Learning and Investment under Demand Uncertainty in Container Shipping*

Jihye Jeon †

June 2021

Abstract

This paper investigates the role of demand uncertainty in explaining cyclical investment fluctuations in the container shipping industry. I develop and estimate a dynamic oligopoly model with learning in which firms choose investment and scrapping. In this model, firms are uncertain about the true parameters in the underlying process for demand, and form and revise their beliefs using available information. Counterfactual analysis reveals that uncertainty about the demand process amplifies investment cycles through (i) leading firms to revise beliefs more drastically as they experience demand fluctuations, and (ii) intensifying strategic incentives among firms.

Keywords: demand uncertainty, learning, dynamic games, investment, shipping.

JEL codes: D83, L10, L92.

*I am indebted to my advisors John Asker, Luis Cabral, Robin Lee, and Ariel Pakes for their guidance and support. I am thankful to Victor Aguirregabiria, David Backus, Allan Collard-Wexler, Francesco Decarolis, Michael Dickstein, Robin Greenwood, Myrto Kalouptsidi, Kei Kawai, Peter Newberry, Marc Rysman, and Laura Veldkamp as well as numerous seminar participants for helpful comments and suggestions. Furthermore, I thank Sharon Cox for assistance in acquiring data and her insights into the industry and Yajie Tang for excellent research assistance. Financial support from the Center for Global Economy and Business is gratefully acknowledged. All errors are my own.

†Department of Economics, Boston University; jjeon@bu.edu.

1 Introduction

In many capital-intensive industries such as oil, shipping, and chemical industries, investment is highly pro-cyclical, even though the cost of investment can be substantially higher during booms. In these industries, firms often invest in long-lived capital while facing considerable uncertainty about future demand. Moreover, the process that dictates the evolution of demand can be complicated or changing over time, creating uncertainty about the demand process itself in addition to uncertainty about what value of demand will be realized from a known process. Although many studies show that uncertainty can amplify investment cycles, there is limited empirical evidence as to its quantitative importance.[1] This paper aims to understand the extent to which uncertainty about the demand process can explain the concentration of investment in boom periods and its welfare consequences in an oligopoly setting with strategic considerations.

During the trade boom in the mid-2000s, container shipping companies ordered a large volume of new ships. This corresponded to the price of these ships reaching a historic high. The downturn in demand following the 2008 crisis resulted in a massive oversupply of ships and, in turn, a sharp decline of shipping rates and profitability. Many industry experts attribute the excess capacity problem to the firms' inability to forecast demand correctly and their over-reaction to recent trends in demand.[2]

> 'The container-shipping industry has been highly unprofitable over the past five years. ... Some of the pain is self-inflicted: as in past cycles, the industry extrapolated the good times and foresaw an unsustainable rise in demand.' (Mckinsey Insights, 2014)[3]

Although anecdotal evidence points to the potentially important role of firm beliefs in driving the investment fluctuations, it is difficult to establish this role empirically due to the lack of direct measurements of beliefs. To overcome this challenge, I adopt the following approach: (i) I use survey data that can inform us indirectly about firm beliefs' about container shipping demand; and (ii) I analyze an oligopoly model of firm learning and investment that allows uncertainty about the demand process. The structural analysis serves two purposes. First, it quantifies the extent

[1] For example, Bernanke (1983) shows that uncertainty increases the value of waiting for new information, thus increasing fluctuations in investment.

[2] The problem is not limited to the 2008 crisis, as suggested by the CEO of one of the largest shipping companies in an interview featured in a *Wall Street Journal* article, "Maersk Line CEO: We Misjudged Container-Shipping Demand", accessed on January 11, 2016 via http://www.wsj.com/articles/SB10001424052702303342104579098680549111434.

[3] "The Hidden Opportunity in Container Shipping," accessed on January 11, 2016. http://www.mckinsey.com/insights/corporate_finance/the_hidden_opportunity_in_container_shipping

to which the uncertainty can explain the observed investment cycles. Second, it helps us understand the mechanisms through which uncertainty affects investment fluctuations. In particular, by incorporating learning in an oligopoly model I can investigate how uncertainty interacts with firms' strategic incentives.

I first leverage the fact that container demand is largely driven by overall trade demand and macroeconomic shocks by adopting auxiliary survey data on forecasts of GDP and trade. I compare these forecasts with expectations about future container trade implied by alternative informational assumptions while being flexible on the specification of the time-series process. I find that incorporating learning leads to agent beliefs that are more consistent with the forecasts of GDP and trade. Although this does not serve as direct evidence of learning, it indicates that incorporating learning may allow us to better approximate beliefs that can explain observed firm behavior.

To measure the quantitative implications of the alternative belief structure, I develop a dynamic oligopoly model of firm investment in the spirit of Ericson and Pakes (1995) that incorporates uncertainty about the demand process. The model captures key features of the industry, but makes several simplifying assumptions that allow me to make progress on relaxing the full-information assumption and identifying the model of firm beliefs.[4] In this model, agents do not know the parameters governing the evolution of demand, but form and revise their expectations based on information available at each decision-making moment. The model allows firms to put heavier weights on more recent observations to capture beliefs that would arise if firms were concerned about structural changes in the underlying process (Evans and Honkapohja (2012)). In each period, firms decide whether to invest in new ships and scrap old ships, and also have the option of borrowing additional capital through chartering.

The informational assumption in the model is a form of behavioral assumption that contrasts with the standard full-information (rational expectations) assumption commonly made in a dynamic oligopoly model.[5]. Under full information, firms know the true distribution of demand, although they may be uncertain about exact future realizations.[6] Adopting the alternative information structure with learning in this paper is important for several reasons. First, conceptually, it incorporates

[4]Section 4 discusses the modeling choices and the motivations in detail. Appendix E presents results from various robustness checks. Appendix F extends the model to incorporate capital constraints and discusses results.

[5]An alternative behavioral approach can be found in Greenwood and Hanson (2015)

[6]Note that the full-information assumption differs from the assumption of perfect foresight under which firms know future realizations exactly.

an additional layer of uncertainty that is potentially important for the present empirical context as evidenced by industry reports, anecdotes, and preliminary analysis based on auxiliary survey data. Recent empirical work shows that this type of uncertainty is important for explaining investment cycles in the macroeconomic context[7]. Incorporating it also allows me to ask how much of the investment fluctuations can be attributed to uncertainty about the demand process versus uncertainty about what demand will be realized from a known process. Doing so in an oligopoly setting sheds light on how uncertainty interacts with strategic incentives. Third, methodologically, it provides a parsimonious and computationally tractable way of approximating agent beliefs that might arise in an environment with a complicated data generating process. It can also be easily implemented in standard dynamic oligopoly models.

The estimation employs firm-level data on capital and investment from 2006 to 2014 as well as price and quantity data for container services from 1997 to 2014. One of the challenges in estimating a learning model is that the researcher does not directly observe agents' beliefs, requiring a simultaneous identification of information and model parameters. An important empirical strategy addressing this issue involves using commonly unavailable data on investment costs and scrap values to pin down the learning process.[8] The typical approach recovers unobserved objects such as investment costs, entry costs, and exit values implied by observed firm decisions while imposing a full-information assumption.[9] In contrast, I use data on ship sales prices and demolition prices to calibrate the investment cost and scrap value. This allows me to recover the model of firm beliefs that can rationalize observed firm behavior given the primitives.

The model estimates are consistent with agents placing a 45% weight on a 10-year-old observation relative to the most current one.[10] Removing demand process uncertainty by endowing agents with knowledge about demand parameters lowers total investment by 17% and investment volatility by 22%. More importantly, it reallocates investment across time, leading firms to withhold investment during demand boom years and suffer less from overcapacity when faced with downturns in demand.

[7] See, for example, Fajgelbaum et al. (2017) and Kozlowski et al. (2016)

[8] This approach is similar to that of Hortacsu and Puller (2008) in the underlying logic. Hortacsu and Puller (2008) use marginal cost data to quantify how much firms' actual bidding deviates from the optimal bidding predicted by their theoretical benchmark for the Texas electricity spot market.

[9] For example, Ryan (2012) and Collard-Wexler (2013) adopt this approach.

[10] This estimate is very close to those in previous studies that estimate a constant-gain learning model based on aggregate survey data such as the Survey of Professional Forecasts or micro data on expectations (e.g. Malmendier and Nagel (2016), Milani (2007), and Orphanides and Williams (2005)). Doraszelski et al. (2018) also find that firms weight recent play disproportionately when forming expectations about competitors' play.

This has a substantial impact on welfare: it increases producer surplus by 85%, while resulting in only a small decrease in consumer surplus.

Through counterfactual simulations with respect to demand volatility, I find that an increase in demand volatility suppresses investment, consistent with findings in previous empirical studies (e.g. Collard-Wexler (2013)).[11] However, I find that the increase in investment volatility in response to increased demand volatility is much higher in the presence of learning. This finding sheds light on the additional informational channel through which demand volatility affects investment: under learning, large fluctuations in demand lead firms to revise their beliefs more frequently and more drastically, which in turn amplifies boom-bust investment cycles.

In order to highlight the effects of strategic incentives and how they interact with agent beliefs, I consider counterfactuals pertaining to competition and learning. Strategic considerations in oligopolistic competition may lead to excessive investment.[12] And agent beliefs may become more highly correlated with recent realizations of demand and more volatile under learning, which can in turn reinforce the strategic incentives. For example, firms generally have greater incentives to steal business from and preempt rivals in periods of high demand due to higher profitability, but under learning strong demand also makes them also more optimistic about the future, further raising these incentives. This can lead to over-investment and investment being concentrated in periods of high demand, which can negatively impact welfare, especially if strong demand coincides with high costs of investment.

I perform a counterfactual experiment whereby the top two firms merge, which eliminates competition and strategic considerations between them. Under the merger, total investment from 2006 to 2014 decreases by 7.5%. The merger leads to a producer surplus increase of $14 billion and a consumer surplus decrease of $1 billion, resulting in a total welfare gain in this setting.[13] I also find that allowing coordination among firms through the merger lowers investment rates and mutes boom-and-bust investment cycles.[14] Furthermore, I compare the results from the merger counter-

[11]Regardless of the presence of learning, if firm profits are concave in demand, increased volatility will decrease investment. The concavity of the profit function may arise from non-linearity in marginal costs, for example.

[12]See, for example, Mankiw and Whinston (1986) and Spence (1977) for theoretical work on this issue.

[13]There are two caveats for the consumer surplus figure. First, it is calculated only with respect to the Asia-Europe market. Second, the magnitude is likely a lower bound since it does not take into account the price competition channel as I maintain the assumption of marginal cost pricing. The model does capture strategic considerations in investment decisions.

[14]US antitrust policy prohibits firms in the same business from colluding on investment decisions, while Japan allows cooperation among rivals along this dimension. O'brien (1987) argues that Japan's support for coordinated

factual under learning and full information. I find that the effects of the merger on both the volume and volatility of investment are higher under learning compared to full information, which suggests that the interaction between learning and dynamic strategic incentives amplifies investment cycles. During high demand periods in which firms have greater strategic incentives to steal business from and preempt rivals, learning also leads firms to collectively become more optimistic, which reinforces strategic incentives.

1.1 Related Literature

This paper builds on a body of literature that studies uncertainty and agents' beliefs in a learning framework. At the 2000 Ely Lecture, Hansen (2007) argued that the rational expectations approach endows agents with too much information and advocated putting econometricians and economic agents on comparable footing.[15] This paper also relates to the literature connecting uncertainty and cyclical investment (e.g. Bernanke (1983), Pindyck (1991), and Dixit (1994)). A few recent papers empirically investigate how uncertainty impacts investment and business cycles (e.g. Fajgelbaum et al. (2017) and Kozlowski et al. (2016)).[16] With much of the literature focusing on aggregate outcomes, this paper contributes to the literature by analyzing implications for firm-level decisions and within-industry investment cycles.

In the area of learning, empirical IO studies have predominantly explored learning about firms' private information (e.g. Jovanovic (1982)), learning about a new technology and spillovers across firms (e.g. Covert (2014) and Hodgson (2018)), or consumers' learning about values of experience goods through experimentation (e.g. Dickstein (2011)). Doraszelski et al. (2018) examine learning about competitors' strategy and demand elasticity parameters in the context of the UK electricity market.

This paper makes a methodological contribution to the literature on the structural analysis of industry dynamics (see Doraszelski and Pakes (2007) for an overview of this literature). Theoretical work includes Hopenhayn (1992) and Ericson and Pakes (1995), and recent empirical papers that

decision-making in investment is partially responsible for the country's success in the steel industry.

[15] Studies that adopt this approach include Cogley and Sargent (2005) that study the role of the Federal Reserve's changing beliefs in monetary policy and Orlik and Veldkamp (2014) that show that uncertainty shocks are countercyclical through a learning model.

[16] Fajgelbaum et al. (2017) develop a model of social learning where uncertainty about fundamentals discourages investment. Kozlowski et al. (2016) show how even transitory shocks can produce persistent effects based on a model in which agents re-estimate the distribution of shocks as they observe new shocks.

adopt the "full-solution approach" similar to this paper include Benkard (2004), Goettler and Gordon (2011), and Igami (2017). The typical approach in modeling beliefs in relation to investment decisions in a dynamic oligopoly setting is to adopt full-information assumptions (e.g. Ryan (2012) and Collard-Wexler (2013)). This paper incorporates learning as a belief-formation process in a dynamic oligopoly framework in order to capture agents' changing beliefs and information sets. It shows that this extra dimension of uncertainty about the underlying model parameters can be crucial in understanding firm behavior in a volatile environment. In particular, the paper shows that allowing such uncertainty creates a new informational channel through which demand fluctuations affect investment, contributing to the body of empirical studies that quantify the effect of demand uncertainty on investment (e.g. Collard-Wexler (2013) and Kellogg (2014)). The framework also makes it possible to study the interaction between firms' information and strategic incentives and its impact on firm investment.

This paper complements empirical studies on the shipping industry including Kalouptsidi (2014), Kalouptsidi (2018), Brancaccio et al. (2019), and Greenwood and Hanson (2015). Kalouptsidi (2014) studies investment cycles in the bulk shipping industry.[17] Kalouptsidi employs a fully rational model and uses second-hand ship prices to identify the value of owning a ship non-parametrically. As the second-hand prices already reflect sellers' and buyers' beliefs about future demand, the author indirectly incorporates firms' beliefs in the estimation of the value of owning ships. By contrast, this study models firms' forecasting process explicitly. This approach will be useful in cases where the industry does not have an active second-hand market or the second-hand market suffers from selection problems.[18] Understanding how firms form expectations is interesting in its own right as well.

Greenwood and Hanson (2015) introduce an alternative and complementary set of behavioral assumptions by considering biases in persistence in earnings and long-run endogenous supply responses by rivals to explain bulk shippers' investment behavior. In the behavioral model, as in

[17] Although bulk and container shipping industries share many similar characteristics, there is stark difference in terms of market power with much higher concentration in the container shipping industry. Kalouptsidi (2014) assumes that each firm owns one ship only and develops a competitive model of the bulk-shipping industry. Also, container shippers operate according to fixed schedules, whereas bulk shippers operate on-demand services much like taxis.

[18] Adverse selection may arise in the second-hand market if sellers privately observe the quality of the goods. If there is selection, the quality of goods traded in the second-hand market may differ from the quality of goods currently owned by firms. In this case, estimating the value of owning the goods from second-hand prices will lead to biased estimates.

the full-information model, agents' perception about the demand process stays fixed (although the perception is allowed to be different from the true process). By contrast, the learning model incorporated in this paper allows uncertainty about the process itself. This is motivated by the fact that agents may perceive the process to change over time especially in the container shipping industry with a relatively short history with changes in demand, technology, and regulations. This approach also allows me to separately quantify the impact of uncertainty about the process and can also be easily incorporated into standard dynamic oligopoly models deployed in the IO literature.

The remainder of the paper is organized as follows. Section 2 describes the industry and the data. Section 3 presents suggestive evidence of firm uncertainty based on forecasts of GDP and trade. Section 4 presents the dynamic model of investment with learning for the shipping industry. Section 5 describes the estimation procedure and discusses estimation results. Section 6 discusses counterfactual experiments. Section 7 concludes.

2 Industry and Data

2.1 Container Shipping Industry

The container shipping industry's core activity is the transportation of containerized goods over sea according to fixed schedules between named ports. The containers come in two standard dimensions (the twenty-foot dry-cargo container (TEU) or the forty-foot dry-cargo container (FEU)), which makes it easier to load, unload, and stack the cargo. The container ships transport a wide range of consumer goods and intermediate goods such as electronics, machinery, textiles, and chemicals. Container trade accounts for over 15% of global seaborne trade by volume and over 60% in value (Stopford (2009)).

Container shipping is a capital-intensive industry in which companies invest in capital by purchasing vessels. The price of building a ship fluctuates depending on the conditions of the shipbuilding and shipping markets at the time of the order, including freight rates, the strength of trade demand, the size of the order book, and expectations.[19] Firms can also scrap old ships that cannot be operated profitably. The demolition prices depend on the demand for scrap metal and

[19]The construction of new ships occurs at shipyards. There are approximately 300 major shipyards and many smaller ones globally.

the availability of ships for scrap.

Container carriers rely on chartered vessels in addition to their own vessels, which are leased out by third parties. Chartered vessels account for approximately 50% of the total container ship capacity operated by the largest 20 firms. The majority of charter contracts for container ships are time charters that involve the hiring of a vessel for a specific period of time with the average contract length of 7-10 months (Reinhardt et al. (2012)). The charterer has operational control of the ships, while the ownership and management of the vessel remain in the hands of the shipowner. This paper focuses on the investment decisions of ship operators as opposed to non-operators due to limited data on individual non-operators' investment. Nevertheless, my model takes into account how charter rates will change with the demand conditions and the individual and aggregate operator-owned volumes. The behavior of non-operating ship owners and the role of the rental market in mitigating volatility faced by the ship operators would be a fruitful area for future research.

The industry is vulnerable to sharp swings in global trade demand, but it is hard for firms to respond quickly to supply-demand imbalances in the short run. There is a gap between the time of placing a new order and the time of receiving the ordered ships due to time-to-build ranging from two to four years. Moreover, whereas bulk shippers can easily move their idle ships into lay-up, container shippers are limited in this respect due to their pre-announced schedules (Stopford (2009)). Since the shipping rates depend directly on the supply of ships relative to demand, the ability to make correct forecasts about future demand and invest accordingly is important in this industry.[20]

Investment is extremely volatile and is highly correlated with the price of new ships as shown in figure 2.[21] Although the price is on average 42% higher compared to the 2009-2014 period, the volume of new orders is higher by more than 60% in the 2006-2008 period.

[20]Freight cost is the most important criterion for customers, although other factors such as transit time, schedule reliability, and frequency of departure matter as well (Reinhardt et al. (2012)).

[21]The prices of building a new ship and the number of ships in the industry order book are available by size category (2500 TEU, 3700 TEU, 6700 TEU, 8800 TEU, 10000 TEU, and 13500 TEU). I first obtain per TEU shipbuilding prices for each size category and construct the weighted average of these prices. The average scrap value is constructed in a similar way.

2.2 Data

This project uses two main datasets on the container shipping industry. The first combines data collected from two sources: MDS Transmodal, a U.K.-based research company, and Clarksons Research, a U.K.-based ship-brokering and research company. This dataset covers quarterly information from 2006 to 2014. The key information includes: (1) volumes and prices of container trade by trade route; (2) firm-level information on the number and the capacity of ships that each firm owns, charters, and has in its order book as well as the capacity deployed in each of the routes the firm operates on; and (3) industry-level charter rates, scrap prices, and shipbuilding prices.

Estimating firms' beliefs for the sample period from 2006 to 2014 requires historical price and quantity data that extend further back than 2006, ideally from the inception of the industry. The first dataset on firm-level investment and capital is therefore supplemented with the historical price and quantity data from the *Review of Maritime Transport* published by the United Nations that goes back to 1997.[22] It contains information on the average freight rates and cargo flows on major routes. The volume of trade is available annually in this dataset, although the price level is available at the quarterly level. The quarterly volume of container trade is imputed based on the data on the value of trade by origin-destination pair from the IMF Direction of Trade Statistics database.[23]

The analysis focuses on major trade routes that together account for approximately 55% of all interregional container trade by trade volume and approximately 60% of deployed ship capacity (see Figure 9 in Appendix B for the average prices on these trade routes from 1997 to 2014). Shipping firms can adjust their capacity across different routes relatively easily and the network of services is constantly changing to meet the needs of trade (Stopford (2009)). For this reason, Kalouptsidi (2014) considers the shipping market to be a global one and includes a single demand component in her model.

In my application, I account for demand in the Asia-Europe (A-E) market separately from demand in other major markets. A-E demand was the main driver of the building boom in the mid-2000s with many ships were built specifically for this market in this period.[24] Moreover, many

[22] Although this is roughly the start date of the official public data on the aggregate price and quantity of container trade, firms may have longer historical data and use them in forming expectations. Section 5.4 discusses my empirical strategy in estimating firms' beliefs given the truncated nature of the price and quantity data.

[23] The imputation assumes that the quarterly container trade volume is proportional to the value of trade in each year.

[24] For example, Maersk ordered eight E-class container ships (of size 14,770 TEU) from 2006 to 2008 all of which

of these newly built ships could not be easily redeployed to other markets due to size restrictions. The A-E traffic mainly goes through the Suez Canal, where restriction on ship size is approximately 18,000 TEU. The traffic on the next largest trade route – the Asia-North America route – transits the Panama Canal (if going to or from the East Coast of North America) which has a smaller size restriction of approximately 5,000 TEU (expanded to approximately 14,000 TEU in June 2016) and West Coast ports are also limited in their ability to handle large ships. Over 90% of ships in the orderbook in my sample were above the 5,000 TEU size restriction, indicating that many of these newly ordered ships are too large to fit through the Panama Canal.

The analysis focuses on firms that deployed over 80,000 TEU of ships quarterly on the Asia-Europe route on average in the 2006 to 2014 period. These firms account for more than 95 % of the total capacity of ships deployed in the Asia-Europe market. This results in a quarterly panel of 17 firms from 2006 to 2014. There is no entry into or exit by these firms during this period. Table 1 provides the summary statistics for this dataset. On average, firms in the sample own 300,000 TEU in capacity, charter 310,000 TEU, and have an order book of 180,000 TEU.

The market structure is more concentrated compared to the bulk shipping industry with more than 40% of total capacity concentrated in the top three firms.[25] Nevertheless, the industry is considered relatively unconcentrated based on the Herfindahl index that is below 1000 (see Figure 10 in Appendix B for the distribution of firm size based on owned capacity).

3 Preliminary Analysis

In this section, I explore suggestive evidence of firm learning and uncertainty based on auxiliary forecast data. The idea is that although direct survey data on firms' beliefs about container shipping demand are not available, I can compare beliefs about closely-related objects such as GDP and overall trade volume, with beliefs about container trade implied under different informational assumptions. Various industry reports and experts confirm that the industry indeed relies on these forecasts when making projections about the container trade industry.[26]

were intended to be operated on the A-E route.

[25] Kalouptsidi (2014) shows that the largest fleet share is 3% for Handysize bulk carriers.

[26] The United Nations ESCAP uses projections about GDP growth in their study of container trade growth and writes that "[g]rowth in the container trade is ultimately driven by economic growth. An underlying assumption of this study is that, for the next decade at least, the structural relationships between the growth in container trade and economic growth will remain basically unchanged" (accessed on July 18, 2019. https://www.unescap.

Given the availability of rich forecast data and its importance in understanding oversupply with many of newly built ships designated directly for this market, I focus on the Asia to Europe market for this exercise. For GDP forecasts, I use the Survey of Professional Forecasters published by the European Central Bank reports. The survey not only reports mean forecasts, but also asks each forecaster to allocate subjective probabilities to ranges of possible outcomes with a width of 0.5 percentage point.[27] This allows me to construct a variance measure of the forecasts by computing an individual forecaster's variance and taking the mean across forecasters.[28] I also obtain the import volume forecasts from the OECD Economic Outlook. I include 14 European countries that are consistently in the data and construct one-year-ahead forecasts on year-on-year growth in imports for this region.[29] Figure 12 in Appendix B shows the forecast data. It is notable that the variance of the forecast jumps in 2008 and maintains the high level through 2014, which already suggests that a full-information model with constant volatility will not be able to match the forecasts successfully.

In order to obtain expectations about container trade implied under full information, I estimate the the process for container trade volume. Under this assumption, estimation involves using the full sample of data from 1997 to 2014, that is, as much data as available to the researcher. To be as flexible as possible on the specification of the process, I consider a general class of time-series models, or an autoregressive integrated moving average ($ARIMA$) process, with varying values for the order of the autoregressive part, the order of the moving average part, and the degree of differencing. I put a time-trend when the degree of differencing is zero. I also explore specifications with time-varying volatility in which the error terms follow a GARCH process in addition to constant-volatility specifications with normally distributed errors. Appendix C.1 provides full details on the specifications and parameter values that are explored including a behavioral model.

I evaluate these candidate specifications following a standard approach based on the Akaike

org/sites/default/files/pub_2398_ch3.pdf). The UNCTAD's Development in International Seaborne Trade writes, "UNCTAD analysis is pointing to continued growth in world seaborne trade that hinges on the continued improvement of the global economy. In line with projected growth in world gross domestic product (GDP), UNCTAD expects global maritime trade to grow by another 4 percent in 2018" (accessed on July 18, 2019. https://unctad.org/en/PublicationChapters/rmt2018ch1_en.pdf).

[27]For example, forecasters are asked to assign a probability to real GDP rising between 0.0% and 0.4%, 0.5% and 0.9%, and so on. I use two-year-ahead forecasts for GDP because there is substantial bunching in the forecasters' probabilities in end bins for one-year forecasts. The bunching makes it difficult to construct variance estimates.

[28]For example, Abel et al. (2016) and Bowles et al. (2007) use the same measure as part of their measure of uncertainty in output growth.

[29]The included countries are Belgium, Czech Republic, Finland, France, Germany, Iceland, Ireland, Italy, Netherlands, Portugal, Spain, Sweden, Switzerland, and United Kingdom.

information criterion (AIC)/ Bayesian information criterion (BIC).[30] The values of the criteria for each specification are listed in table 10. The simplest model of $AR(1)$ with constant volatility minimizes both the AIC and BIC, but I also report results with time-varying volatility, specifically $GARCH(0,1)$.[31]

To evaluate expectations implied by a learning model, I focus the simplest $AR(1)$ process and non-Bayesian learning, as Bayesian learning requires additional assumptions about priors. In this model, since agents do not the know parameters of the model (the slope, constant, and variance of the error), they use past observations to estimate these parameters and form their expectations based on them. The estimation then amounts to fitting the $AR(1)$ process using historical data including observations up to the most current period, that is, $\{Q_\tau\}_{\tau=1}^{t}$ at each t where Q_τ denotes the logged quantity of container trade at τ.[32]

Table 2 reports the correlation between the forecasts of GDP and trade and the forecasts of container trade implied by models with and without learning. The learning model is able to generate beliefs that are highly correlated with both the GDP and trade forecasts. The variance of GDP forecasts is especially matched well with the correlation coefficient of 0.9. By contrast, the full-information model cannot match this pattern even when the volatility of the error term is allowed to vary. The correlation is 0.17 under the GARCH process and zero by construction under the constant-volatility specification. The mean container trade forecasts generated by full-information models are also less consistent with the GDP or trade forecasts. The correlation is negative for the GDP forecast as full-information models predict that the growth rate is higher during periods of weak demand.[33] This is true even when I add richness to the specification through adding more lags or moving average terms.

While the results from this exercise alone do not provide direct evidence of firm uncertainty or reveal which exact model and information firms are using to form beliefs, they suggest that incorpo-

[30] The AIC and BIC provide selection criteria for the model that better fits the data penalizing the number of parameters. The AIC imposes a weaker penalty on the number of parameters. The AIC is given by $AIC = (2k - 2\ln(\hat{L}))/N$ where k is the number of parameters in the model, L is the maximized value of the likelihood function of the model, and N is the number of observations. The BIC is given by $BIC = (\ln(N)k - 2\ln(\hat{L}))/N$.

[31] A model with the lowest number of parameters may be favored due to the relative short time-series. As a robustness check, I repeat the exercise with more lags ($p = 2, 3, 4$), and find that this does not change the correlation estimates significatly. The $GARCH(1,1)$ specification results in a non-significant, negative GARCH coefficient.

[32] In the structural estimation section, I explore various alternative specifications of learning including Bayesian learning. Here, I consider the simplest form of adaptive learning with the same weight given to all past observations.

[33] This is because the AR(1) process has the mean-reversion property. So with the constant parameter estimates as in the full-information model, the expected growth rate is larger when current demand is lower.

rating learning might help us in approximating firm beliefs in this environment. This is important because firm beliefs are key inputs when considering investment decisions or other dynamic decisions such as entry and exit, but they are often hard to capture especially in a volatile environment with potential structural changes. Moreover, a learning model provides a unified information structure with a minimal departure from the rational expectations benchmark that can explain investment fluctuations observed in the data. It is also parsimonious, providing computational tractability as discussed in more detail in section 4.

4 Model

Motivated by the finding from the previous section, I propose a dynamic model of firm investment that allows uncertainty about the demand process. The model builds on the dynamic oligopoly framework developed by Ericson and Pakes (1995) and the learning literature in macroeconomics. Firms' beliefs about demand change over time as firms re-estimate the parameters of the demand process using up-to-date information available to them. In each period, a firm decides whether to invest in new ships and whether to scrap existing ships based on its own capital and order-book levels, and rivals' aggregate capital and order-book levels as well as its beliefs about future demand. In the product market competition stage, firms decide on how much capacity to charter (lease from a third-party chartering company) and how much capacity to deploy in each market. I make several simplifying assumptions motivated by institutional details and patterns in the data, which I will discuss throughout this section. I start by describing the environment in 4.1, and demand for container shipping services and period profits in section 4.2. Section 4.3 describes the model of firm beliefs, and section 4.4 presents firms' dynamic problem. Section 4.5 provides a definition of equilibrium.

4.1 Environment

Time is discrete with an infinite horizon and is denoted by $t \in \{0, 1, 2, ...\}$. There are n incumbent firms and the set of incumbent firms is denoted by $N = \{1, 2, ..., n\}$. Firms are heterogeneous with respect to their firm-specific state, $x_{it} = (k_{it}, b_{it})$, where k_{it} is the capacity of ships owned by firm i

and b_{it} is the backlog, or the capacity of firm i's order book.[34] Based on the discussion in section 2.2, the model takes into account demand in the Asia-Europe and other markets separately.[35] Then, the underlying industry state is $s_t = ((x_{it})_i, d_t)$ where $(x_{it})_i$ is the list of all incumbents' firm-specific states and $d_t = (z_t, \tilde{z}_t)$ includes the demand states of the Asia-Europe market and the outside market. Note that I still allow firms' values and actions to depend on the parameters summarizing firms' beliefs (through indexing the value function with these parameters).[36]

The timing of events is as follows: (1) Firms observe their current state as well as their private cost shocks associated with investing and scrapping. They update their beliefs about demand. (2) Firms make investment and scrapping decisions. (3) Firms choose how much capacity to charter and how much capacity to deploy in the Asia-Europe market and the outside market. They receive period profits. (4) The dynamic decisions are implemented and the delivery and depreciation outcomes are realized. The industry evolves to a new state.

Computing a Markov perfect equilibrium (in which each incumbent firm follows a Markov strategy that is optimal when all competitors follow the same strategy) for this setting is infeasible due to the curse of dimensionality. As the number of incumbent firms grows, the number of states grows more than exponentially.[37] To address this challenge, I consider an alternative equilibrium concept that can be viewed in the context of the moment-based Markov Equilibrium (MME) of Ifrach and Weintraub (2016), or more broadly the experience-based equilibrium (EBE) of Fershtman and Pakes (2012).

In MME, firms keep track of and condition their strategies on the detailed state of strategically important firms (dominant firms) and a few moments of the distribution describing non-dominant

[34] The owned capacity space denoted by $\mathcal{K}$ is discretized into 19 points such that $\mathcal{K} = \{k_0, k_1, k_2, ..., k_{18}\}$ and the order book capacity space denoted by $\mathcal{B}$ into seven points such that $\mathcal{B} = \{b_0, b_1, ..., b_6\}$. $\mathcal{K}$ and $\mathcal{B}$ are both discretized in 100,000 TEU increments such that $k_0 = 0$ TEU, $k_1 = 100,000$ TEU, and so on, and $b_0 = 0$ TEU, $b_1 = 100,000$ TEU, and so on. The computational constraints limit us to a coarse discretization of the state space as in many papers estimating a dynamic oligopoly model (see, for example, Benkard (2004) and Collard-Wexler (2013)). An earlier version of the paper uses a coarser state space with only 15 points for k_{it} and 5 points for b_{it} and produces similar results.

[35] To reiterate, although it is natural to consider a shipping market to be a global one, since firms can reposition ships across markets relatively easily, the geographical features (due to the size restrictions on the Suez and Panama Canals) make it difficult for firms to move large ships positioned in the Asia-Europe market to other markets. Moreover, aggregating multiple routes in the outside market helps reduce the size of the state space.

[36] The belief parameters are not included as state variables, since firms use their current beliefs as the best forecasts of future demand, implying that the beliefs are fixed objects from the firms' current perspective. This assumption and its implications are discussed in more detail in section 4.4.

[37] There are 17 active firms in my application. Even a simple specification with a single state variable that can take up to five different values would result in over one billion of states.

firms' states, instead of the detailed state of all incumbents. The main difference between MME and oblivious equilibrium (OE) introduced by Weintraub et al. (2008) is that MME relaxes the assumption that firms believe that the average industry state holds at any time. My application allows firms to keep track of their own firm-specific states, the sum of all incumbents' states, and the aggregate demand states. Firms' strategies thus depend on the firm-specific state, $x_{it} = (k_{it}, b_{it})$, and the *moment-based* industry state defined as $\hat{s}_t = (\sum_i x_{it}, d_t)$.[38] In appendix E.2, I consider a version that allows richer information by adding a dominant firm's state into the moment-based industry state and show that the model predictions are robust to this change.

4.2 Period Profit

Each market has two directional routes (e.g. Asia to Europe and Europe to Asia in the Asia-Europe market). Firms face constant elasticity demand in each route they operator in:

$$\log Q_{jt} = z_{jt} + \alpha_1 \log P_{jt} \tag{1}$$

where z_{jt} denotes the demand state, P_{jt} the price, and Q_{jt} the quantity on route j at time t.

In each period, firms choose (a) how much capacity to charter (h_{it}), and (b) how much capacity to allocate to the Asia-Europe market ($\bar{q}_{it}$) and the outside market ($\tilde{q}_{it}$) given their present state.[39] In other words, a firm chooses how much of its total capacity to allocate to the Asia-Europe market or the outside market where the total capacity is determined as the sum of its chartered and owned capacity (i.e. $k_{it} + h_{it} = \bar{q}_{it} + \tilde{q}_{it}$).

The marginal cost of providing services on a route is linearly increasing in quantity up to the

[38] MME strategies are not necessarily optimal, however; there may be a profitable unilateral deviation to a strategy that depends on the detailed state of all firms. This is because the moment-based state may not be sufficient statistics to predict the future evolution of the industry.

[39] I model the chartering and capacity allocation decisions as simultaneous actions instead sequential actions of deployment followed by chartering in order to capture the relative flexibility of chartering and frictions in changing deployment. The average charter contract length is 7-10 months on average, but firms operate under multiple charter contracts at once such that they have many margins through which to adjust chartered capacity in response to demand shocks. In addition, moving ships across markets is not frictionless due to ship size issues and pre-announced schedules.

firms' capacity constraint as given in equation (2).

$$mc(q_{ijt}, \bar{q}_{it}) = \begin{cases} a + \frac{bq_{ijt}}{\bar{q}_{it}} & \text{if } q_{ijt} \leq \bar{q}_{it} \\ \infty & \text{otherwise.} \end{cases} \tag{2}$$

This functional form implies that (i) the marginal cost increases as the firm's quantity gets closer to the firm's full capacity; and (ii) firms with higher capacity have a lower marginal cost of producing the same level of quantity. This assumption is based on three institutional details. First, it becomes increasingly difficult to schedule loading and unloading as the ship reaches its full capacity. Second, firms that deploy greater capacity on the route have a relative cost advantage due to the fact that operating expenses such as crew, insurance, and administration offer scales of economies (Stopford (2009)). Lastly, this functional form allows me to indirectly account for the fact that larger firms tend to have larger ships and thus higher fuel efficiency without having to include the size of ships as a state variable.[40] This functional form is also similar to that of Ryan (2012), where the marginal cost increases as firms operate closer to maximum capacity, but only after a threshold is reached.

I assume that in the product market firms are not withholding capacity strategically given their capacity constraints (although still making positive profits due to the convex costs). This assumption is motivated by two observations from the data. First, a regression analysis suggests that variation in the level of competition across routes does not does not explain variation pricing.[41] Second, the "effective capacity" (the total capacity of ships firms make available on the market as a share of the total capacity of ships they have) remains relatively stable over time even in the face of huge excess capacity during the post-crisis period (see Figure 11 in Appendix B). The share stays at 92% on average in 2006-2014 with a dip to 87% in 2009 -2010. Given the assumption, the supply

[40] The correlation coefficient between the capacity on the A-E market and ship size is 0.83.

[41] A regression of prices on market competitiveness suffers from the classic problem that market structure is endogenous. To address this problem, I instrument the number of firms on the route from region A to region B with the number of firms on routes that are "neighbors" to region A. Note that demand for route A to B is largely driven B's demand for imports from A. But firms serving this route have to serve route B to A as well. Therefore, an increase in firms operating routes connecting A and other regions will increase the number of firms on route B to A (, hence, route A to B) for reasons that are unrelated to demand and cost conditions of route A to B. I find that the coefficient on the number of firms is negative in the OLS estimation, but becomes insignificant in the IV estimation. The coefficients are reported in table 9 in Appendix B.

curve for route j is given as the horizontal sum of all firms' supply curves as follows:

$$P_{jt} = a + \frac{bQ_{jt}}{\bar{Q}_t} \quad \text{if } Q_{jt} \leq \bar{Q}_t \tag{3}$$

where $\bar{Q}_t = \sum_i \bar{q}_{it}$. The price in the Asia-Europe market is determined by the intersection of the demand curve given in equation (1) and the supply curve given in equation (3).

The period profit is the sum of profits from providing shipping services on the Asia to Europe and the Europe to Asia routes plus the profit from the outside market minus the charter cost and the fixed cost of capital:

$$\begin{aligned}\pi(x_{it}, \hat{s}_t) = \max_{\bar{q}_{it}, h_{it}} \Bigg\{ & \underbrace{\Big(\sum_{j \in \{1,2\}} P_{jt} q_{ijt} - c(q_{ijt}, \bar{q}_{it}) \Big)}_{\text{Profit from A-E market}} \\ & + \underbrace{R(\tilde{q}_{it}, \tilde{Q}_t, \hat{s}_t)}_{\text{Profit from outside market}} - \underbrace{CC(h_{it}, x_{it}, \hat{s}_t)}_{\text{Charter cost}} - \underbrace{FC \cdot k_{it}}_{\text{Fixed cost of capital}} \Bigg\}\end{aligned} \tag{4}$$

where FC is the fixed cost of holding one unit of capital, R is the profit from the outside market, CC is the charter cost, and $\tilde{q}_{it}$ is the capacity deployed in the outside market. The fixed cost of holding ships includes all costs that do not vary with the output level (or how full the ships are) such as docking fees, maintenance costs, canal dues, and port charges. I do not explicitly model the chartering market and the product market competition in the outside market but account for them in a reduced-form way.[42] The detailed specification of the reduced-form functions for the charter cost and the outside-market profit is given in section 5.2.

4.3 A Model of Firms' Beliefs about Demand

I consider classes of learning models most widely used in explaining macroeconomic fluctuations including adaptive learning (which is a frequentist approach to learning and is sometimes referred to as least-squared learning) and Bayesian learning. I use the adaptive learning model as a benchmark for several reasons. It does not require the estimation of prior beliefs, which can be challenging with

[42] It is hard to infer capacity allocated to each small market because of the granularity of the data. Thus, explicitly modeling product market competition in the outside market will require many strong assumptions about how capacity is allocated across all individual markets. Similarly, I do not explicitly model the charter market due to limited data on individual non-operator ship owners.

the relatively short time-series data. More importantly, the adaptive framework has a parameter that allows agents to put heavier weights on more recent observations, which help match the investment fluctuations and the correlation between investment and demand observed in the data better than Bayesian learning or other alternative models.[43] Evans and Honkapohja (2012) show that this is a natural way to form expectations if agents were concerned about the possibility of structural changes, and several empirical papers have used this model to empirically explain macroeconomic fluctuations (see, for example, Milani (2007)). Appendix C.3 presents results under the alternative model with Bayesian learning.

Under adaptive learning, agents form expectations about demand based on information available to them in each period. They operate like econometricians who estimate the parameters of the model based on the best information at their disposal and make forecasts using their estimates. Agents contemplate a first-order autoregressive model for demand in the Asia-Europe market, denoted by z_t, as the following:

$$\begin{aligned} z_t &= \rho^0 + \rho^1 z_{t-1} + \omega_t \\ &= \rho y_t + \omega_t \end{aligned} \tag{5}$$

where $\omega_t \sim N(0, \sigma^2), \rho = [\rho^0, \rho^1]'$, and $y_t = [1, z_{t-1}]'$ (see section 2.2 for the discussion on why this market is modeled separately from other markets). Similarly, the model for demand in other major markets, henceforth called the "outside market" ($\tilde{z}_t$), is given as:

$$\begin{aligned} \tilde{z}_t &= \tilde{\rho}^0 + \tilde{\rho}^1 \tilde{z}_{t-1} + \tilde{\omega}_t \\ &= \tilde{\rho}' \tilde{y}_t + \tilde{\omega}_t \end{aligned} \tag{6}$$

where $\tilde{\omega}_t \sim N(0, \tilde{\sigma}^2), \tilde{\rho} = [\tilde{\rho}^0, \tilde{\rho}^1]'$, and $\tilde{y}_t = [1, \tilde{z}_{t-1}]'$.[44] Agents are uncertain about the parameters in the demand model, $\{\rho^0, \rho^1, \sigma, \tilde{\rho}^0, \tilde{\rho}^1, \tilde{\sigma}\}$. Thus, they revise their expectations by re-estimating

[43] There are alternative models of firm beliefs that are not explored in this paper such as the Cognitive Hierarchy model (Camerer et al. (2004)) and the Level-k model (Costa-Gomes and Crawford (2006); Crawford and Iriberri (2007)). This type of models would be more appropriate for empirical settings in which there is a clear hierarchy of beliefs among firms that can be established empirically, for example, due to the heterogeneity in access to information or technologies across firms.

[44] Allowing correlation in demand in these two markets is straightforward and does not change the results quantitatively and qualitatively.

these parameters in each period based on demand realizations up to time t, $\{z_\tau, \tilde{z}_\tau\}_{\tau=0}^t$. At each t, firms' beliefs about demand can be described by the estimates of the $AR(1)$ parameters, denoted as $\eta_t = (\rho_t^0, \rho_t^1, \sigma_t, \tilde{\rho}_t^0, \tilde{\rho}_t^1, \tilde{\sigma}_t)$.

Firms are assumed to have homogenous beliefs about the aggregate demand. The prices and volumes of container trade are public information periodically published in trade journals and other publications. Moreover, swings in global trade demand common to all firms are the main source of demand shocks in this industry.[45] The model also assumes that agents use their current beliefs in forecasting demand (see Appendix E.1 for an approach to relax this assumption and a discussion of implications for estimation results as well as challenges in implementation).[46] This assumption has two behavioral interpretations. The first interpretation is that agents believe current beliefs to be the correct or best forecasts for future demand. The alternative interpretation is that agents use current beliefs in forecasting as these approximate future beliefs.

Let $Y_t = [y_0, y_1, ..., y_t]'$ and $R_t = \frac{Y_t' Y_t}{t}$. The expectations at time t regarding the Asia-Europe market demand under adaptive learning can be written recursively as

$$\rho_t = \rho_{t-1} + \lambda_t (R_t)^{-1} y_t \left(z_t - \rho_{t-1}' y_t\right) \tag{7}$$

$$R_t = R_{t-1} + \lambda_t (y_t y_t' - R_{t-1}) \tag{8}$$

where λ_t is the weight parameter that governs how responsive the estimate revisions are to new data (Evans and Honkapohja (2012)). Figure 8 in appendix B plots the relationship between relative weights placed on observations and the value of λ_t. If $\lambda_t = \frac{1}{t}$, agents put equal weight on all observations in their information set. If λ_t is some constant between 0 and 1, weights geometrically decline with the age of observation such that agents assign heavier weights to more recent observations. A larger value of λ leads to heavier discounting of older observations. For

[45] On a practical level there are no publicly available data that provide information on firm-level demand to my knowledge, which would be necessary to allow heterogenous firm beliefs. Nevertheless, heterogeneity in firms' beliefs would arise if firms experienced different demand shocks, for example, through different customer pools. How firms form heterogeneous beliefs and how they affect firm decisions and industry dynamics are interesting topics of study for future work.

[46] This implies that agents do not internalize the possibility of learning in the future. In the context of this paper, since information about the aggregate trade demand is exogenous to agents' actions, there is no room for experimentation regardless of the assumption about learning. In contrast, suppose information is endogenous to agents' decisions, for example, because agents are making consumption decisions for experience goods for which quality is difficult to observe in advance. In this case, the assumption of myopic learning rules out experimentation, whereas allowing agents to internalize learning in the future may encourage experimentation.

example, when $\lambda = 0.03$, agents put a 30% weight on a 10-year-old observation relative to the most current observation, while when $\lambda = 0.02$, they put a 45% weight on a 10-year-old observation.

4.4 Firms' Dynamic Problem

Firms make an investment decision ($\iota_{it} \in \{0, 1\}$) and a scrapping decision ($\delta_{it} \in \{0, 1\}$) in order to maximize expected discounted profits.[47] I denote the strategy profile as $\mu_{it} = (\iota_{it}, \delta_{it})$. Each investing firm pays an investment cost. The investment cost consists of a part common to all firms that is a function of the aggregate state, $\kappa(\hat{s}_t)$, and a privately observed part of the cost, $\varepsilon^{\iota}_{it} \sim N(0, (\sigma^{\iota})^2)$. If a firm decides to scrap its ships, the firm receives a scrap value. The scrap value is the sum of the value common to all firms, $\phi(\hat{s}_t)$, and an iid private value distributed as $\varepsilon^{\delta}_{it} \sim N(0, (\sigma^{\delta})^2)$. I assume that there is also a fixed rate (ζ) at which involuntary scrapping occurs when the ships become inoperable in which case the firm also receives the same scrap value.[48] I denote as $\nu(\delta_{it}, x_{it})$ the expected amount of capital reduction from voluntary and involuntary scrapping before the realization of the depreciation outcome such that $\nu(\delta_{it}, x_{it})$ is 1 if $\delta_{it} = 1$, and ζk_{it} otherwise.

The value function of a firm before observing its private shocks can be written as

$$V^{\eta_t}(x_{it}, \hat{s}_t) = E\Big[\max_{\iota_{it}, \delta_{it}} \pi(x_{it}, \hat{s}_t) - \iota_{it}\left(\kappa(\hat{s}_t) + \varepsilon^{\iota}_{it}\right) + \nu(\delta_{it}, x_{it})\left(\phi(\hat{s}_t) + \varepsilon^{\delta}_{it}\right)$$
$$+\beta E\left[V^{\eta_t}(x_{it+1}, \hat{s}_{t+1} | x_{it}, \hat{s}_t)\right]\Big]$$

where η_t is the vector of parameters summarizing firms' beliefs in period t about future demand. The first expectation is over $(\varepsilon^{\iota}_{it}, \varepsilon^{\delta}_{it})$ and the second expectation is over the evolution of the state $(x_{it}, \hat{s}_t)$. The value function is a function of η_t as it depends on how firms perceive the demand state to evolve. Note that the problem is still stationary due to the assumption that firms use the

[47] Firms are restricted to invest and/or scrap up to only one unit (100,000 TEU) per period due to the discretization of the state space and computational burden as described in section 4.1. In the data there are no observations of a capital reduction by more than one unit and there are only three instances of an investment of more than one unit. Capping the maximum investment level to one unit for each firm reduces the action space, thus significantly alleviating the computational burden.

[48] Ship scraps are valuable primarily for their steel, and thus are priced based on the amount tonnage rather than the age or functionality of the ships. If a firm scraps its vessels, there is no involuntary scrapping in the same period such that the maximum reduction in k_{it} is one unit. This assumption is made because the data do not provide any observations of a capital reduction by more than one unit. The interpretation of this assumption can be that when a firm decides to scrap its vessels, it chooses the oldest vessels that are about to deprecate on their own. This assumption can be easily relaxed.

current period's beliefs (η_t) in forecasting future demand (see Appendix E.1 for the discussion on relaxing this assumption.) This means that the continuation value is a function of η_t only, and not η_{t+1}, η_{t+2}, and so on. Firms pick the action that maximizes the net present value such that

$$(\iota_{it}^*, \delta_{it}^*) = \arg\max_{\iota_{it}, \delta_{it}} \pi(x_{it}, \hat{s}_t) - \iota_{it} \left(\kappa(\hat{s}_t) + \varepsilon_{it}^{\iota}\right) + \nu(\delta_{it}, x_{it}) \left(\phi(\hat{s}_t) + \varepsilon_{it}^{\delta}\right) \\ + \beta E \left[V^{\eta_t}(x_{it+1}, \hat{s}_{t+1} | x_{it}, \hat{s}_t)\right]$$

The current model does not allow for persistent heterogeneity in the investment costs and scrap values across firms. The transaction-level pricing data on investment and demolition indicate that there is no significant firm heterogeneity at least in the observed transaction prices of investment and scrapping. The model incorporates firm heterogeneity in other areas, however, since it may be important given the persistent concentration of market power. First, the cost of chartering ships from a third party is allowed to depend on firm size, as larger firms may have greater bargaining power. Second, the marginal cost of production depends on the capacity of the firm's deployed ships as described in section 4.2.

State Transitions

When a firm invests, the order book capacity increases by one unit when there is no delivery at t and stays constant if there is delivery. A firm's own capacity is determined by scrapping decision, and depreciation and delivery outcomes. The transition of the firm-specific state is described as:

$$k_{it+1} = k_{it} + \tau_{it} - \min(\delta_{it} + \psi_{it}, 1)$$
$$b_{it+1} = b_{it} + \iota_{it} - \tau_{it}$$

where τ_{it} is delivery and ψ_{it} is depreciation. The probability of delivery is a linear function of the firm's order-book capacity such that the delivery occurs with the probability of ξb_{it} for some constant ξ. Similarly, the probability of depreciation is ζk_{it} such that it linearly increases in the capital stock. The perceived evolutions at time t of the aggregate demand states for the Asia-Europe market and the outside market follow first-order autoregressive processes as follows:

$$z_t = \rho_t^0 + \rho_t^1 z_{t-1} + \omega_t$$
$$\tilde{z}_t = \tilde{\rho}_t^0 + \tilde{\rho}_t^1 \tilde{z}_{t-1} + \tilde{\omega}_t$$

where $\omega_t \sim N(0, \sigma_t^2)$ and $\tilde{\omega}_t \sim N(0, \tilde{\sigma}_t^2)$.[49] This process is described in more detail in section 4.3. The parameters in the $AR(1)$ model, $\eta_t = (\rho_t^0, \rho_t^1, \sigma_t, \tilde{\rho}_t^0, \tilde{\rho}_t^1, \tilde{\sigma}_t)$, summarize the beliefs about the evolution of future demand at time t. How firms update these beliefs as they receive new information is described in section 4.3.

4.5 Equilibrium

The value function can be re-written as the perceived value of a firm using moment-based strategy μ' in response to all other firms following strategy μ:

$$\hat{V}^{\eta}_{\mu',\mu}(x, \hat{s}) = E\left[\pi(x, \hat{s}) - \iota\left(\kappa(\hat{s}) + \varepsilon^{\iota}\right) + \nu(\delta, x)\left(\phi(\hat{s}) + \varepsilon^{\delta}\right) + \beta E_{\mu',\mu}\hat{V}^{\eta}(x', \hat{s}'|x, \hat{s})\right].$$

The definition of an equilibrium is then given as follows.

Definition An equilibrium comprises an investment and scrapping strategy μ that satisfies the following conditions:

(a) Firm strategies satisfy the optimality condition:

$$\sup_{\mu' \in \mathcal{M}} \hat{V}^{\eta}_{\mu',\mu}(x, \hat{s}) = \hat{V}^{\eta}_{\mu}(x, \hat{s}) \quad \forall (x, \hat{s}) \in \mathcal{X} \times \hat{\mathcal{S}}.$$

(b) The perceived transition kernel is given by:

$$\hat{P}_{\mu} = \Phi P_{\mu}$$

where $\hat{P}_{\mu}$ is the transition kernel of the moment-based state when firms use strategy μ,

[49] I explore alternative specifications including a case in which the errors in the AR(1) processes follow heavier-tailed t-distributions and a case in which correlation between demand in the Asia-Europe market and demand in the outside market is allowed. Main results are robust to these alternative specifications.

P_μ is that of the underlying state, and Φ is an operator such that $\hat{P}_\mu$ approximates the process of the moment-based state, P_μ.[50]

The equilibrium is computed using an algorithm based on value-function iteration. Appendix D describes the algorithm in detail.

5 Estimation and Empirical Results

The estimation of the dynamic model of investment with learning proceeds as follows. First, I estimate demand for shipping services to recover the elasticity of demand and demand states. Second, I estimate parameters governing static competition, including the marginal cost of production, the charter cost, and the outside-market profit, which are used to compute period profits. Third, I calibrate the investment cost and the scrap value based on the pricing data of new and scrapped ships as well as other model primitives such as the delivery and depreciation processes. Fourth, I discuss the empirical implementation of the learning model. Finally, I estimate the dynamic model through the method of simulated moments.

5.1 Estimating Demand for Shipping Services

The goal of this section is to estimate the price elasticity of demand and to construct demand states for the Asia-Europe market and the outside market.[51] The empirical analogue of the constant elasticity demand model in equation (1) is:

$$\log Q_{jt} = \alpha_0 + \alpha_1 \log P_{jt} + \alpha_2 W_{jt} + \varepsilon_{jt} \tag{9}$$

where j is an indicator for trade routes, Q_{jt} is the amount of container shipping services in terms of TEU, P_{jt} is the average price per TEU, and W_{jt} is a demand shifter. I estimate equation (9) using instrumental variables regression in order to correct for the endogeneity of prices. Price is instrumented with the average size and age of ships and the proportion of ships that are over 20 years old. The size of ships is one of the key determinants of cost efficiency as larger ships require

[50] In practice, the moment-based industry state's evolution is defined as the long-run average of observed transitions from the moment-based state in the current period to the moment in the next period consistent with strategy μ.

[51] This section follows the demand estimation of Kalouptsidi (2014) closely.

less fuel per TEU on average. The age of ships matters as well, since older ships tend to require higher maintenance costs. Log GDP for the destination area is used as a demand shifter.

The estimation uses data from six major trade routes from 2001:Q2 to 2014:Q4.[52] The demand parameters are identified by the time-series variation as well as the cross-sectional variation across six different routes in the data along with the constant elasticity functional form assumption. In particular, since ships have to travel back and forth on the two routes in each market they serve, two routes in the same market (e.g. Asia to Europe and Europe to Asia) have the same level of supply while facing different demand shocks, which helps the identification of the demand parameters.

The price elasticity of demand is estimated to be -3.89 (see Table 12 in Appendix C.2 for detailed results). This implies that a change in price from \$1510 per TEU to \$1360 per TEU would result in a change in quarterly quantity demanded of approximately 0.92 million TEU on the Asia to Europe route. Stopford (2009) explains that container trade is price elastic because lowering prices encourages the substitution of cheap foreign substitutes for local products. Moreover, other transportation modes are available, such as road and rail transportation and air freight. Kalouptsidi (2014) estimates demand for bulk shipping to be more elastic at -6.17.

Given the elasticity of demand estimates, I construct the demand state for each trade route (z_{jt}) as the intercept of the demand curve:

$$z_{jt} = \hat{\alpha}_0 + \hat{\alpha}_2 W_{jt} + \hat{\varepsilon}_{jt} \tag{10}$$

where $\{\hat{\alpha}_0, \hat{\alpha}_2\}$ are parameters estimated from the regression and $\hat{\varepsilon}_{jt}$ is the residual. Finally, I construct aggregate demand states for the Asia-Europe market and the outside market from the route-level demand states. For the Asia-Europe market, I take the demand state for the Asia to Europe direction. Since the container trade volume is less than half on the opposite Europe to Asia direction, firms' investment and capacity deployment decisions in the market are mostly dictated by the trade demand on the Asia to Europe direction. For the outside market, I take the sum of the demand states in the non-Asia-Europe routes. Figure 1 plots the demand states for 1997 to 2014 for the Asia-Europe and the outside markets. There is a large drop in demand in both markets

[52] Although the price, quantity, and GDP data are available from 1997, the instruments are available starting from 2001:Q2. The included trade routes are Asia to Europe, Europe to Asia, Asia to North America, North America to Asia, Europe to North America, and North America to Europe.

around 2008-2009.

5.2 Estimating the Profit Function

The second step of the estimation is to construct period profits by estimating the marginal cost, charter cost, and outside market profit functions. Firms' capacity deployment decisions yield a supply curve that, along with the demand curve, determines the equilibrium prices and quantities for the Asia-Europe market. The marginal cost of providing container shipping services is specified in equation (2), which serves as the basis for the maximum likelihood estimation of the cost parameters (a, b).

The outside market profit and the charter cost functions are specified in a reduced-form way as:

$$R(\tilde{q}_{it}, x_{it}, \hat{s}_t) = \tilde{q}_{it}\left(r_0 + r_1\tilde{z}_t + r_2\tilde{Q}_t\right)$$

$$CC(h_{it}, x_{it}, \hat{s}_t) = h_{it}(\gamma_0 + \gamma_1 z_t + \gamma_2 k_{it} + \gamma_3 K_t).$$

where $K_t = \sum_i k_{it}$. The profit from each unit of capacity deployed in the outside market is allowed to depend on the total deployed capacity in the outside market $(\tilde{Q}_t)$ since higher supply may lead to fiercer price competition and lower profit. The charter cost depends on firm-level own capacity (k_{it}). Charter cost is also allowed to depend on the total capacity owned by operator (K_t) as it is likely to affect demand for chartering.[53]

The estimation of the charter cost and outside market profit functions is based on firms' static profit maximization problem.[54] Given the demand estimates, I estimate these objects via maximum likelihood based on the first-order conditions with respect to the capacity deployed on the Asia-Europe route $(\bar{q}_{ijt})$ and the chartering decisions (h_{it}), respectively. The variations in capacity deployment and charter decisions across different firm types and across time along with the first-order conditions and the functional form assumptions provide identification for these parameters.

Table 13 in Appendix C.2 reports estimates of the profit function parameters. The coefficients on

[53] I estimate an alternative specification that allows the charter rate to depend on the outside market demand, and find that the coefficient is not significantly different from zero.

[54] I use implied charter costs from observed chartering decisions instead of calibrating charter costs using the available data due to several limitations to directly using the observed charter rates. Since the rates are negotiated between the operator and the ship owner, larger operators may get discounts even though we only observe the average rates. Moreover, the observed charter rate only includes the transferred amount from the ship operator to the owner, and does not reflect other costs associated with operating chartered vessels such as port costs and crew costs.

the Asia-Europe market demand state in the outside market profit and charter cost functions (r_1 and γ_1) are positive. This implies that stronger demand leads to higher outside market profits as well as higher charter costs. The estimates also show that when there is more aggregate deployed capacity in the outside market, firms earn less from that market on average, which captures competitive effects. In addition, larger firms tend to face lower charter costs, and an increase in total industry capacity owned by ship operators lowers charter costs. Finally, the sign on the marginal cost parameter b is positive, capturing the fact that the cost increases as the firm operates closer to its capacity constraint. The estimates also suggest that there is substantial cost heterogeneity across firm size. For example, the marginal cost at $q_{it} = 100,000$ TEU is 38% higher for firms with maximum capacity of 200,000 TEU than for firms with a constraint of 400,000 TEU.[55]

5.3 Estimating Other Model Primitives

Data on investment costs and scrap values that are typically unavailable in other settings allow a flexible specification and estimation of the model of firm beliefs. I use industry-level shipbuilding and demolition price data to estimate the investment cost and the scrap value, respectively, as functions of the industry state variables (industry owned-ship and order-book capacities, and demand states for the Asia-Europe and outside markets) via least squares. Figure 3 compares investment costs and scrap values observed in the data to predicted values obtained from the regression (see Table 14 in Appendix C.2 for the detailed estimates).

The delivery process of newly ordered ships and the depreciation process of existing ships are also estimated separately from the estimation of dynamic parameters. The mean delivery rate is estimated based on a simple regression of delivery on the firm's order-book size with no constant.[56] For the depreciation process, I set an exogenous rate. This is because the data do not differentiate between depreciation and the scrapping of ships that can still be operated physically. Thus, the depreciation rate and the distribution of the private shocks to the scrap value can not be separately identified. The depreciation rate, ζ, is set such that the average age at which ships naturally depreciate is 20 years.[57]

[55] Note, however, that the marginal cost of operating at $x\%$ of a firm's capacity constraint is constant regardless of the firm size by construction.

[56] The current formulation assumes that the delivery rate depends solely on the firm's own order-book size, since the industry order-book size does not have a statistically significant effect on the delivery rate.

[57] Although historically the lifespan of container ships was 25 to 30 years, it has fallen in recent years especially for

5.4 Empirical Implementation of the Learning Model

This section discusses the implementation of the learning model described in section 4.3 and presents expectations about demand implied by the model. The truncated nature of the price and quantity data for container trade poses a challenge in implementing the learning model. An agent's information set in each period includes all observations from the past. However, although firms may have access to observations from the inception of the industry, the researcher may not. This problem arises in most empirical settings when dealing with a learning model. In my particular setting, data on prices and quantities for major trade routes are available starting from 1997, although the first international voyage dates back to 1966.

Given this challenge, I explore two alternative methods of empirically implementing an adaptive learning model: the truncation approach and the imputation approach. Based on the model fit, I adopt the former. Appendix E.3 discusses details of the imputation approach and the pros and cons of the two approaches. It also demonstrates that beliefs are robust to these alternatives.

The approach I adopt entails setting the initial period of the information set as the start date of the data. I consider the weight parameter $\lambda_t = \frac{1}{t}$ as well as $\lambda \in (0, 0.04]$ in increments of 0.0025.[58] If $\lambda_t = \frac{1}{t}$, equal weights are applied to all past observations. In practice, the estimation procedure under this parameter value amounts to applying least squares to estimate equation (5) for each period separately. The regression at period t uses demand-state data from the first period to period t, or $\{z_\tau, \tilde{z}_\tau\}_{\tau=0}^{t}$ (see section 5.1 for steps of recovering these states). For cases of a constant λ_t, weighted least squares are applied where the weight on an observation from the period τ is given by $(1 - \lambda_t)^{t-\tau}$.

Figure 4 shows firms' demand parameter estimates from 2000 to 2014 under adaptive learning with $\lambda_t = 0.02$ for the Asia-Europe market (see Figure 13 in Appendix C.2 for the outside market). The estimates in the shaded area are for 2006 to 2014, which will be used in the estimation of the dynamic model. The estimate of the persistent parameter ρ_t^1 rises from 2006 to 2007 and shows a general downward trend thereafter. The variance parameter σ_t spikes in early 2009 and stays high throughout the end of the sample period.

larger ships. *Vesselvalues* reports that the the average age of all sizes of container ships sold for scrap was around 22 years old and the average age at which a Post-Panamax container ship was sold for scrap was around 19.5 years.

[58]Orphanides and Williams (2005) suggest that a constant gain parameter in the range between 0.01 and 0.04 matches the data on expectations well.

Under adaptive learning, the degree to which the parameter estimates react to recent events grows as agents put more weight on recent observations (as shown in Figure 14 in Appendix C.2). For example, the degree to which σ_t jumps around 2009 is the smallest in the case where agents weigh all past demand realizations equally ($\lambda_t = 1/t$). When λ_t is a constant, the larger λ_t, the larger the jump in σ_t around 2009. Similarly, the larger the fall in the persistence parameter ρ^1 in the post-2008 period, the larger λ_t becomes. It is this variation in beliefs and the variation in the data on investment and scrapping around demand shocks that identify the model of firm beliefs. The identification of λ_t is discussed in more detail in section 5.5.

5.5 Estimating the Dynamic Model of Investment with Learning

The last and most computationally intense step of the estimation entails estimating the model of firm beliefs and the dynamic parameters. The typical empirical strategy of estimating a dynamic game of investment is to recover objects such as investment costs, entry costs, and exit values by searching for parameters that minimize the distance between actions observed in the data and those that the parameters imply (e.g. Collard-Wexler (2013) and Ryan (2012)). This paper instead employs data on shipbuilding and demolition prices to estimate investment costs and scrap values as described in section 5.3, which opens up the possibility of identifing the model of firms beliefs. Although the application is different, the underlying logic of this approach is similar to that of Hortacsu and Puller (2008) in which the authors use marginal cost data to quantify how much firms' bidding deviates from the optimal bidding benchmark.

I employ the method of simulated moments (MSM) to estimate the dynamic model, which minimizes a distance criterion between key moments from the actual data and the simulated data. Let θ denote the vector of dynamic and belief parameters such that $\theta = (\sigma^\iota, \sigma^\delta, FC, \lambda_t)$. I solve for an equilibrium of the dynamic investment model and obtain the optimal investment policy function for each candidate parameter vector.[59] Using equilibrium strategies obtained in the previous step, I simulate the equilibrium path for the 2006 to 2014 period $S = 1000$ times. And from these paths,

[59] Recently, empirical techniques have been proposed to estimate the dynamic industry equilibrium without having to solve for an equilibrium (e.g. Aguirregabiria and Mira (2007), Bajari et al. (2007), Pakes et al. (2007)). The first stage of this approach entails recovering firms' policy functions by regressing observed actions on observed state variables. The second stage involves estimating structure parameters that make these policies optimal. This approach relies on flexible functional forms in the first step, so the data requirement is too high given the global nature of my data set. I use a full-solution method instead, which involves solving the model at every guess of the parameter, but is more efficient.

I obtain the simulated moments as follows:

$$\Gamma(\theta) = \frac{1}{S} \sum_{s=1}^{S} \Gamma_s(\theta).$$

I search for the parameter vector that minimizes the weighted distance between the data and simulated moments given as:

$$f(\theta) = \left(\Gamma^d - \Gamma(\theta)\right)' W \left(\Gamma^d - \Gamma(\theta)\right). \tag{11}$$

where Γ^d is the set of data moments.[60]

The moments used in the estimation include the average investment before and after 2008, the volatility of investment, the correlation in demand and investment, and the aggregate capacity of owned and backlogged ships. Table 3 lists these moments and compares the data moments and simulated moments under the parameter estimates.

The results (reported in Table 4) indicate that the weighting parameter estimate is $\lambda_t = 0.02$. I will refer to the adaptive learning model with $\lambda_t = 0.02$ as the *baseline learning model* in the rest of the paper. This implies that agents put approximately 45% weights on a 10-year-old observation compared to the most recent observation. This estimate is very close to the values that previous studies in macroeconomics have estimated based on aggregate survey data such as the Survey of Professional Forecasts or micro data on expectations. For example, Malmendier and Nagel (2016), Milani (2007), and Orphanides and Williams (2005) estimate the constant-gain parameter (λ_t) to be 0.0175, 0.0183, and 0.02, respectively, with respect to expectations about macroeconomic conditions and monetary policy. Figure 5 shows that the baseline learning model does well at predicting the investment boom in 2007 and the plunge in investment in 2009.

The fixed cost of holding one unit of capital (100,000 TEU) in one quarter is estimated to be \$25 million, which is approximately 36% of the period's profit from one unit of capital (where the period profit is the sum of profits from the Asia-Europe market and the outside market minus the charter cost and does not include the investment cost and scrap value). This fixed cost includes all costs

[60] The search is done over grids of $(\sigma^{\iota}, \sigma^{\delta}, FC, \lambda_t)$. The grids for σ^{ι} and σ^{δ} are in increments of 0.005 and the grid for FC is in increments of \$50/TEU. The candidate belief parameter values include $\lambda_t = \frac{1}{t}$ and $\lambda_t = \{0.025, 0.05, 0.075, ..., 0.4\}$. I use the inverse of the variance-covariance matrix of the simulated moments as the weighting matrix (W).

that owning and operating ships impose, regardless of the production level, such as maintenance costs, canal dues, and port charges. It also includes the cost of labor needed in the operation of the ships regardless of how full they are.

The identification relies on a revealed-preference argument. I have recovered the values of benefits and costs of each of the options that the firm faces–investment, scrapping, and staying for each state in the state space as described in section 5.3. As a result, given these values, firms' choices in various states observed in the data reveal their expectations about future demand.

More concretely, the estimation relies on the variation in firms' beliefs across different weighting parameter values and the variation in firm behavior across time and firms observed in the data. As firms discount older observations more heavily, their beliefs become more responsive to recent shocks. This will amplify the effect of recent demand shocks on investment, which will increase the correlation between demand and investment. The left panel of Figure 6 illustrates this relationship by plotting the comparative statics of the correlation of demand and investment for different values of λ_t with all other parameters fixed at each estimated value. Similarly, when λ_t increases and firms revise their beliefs more dramatically in response to demand shocks, investment becomes more volatile, as illustrated in the right panel of Figure 6.

In principle, the parameters are identified by both time-series and cross-sectional variations. In this dataset, the main identification is coming from the time-series variation in investment and scrapping as well as investment costs and scrap values, and it is valuable to observe a boom and a bust in my sample period. The shipping industry provides a great setting in that it is exposed to large exogenous fluctuations in demand resulting from cycles in world trade.

In the sample period that this paper focuses on, the financial market also experienced an upheaval. Therefore, one might worry that omitting information on credit market conditions might bias the main results of the paper. In appendix F, I incorporate credit market frictions in the form of collateral constraints in the model and discuss implications.

6 Counterfactual Analysis

I first simulate a model of full information which removes uncertainty about the parameters in the demand process to quantify the effect of the uncertainty. Through the next two counterfactual experiments, I seek to understand the mechanisms through which learning affects firm investment behavior. First, I address the long-standing question on the effect of demand volatility on investment. By applying the learning framework I shed light on the informational channel through which demand fluctuations affect investment. Second, I conduct counterfactuals with respect to competition and industry consolidation by simulating the industry under a merger of the top two firms. This exercise helps us understand how the interaction between strategic incentives and agent beliefs affects investment.

6.1 Uncertainty about the Demand Process

To remove uncertainty about the demand process, I consider a model in which agents know the parameters of the demand process, and the only uncertainty is about what value of demand will be realized due to the variance in the process. This is a naive model that is expected to perform poorly in matching patterns in the data, but the comparison will help us understand the role of information.

The model governing the evolution of demand is given by equations (5) and (6), as in the adaptive learning model. I also consider a specification with time-varying volatility.[61] In this version, demand is assumed to follow the same $AR(1)$ process, but the error terms are assumed to follow a $GARCH(1,1)$ process such that the current period's variance depends on the last period's realized error and variance:

$$\sigma_t^2 = a_0 + a_1\omega_{t-1}^2 + b_1\sigma_{t-1}^2$$

where ω_{t-1} is the realized error in period $t-1$.

In the full-information model, the parameters in the demand model are known to the agents. Then, estimating beliefs under this model involves estimating the demand process using the full sample of data or as much data as are available to the researcher. Beliefs implied by the full-information model are presented in Figure 7 (see table 16 in appendix C.3 for the estimates of the GARCH parameters). Compared to the baseline learning model, the persistent parameter (ρ^1)

[61] As an alternative model of stochastic volatility, I consider a Markov regime-switching model where the variance and persistence parameter are no longer a constant but can take on one of two values. It is rejected that there are two regimes. The specification in which only the variance can switch between two values produces similar results as the GARCH model.

is lower in the pre-2008 period and higher afterward. The jump in the variance around 2009 is substantially larger under the full information model with time-varying volatility compared to the learning model. But the hike is more short-lived and the variance is more volatile, while in the learning model the variance remains high throughout the end of the sample period.

The second column of Table 5 shows the simulated moments under the full-information model. Primitives are re-estimated for each alternative model of beliefs.[62] One of the striking features that arise when agents are endowed with information about the demand process is that the correlation between demand and investment becomes negative such that firms restrain from investing in the peak demand season of 2006-2007 and invest more heavily after 2008. There are two forces driving this prediction. First, under full information firms' beliefs about the underlying demand process remain constant even if they receive a series of high demand draws unlike in the learning model. Hence, due to the mean-reverting property of the autoregressive process, firms expect slower growth in high demand periods. Second, there is an a positive relationship between demand and the shipbuilding prices, which further discourages investment during high demand periods. The results also suggest that investment is less volatile under full information, suggesting that fluctuations in agent beliefs arising from learning amplify cycles of investment.

Surprisingly, allowing time-varying volatility under full information does not have a substantial effect on the level and the timing of investment. The correlation between investment and demand is only slightly less negative compared to the model with constant volatility. This finding provides some insights into the role of firms' beliefs about demand. It shows that although changes in volatility in demand may partially account for the observed boom and bust cycle of investment, the main force driving the cycle was the changes in the level of demand forecasts over time.

6.2 Demand Volatility

Demand volatility can affect investment in several ways. First, as real options theory predicts, an increase in demand volatility raises the cost of investment, since once a firm makes an investment it cannot disinvest should market conditions change adversely. Second, an increase in demand volatility may also increase the volatility of investment costs. Finally, the presence of learning opens up an

[62] Results are not qualitatively different when I fix the primitives at the levels recovered under the baseline learning model.

additional channel through which demand fluctuations affect investment, since increased demand volatility makes agents revise their expectations more often and more drastically.

To quantify the effect of demand volatility, I conduct the following counterfactual simulations. I simulate two sets of demand series for 2006 to 2014–one with high volatility and the other with low volatility. In the high volatility case, the variances in the demand processes for the Asia-Europe and outside markets are doubled from the estimates based on the full sample of data. In the low-volatility case, the variances are halved from the estimates. The remaining parameters and the demand realizations prior to 2006 are set to the estimated levels.

Table 6 shows simulation results for the high- and low- volatility cases under learning and full information, respectively. An increase in demand volatility has a negative effect on investment, which is consistent with findings in previous studies such as Bloom (2009) and Collard-Wexler (2013). Going from low to high volatility reduces investment by 8.6% under learning. This suggests that the value function is concave with respect to demand. If the value function is concave, lower volatility in demand raises the expected value of owning a ship, increasing the average level of investment. In addition, an increase in demand volatility also increases the volatility of investment, as higher demand volatility leads to more volatile shipbuilding prices.

Note that the learning model and the full-information model yield qualitatively and quantitatively different predictions about investment patterns: under learning higher demand volatility generates larger investment boom-and-bust cycles that are more highly correlated with demand cycles. First, the increase in the volatility of investment in response to an increase in demand volatility is higher under learning. This is because when learning is present higher demand volatility also leads to larger changes in firms' expectations about future demand, which further increases the volatility of investment. Under full information, the effect is solely through changes in the shipbuilding price. Second, there is a higher correlation between demand and investment under learning, since learning generates agent beliefs that are more correlated with demand. By contrast, the increase in demand volatility has an insignificant effect on the correlation between demand and investment under full information.

6.3 Strategic Incentives in Investment Decisions

In this section, I study the effects of strategic incentives and industry consolidation as well as how these effects interact with agent beliefs. To deal with the recent excess capacity in the industry, container shipping firms have increasingly moved towards consolidation. In July 2014 Maersk Line and MSC–the world's two biggest container-shipping companies–formed an alliance named 2M, which akin to a code-sharing deal between airlines, is meant to help firms cut costs by using each other's ships and port facilities and reduce competition. More firms are planning mergers and acquisitions as well. Cosco and CSCL, the sixth- and seventh-largest carriers by operated fleet capacity, have proposed a merger. And CMA-CGM has proposed the acquisition of APL.

On one hand, increased consolidation may hurt consumers through reduced competition. On the other hand, there are potential sources of efficiency gains on the producers' side, which makes the final direction of the welfare change ambiguous. In particular, consolidation may reduce the business-stealing effect and preemption motives that can lead to the capital levels higher than the socially optimal level. Mankiw and Whinston (1986) show that the business stealing effect can result in socially inefficient levels of entry when there are fixed costs of entry. Also, many theoretical studies predict that strategic incentives can lead to excess capacity, since firms may use investment as a commitment to deter entry or expansion of rivals (e.g., Spence (1977)).

My model incorporates these strategic incentives that arise in dynamic decisions such as the business-stealing and preemption motives.[63] First, the business-stealing effect arises since a firm's investment in an extra unit of capacity has a negative effect on the price and competitors' profitability, and this negative effect of increasing one's own capital is internalized by all incumbents in the market. Second, as the volume of the industry order book grows and shipyards get closer to their full capacity, the price of building a new ship increases. This generates dynamic incentives for firms to preemptively commit to investment before others do when they expect strong demand.

In order to quantify the effects of reduced competition and strategic interaction that would result from industry consolidation, I consider a merger between top two firms that jointly account for over

[63] An important caveat is that I only account for strategic incentives that arise in dynamic competition, but not the incentives that arise in static competition due to the assumption of marginal cost pricing in the product market.This assumption has been made based on two facts. First, the industry is relatively unconcentrated with the Herfandahl index below 1000. Second, the route's level of competitiveness measured by the number of active firms does not lead to changes in the price.

35% of total capacity. Then, to understand how strategic effects interact with agent learning, I compare results from the merger counterfactuals under the assumptions of full information and learning. Firms have greater strategic considerations during periods of strong demand with high profitability. But, under uncertainty about the demand process, when firms experience a series of strong demand, they also revise their beliefs about future demand upward, which reinforces strategic incentives.

The increased size of the merged firm would not only reduce strategic interaction, but would also result in changes in costs, bargaining power with the charterer, etc. Therefore, to disentangle the effect of strategic incentives from the effect arising from a change in the firm size distribution, I assume that the merged firm maximizes joint profits, but still operates two plants, which keeps the investment and chartering costs comparable to the case with no merger.

As shown in Table 7 Panel A, reducing competition externalities through a merger has a substantial effect on investment: from 2006 to 2014, investment drops by 7.5%. Investment falls heavily for the merging firms by 40%, but also falls for non-top-two firms by 2.5%. In addition to the level of investment, the timing of investment changes as well. In particular, the volatility of investment and the correlation between investment and demand decrease under the merger case. Compared to the baseline case, investment is relatively less concentrated in times of strong demand when the price of investment is also very high (see Figure 17 in Appendix C.4 for yearly investment under different scenarios). These results suggest that strategic interaction among firms not only is responsible for raising investment as many theoretical studies suggest, but also raises the volatility of investment. In terms of welfare, this means that the producer surplus almost doubles while consumer surplus drops only slightly by 1%.

I compare results from the merger counterfactuals under full information and learning in panel B of Table 7. The results show that, compared to the learning model, the full-information model underestimates changes in investment resulting from the merger, and thus underestimates welfare changes, especially the producer surplus gain. For example, the predicted decrease in investment resulting from the merger is approximately 14,000 TEU under the learning model, compared to 9,000 TEU under the full-information model. The underestimation of the changes in investment volatility is also dramatic (25,000 TEU compared to 4,000 TEU). The fact that the effect of strategic incentives is greater under learning sheds light on the relationship between the competitive forces and firm beliefs. Strong demand for shipping raises firms' strategic incentives, for example, to preemptively commit to investment and to steal business from others. But under learning strong demand also makes agents more optimistic, which amplifies these strategic incentives. These results also suggest that if a regulator used the full-information model in evaluating this merger, he may underestimate the welfare gains.

7 Conclusion

Information can play an important role when firms make long-term capital investments while facing large fluctuations and uncertainty in demand. This paper empirically examines the role of information in generating investment boom-and-bust cycles and overcapacity in the context of the container shipping industry. I develop a dynamic oligopoly model of investment that incorporates an additional layer of uncertainty – uncertainty about the aggregate demand process. In this model, agents form expectations about demand using the best information available to them in each period, and use their changing forecasts in their investment and scrapping decisions. This learning framework provides a computationally tractable way to incorporate more sophisticated information sets. This is useful particularly for settings that are highly volatile and complicated with potential structural breaks such that it is hard for economic agents and researchers to observe and estimate the underlying process of interest.

A key empirical strategy of the paper is to employ data on shipbuilding and demolition prices, which allows me to identify the model of firm beliefs. I find that the uncertainty about the demand process amplifies investment cycles and raises the correlation between investment and demand. The counterfactual analysis reveals the mechanisms through which learning affects investment cycles. I find that under learning higher demand volatility leads to more frequent and larger revisions of expectations about demand, amplifying the magnitude of investment cycles. In addition, by simulating policies that reduce competition among firms, I find that learning strengthens firms' strategic incentives, which also amplifies investment cycles.

References

Abel, Joshua, Robert Rich, Joseph Song, and Joseph Tracy, "The measurement and behavior of uncertainty: evidence from the ECB Survey of Professional Forecasters," *Journal of Applied Econometrics*, 2016, *31* (3), 533–550.

Aguirregabiria, Victor and Pedro Mira, "Sequential estimation of dynamic discrete games," *Econometrica*, 2007, pp. 1–53.

Bajari, Patrick, Jonathan Levin et al., "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 2007, *75* (5), 1331–1370.

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

Bernanke, Ben S, "Irreversibility, Uncertainty, and Cyclical Investment," *The Quarterly Journal of Economics*, 1983, *98* (1), 85–106.

Bloom, Nick, "The Impact of uncertainty shocks," *Econometrica*, 2009.

Bowles, Carlos, Roberta Friz, Veronique Genre, Geoff Kenny, Aidan Meyler, and Tuomas Rautanen, "The ECB survey of professional forecasters (SPF)-A review after eight years' experience," *ECB occasional paper*, 2007, (59).

Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papageorgiou, "Geography, Transportation, and Endogenous Trade Costs," *Econometrica, Forthcoming*, 2019.

Camerer, Colin F, Teck-Hua Ho, and Juin-Kuan Chong, "A cognitive hierarchy model of games," *The Quarterly Journal of Economics*, 2004, *119* (3), 861–898.

Cogley, Timothy and Thomas J Sargent, "The conquest of US inflation: learning and robustness to model uncertainty," *Review of Economic dynamics*, 2005, *8* (2), 528–563.

Collard-Wexler, Allan, "Demand Fluctuations in the Ready-Mix Concrete Industry," *Econometrica*, 2013, *81* (3), 1003–1037.

Costa-Gomes, Miguel A and Vincent P Crawford, "Cognition and behavior in two-person guessing games: An experimental study," *American Economic Review*, 2006, *96* (5), 1737–1768.

Covert, Thomas, "Experiential and social learning in firms: the case of hydraulic fracturing in the Bakken Shale," *Available at SSRN 2481321*, 2014.

Crawford, Vincent P and Nagore Iriberri, "Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions?," *Econometrica*, 2007, *75* (6), 1721–1770.

Dickstein, Michael J, "Efficient provision of experience goods: Evidence from antidepressant choice," *Job Market Paper, Harvard University*, 2011.

Dixit, Avinash K, *Investment under uncertainty*, Princeton university press, 1994.

Doraszelski, Ulrich and Ariel Pakes, "A framework for applied dynamic analysis in IO," *Handbook of industrial organization*, 2007, *3*, 1887–1966.

_ , Gregory Lewis, and Ariel Pakes, "Just starting out: Learning and equilibrium in a new market," *American Economic Review*, 2018, *108* (3), 565–615.

Ericson, Richard and Ariel Pakes, "Markov-perfect industry dynamics: A framework for empirical work," *The Review of Economic Studies*, 1995, *62* (1), 53–82.

Evans, George W and Seppo Honkapohja, *Learning and expectations in macroeconomics*, Princeton University Press, 2012.

Fajgelbaum, Pablo D, Edouard Schaal, and Mathieu Taschereau-Dumouchel, "Uncertainty traps," *The Quarterly Journal of Economics*, 2017, *132* (4), 1641–1692.

Fershtman, Chaim and Ariel Pakes, "Dynamic games with asymmetric information: A framework for empirical work*," *The Quarterly Journal of Economics*, 2012, p. qjs025.

Goettler, Ronald L and Brett R Gordon, "Does AMD spur Intel to innovate more?," *Journal of Political Economy*, 2011, *119* (6), 1141–1200.

Greenwood, Robin and Samuel G Hanson, "Waves in ship prices and investment," *The Quarterly Journal of Economics*, 2015, *55*, 109.

Hansen, Lars Peter, "Beliefs, doubts and learning: Valuing macroeconomic risk," *The American Economic Review*, 2007, *97* (2), 1–30.

Hodgson, Charles, "Information Externalities, Free Riding, and Optimal Exploration in the UK Oil Industry," Technical Report, Working Paper 2018.

Hopenhayn, Hugo A, "Entry, exit, and firm dynamics in long run equilibrium," *Econometrica: Journal of the Econometric Society*, 1992, pp. 1127–1150.

Hortacsu, Ali and Steven L Puller, "Understanding strategic bidding in multi-unit auctions: a case study of the Texas electricity spot market," *The RAND Journal of Economics*, 2008, *39* (1), 86–114.

Ifrach, Bar and Gabriel Y Weintraub, "A Framework for Dynamic Oligopoly in Concentrated Industries," 2016.

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

Jovanovic, Boyan, "Selection and the Evolution of Industry," *Econometrica: Journal of the Econometric Society*, 1982, pp. 649–670.

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

— , "Detection and impact of industrial subsidies: The case of Chinese shipbuilding," *The Review of Economic Studies*, 2018, *85* (2), 1111–1158.

Kellogg, Ryan, "The effect of uncertainty on investment: evidence from Texas Oil Drilling," *The American Economic Review*, 2014, *104* (6), 1698–1734.

Kozlowski, Julian, Laura Veldkamp, and Venky Venkateswaran, "The tail that wags the economy: Belief-driven business cycles and persistent stagnation," 2016.

Malmendier, Ulrike and Stefan Nagel, "Learning from inflation experiences," *The Quarterly Journal of Economics*, 2016, *131* (1), 53–87.

Mankiw, N Gregory and Michael D Whinston, "Free entry and social inefficiency," *The RAND Journal of Economics*, 1986, pp. 48–58.

Milani, Fabio, "Expectations, learning and macroeconomic persistence," *Journal of monetary Economics*, 2007, *54* (7), 2065–2082.

O'brien, Patricia A, "Coordinating market forces: the anatomy of investment decisions in the Japanese steel industry, 1945-1975," *Business and Economic History*, 1987, pp. 205–209.

Orlik, Anna and Laura Veldkamp, "Understanding uncertainty shocks and the role of black swans," Technical Report, National Bureau of Economic Research 2014.

Orphanides, Athanasios and John C Williams, "The decline of activist stabilization policy: Natural rate misperceptions, learning, and expectations," *Journal of Economic dynamics and control*, 2005, *29* (11), 1927–1950.

Pakes, Ariel, Michael Ostrovsky, and Steven Berry, "Simple estimators for the parameters of discrete dynamic games (with entry/exit examples)," *The RAND Journal of Economics*, 2007, *38* (2), 373–399.

Pindyck, Robert S, "Irreversibility, Uncertainty, and Investment," *Journal of Economic Literature*, 1991, *29* (3), 1110–48.

Reinhardt, Forest L, Ramon Casadesus-Masanell, and Frederik Nellemann, "Maersk Line and the Future of Container Shipping," 2012.

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

Spence, A Michael, "Entry, capacity, investment and oligopolistic pricing," *The Bell Journal of Economics*, 1977, pp. 534–544.

Stopford, M, "Maritime Economics, 3rd," 2009.

Weintraub, Gabriel Y, C Lanier Benkard, and Benjamin Van Roy, "Markov perfect industry dynamics with many firms," *Econometrica*, 2008, *76* (6), 1375–1411.

Figures and Tables

Figure 1: Demand States

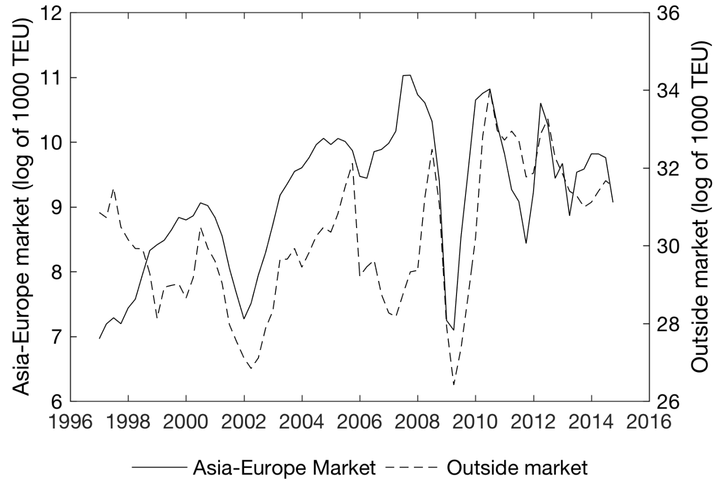

Figure 2: Total Investment and Investment Costs

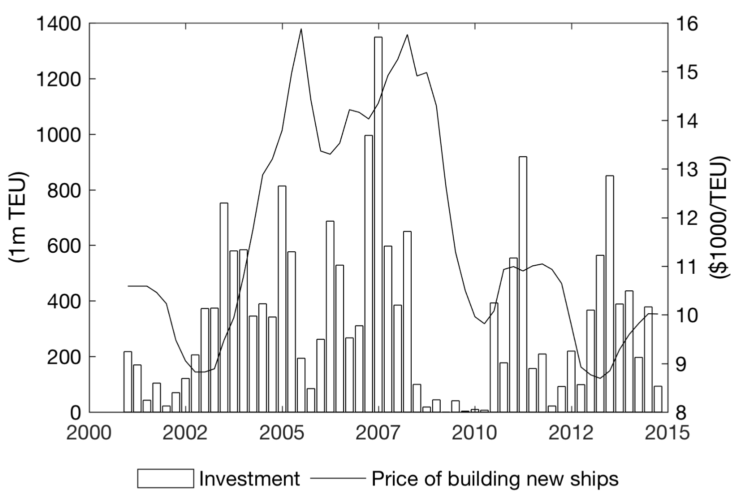

Notes: This figure shows the volume of new orders (left axis) and the average price of building new ships (right axis) from 2001:Q1 to 2014:Q4.

Figure 3: Predicted Investment Costs and Scrap Values

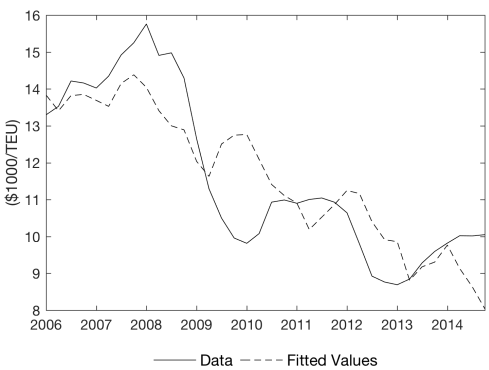

(a) Investment costs

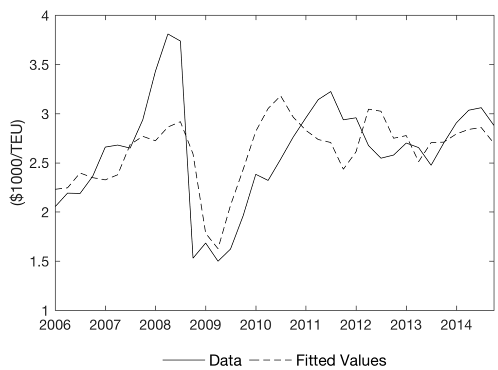

(b) Scrap values

Notes: The left panel shows the average shipbuilding price observed in the data and the predicted shipbuilding price from the regression of the shipbuilding price on the industry state variables. The right panel shows the average scrap value and the predicted scrap value.

Figure 4: Beliefs under Learning for the Asia-Europe Market

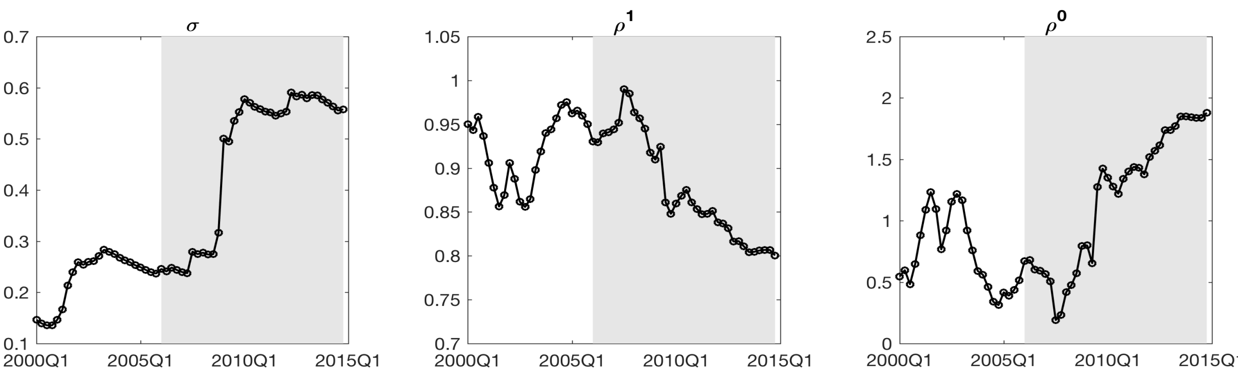

Notes: This figure shows firms' beliefs about demand in the Asia-Europe market for 2000:Q1 to 2014:Q4 under adaptive learning with $\lambda_t = 0.02$. The beliefs are summarized by the three parameters, $\{\sigma_t, \rho_t^0, \rho_t^1\}$, in the $AR(1)$ process as given in equation (5). Beliefs for 2006-2014 in the shaded area are used in the main analysis.

Figure 5: Model Fits

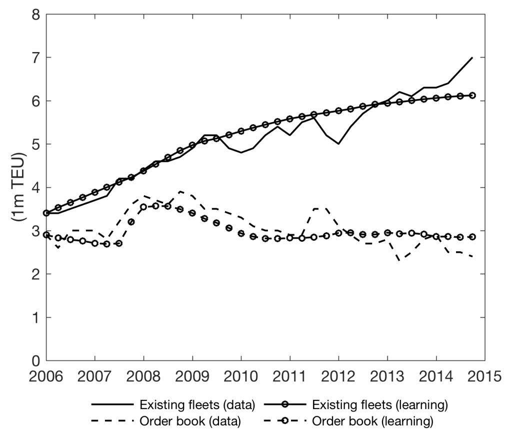

(a) Capacity of owned ships and order book

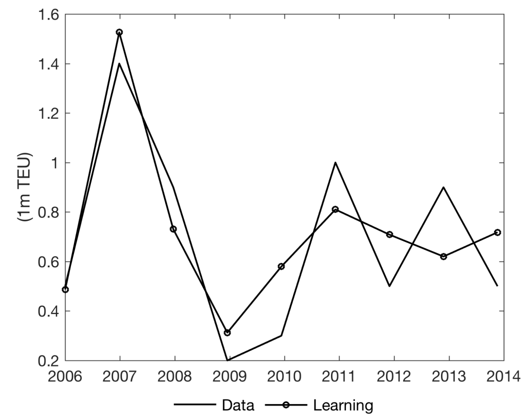

(b) Yearly investment

Notes: The left panel shows the industry evolution simulated under the baseline learning model (adaptive learning with $\lambda_t = 0.02$) and the industry evolution in the data. The right panel shows yearly investment simulated under the baseline learning model and observed in the data, respectively. The simulated moments are based on 1000 equilibrium paths.

Figure 6: Comparative Statics for the Belief Parameter

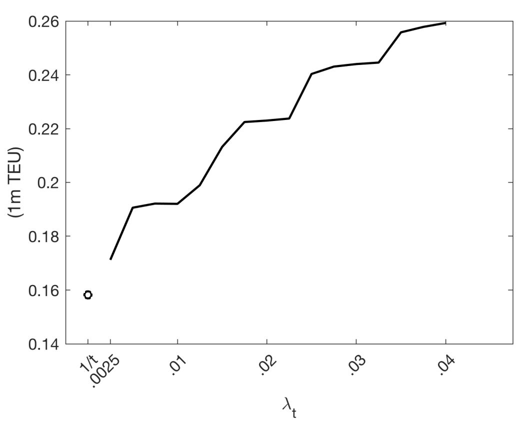

(a) Correlation between demand and investment

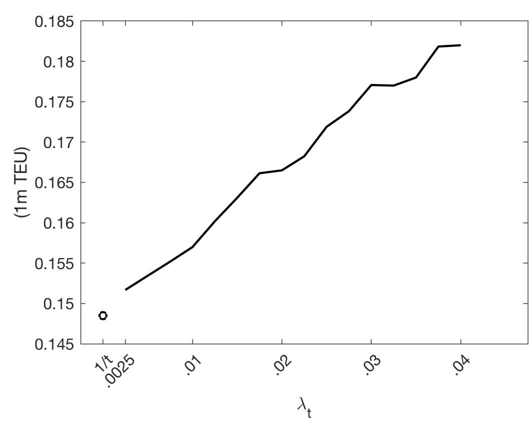

(b) Volatility of investment

Figure 7: Beliefs under Full Information Models of Beliefs for the Asia-Europe Market

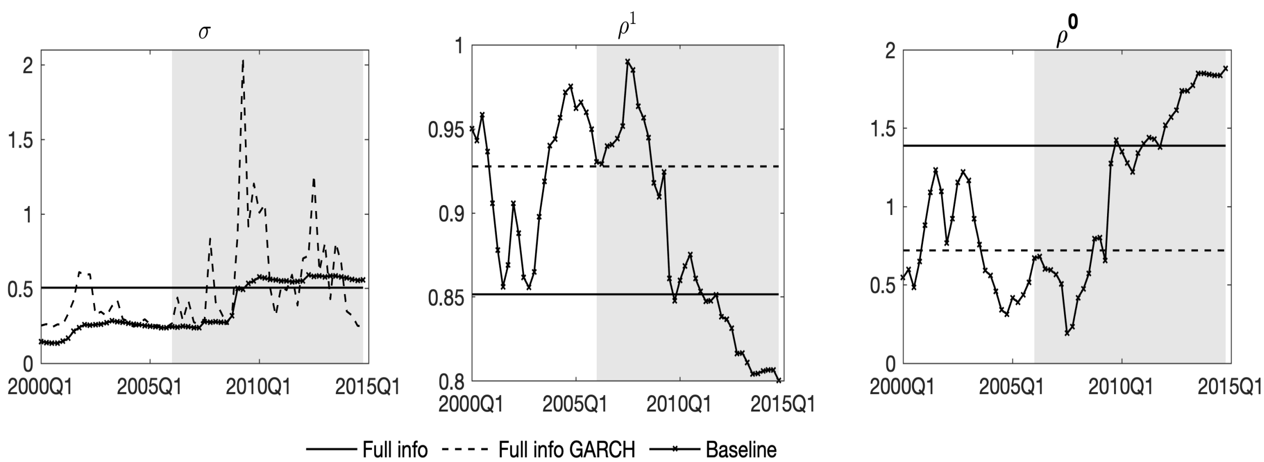

Notes: This figure shows firms' beliefs about demand in the Asia-Europe market for 2000:Q1 to 2014:Q4 under full-information models. Beliefs for 2006-2014 in the shaded area are used in the main analysis. For the full-information GARCH model, the leftmost panel includes an inferred conditional standard deviation of ω_t.

Table 1: Descriptive Statistics

		Mean	Std. Dev.	Min	Max
Industry-level data (2006-2014)					
Shipbuilding price ($1000/TEU)		11.62	2.22	8.69	15.76
Scrap price ($1000/TEU)		2.62	0.55	1.50	3.81
Market-level data (1997-2014)					
Asia to Europe	Quantity (1 million TEU)	2.37	1.10	0.70	3.98
	Price ($1000/TEU)	1.51	0.28	0.80	2.09
Europe to Asia	Quantity (1 million TEU)	1.08	0.39	0.51	1.76
	Price ($1000/TEU)	0.78	0.10	0.57	1.07
Asia to North America	Quantity (1 million TEU)	2.57	0.78	1.12	3.92
	Price ($1000/TEU)	1.67	0.21	1.27	2.20
North America to Asia	Quantity (1 million TEU)	1.20	0.41	0.63	2.14
	Price ($1000/TEU)	0.89	0.16	0.68	1.43
Europe to North America	Quantity (1 million TEU)	0.78	0.15	0.48	1.05
	Price ($1000/TEU)	1.32	0.16	0.93	1.77
North America to Europe	Quantity (1 million TEU)	0.55	0.14	0.32	0.76
	Price ($1000/TEU)	0.99	0.21	0.67	1.60
Firm-level data (2006-2014)					
Capacity of owned ships (1 million TEU)		0.30	0.25	0.04	1.47
Capacity of ships in order book (1 million TEU)		0.18	0.13	0.00	0.64
Capacity of chartered ships (1 million TEU)		0.31	0.29	0.01	1.55
Capacity of ships deployed in Asia-Europe market (1 million TEU)		0.22	0.19	0.04	0.99

Notes: There are 36 industry-level, 216 market-level, and 612 firm-level observations. In addition to the Asai-Europe route, other routes include Asia to North America, North America to Asia, North America to Europe, and Europe to North America routes.

Table 2: Correlation between GDP/ Trade Forecasts and Container Trade Volume Forecasts

	Full information		Learning
	Constant volatility	Time-varying volatility	
GDP - Mean	-0.38	-0.19	0.66
	(0.16)	(0.17)	(0.13)
GDP - Variance	.	0.17	0.90
	.	(0.17)	(0.08)
Trade - Mean	0.06	0.18	0.42
	(0.25)	(0.25)	(0.23)

Notes: Standard errors are in parentheses.

Table 3: Data and Simulated Moments

	Data moments	Simulated moments
Average investment in 2006-2008 (1 million TEU)	0.23	0.23
		(0.03)
Average investment in 2009-2014 (1 million TEU)	0.14	0.15
		(0.02)
Total capacity of owned ships (1 million TEU)	5.12	5.15
		(0.27)
Total capacity in the order book (1 million TEU)	3.01	2.98
		(0.14)
Correlation between demand and investment	0.19	0.22
		(0.12)
Volatility of investment (1 million TEU)	0.17	0.17
		(0.03)

Notes: This table compares moments observed in the data and moments simulated under the estimated parameters. The simulated moments are computed based on 1000 series of equilibrium paths. Standard deviations are in parentheses.

Table 4: Dynamic Parameter Estimates

λ_t	0.02 (0.005)
σ^ι (1 billion US dollars)	0.275 (0.055)
σ^δ (1 billion US dollars)	0.43 (0.092)
FC (1 billion US dollars)	0.025 (0.0051)

Notes: This table shows estimates of dynamic parameters. λ_t is the weighting parameter in the adaptive learning model that governs how heavily agents discount older observations when forming expectations about demand. σ^ι is the standard deviation of the i.i.d. shock around the investment cost of building 100,000 TEU and σ^δ around the scrap value. FC is the fixed cost of holding 100,000 TEU of capacity. Standard errors are in parentheses.

Table 5: Full Information Counterfactuals

	Baseline	Full info	Full info GARCH
Average investment in 2006-2008 (1 million TEU)	0.23	0.17	0.17
	(0.03)	(0.03)	(0.03)
Average investment in 2009-2014 (1 million TEU)	0.15	0.21	0.22
	(0.02)	(0.02)	(0.03)
Total capacity of owned ships (1 million TEU)	5.15	5.05	5.17
	(0.27)	(0.27)	(0.31)
Total capacity in the order book (1 million TEU)	2.98	3.06	3.04
	(0.14)	(0.15)	(0.13)
Correlation between demand and investment	0.22	-0.24	-0.21
	(0.12)	(0.16)	(0.17)
Volatility of investment (1 million TEU)	0.17	0.14	0.15
	(0.03)	(0.02)	(0.02)

Notes: Standard deviations are in parentheses. Primitives are re-estimated for each alternative model of beliefs.

Table 6: Demand Volatility Counterfactuals

Model	Learning		Full info	
Volatility	High	Low	High	Low
Investment (100,000 TEU)	1.48	1.62	1.35	1.57
Volatility of investment (100,000 TEU)	0.75	0.45	0.47	0.31
Corr. between demand and investment	0.04	0.02	-0.12	-0.13
Consumer surplus (1 billion US dollars)	112.60	85.30	113.27	84.11
Producer surplus (1 billion US dollars)	24.59	33.28	26.84	35.08
Total surplus (1 billion US dollars)	137.19	118.58	140.12	119.18

Notes: This table shows results from demand volatility counterfactuals. The owned capacity, order book, and investment are reported as the average over time, and the welfare measures as the sum over the entire period. Consumer surplus is calculated with respect to the Asia-Europe market only.

Table 7: Merger Counterfactuals

Panel A: Industry Outcomes and Welfare	
	Merger(%Δ)
Owned capacity (1 million TEU)	5.02 (-2.53)
Orderbook (1 million TEU)	2.78 (-6.80)
Investment (1 million TEU)	0.17 (-7.50)
Volatility of investment (1 million TEU)	0.15 (-14.66)
Correlation between investment and demand	0.15 (-30.46)
Consumer surplus (1 bil. US dollars)	81.69 (-1.13)
Producer surplus (1 bil. US dollars)	28.85 (93.80)
Total surplus (1 bil. US dollars)	110.54 (13.36)
Investment by top two firms (1 million TEU)	0.01 (-40.12)
Investment by other firms (1 million TEU)	0.15 (-2.47)
Owned capacity of top two firms (1 million TEU)	1.43 (-5.59)
Owned capacity of other firms (1 million TEU)	3.59 (-1.25)
Producer surplus of top two firms (1 bil. US dollars)	25.88 (105.15)
Producer surplus of other firms (1 bil. US dollars)	2.97 (20.85)

Panel B: Welfare Changes under Learning and Full-Information Models		
	Merger	
	Learning	RE
Δ in investment (1 million TEU)	-0.014	-0.009
Δ in investment volatility (1 million TEU)	-0.025	-0.004
Δ in consumer surplus (1 bil. US dollars)	-0.94	-0.46
Δ in producer surplus (1 bil. US dollars)	13.96	10.04
Δ in total surplus (1 bil. US dollars)	13.03	9.58

Notes: Panel A shows results from the merger simulations over the sample period (2006:Q1-2014Q4) with the percentage changes from the case of no merger in parentheses. The owned capacity, order book, and investment are reported as the average over time, and the welfare measures as the sum over the entire period. Panel B compares changes predicted by the learning model and the full-information model. Consumer surplus is calculated with respect to the Asia-Europe market only.

Appendix

A Simple Dynamic Oligopoly Model with Learning and Merger Counterfactuals

In order to demonstrate the key channels through which uncertainty and strategic interaction affect firm behavior, I analyze a simple duopoly model of firm investment allowing uncertainty about the demand process. This model retains the main framework of the model presented in section 4, while abstracting from certain features that are not central to understanding the key forces of interest, such as chartering and endogenous investment costs. Using this model I simulate a merger of the two firms under learning and under full information, respectively. Comparing results from these two simulations will shed light on whether uncertainty about the demand process accentuates the role of strategic interaction in investment and welfare.

In this model, time is discrete with an infinite horizon and is denoted by $t \in \{0, 1, 2, ...\}$. There are two incumbent firms denoted by $i \in \{1, 2\}$. Firms are heterogeneous with respect to their firm-specific state, $x_{it} = (k_{it}, b_{it})$, where k_{it} is the capacity of ships owned by firm i and b_{it} is the backlog of firm i. The backlog is capped to one unit of capital and there is one period of time to build. Firm profit also depends on market-level demand d_t. The state for firm i is then given by $s_{it} = (x_{it}, x_{-it}, d_t)$.

The timing of events is as follows: (1) Firms observe their current state as well as their private cost shocks associated with investment. They update their beliefs about demand. (2) Firms make investment and scrapping decisions. (3) The industry evolves to a new state.

Agents consider a first-order autoregressive process for the evolution of demand given as

$$d_t = \rho^0 + \rho^1 d_{t-1} + \omega_t.$$

where $\omega_t \sim N(0, \sigma^2)$. They revise expectations with respect to the evolution of the demand state based on adaptive learning in the same way as described in section 4.3. That is, in each period t, firms re-estimate parameters $\eta_t = \{\rho_t^0, \rho_t^1, \sigma_t\}$ based on information available to them, which is the history of demand realizations $\{d_0, d_1, ..., d_t\}$.

In each period, firms choose whether to invest in one unit of capital ($\iota_{it} \in \{0,1\}$) and whether to scrapping one unit ($\delta_{it} \in \{0,1\}$). Firms are not allowed to both invest and scrap in the same period. A firm pays an investment cost, κ, if it decides to invest, and receives a scrap value, ϕ, if it decides to scrap. There is also a private cost shock associated with each action, $\epsilon_{it} = \{\epsilon_{it}^0, \epsilon_{it}^1, \epsilon_{it}^2\}$. The private shocks follow a type 1 extreme value distribution.

The value function of a firm before observing its private shock can be written as

$$\begin{aligned} V^{\eta_t}(s_{it}) = & E\Big[\max\{\pi(s_{it}) + \epsilon_{it}^0 + \beta E[V^{\eta_t}(s_{it+1}|s_{it}, \iota_{it}=0, \delta_{it}=0)] \\ & , \pi(s_{it}) - \kappa + \epsilon_{it}^1 + \beta E[V^{\eta_t}(s_{it+1}|s_{it}, \iota_{it}=1, \delta_{it}=0)] \\ & , \pi(s_{it}) + \phi + \epsilon_{it}^2 + \beta E[V^{\eta_t}(s_{it+1}|s_{it}, \iota_{it}=0, \delta_{it}=1)]\}\Big] \end{aligned}$$

The period profit is given as follows

$$\pi(x_{it}, x_{-it}, d_t) = \frac{k_{it}}{\sqrt{k_{it} + k_{-it}}} d_t - a k_{it}. \tag{12}$$

Simulations of Mergers and Uncertainty

This model has at least two sources of strategic considerations in firms' investment decisions. First, there is a business stealing incentive, which arises because a firm's investment in an extra unit of capacity negatively affects its rival's profit and is internalized by both firms as seen in equation (12). Second, a firm has preemption motives in that it invests earlier and more than they would in the absence of competition in order to discourage its rival's investment.

I simulate a merger which removes these strategic considerations between the firms under learning and full information, respectively. I use the demand realizations that I compute from the data and use in the estimation as in section 5.1. Panel A of table 8 shows calibrated parameters and panel B shows results from the merger simulations. Comparing the first and third columns of panel B shows that uncertainty increases both the volume and volatility of investment as well as the total capacity. Comparing the first and second columns shows that strategic incentives or competition have the similar effect. Lastly, comparing the third and and sixth columns shows that the effect of competition is more pronounced under learning. In particular, investment becomes more highly procyclical under learning. During periods of high demand firms revise their beliefs about future demand upward, which amplifies strategic incentives and raises investment. By contrast, weak demand lead firms to revise their beliefs downward, resulting in suppressed investment or higher scrapping.

Table 8: Merger Simulations with a Dynamic Duopoly Model

Panel A: Calibrated parameters						
Discount rate	β	0.9				
Mean investment cost	κ	8				
Mean scrap value	ϕ	.1				
Fixed cost	a	0.04				
Panel B: Simulation results						
	Learning			Full Information		
	Duopoly	Monopoly	%Δ	Duopoly	Monopoly	%Δ
Total capacity	10.28	7.73	-0.25	9.21	7.47	-0.19
	(0.98)	(0.53)		(0.48)	(0.43)	
Total investment	0.24	0.13	-0.47	0.19	0.11	-0.45
	(0.10)	(0.05)		(0.05)	(0.04)	
Volatility of investment	0.46	0.34	-0.27	0.42	0.32	-0.25
	(0.09)	(0.10)		(0.07)	(0.08)	
Correlation between investment and demand	0.16	0.09	-0.46	0.07	0.06	-0.18
	(0.14)	(0.17)		(0.17)	(0.16)	

Notes: Standard deviations are in parentheses. The simulation is based on 100 sample paths.

B Additional Tables Figures

Table 9: Regression of Price per TEU on Competition

	(1)	(2)
Number of Firms	-0.007***	0.003
	(0.002)	(0.005)
Route FE	Yes	Yes
Constant	1.50***	1.14***
	(0.07)	(0.19)
Observations	1080	1080
R^2	0.890	0.886

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.001$

Notes: The dependent variable is price of container shipping services per TEU at the route level where a route connects two regions. In the second column, the number of firms on route from region A to region B has been instrumented with the number of firms operating routes connecting region A and other regions. The F-statistics from the first stage regression is 135.43.

Figure 8: Weights on Observations under Adaptive Learning

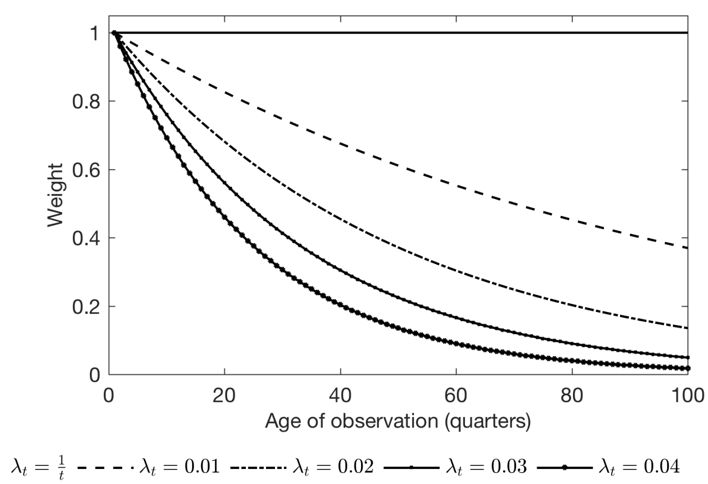

Nots: This figure plots weights that are applied to observations for different values of λ_t in the adaptive learning model where λ_t is the weight parameter that governs how responsive estimate revisions are to new data (see equation (7)).

Figure 9: Prices on Major Trade Routes

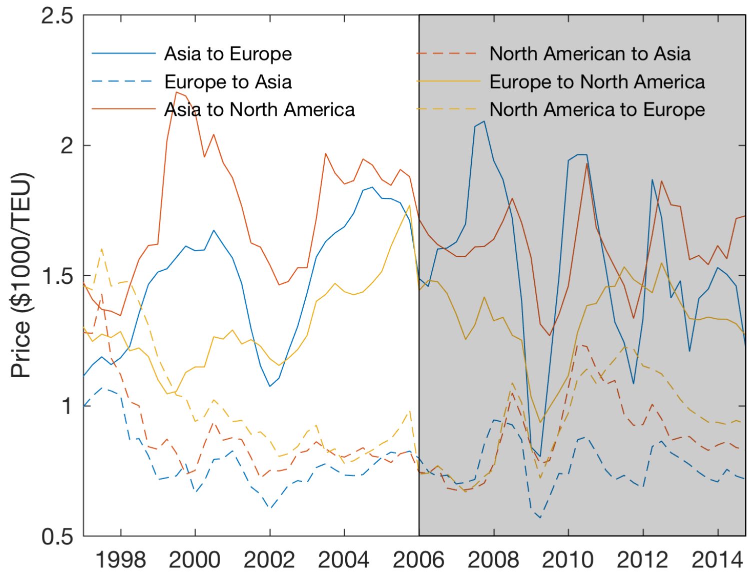

Notes: This figure shows quarterly average prices of shipping a unit of trade goods (TEU) on major container trade routes from 1997 to 2014. The shaded area covers the period from which this paper's main analysis derives, from 2006 to 2014.

Figure 10: The Distribution of Firm Size

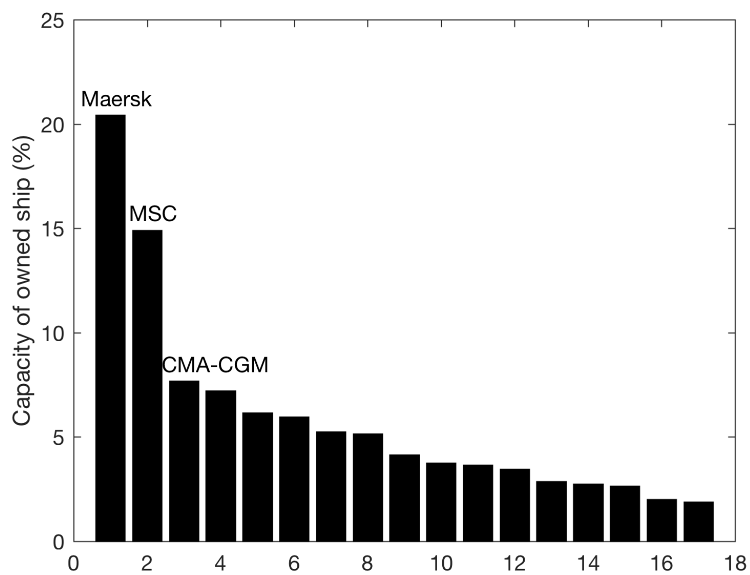

Notes: This figure shows the capacity owned by each firm as a percentage of total industry capacity, where the capacity is averaged over the period from 2006:Q1 to 2014:Q4.

Figure 11: Deployed Ship Capacity as a Share of Total Ship Capacity

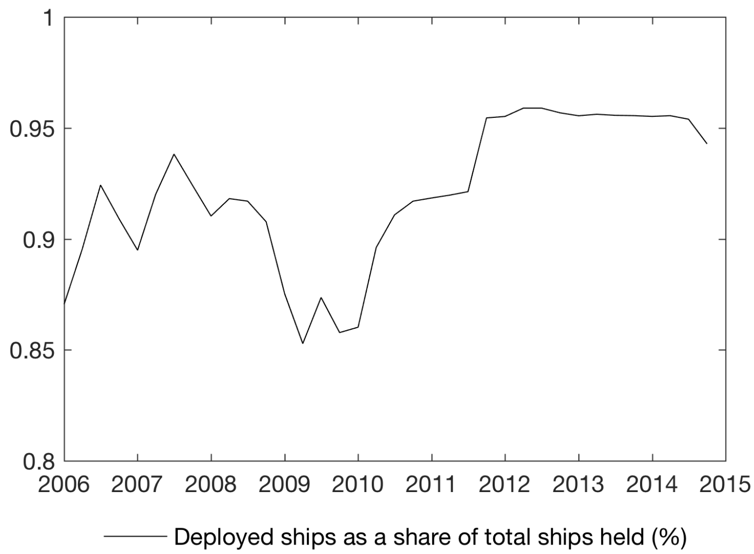

Figure 12: Forecasts on Year-On-Year Growth (%) in GDP & Imports for Europe

(a) Mean

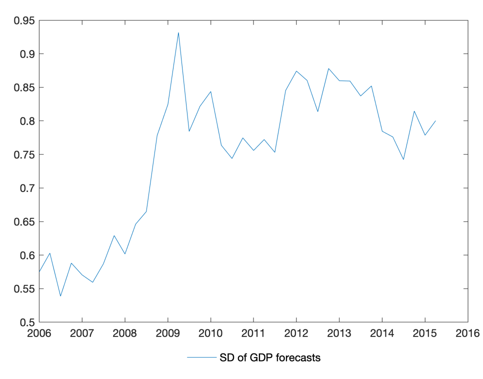

(b) SD

Notes: The left panel plots two-year-ahead forecasts for GDP in the Euro area and one-year-ahead forecasts for imports in selected European countries. The right panel plots the standard deviation of the GDP forecasts.

C Detailed Estimation Results

C.1 Preliminary Analysis

I consider the $ARIMA(p, d, q)$ model given by

$$X_t - \alpha_1 X_{t-1} - \cdots - \alpha_p X_{t-p} = \varepsilon_t + \theta_1 \varepsilon_{t-1} + \cdots + \theta_q \varepsilon_{t-q}$$

where X_t is the d-degree differenced data on the logged quantity of container trade in TEU, p is the order of the autoregressive part, and q is the order of the moving average part. I explore a wide range of parameter values including $0 \leq p \leq 4, 0 \leq d \leq 1, 0 \leq q \leq 4, 0 \leq p' \leq 1, 0 \leq q' \leq 1$, and also consider putting a time-trend when the degree of differencing (d) is zero. To accommodate the volatile nature of the market, I explore specifications with time-varying volatility in which the error terms follow a $GARCH(p', q')$ process in addition to specifications with constant volatility in which the error terms, ε_t, have a normal distribution with zero mean. I evaluate these candidate specifications based on the Akaike information criterion (AIC)/ Bayesian information criterion (BIC).Table 10 lists all candidate specifications considered and the AIC and BIC criteria.

Table 10: The AIC and BIC for different specifications of the container trade quantity process

$ARIMA$ (p, d, q)	$GARCH$ (p', q')	AIC	BIC
(1,0,0)	(0,0)	-2.79	-2.69
(2,0,0)	(0,0)	-2.75	-2.62
(3,0,0)	(0,0)	-2.72	-2.55
(4,0,0)	(0,0)	-2.69	-2.49
(1,0,1)	(0,0)	-2.67	-2.51
(1,0,2)	(0,0)	-2.63	-2.43
(2,0,1)	(0,0)	-2.64	-2.45
(2,0,2)	(0,0)	-2.62	-2.40
(1,1,1)	(0,0)	-2.70	-2.57
(1,1,2)	(0,0)	-2.67	-2.51
(2,1,1)	(0,0)	-2.67	-2.51
(2,1,2)	(0,0)	-2.75	-2.59
(1,0,0)	(0,1)	-2.74	-2.59

Notes: The AIC is given by AIC $= (2k - 2\ln(\hat{L}))/N$ where k is the number of parameters in the model, L is the maximized value of the likelihood function of the model, and N is the number of observations. The BIC is given by BIC $= (\ln(N)k - 2\ln(\hat{L}))/N$.

I also repeat the exercise in section 3 with a behavioral model similar to Greenwood and Hanson (2015). That is, I assume that the logged quantity of container trade follows an $AR(1)$ process

$$X_t - \bar{X} = \rho_1(X_t - \bar{X}) + \varepsilon_t$$

with $\rho_1 \in [0, 1)$ and $Var[\varepsilon_t] = \sigma_\varepsilon^2$. I estimate this model using the full sample of the data from 1997 to 2014 for the Asia to Europe route and firms to over-extrapolate shocks in container trade by letting them use $\rho_k \in [\rho_1, 1)$ as the coefficient on the $(X_t - \bar{X})$ term in their forecasts. The results as shown in table 11 suggest that beliefs from this model are more highly correlated with the GDP and trade forecasts compared to the full-information model, but less highly correlated compared to the learning model.

Table 11: Correlation between GDP/ Trade Forecasts and Container Trade Volume Forecasts

	Full information		Learning	Over-extrapolation allowed
	Constant volatility	Time-varying volatility		
GDP - Mean	-0.38	-0.19	0.66	0.08
	(0.16)	(0.17)	(0.13)	(0.17)
GDP - Variance	.	0.17	0.90	.
	.	(0.17)	(0.08)	.
Trade - Mean	0.06	0.18	0.42	0.30
	(0.25)	(0.25)	(0.23)	(0.24)

Notes: Standard errors are in parentheses.

C.2 Firm Beliefs and Other Primitives

This section presents detailed results from the empirical implementation of the learning model in section 5.4 and the first three steps of the estimation described in sections 5.1 to 5.3.

Figure 13: Beliefs under Learning for the Outside Market

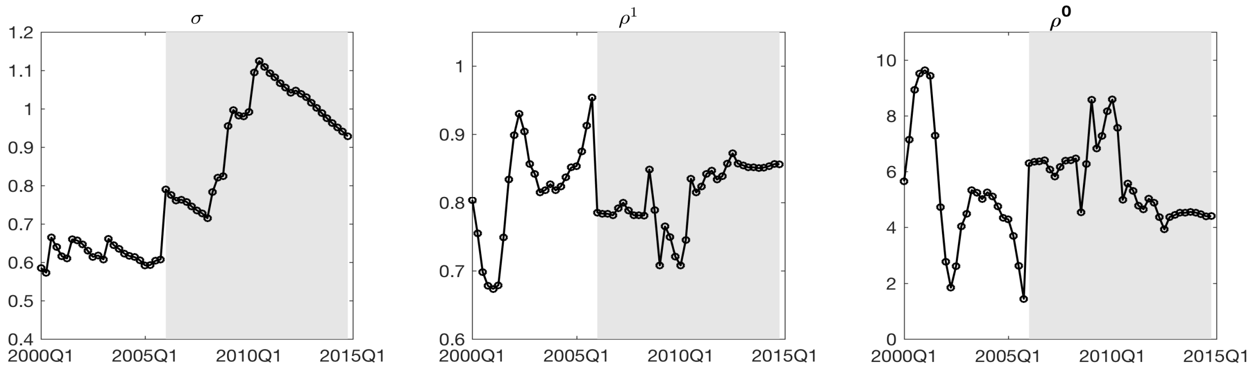

Notes: This figure shows firms' beliefs about demand in the outside market for 2000:Q1 to 2014:Q4 under adaptive learning with $\lambda_t = 0.02$. The beliefs are summarized by the three parameters, $\{\tilde{\sigma}_t, \tilde{\rho}_t^0, \tilde{\rho}_t^1\}$, in the AR(1) process as given in equation (6). Beliefs for 2006-2014 in the shaded area are used in the main analysis.

Figure 14: Beliefs under Learning with Different Weighting Parameters for the Asia-Europe Market

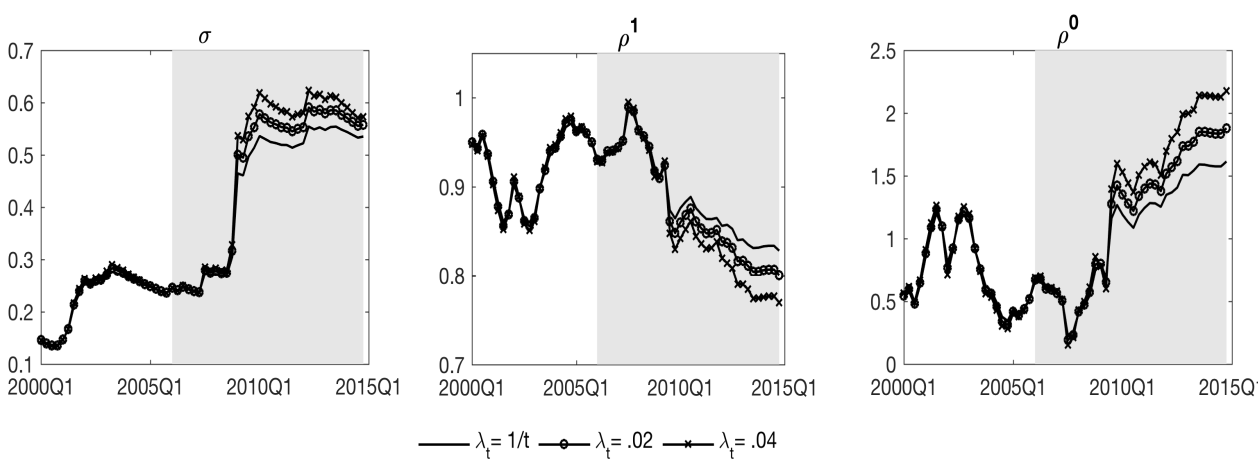

Notes: This figure shows firms' beliefs about demand in the Asia-Europe market for 2000:Q1 to 2014:Q4 under adaptive learning for different values of λ_t. The beliefs are summarized by the three parameters, $\{\sigma_t, \rho_t^0, \rho_t^1\}$, in the AR(1) process as given in equation (**??**). Beliefs for 2006-2014 in the shaded area are used in the main analysis.

Table 12: IV Regression Results for Demand for Container Shipping

Dependent Variable	First stage Log price	Second stage Log quantity
Size of owned ships (1000 TEU)	-0.13**	
	(0.06)	
Age of owned ships (year)	0.03	
	(0.03)	
Fraction of 20+ y.o. ships	-0.02*	
	(0.01)	
Log GDP	0.44***	2.73***
	(0.12)	(0.53)
Log price		-3.89**
		(1.87)
Route FE	Yes	Yes
Constant	-6.27***	-32.66***
	(1.79)	(7.48)
R^2	0.83	0.11

Notes: Standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.001$.

Table 13: Estimates of the Profit Function Coefficient Parameters

Marginal cost		
a	0.265	(0.011)
b	1.750	(0.024)
Outside market profit (R)		
r_0	-1.238	(0.177)
r_1	0.089	(0.006)
r_2	-0.117	(0.008)
Charter cost (CC)		
γ_0	0.206	(0.096)
γ_1	0.087	(0.007)
γ_2	-0.084	(0.021)
γ_3	-0.064	(0.009)

Notes: This table reports estimates of the parameters in the marginal cost, outside market profit, and charter cost functions. The unit of the aggregate deployed capacity ($\tilde{Q}_t$) in the outside market profit function; as well as the firm-level owned capacity (k_{it}) and the aggregate owned capacity (K_{it}) in the charter cost function, is 1 million TEU. Standard errors for the estimates are in parentheses.

Table 14: Estimates of the Investment Cost and Scrap Value

	Investment cost ($1000/TEU)	Scrap value ($1000/TEU)
Total capacity of owned ships	-1.35***	0.11
(1 million TEU)	(0.35)	(0.12)
Total capacity in order book	1.12**	0.06
(1 million TEU)	(0.54)	(0.19)
Demand state: A-E market	0.50	0.25**
	(0.31)	(0.11)
Demand state: outside market	-0.16	0.08
	(0.17)	(0.06)
Constant	15.09**	-3.17*
	(4.81)	(1.68)
R^2	0.69	0.38

Notes: This table reports coefficient estimates in the investment cost and scrap value functions. Standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.001$.

C.3 Additional Results from Alternative Models of Firm Beliefs

This section presents additional results from alternative models of firm beliefs including a Bayesian learning model.

Bayesian Learning

Under Bayesian learning, each firm starts with prior beliefs about the parameters of the model. Then, based on its information set, $\{z_\tau, \tilde{z}_\tau\}_{\tau=0}^{t}$, the firm updates its beliefs about the parameters in the demand process, $(\rho_t^0, \rho_t^1, \sigma_t, \tilde{\rho}_t^0, \tilde{\rho}_t^1, \tilde{\sigma}_t)$. The $AR(1)$ coefficients for the Asia-Europe market, $\rho = [\rho^0, \rho^1]$, have normal priors given by $\rho_0 \sim N(\mu_0, \Sigma_0)$. The prior of σ^2 follows an inverse Gamma distribution. Then, the posterior distribution $\rho_t \sim N(\mu_t, \Sigma_t)$ has the mean and the variance given by

$$\mu_t = \Sigma_t \left(\Sigma_0^{-1}\mu_0 + \sigma^{-2}(Y_t' Z_t)\right)$$
$$\Sigma_t = \left(\Sigma_0^{-1} + \sigma^{-2}(Y_t' Y_t)\right)^{-1}.$$

The beliefs are defined similarly for the outside market.

The first three years of the price and quantity data (1997:Q1-1999Q4) are used in the estimation of the prior beliefs, although I explore alternative estimations of the priors.[64] I start from diffuse priors and apply the Gibbs sampling methods. The estimates of the prior are presented in table 15. In the first quarter of 2000, firms start with the prior beliefs about the parameters and revise their beliefs using Bayesian updating in each period based on newly realized data. I apply the Gibbs sampling techniques to estimate the posterior beliefs.

Figure 15 shows posterior beliefs under Bayesian learning. Compared to the baseline model of adaptive learning, the degree to which firms' beliefs react to new data is smaller under Bayesian learning. This is because there are lower weights placed on new data under Bayesian learning, as agents assign positive weights on their prior beliefs. Consequently, although the timing of the investment boom and bust predicted under Bayesian learning is consistent with the data, the magnitudes of the rise and the fall in investment are smaller than observed in the data.

[64] I explore using the full sample from 1997 to 2014 for estimating the priors with the same updating for posterior beliefs, in which case the posterior beliefs remain almost identical in this case. The choice of the sample period for prior beliefs is thus expected to have little impact on empirical results.

Table 15: Moments of the Prior Distributions under Bayesian Learning

Asia-Europe market		
ρ^0	ρ^1	σ
0.51	0.95	0.17
(0.63)	(0.08)	(0.02)
Outside market		
$\tilde{\rho}^0$	$\tilde{\rho}^1$	$\tilde{\sigma}$
8.11	0.72	0.68
(5.19)	(0.17)	(0.27)

Notes: This table shows the estimated means and standard deviations (in parentheses) of the prior distributions of $AR(1)$ parameters. The estimation is based on data from 1997:Q1 to 1999:Q4.

Figure 15: Beliefs under Bayesian Learning

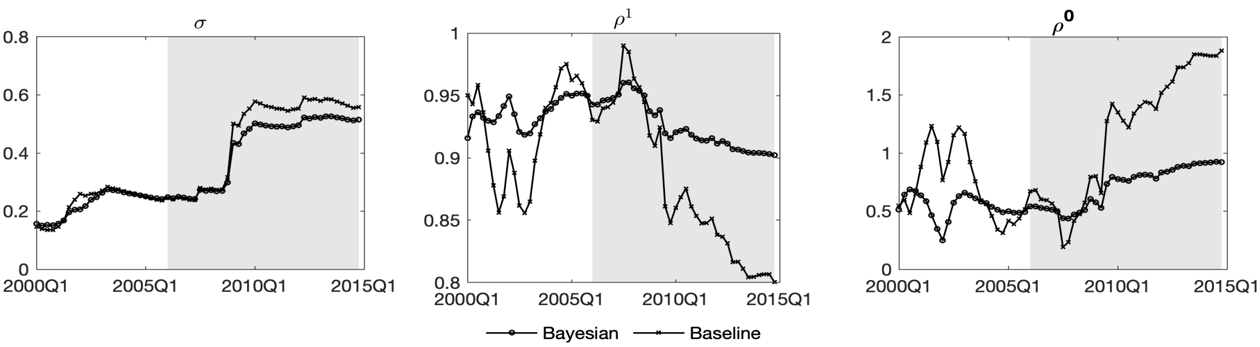

(a) Asia-Europe market

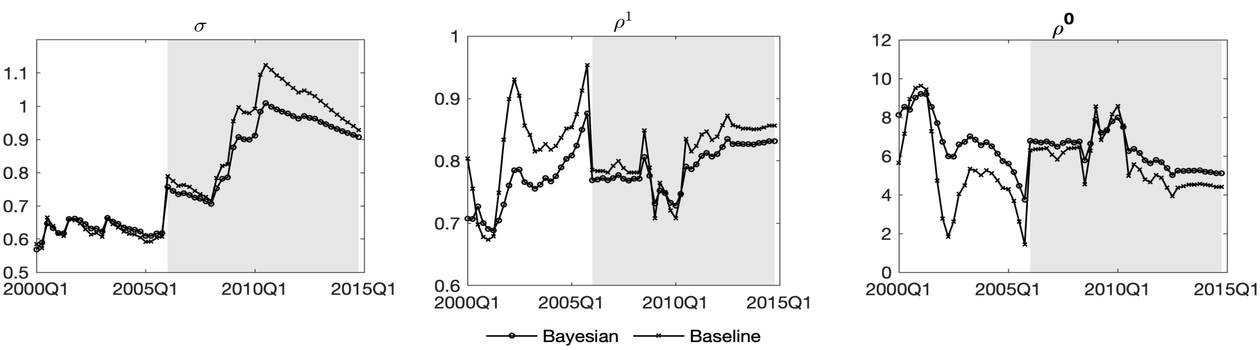

(b) Outside market

Notes: This figure shows firms' beliefs about demand in the outside market for 2000:Q1 to 2014:Q4 under Bayesian learning. Beliefs for 2006-2014 in the shaded area are used in the main analysis.

Full Information Model

Figure 16: Beliefs under Full Information for the Outside Market

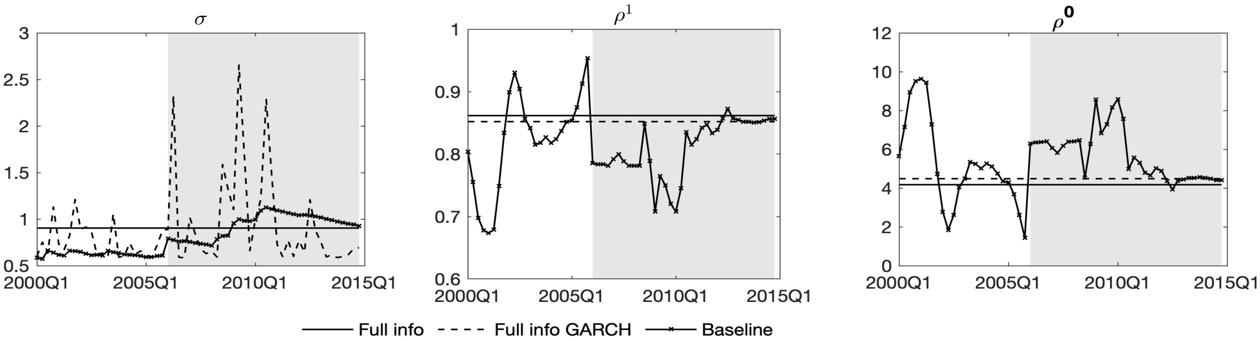

Notes: This figure shows firms' beliefs about demand in the outside market for 2000:Q1 to 2014:Q4 under full information models. Beliefs for 2006-2014 in the shaded area are used in the main analysis.

Table 16: Estimates of the Time-Varying Volatility Models

Asia-Europe market		
a_0	a_1	b_1
0.046	0.83	0.17
(0.015)	(0.28)	(0.16)
Outside market		
$\tilde{a}_0$	$\tilde{a}_1$	
0.34	0.73	
(0.11)	(0.25)	

Notes: Standard errors are in parentheses.

C.4 Additional Results from Counterfactual Simulations

Figure 17: Yearly Investment under the Merger Cases

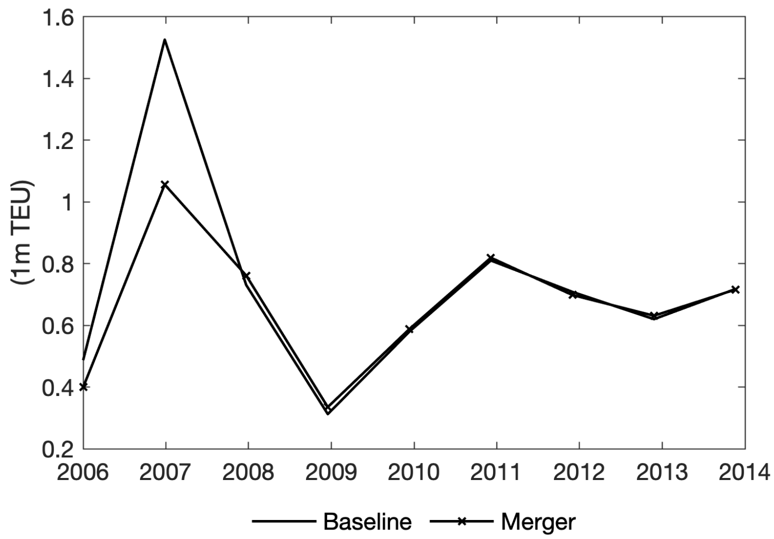

Notes: The simulations are based on 1000 equilibrium paths.

D Computation

To compute strategies under MME for the model described in section 4.4, I adopt a computational algorithm that is analogous to the standard value function iteration algorithm except for an extra simulation step. Because the transition of the moment-based industry state $\hat{s}$ may not be Markov, a simulation step is used to generate the Markov approximation of the transition of this state. The algorithm starts with a choice-specific value function that maps from the set of state-action pairs to values denoted as $W^{\eta}(\mu, x, \hat{s})$. It contains expected values of different actions prior to drawing random costs of investing and scrapping given beliefs about demand η. Then, based on a simulation run in which firms play optimal strategies implied by these choice-specific values, the algorithm constructs the perceived transition kernel $\hat{P}_{\mu}[m'|\hat{s}]$. The next step updates the values and strategies using the best response against the current strategy and the perceived transitions kernel. Finally, equilibrium conditions are checked based on the norm of the distance between the values in the memory and the updated values. A more detailed description of the algorithm is provided as follows:

1. Initialize $W^{\eta}(\mu, x, \hat{s})$ for all $(\mu, x, \hat{s}) \in \mathcal{M} \times \mathcal{X} \times \hat{\mathcal{S}}$, and optimal strategies, μ^*, that W^{η} implies.

2. Simulate a sample path of $\{\hat{s}_t\}_{t=1}^{T}$ for large T based on μ^*. Calculate the empirical frequencies of industry state $h(\hat{s}) = \frac{1}{T} I\{\hat{s}_t = \hat{s}\}$ for all $\hat{s} \in \hat{\mathcal{S}}$. Calculate the empirical transition kernel as

$$\hat{P}_{\mu}[m'|\hat{s}] = \frac{\sum_{t=1}^{T} I\{\hat{s}_t = \hat{s}, m_{t+1} = m'\}}{\sum_{t=1}^{T} I\{\hat{s}_t = \hat{s}\}}.$$

3. Calculate the new values for each state-action pair $(\mu, x, \hat{s})$ as:

$$\tilde{W}^{\eta}(\mu, x, \hat{s}) = \pi(x, \hat{s}) - \iota\kappa(\hat{s}) + \nu(\delta, x)\phi(\hat{s}) + \beta E_{a,\mu}\left[V^{\eta}(x', \hat{s}'|x, \hat{s})\right]$$

   and obtain the new best response $\tilde{\mu}^* = \arg\max_{\mu} W(\mu, x, \hat{s}|\mu, \mu^*)$ for all $(x, \hat{s}) \in \mathcal{X} \times \hat{\mathcal{S}}$.

4. Calculate the following norm: $\max_{x,\mu} \sum_{s\in\hat{\mathcal{S}}} |\tilde{W}^{\eta}(\mu, x, \hat{s}) - W^{\eta}(\mu, x, \hat{s})| h(\hat{s})$.

5. If the norm is greater than ε, update the values and the strategy profile with $\tilde{W}$ and $\tilde{\mu}^*$ and repeat steps 2-5.

E Robustness

E.1 Relaxing the Myopic Learning Assumption

The main specification of this paper assumes that agents do not internalize the possibility of learning in the future and use their current beliefs in their forecasts. This implies that although the parameters $\{\rho_t^0, \rho_t^1, \sigma_t, \tilde{\rho}_t^0, \tilde{\rho}_t^1, \tilde{\sigma}_t\}$ that summarize agents' current beliefs are still state variables, they are not "active" in the sense that they stay fixed over time, as forecasting demand for all future periods requires only the current beliefs. Therefore, I can solve the model while fixing the belief parameters at the levels implied by demand realizations observed in the data and the specified learning model.

In contrast, if I allow agents to forecast using beliefs that change as they receive new draws of demand, the belief parameters now become state variables that evolve stochastically depending on the realizations of demand. Therefore, implementing this requires solving the model for each point on a fine grid of belief parameter values in addition to values of all other parameters.

Since fully relaxing the myopic learning assumption is computationally infeasible, I consider partially relaxing it by allowing agents to internalize future learning for one period ahead. That is, at time t agents use their current beliefs in predicting the distribution of demand for $t+1$ and use their updated beliefs based on demand realized at $t+1$ to predict demand from $t+2$ onwards. Although still restrictive, this exercise will be informative, especially because due to discounting the effect should be the strongest for $t+1$ and subside over time. I solve this model for the baseline case where all other parameters are held at their estimated values.

Table 17 shows simulated moments under this new specification (referred to as "non-myopic learning") in comparison with the moments under the baseline specification. The results show that the correlation between demand and investment is higher under non-myopic learning while other moments are almost identical under the two specifications. That is, when firms internalize the possibility of learning in the future, their recent demand draws, and thus their current beliefs have stronger positive effects on investment. For example, when firms receive favorable demand draws, their current beliefs are revised upward, but this also makes them believe that future draws will be more favorable and they will continue to be more optimistic relative to the myopic learning case. Thus, under non-myopic learning the effect of recent demand on investment is further reinforced.

Table 17: Simulated Moments under Non-Myopic Learning

	Non-myopic learning	Baseline
Average investment in 2006-2008 (1 million TEU)	0.24	0.23
	(0.03)	(0.03)
Average investment in 2009-2014 (1 million TEU)	0.15	0.15
	(0.02)	(0.02)
Total capacity of owned ships (1 million TEU)	5.16	5.15
	(0.27)	(0.27)
Total capacity in the order book (1 million TEU)	3.00	2.98
	(0.14)	(0.14)
Correlation between demand and investment	0.27	0.22
	(0.11)	(0.12)
Volatility of investment (1 million TEU)	0.17	0.17
	(0.02)	(0.03)

Notes: The simulated moments are computed based on 1000 series of equilibrium paths. Standard deviations are in parentheses.

E.2 Adding a Dominant Firm's State in the Moment-Based State

The moment-based Markov equilibrium as proposed by Ifrach and Weintraub (2016) allows firms to keep track of the detailed state of dominant firms (strategically important firms) as well as moments describing the state of fringe firms as their moment-based industry state. In my application, firms' industry states are further reduced to the the sum of states of all firms However, MME strategies may not be optimal (i.e. there may be a profitable unilateral deviation to a strategy that depends on more detailed information), if moments do not summarize all payoff-relevant information. In order to investigate how robust equilibrium strategies are to changes in the moment-based industry state, I consider a version in which richer information is allowed in the industry state and compare model predictions and values to the baseline case.

In particular, firms condition their strategy on the firm-specific state of the largest firm (the dominant firm) in addition to the states in the baseline case including their own firm-specific state, the sum of all firms' states, and demand states. In one version, the dominant firm's capital, denoted as k_1 is included in the information set and in the other version, the dominant firm's order book, b_1. Let $\hat{s}'$ denote the new industry state and let μ' and $\hat{V}'$ denote the optimal strategy and the value of the new game based on $\hat{s}'$ as the industry state. The difference in the values of the baseline model and the model that includes the dominant firm's state for each underlying state s is defined as:

$$\Delta_{\mu'}(x,s) = \frac{V'^{\eta}_{\mu',\mu}(x,\hat{s}') - \hat{V}^{\eta}_{\mu}(x,\hat{s})}{\hat{V}^{\eta}_{\mu}(x,\hat{s})}.$$

The expected value of this deviation is computed as the weighted average through a simulation where the weights come from simulations based on the baseline model, or $\hat{V}$. Table 18 shows that model predictions stay robust when either of the dominant firm states is added. The average difference in the values is not significantly different from zero for both cases.

Table 18: Adding a Dominant Firm's State in the Moment-Based State

Panel A: Simulated moments						
	Baseline		Model with dominant firm's capital state		Model with dominant firm's order book state	
Average investment in 2006-2008 (1 million TEU)	0.23	(0.03)	0.23	(0.03)	0.23	(0.03)
Average investment in 2009-2014 (1 million TEU)	0.15	(0.02)	0.15	(0.02)	0.15	(0.02)
Total capacity of owned ships (1 million TEU)	5.15	(0.27)	5.13	(0.28)	5.14	(0.28)
Total capacity in the order book (1 million TEU)	2.98	(0.14)	2.98	(0.14)	2.99	(0.14)
Correlation between demand and investment	0.22	(0.12)	0.21	(0.12)	0.22	(0.12)
Std. dev. in investment (1 million TEU)	0.17	(0.03)	0.17	(0.03)	0.17	(0.03)
Panel B: Average difference in values						
All firms (%)			-0.35	(0.46)	-0.42	(0.54)
Dominant firm (%)			-0.18	(0.22)	-0.21	(0.26)
Fringe firms (%)			-0.36	(0.48)	-0.43	(0.56)

Notes: Standard errors are in parentheses.

E.3 Implementation of the Adaptive Learning Model

I explore and compare two approaches in implementing the learning model. The truncation approach entails setting the initial period of the information set as the start date of the data. This method is straightforward to implement and is appropriate if firms also do not have access to information beyond the data available to the researcher. However, bias can arise if the agents' information set includes observations extending further back than the start date of the data. The bias would be mitigated if agents discount older observations more heavily when forming expectations.

The imputation approach employs external data that provide information about the missing data. This approach is appealing if agents indeed use a longer historical dataset in forming expectations than observed by the researcher, and the researcher has access to the external data that provide a good approximation to these data. Bias can arise, however, from the imputation process depending on the quality and scope of the external data. In the context of this paper, for example, one could consider using international trade data to proxy demand for container shipping.

The imputation approach is implemented as follows. I set the start date for firms' information as the second quarter of 1966, which is the date of the first international container voyage. Then, I employ quarterly data on the value of trade by origin-destination pair from the IMF Direction of Trade Statistics database to impute the missing data on demand states from 1966:Q2-1996:Q4.[65] Finally, I estimate the beliefs using the imputed longer time-series data in the same way as the truncation approach.

Figure 18 compares beliefs for the Asia-Europe market under the truncation and imputation approaches. The beliefs are closer to one another, especially for the period of the main analysis, from 2006 to 2014. The model fits under the two approaches are also close to one another, although they are better under the truncation approach, especially for the correlation between demand and investment as shown in Table 19.

Table 19: Data Moments and Simulated Moments under the Truncation and Imputation Approaches

	Data	Truncation	Imputation
Average investment in 2006-2008 (1 million TEU)	0.23	0.23	0.22
Average investment in 2009-2014 (1 million TEU)	0.14	0.15	0.16
Total capacity of owned ships (1 million TEU)	5.09	5.15	5.17
Total capacity in the order book (1 million TEU)	3.07	2.98	2.98
Correlation between demand and investment	0.19	0.22	0.26
Std. dev. in investment (1 million TEU)	0.17	0.17	0.18

Notes: This table compares moments observed in the data and moments simulated under the truncation and imputation approaches of the baseline learning model.

[65] To translate the value of trade to the quantity of container trade, the demand state for the 1997-2014 period was regressed on the de-trended value of trade. Then, the demand states for periods with missing data are constructed as predicted values from the regression. For the 1997-2014 period, actual demand states are used.

Figure 18: Beliefs under Adaptive Learning Based on Two Alternative Approaches

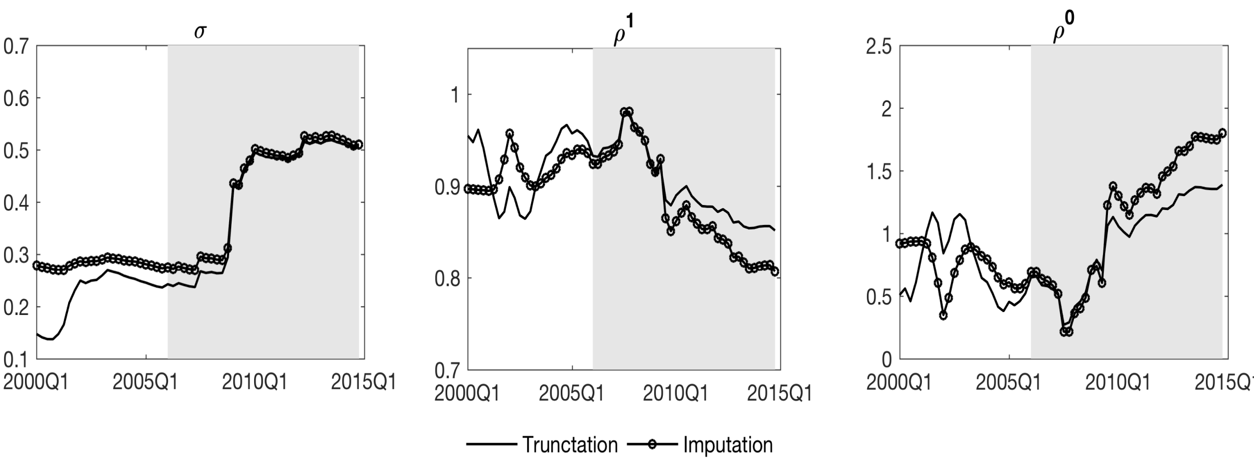

Notes: This figure shows firms' beliefs about future demand under adaptive learning estimated with the truncation approach and the imputation approach, respectively, for the case of $\lambda_t = 0.02$. Beliefs for 2006-2014 in the shaded area are used in the main analysis.

F Credit Market Conditions

In the sample period that this study focuses on (2006-2014), the financial market also experienced an upheaval along with the international trade market. Therefore, one might worry that omitting information on credit market conditions might bias the main results of the paper. I address this concern in two ways. First, I simulate a model that takes into account credit market frictions in the form of collateral constraints. Second, I employ auxiliary data on companies' financials information and examine evidence of financial constraints.

F.1 A Model with Credit Market Frictions

Vessels either currently in construction or held by the company typically serve as collateral for bank loans. Therefore, the dip in the value of ships and the resulting tightening of credit can be another factor contributing to the sharp decline of investment in new ships during the financial crisis. To incorporate this idea, I allow financial frictions in the form of collateral constraints. In particular, I model a borrowing constraint that depends on the resale value of the firm's current stock of capital. The dynamic problem firms face with such financial frictions can be written as

$$\begin{aligned} V^{\eta_t}(x_{it}, \hat{s}_t) = & \max_{\iota_{it}, \delta_{it}} \pi(x_{it}, \hat{s}_t) - \iota_{it}\left(\kappa(\hat{s}_t) + \varepsilon^{\iota}_{it}\right) + \nu(\delta_{it}, x_{it})\left(\phi(\hat{s}_t) + \varepsilon^{\delta}_{it}\right) \\ & + \beta E\left[V^{\eta_t}(x_{it+1}, \hat{s}_{t+1} | x_{it}, \hat{s}_t)\right] \\ & \text{s.t. } k_{it}\chi(\hat{s}_t) \geq d\iota_{it}\left(\kappa(\hat{s}_t) + \varepsilon^{\iota}_{it}\right) \end{aligned}$$

where χ is the resale value of one unit of ships that depends on the state $\hat{s}_t$ and d is a parameter determining tightness of the collateral requirement.

To simulate this model with credit constraints, I first supplement my data with the data on the price of second-hand ships. I estimate the resale value as a linear function of the aggregate state variables with the estimates presented in table 20. The parameter d is calibrated to be 140%, which requires that the existing ships be valued at least 1.4 times the value of the purchased ships. I use this value, which likely overestimates the tightness of the credit constraint, because the goal of the excise is to learn about the role of financial frictions.[66] I simulate the industry imposing the credit

[66] The value is likely an overestimate because the required minimum value-to-loan ratio is generally between 120% and 140%. Moreover, firms typically use vessels under construction that they are purchasing as the main collateral.

constraint starting from the third quarter of 2008, during which both investment and the value of ships collapse.[67]

Table 20: Estimates of the Resale Value

	Resale value ($1000/TEU)
Total capacity of owned ships	-3.16***
	(0.63)
Total capacity in order book	3.11**
	(1.00)
Demand state: A-E market	0.25
	(0.56)
Demand state: outside market	0.06
	(0.30)
Constant	13.38
	(8.80)
R^2	0.74

Notes: This table reports coefficient estimates in the resale value function. Standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.001$.

Table 21 presents a comparison between the computed equilibrium under the baseline model without the financial frictions and that under the model with the frictions. The results suggest that while introducing endogenous financial frictions have distributional effects on investment and capital, the main results from the model are robust. The total amount of capital and the correlation between investment and demand remain almost identical, and the total amount of investment falls by 5%. There are asymmetric effects on the firms. The amount of investment decreases for the smaller firms that are more constrained due to their low capital stock, while it increases slightly for the largest two firms.

F.2 Credit Constraints and Investment

Using data from Compustat on firms' debts and liabilities, I regress investment levels on state variables and variables relating to the firm's credit constraints including long-term debt and debt in current liabilities.[68] If financial constraints were the main determinants of investment, we would

Existing capital is used as additional collateral when the credit market is extremely tight.

[67] I assume that the change in the credit market came as a surprise, in that firms did not expect the change prior to the third quarter of 2008.

[68] 281 company-quarter-level observations on company financials are available out of 612 observations used in the main analysis. There is, however, substantial variation on the magnitude of debts across firms in the data. The average firm-level long-term debt over the sample period varies from 0.06 million dollars for UASC to 4.3 billion dollars for Hyundai.

Table 21: Results from the Model with Credit Market Frictions

	Baseline	With frictions
Total capacity of owned ships (1 million TEU)	51.47	51.17
	(2.72)	(2.67)
Total capacity in the order book (1 million TEU)	29.81	29.06
	(1.41)	(1.54)
Correlation between demand and investment	0.22	0.23
	(0.12)	(0.11)
Volatility of investment (1 million TEU)	1.73	1.70
	(0.26)	(0.26)
Total investment (1 million TEU)	64.92	61.18
	(6.31)	(5.84)
Total investment by top two firms (1 million TEU)	0.87	0.88
	(0.25)	(0.26)
Total investment by other firms (1 million TEU)	5.62	5.24
	(0.58)	(0.55)

expect that firms that hold a higher amount of debt (thus facing harsher credit constraints) would withhold investment to a greater extent. The regression results presented in Table 22, nonetheless, suggest that debt levels do not have statistically significant effects on firms' investment.

Table 22: Regression of Investment on Debt-Related Variables

Dependent variable: Investment (1000 TEU)		
Owned-ship capacity (1000 TEU)	-.037	(.027)
Order-book capacity (1000 TEU)	-.024	(.017)
Aggregate owned-ship capacity (1000 TEU)	.012	(.01)
Aggregate order-book capacity (1000 TEU)	-.015**	(.0064)
Demand state (Asia to Europe)	1.1	(2.3)
Demand state (outside market)	.06	(1.3)
Chartered ship capacity (1000 TEU)	-.025	(.024)
Aggregate chartered ship capacity(1000 TEU)	-.019*	(.011)
Deployment in Asia-Europe market (1000 TEU)	.087**	(.043)
Aggregate deployment in Asia-Europe market (1000 TEU)	.019**	(.0078)
Long-term debt (1 billion US dollars)	.00079	(.002)
Debt in current liabilities (1 billion US dollars)	-.0019	(.0029)
Constant	-11	(38)
Observations	281	
R^2	0.076	

Notes: Standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.001$.