innovation step size. Specifically, we simulate scenarios where the incremental vessel size that carriers can upgrade by is varied, while keeping the maximum vessel size fixed. This analysis demonstrates how differences in innovation step size can drive the industry toward

[37] Middle-level demand corresponds to the average demand shifter estimated from the data. High-level demand assumes a 25% increase, while low-level demand assumes a 25% decrease.

[38] In the high-demand scenario, both the 18,000 TEU and 24,000 TEU maximum vessel size cases achieve a total capacity of 72,000 TEU (Figure 14, right panel). However, their long-run performance in terms of freight price dynamics differs, particularly under price volatility. For more details, see Appendix E.

Figure 12: Different Maximum Vessel Size: Price Across Different Demand Level

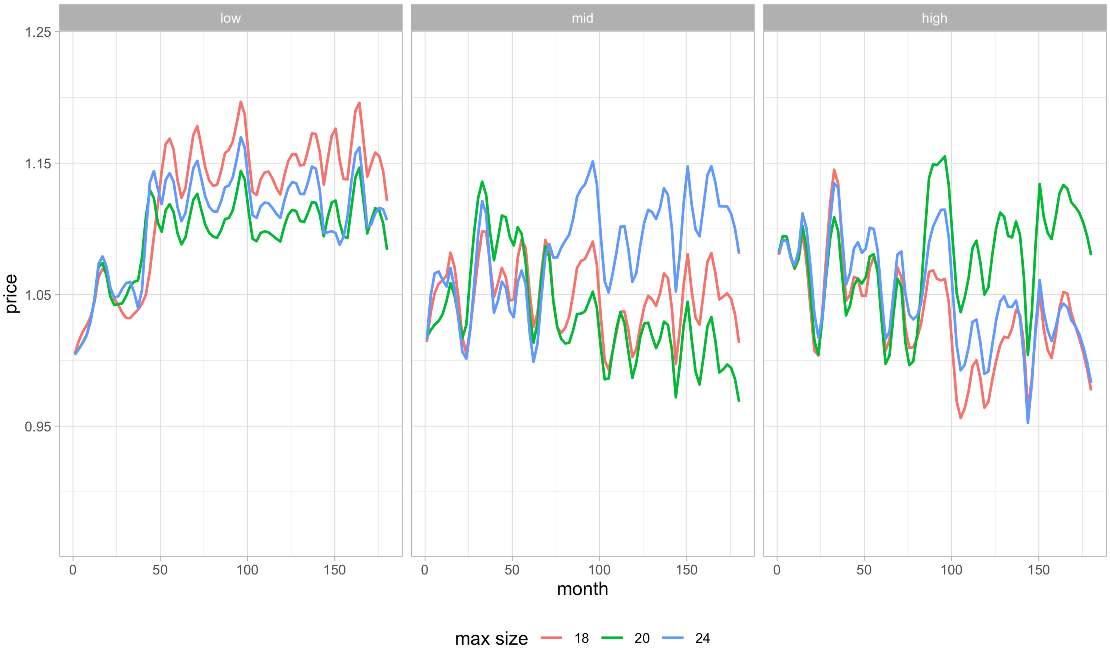

Note:

distinct equilibria. The non-stationary setting of our model is particularly advantageous here, as it allows us to articulate how the industry evolves toward long-run equilibrium.

In this simulation, we fix the maximum vessel size (technology frontier) at 24,000 TEU and compare the industry outcomes under innovation step sizes of 1,500 TEU and 2,000 TEU. Figure 15 illustrates the transition paths for the average vessel size and the number of market players. With a larger innovation step size, the market consolidates more rapidly, reaching a duopoly within 7–8 years. In contrast, the scenario with a smaller step size stabilizes at an oligopoly with three players. Despite the maximum vessel size being set at 24,000 TEU, the industry's average vessel size converges to around 19,000 TEU.

The reason behind these differing equilibria is that a larger step size amplifies the business-stealing payoff, incentivizing carriers to invest and upgrade more aggressively. Figure 16 depicts the probability of investment for a carrier currently operating vessels with an average size of 18,000 TEU. In the case with a step size of 1,500 TEU, the investment probability

Figure 13: Different Maximum Vessel Size: Transition Path Across Different Demand Level

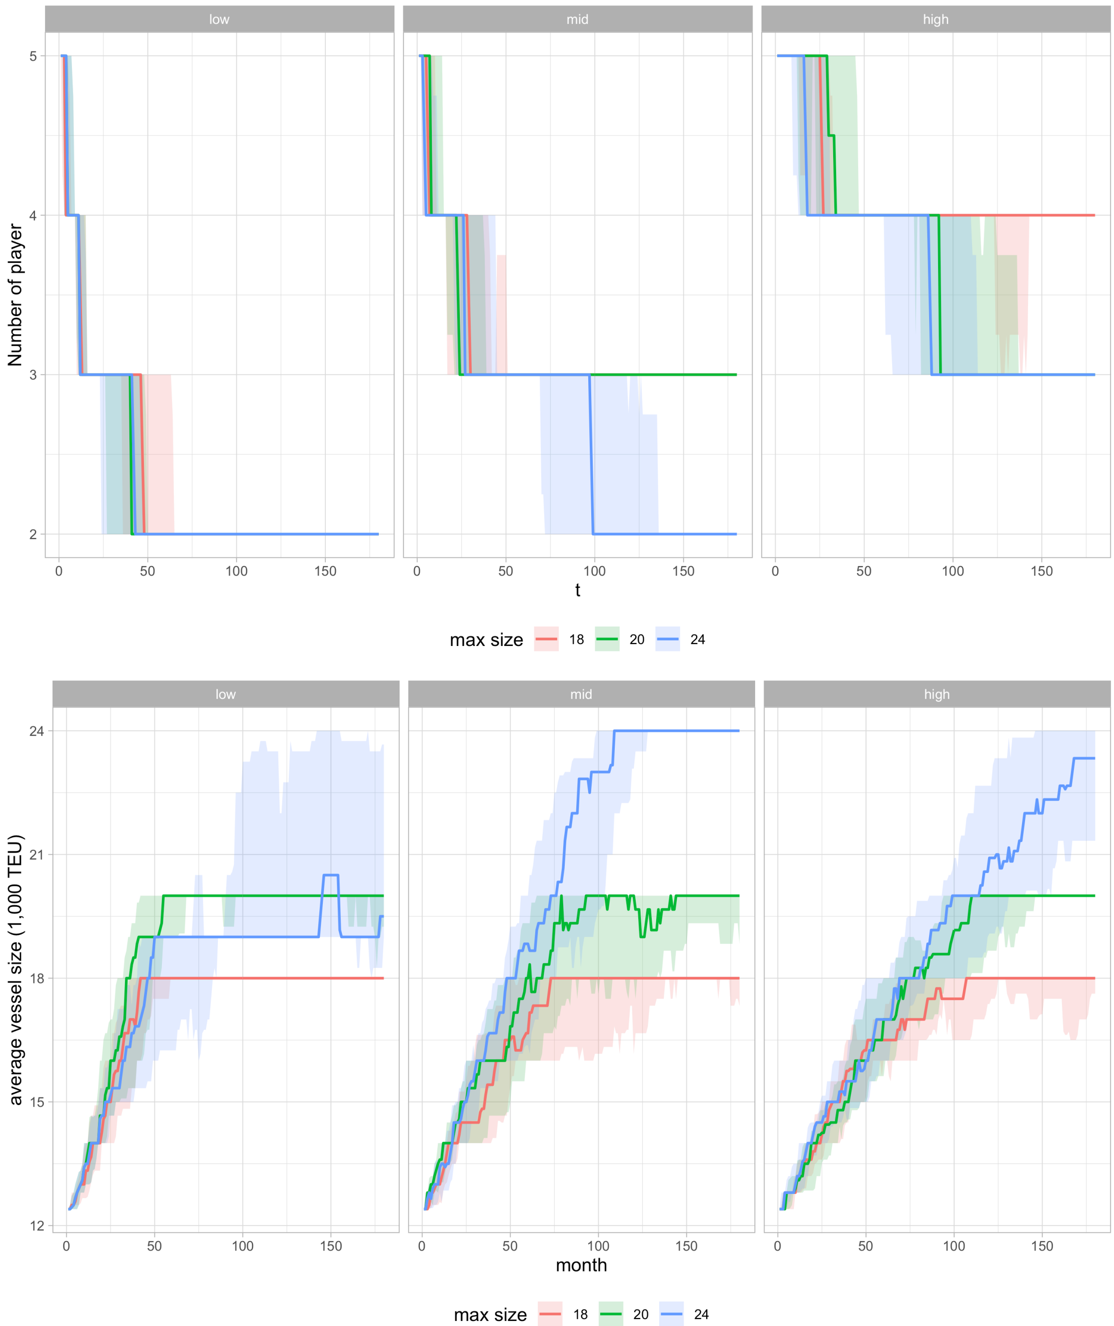

Note:

Figure 14: Different Maximum Vessel Size: Total Capacity Across Different Demand

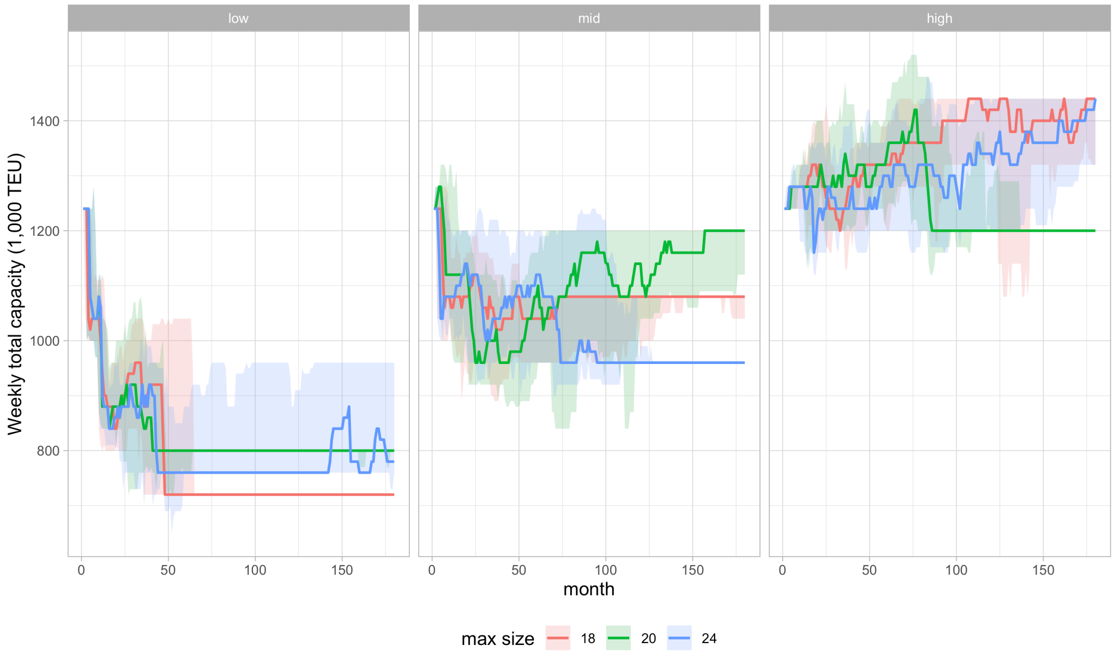

Note:

is significantly lower than that of the 2,000 TEU scenario. Specifically, if the carrier is a market leader, its probability of upgrading is nearly zero.[39] This explains why, in the 1,500 TEU step-size scenario, the industry's average vessel size converges to approximately 19,500 TEU (see the red line in the lower panel of Figure 15).

Innovation step size also affects price dynamics. As shown in Figure 17, prices are lower during the transition phase (0–72 months) in the scenario with a larger innovation step size, primarily due to more aggressive investment. However, this aggressive investment accelerates market consolidation, leading to reduced total capacity and higher long-run prices (Figure 26, blue line). This highlights a trade-off in the welfare impact of innovation step size between the transition phase and the long-run equilibrium.

We also explored the effects of varying initial industry states on transition paths and

[39] As shown in Figure 16, the probability of investment increases as the carrier lags further behind the industry average. This aligns with our earlier findings that market followers have stronger incentives to invest compared to market leaders.

Figure 15: Different Step Sizes: Transition Path

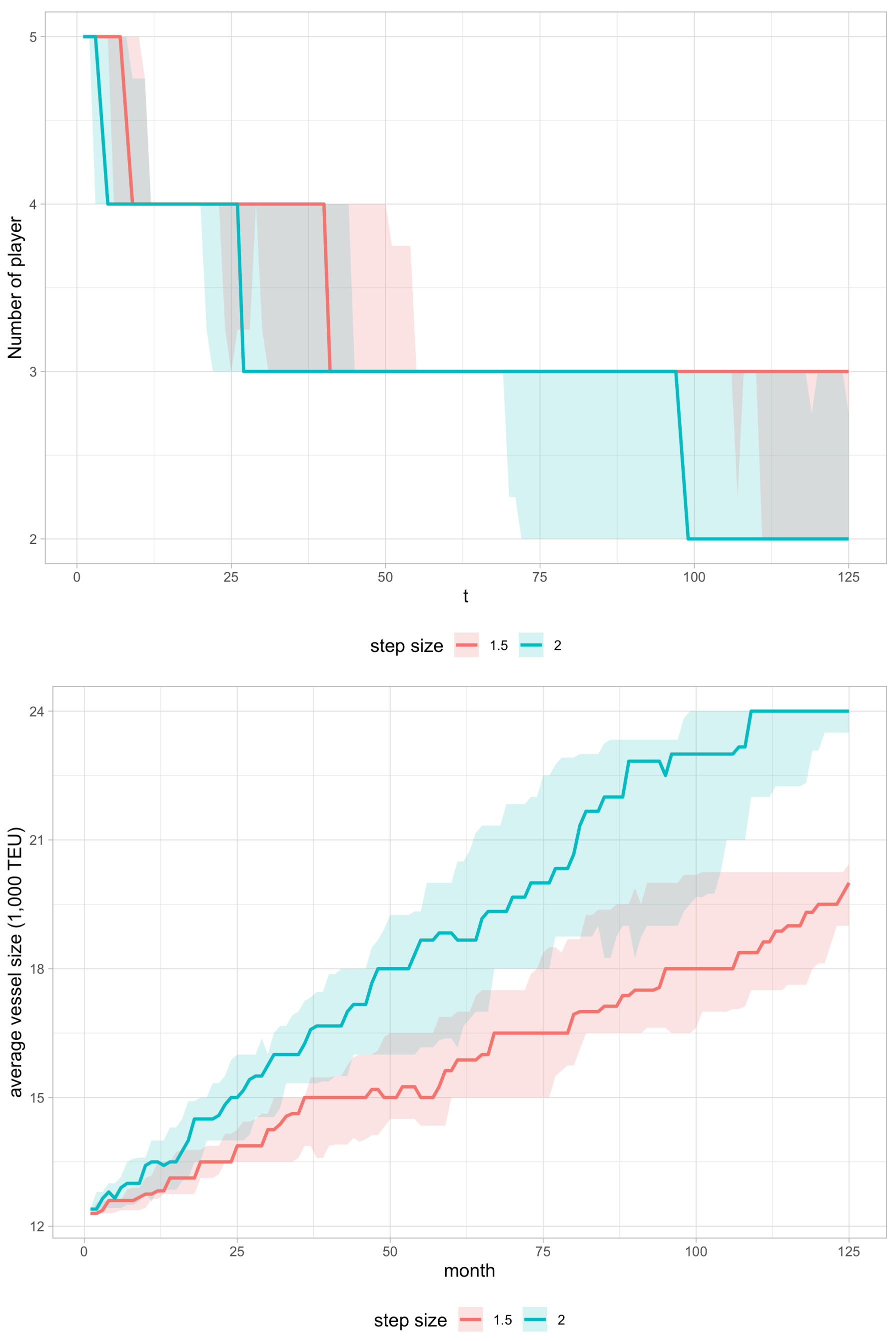

Note: The top panel shows the number of players in the market over time, and the bottom panel shows the average vessel size.

Figure 16: Different Step Sizes: Incentive to Invest

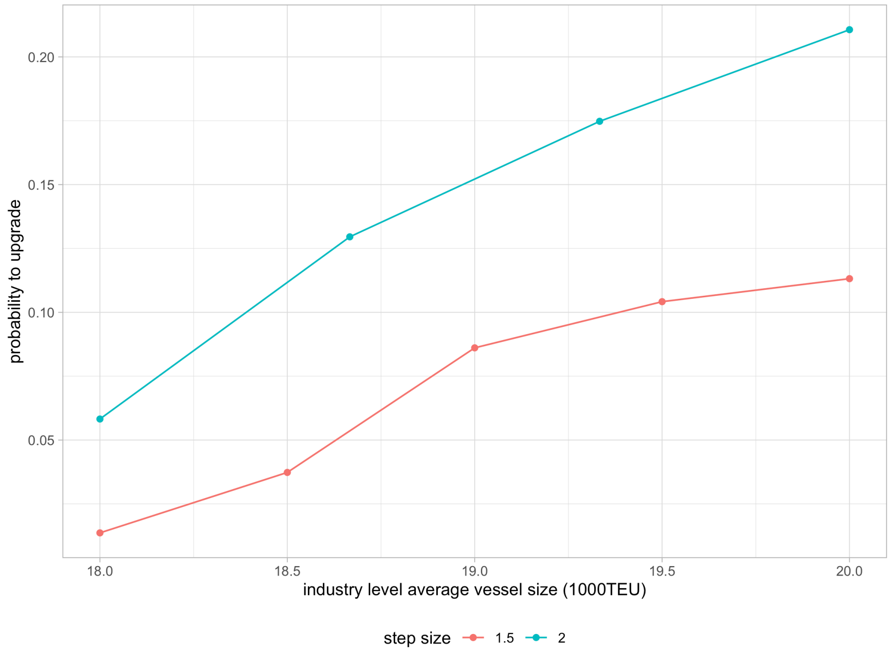

Note: Investment probabilities for a carrier currently operating at 18,000 TEU.

equilibria. While differences in initial conditions primarily influence the transition path, the long-run equilibrium remains unaffected. For further details on this analysis, see Appendix F.

6 Conclusion

This paper examines the competitive effects of technological innovation in the container shipping industry, with a particular focus on the relationship between increasing vessel size, market structure, and welfare. Our analysis highlights that while advancements in shipping technology, such as the trend toward larger vessels, generate significant economies of scale and reduce operational costs, they also exacerbate market concentration. This trade-off between efficiency and competition raises important questions about whether the pursuit of ever-larger vessels has gone too far from a welfare perspective.

Using a dynamic oligopoly framework, we show that the impact of technological ad-

Figure 17: Different Step Sizes: Price

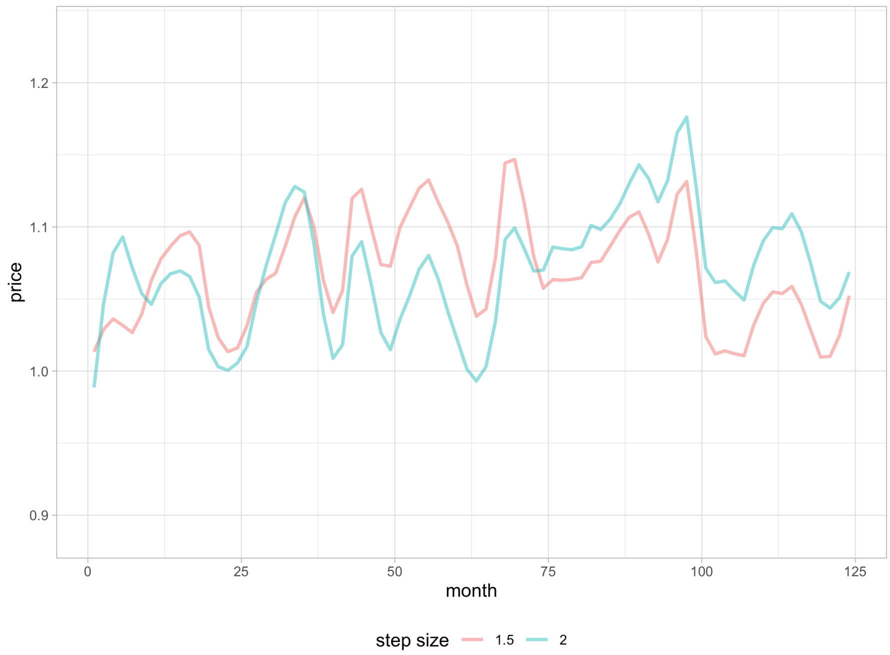

Note: Price dynamics under different innovation step sizes.

vancements depends critically on both the industry's technological frontier and the pace of innovation. Larger vessels reduce costs but lead to greater market consolidation, which can erode consumer benefits through higher freight prices. Our counterfactual analysis reveals that the welfare-optimal technology frontier lies at an average vessel size of around 20,000 TEU under current demand conditions, as further increases in vessel size risk over-consolidating the market. Moreover, we find that the speed of innovation plays a crucial role in shaping industry equilibrium, with smaller innovation steps fostering competition and larger steps driving consolidation. These findings have significant policy implications, particularly for the design of infrastructure investments and regulatory frameworks. Policymakers must balance the pro-competitive effects of technological advancements with the risk of market concentration, ensuring that the benefits of innovation are distributed equitably across the economy.

References

Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt (2005). Competition and innovation: An inverted-u relationship. *The quarterly journal of economics 120*(2), 701–728.

Asturias, J. (2020). Endogenous transportation costs. *European economic review 123*, 103366.

Benkard, C. L. (2004). A dynamic analysis of the market for wide-bodied commercial aircraft. *The Review of Economic Studies 71*(3), 581–611.

Bernhofen, D. M., Z. El-Sahli, and R. Kneller (2016). Estimating the effects of the container revolution on world trade. *Journal of International Economics 98*, 36–50.

Berry, S. and A. Pakes (1993). Some applications and limitations of recent advances in empirical industrial organization: Merger analysis. *The American Economic Review 83*(2), 247–252.

Brancaccio, G., M. Kalouptsidi, and T. Papageorgiou (2020). Geography, transportation, and endogenous trade costs. *Econometrica 88*(2), 657–691.

Brancaccio, G., M. Kalouptsidi, and T. Papageorgiou (2024). Investment in infrastructure and trade: The case of ports. Technical report, National Bureau of Economic Research.

Brooks, L., N. Gendron-Carrier, and G. Rua (2018). The local impact of containerization. *Finance and Economics Discussion Series 45*.

Coşar, A. K. and B. Demir (2018). Shipping inside the box: Containerization and trade. *Journal of International Economics 114*, 331–345.

Fuchs, S. and W. F. Wong (2024). Multimodal transport networks.

Ganapati, S., W. F. Wong, and O. Ziv (2024). Entrepot: Hubs, scale, and trade costs. *American Economic Journal: Macroeconomics 16*(4), 239–278.

Garg, S. and S. Saxena (2023). Dynamic effects of price controls and deregulation policies: Evidence from the indian cement industry.

Goettler, R. L. and B. R. Gordon (2011). Does amd spur intel to innovate more? *Journal of Political Economy 119*(6), 1141–1200.

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

Igami, M. (2017). Estimating the innovator's dilemma: Structural analysis of creative destruction in the hard disk drive industry, 1981–1998. *Journal of Political Economy 125*(3), 798–847.

Igami, M. (2018). Industry dynamics of offshoring: The case of hard disk drives. *American Economic Journal: Microeconomics 10*(1), 67–101.

Igami, M. and K. Uetake (2020). Mergers, innovation, and entry-exit dynamics: Consolidation of the hard disk drive industry, 1996–2016. *The Review of Economic Studies 87*(6), 2672–2702.

Imai, A., E. Nishimura, S. Papadimitriou, and M. Liu (2006). The economic viability of container mega-ships. *Transportation Research Part E: Logistics and Transportation Review 42*(1), 21–41.

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

Kalouptsidi, M. (2014). Time to build and fluctuations in bulk shipping. *American Economic Review 104*(2), 564–608.

Kim, M. (2013). Strategic responses to used goods markets: Airbus and boeing. *Available at SSRN 2267178*.

Marshall, G. and A. Parra (2019). Innovation and competition: The role of the product market. *International Journal of Industrial Organization 65*, 221–247.

Merk, O. (2018). Container ship size and port relocation.

Murray, W. (2016). Economies of scale in container ship costs.

Otani, S. (2024). Industry dynamics with cartels: The case of the container shipping industry. *arXiv preprint arXiv:2407.15147*.

Ryan, S. P. (2012). The costs of environmental regulation in a concentrated industry. *Econometrica 80*(3), 1019–1061.

Schumpeter, J. A. (1942). *Capitalism, socialism and democracy.* routledge.

UNCT, A. (2022). Review of maritime transport 2022. navigating stormy waters.

Veldman, S. (2011). On the ongoing increase of containership size. In *Advances in maritime logistics and supply chain systems*, pp. 201–228. World Scientific.

Wong, W. F. (2022). The round trip effect: Endogenous transport costs and international trade. *American Economic Journal: Applied Economics 14*(4), 127–166.

Yang, D., M. Liu, and X. Shi (2011). Verifying liner shipping alliance's stability by applying core theory. *Research in Transportation Economics 32*(1), 15–24.

A Details On Demand Estimation

The container trade volume and price data from Container Trade Statistics cover 22 major trade routes, corresponding to 11 round-trip lanes: Asia-Europe, Asia-US, Asia-South America, Asia-Indian Subcontinent, Europe-US, Europe-South America, Europe-Indian Subcontinent, US-Indian Subcontinent, South America-Indian Subcontinent, Asia-Australia, and Australia-Europe. To ensure proper identification and control for fixed effects at each node in the shipping network, we include each region as either an origin or a destination in at least two routes. Since our data extends only until February 2024, we set the start date of our sample to November 2023 to maintain a balanced panel before and after the Houthi attacks, which serve as the treatment event.

To account for trade volume heterogeneity across routes, we employ a weighted two-stage least squares approach, using trade volume as weights, thereby ensuring that busier trade lanes have greater representation in the estimation. Additionally, we classify Asia-US and ISC-US routes as part of the treatment group, but only from January 2024 onward. This is because the Houthi attacks primarily affected Asia-Europe trade by extending transit times by approximately two weeks, which increased the capacity required to maintain weekly sailing schedules. This capacity spillover led to shortages on other trade lanes, but with a lagged effect. To account for this transmission delay, we introduce Asia-US and ISC-US routes into the treatment group one month after the attacks. This adjustment highlights the interconnected nature of the global container shipping network.

An additional consideration is the relationship between the estimated price elasticity of demand for container shipping services and the broader trade elasticity documented in the literature. Since shipping costs are additive to the total cost of imports, the two elasticities are conceptually distinct. Specifically, the demand elasticity we estimate, σ, relates to trade elasticity, σ^{trade}, through the equation:

$$\sigma \approx \frac{f}{f+z} \cdot \sigma^{trade},$$

where f represents the freight price and z denotes the cargo value. Assuming that freight rates account for approximately 10–15% of total import costs, our estimated demand elasticity translates to a trade elasticity in the range of 10–12—on the higher end, yet still within the range observed in the trade literature. Moreover, our estimates should be interpreted as an upper bound for short-run demand elasticity. This is due to potential upward bias introduced by the increase in transit time caused by the attack, as longer transit times reduce the effective value of fast shipping. For a more detailed discussion on demand estimation, including adjustments for transit speed valuation, we refer the reader to the author's related work using proprietary data from a logistics freight forwarder.

B Details on Marginal Cost Estimation

We estimate the marginal cost function for carriers in two steps. First, we estimate Equation 10 under the assumption that all carriers share the same marginal cost parameters, b and c.[40] This assumption is primarily driven by data limitations, as our dataset from CTS for 2015–2016 only provides *aggregate* TEU volumes, preventing us from modeling carrier-specific heterogeneity in capacity utilization. However, this simplification allows us to directly link *observed* freight prices to *observed* aggregate capacity utilization ratios, enabling a straightforward estimation of b and c via linear regression.

Figure 18 illustrates the positive relationship between freight rates and capacity utilization, with estimation results summarized in Table 5. The coefficient b is significantly positive, indicating an increasing marginal cost curve. This suggests that as vessel-level capacity utilization nears 100%, the marginal cost of accommodating an additional container rises. This result aligns with operational realities, as fully loaded vessels require longer loading and unloading times, increasing handling costs.

[40]However, the slope of the marginal cost function varies across carriers as it depends on their deployed capacity. See Equation 4 for further details.

Figure 18: Freight price versus effective capacity utilization rate

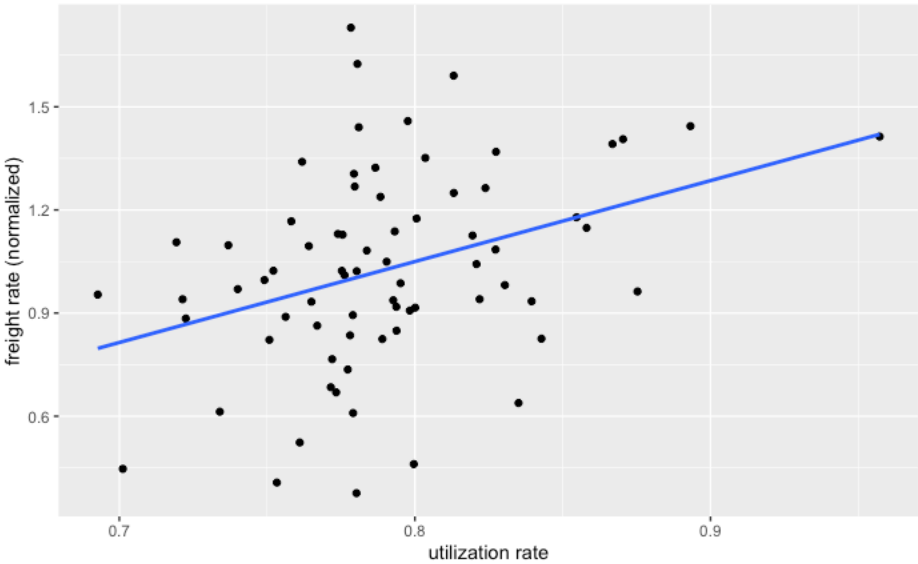

Table 5: Step 1 estimation

	Dependent variable:
	rate_norm
util (b)	2.354***
	(0.729)
Constant (c)	−0.833
	(0.577)
Observations	72
R^2	0.130
Adjusted R^2	0.117
Residual Std. Error	0.272 (df = 70)
F Statistic	10.436*** (df = 1; 70)
Note:	*p<0.1; **p<0.05; ***p<0.01

Figure 19: Capacity Utilization

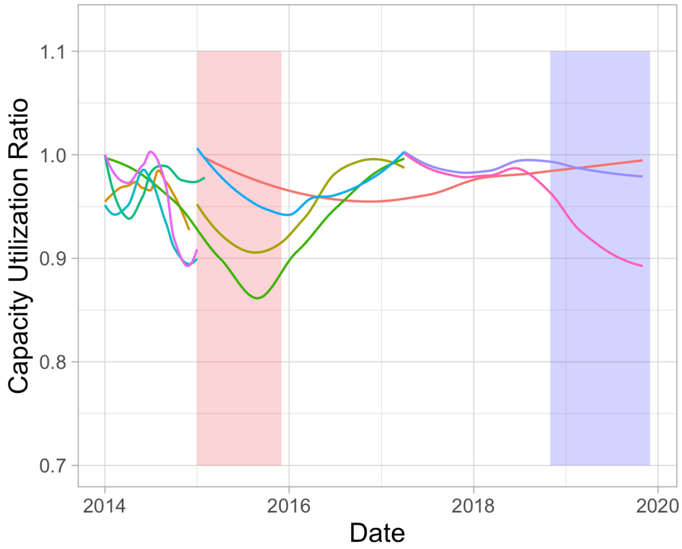

C Static Model Fit

To evaluate the model fit, we divide the sample into an in-sample period (2015, shaded red in Figure 19) and an out-of-sample period (October 2018 to December 2019, shaded blue in Figure 19). The year 2015 is chosen as the in-sample period because it exhibits the greatest variation in capacity utilization across carriers, driven by the downturn in aggregate demand. We estimate the fixed cost parameters, f and α, using variation within the in-sample data and then assess the model's predictive performance on the out-of-sample period in 2019.

The model performs well in fitting the in-sample data (see Figure 20). It successfully captures the cross-sectional variation across carriers, correctly predicting that carriers operating smaller vessels reduce their capacity utilization more drastically than those with larger vessels. Additionally, the model tracks the overall time-series trend, reinforcing the accuracy of its estimated relationship between aggregate demand and capacity utilization. The correlation between model-generated and observed capacity utilization rates is 0.62 (see Figure 21).

For the out-of-sample period, while the model fit is weaker compared to the in-sample

Figure 20: Model Fit: In-sample capacity utilization

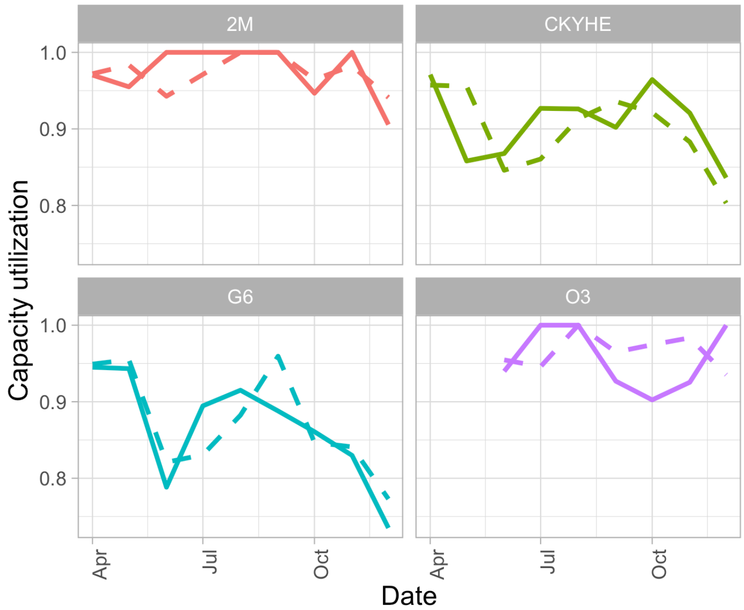

Note: The dotted line is the model predicted capacity utilization and the solid line is the capacity utilization in the data.

period, it still captures key cross-sectional variations across carriers. Specifically, it correctly predicts that THE Alliance, which operates smaller vessels, experiences lower capacity utilization during demand downturns (see Figures 22 and 23).

D Dynamic model fit

We compare the model-generated and empirical transition paths for both the average vessel size (Figure 24) and the number of players in the market (Figure 25). The model fits the observed trajectory of vessel size growth well, closely matching the empirical trend. However, the model predicts industry consolidation—reducing from five to four players—earlier than observed in the data. This discrepancy suggests that either the pace of consolidation in reality is slower due to unmodeled frictions, or that additional strategic considerations, such as coordination challenges or regulatory constraints, may delay firm exits in practice.

To provide intuition on the identification of key dynamic parameters, the investment cost

Figure 21: Model Fit: In-sample capacity utilization dotted plot

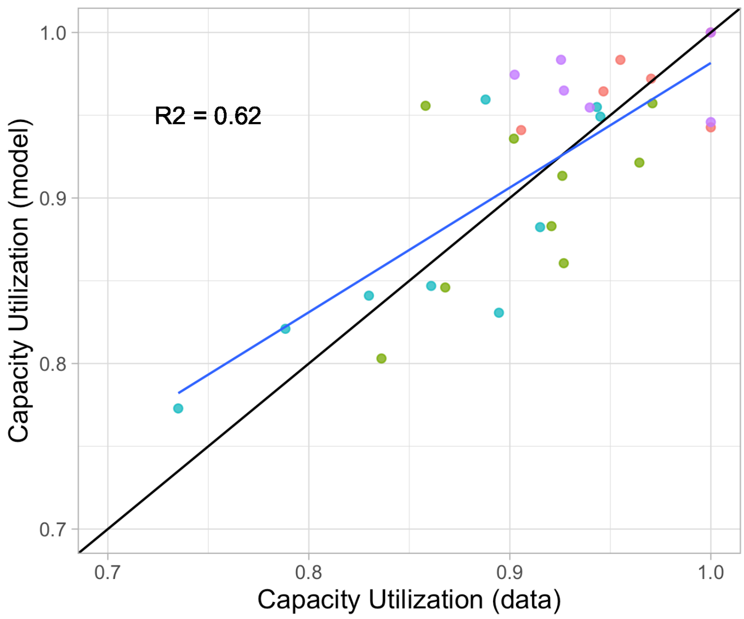

Note: This graph plotted the model generated capacity utilization (y-axis) against the observed capacity utilization in the data (x-axis).

Figure 22: Model Fit: Out-of-sample capacity utilization

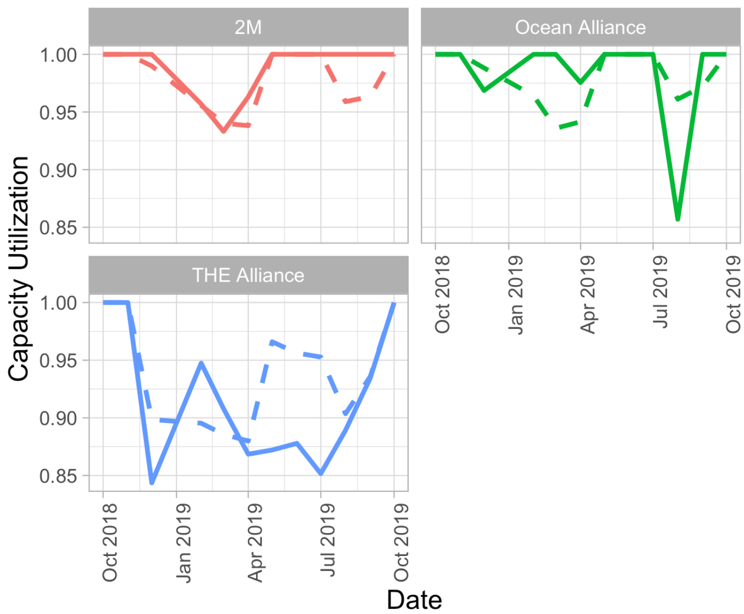

Note: The dotted line is the model predicted capacity utilization and the solid line is the capacity utilization in the data.

Figure 23: Model Fit: Out-of-sample capacity utilization dotted plot

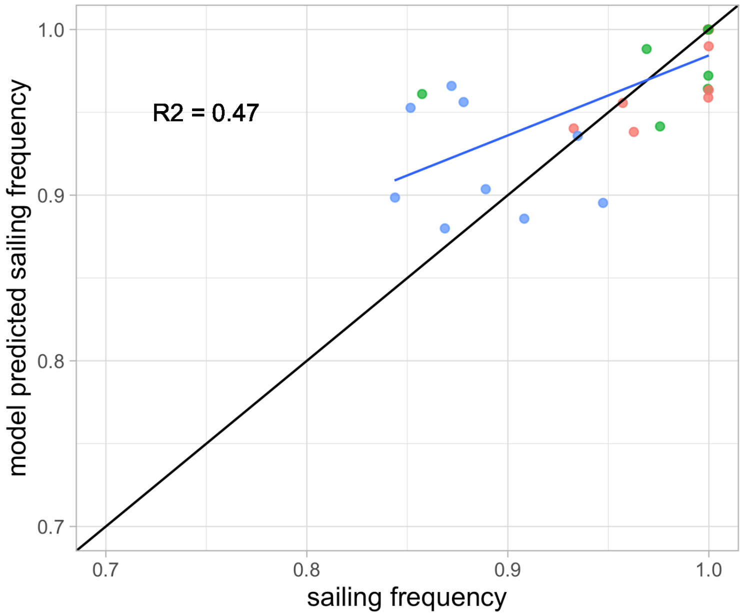

Note: This graph plotted the model generated capacity utilization (y-axis) against the observed capacity utilization in the data (x-axis).

parameter ι primarily governs the speed of vessel size upgrades. Since larger vessels reduce per-unit operating costs, carriers have an incentive to invest, and the empirical trajectory of vessel size growth provides identification for ι. Meanwhile, the maintenance cost parameter (m) and entry cost parameter (ζ) shape the market structure by influencing firm survival and new entry decisions. These parameters are identified through the observed number of active alliances during the sample period, as they determine whether firms find it profitable to remain in the market or whether new entrants can viably compete. The interplay of these costs captures the long-run industry equilibrium, where firms balance the benefits of operating larger vessels against the financial burden of maintaining their fleet and the barriers to entry.

Figure 24: Model Fit: Average vessel size

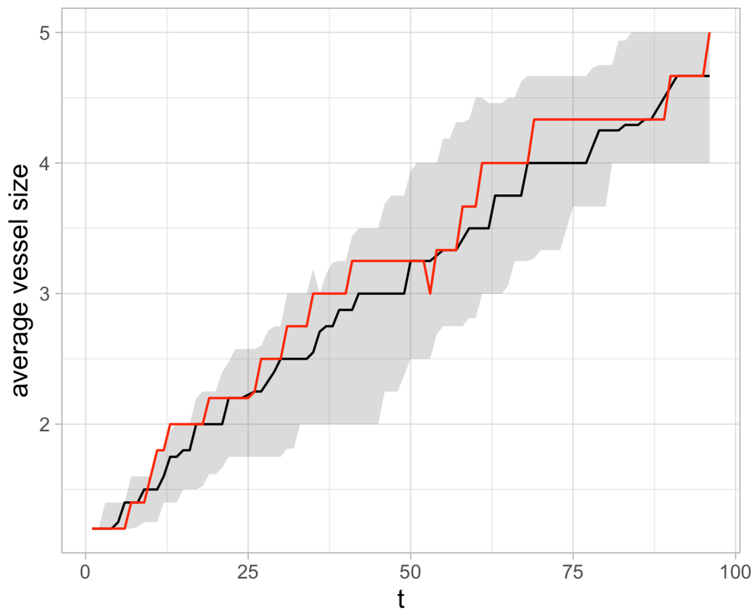

Note: The black line with shaded confidence interval is the model predicted average vessel size and the red line is the vessel size observed in the data.

Figure 25: Model Fit: Number of firms

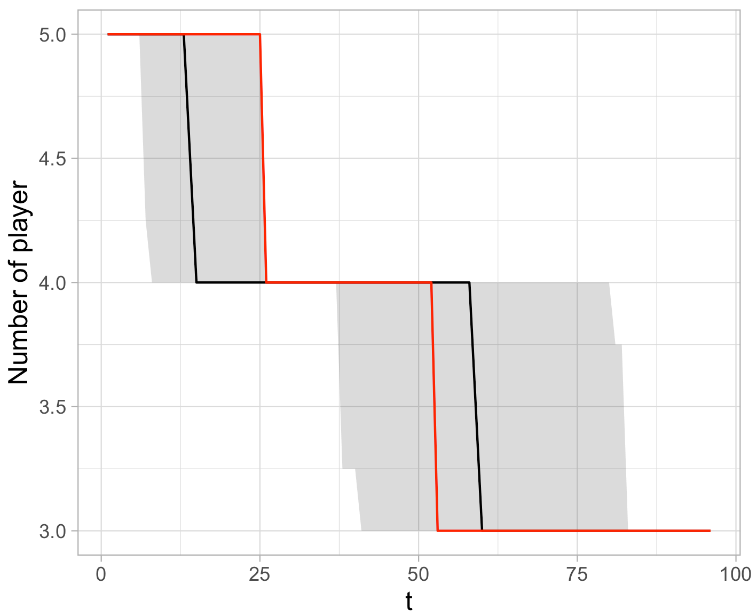

Note: The black line with shaded confidence interval is the model predicted number of players and the red line is the number of player observed in the data.

E Price Dynamics Under The Same Total Capacity

Even when the *total* industry capacity remains constant, differences in the equilibrium composition—specifically, the number of firms and the average vessel size—can have significant implications for freight prices and consumer welfare. This distinction arises because, during demand downturns, carriers do not operate their entire fleet, adjusting their deployed capacity in response to market conditions. Carriers with smaller vessels, on average, tend to reduce their capacity utilization ratios more sharply than those operating larger vessels, as smaller ships are less cost-efficient at lower utilization rates. This variation is central to our earlier identification and estimation of economies of scale.

To further explore this dynamic, we examine how different industry structures—varying in the number of firms and vessel sizes—translate into different pricing patterns. In an equilibrium characterized by fewer firms operating larger vessels, the total *deployed* capacity during downturns is higher than in an equilibrium with more firms operating smaller vessels. This occurs because firms with larger vessels face stronger cost incentives to maintain higher utilization levels, leading to a less pronounced contraction in active capacity. Consequently, the equilibrium structure not only influences the average freight price but also affects price volatility and market stability. In a more consolidated market with larger vessels, prices may be more stable during downturns due to higher deployment persistence, but the reduction in competition could lead to higher prices in the long run. Conversely, a market with more firms and smaller vessels may exhibit greater price fluctuations, as firms adjust their capacity more aggressively in response to demand shocks.

These differences underscore the broader welfare implications of industry consolidation and vessel size growth. While larger vessels reduce costs through economies of scale, the corresponding reduction in market competition may offset these efficiency gains, particularly in periods of demand volatility. Understanding these trade-offs is crucial for policymakers evaluating the competitive effects of technological advancements in the shipping industry.

Figure 26: Different Step Sizes: Total Capacity

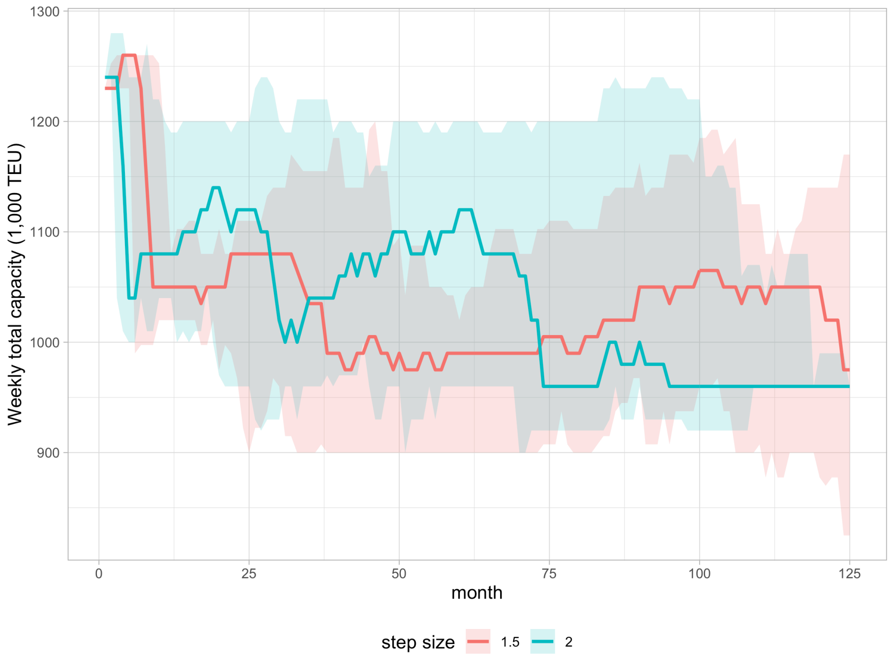

Note: Total capacity under different innovation step sizes.

F Transition Path and Long-run Equilibrium

F.1 More Details on Innovation Step Size

F.2 Initial State

In the main text, we examined how different innovation step sizes influence the transition path and long-run equilibrium of the industry. We now turn our attention to another critical factor: the role of initial industry conditions. Specifically, we analyze how different starting states affect the transition dynamics and the speed of market consolidation. To isolate the effect of initial conditions, we consider two counterfactual cases:

1. **Case 1**: The industry starts with five carriers, each operating an average vessel size of 14,000 TEU.

2. **Case 2**: The industry starts with five carriers but with an uneven vessel size distribu-

tion: three carriers operate vessels of 12,000 TEU, one operates vessels of 16,000 TEU, and one operates vessels of 18,000 TEU.

To ensure that the observed differences arise purely from initial conditions, we hold both the total industry capacity and the innovation step size constant across both cases. The only variation is that Case 1 begins with a symmetric industry structure, while Case 2 starts with a more heterogeneous distribution of vessel sizes.

Figure 27 presents the transition paths of both cases. Despite differences in initial conditions, we find that both scenarios ultimately converge to the same long-run equilibrium. However, the speed of transition differs significantly: the market consolidates much faster in the asymmetric case (Case 2). This result is consistent with our dynamic estimation, which shows that the incentive to invest is much stronger for follower firms than for market leaders. In an industry that starts from an uneven initial state, follower firms invest more aggressively to remain competitive, accelerating the consolidation process. Conversely, when all five firms begin at similar positions, the investment and consolidation dynamics unfold more gradually.

Another notable observation is the self-reinforcing relationship between investment and market consolidation. In the asymmetric case, the initial imbalance between firms leads to faster market consolidation as weaker firms struggle to keep pace with larger competitors. This early consolidation, in turn, strengthens the incentive for firms to invest in larger vessels, further accelerating the transition. This amplification effect underscores how initial asymmetries in market structure can significantly influence industry dynamics.

Figure 28 plots the trajectory of freight prices under the two scenarios. On average, prices are higher in the asymmetric initial state case, as market consolidation occurs more quickly, leading to reduced competition. However, this conclusion does not hold universally. Whether an asymmetric initial state benefits or harms consumer welfare depends on whether the long-run equilibrium itself is more or less competitive. For example, in this case, the asymmetric initial state amplifies the negative welfare impact on consumers by accelerating

Figure 27: Different Initial State: Transition Path

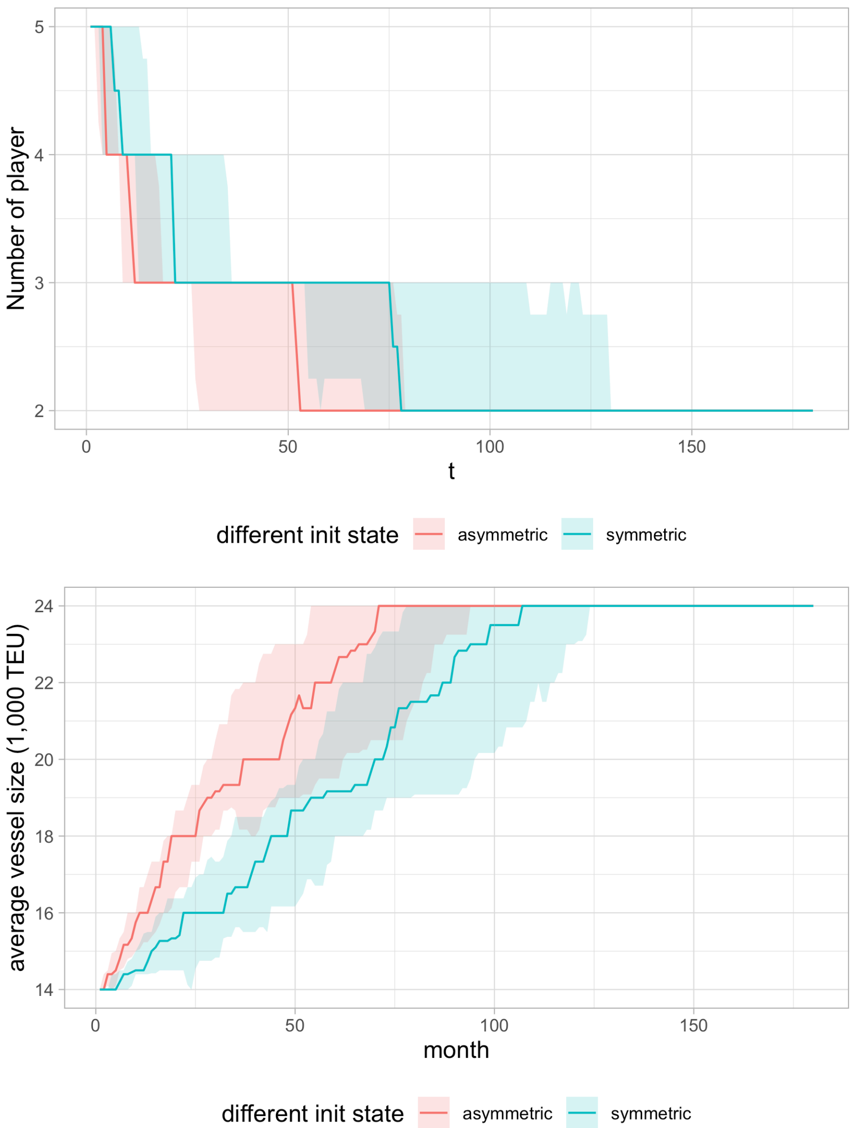

Note: The top panel shows the number of players in the market over time, and the bottom panel shows the average vessel size.

Figure 28: Different Initial State: Price

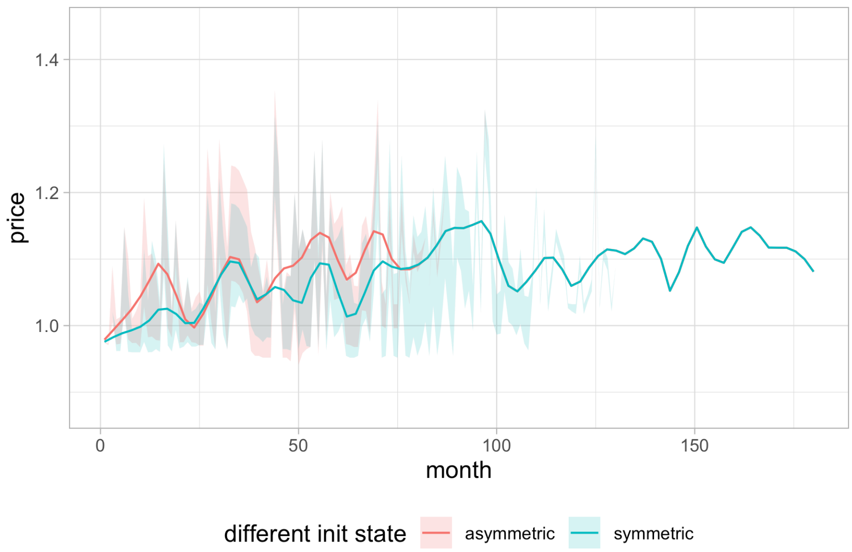

Note: Price dynamics under different initial states.

capacity reduction, as shown in Figure 29.

Thus, the welfare implications of different initial states depend on the nature of the long-run equilibrium. If the ultimate equilibrium supports greater competition and consumer welfare, an asymmetric initial state will expedite these benefits. Conversely, if the equilibrium leads to higher market concentration, an asymmetric initial state will exacerbate its negative effects. The transition path, in essence, serves as an amplifier, magnifying the welfare effects of the final industry structure.

Figure 29: Different Initial State: Total Capacity

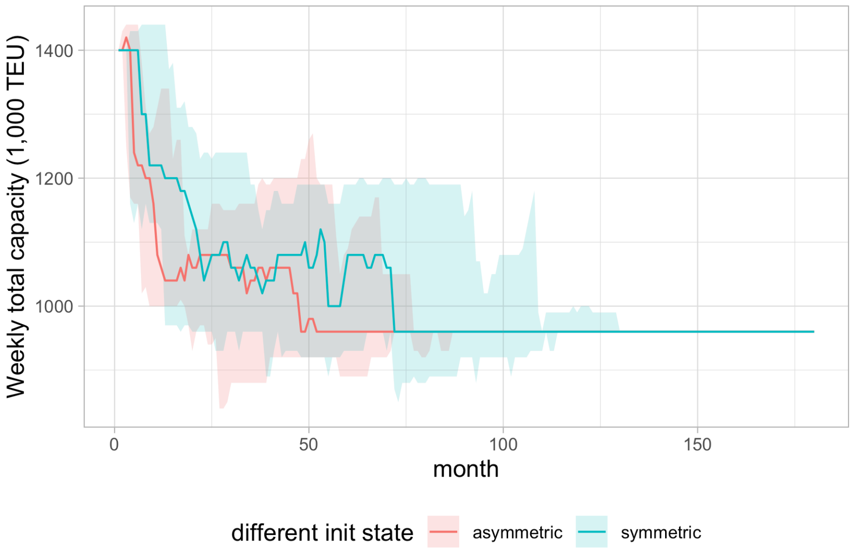

Note: Total capacity dynamics under different initial states.