

Time to Build and Shipping Prices

Myrto Kalouptsidi^{*}

Department of Economics, Yale University

January 14th 2011

Abstract

This paper quantifies the impact of demand uncertainty and time to build on prices, investment and surplus in the world bulk shipping industry. We explore the impact of both construction lags and the lengthening of delivery lags in periods of high investment activity by constructing a dynamic model of ship entry and exit. A rich dataset of second-hand ship sales provides direct information on key dynamic objects of the model and allows for their nonparametric estimation. We are able to estimate a ship's value function directly and use it to recover entry costs, scrap values and per period payoffs. We find that moving from time-varying, to constant, to no time to build reduces the level and volatility prices, while increasing entry and consumer (i.e. shipper) surplus. As the duration of time to build can be steered by shipbuilding subsidies, the findings of this paper provide also a first step in examining the impact of such subsidization programs -prevalent in China, South Korea and Japan- on the shipping industry.

^{*}37 Hillhouse Ave, New Haven, CT 06511. email: myrto.kalouptsidi@yale.edu

[†]I am extremely thankful to my advisors, Steve Berry, Lanier Benkard and Phil Haile for their numerous useful comments and suggestions. I have also gained a lot from discussions and comments from Eduardo Engel, Costas Arkolakis, Chris Conlon, Rick Mansfield. Finally, I would like to thank George Banos, Danae Bati and Alex Hatzis for many discussions on how the shipping industry functions.

1 Introduction

Under uncertainty, time to build can substantially affect investment and prices both in terms of their level and volatility. Time to build magnifies the impact of demand shocks by rendering supply less elastic in the short-run. Many industries are characterized by demand uncertainty and multi-period construction processes and the world bulk shipping industry is an outstanding example. In recent years, the rapid growth of imports of raw materials in developing countries such as China, was accompanied by a large increase in freight rates, shown in Figure 1. This led to an unprecedented surge in new ship orders, shown in Figure 2. Long construction lags inherent in shipbuilding, as well as bottlenecks due to shipyard capacity constraints that often bind, create time to build: a gap between the order and delivery date of a ship. As a result, when the crisis of 2008 hit and the operating fleet suffered from an inordinately low demand level, another 70% of that fleet was scheduled to be delivered between 2008 and 2012.

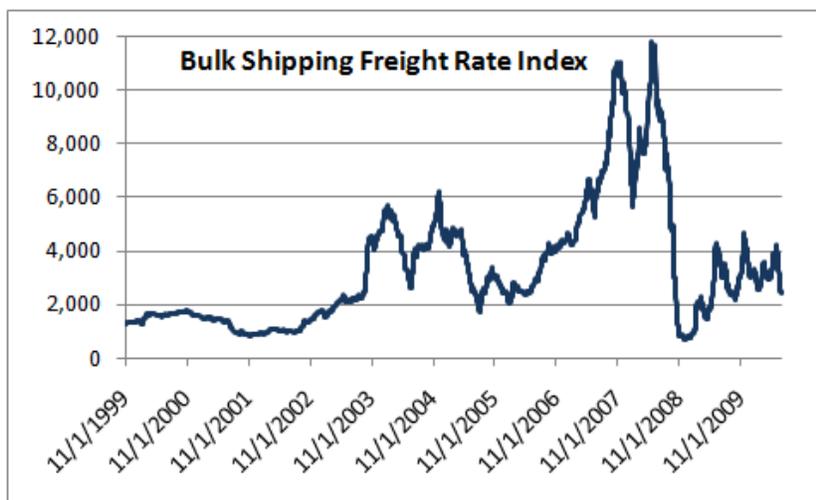


Figure 1: The Baltic Dry Index: daily index based on a weighted average of rates on 20 representative bulk routes. Compiled by the Baltic Exchange. 11/1/1999 = 1334.

It is the goal of this paper to quantify the impact of demand uncertainty and time to build on prices, investment and surplus in the world bulk shipping industry. We explore the impact of both pure construction lags, as well as the lengthening of delivery lags in periods of high investment activity due to capacity constraints. As the duration of time to build can be steered by shipbuilding subsidies, the findings of this paper provide also a first step in examining the impact of subsidization programs in China, South Korea and Japan on the shipping industry.

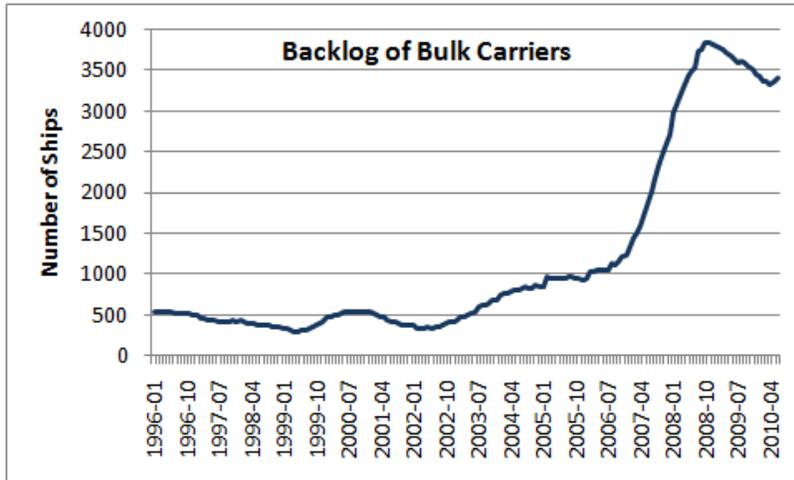


Figure 2: Total number of dry bulk carriers under order at world shipyards.

We construct a dynamic model of ship entry and exit. A rich dataset of second-hand ship sales provides direct information on key dynamic objects of the model and allows for their nonparametric estimation. We find that moving from time-varying to constant to no time to build reduces prices, while increasing investment in new ships and the surplus of industry participants.

Our model builds on the two key features: first, demand for sea transport is inherently uncertain and volatile; second, supply adjusts sluggishly due to entry costs, time to build and convex operating costs of ships. In the model, a firm is a ship. Firms enter, age and endogenously exit. A free entry condition every period determines the number of entrants. Due to time to build, operation begins a number of periods after the entry decision. This number of periods varies over time: it increases due to bottlenecks in production in periods of high investment activity. Once a firm is in the market, it ages and decides whether to exit (demolish the ship) or continue to operate in the market.

We use the second-hand price of a ship to estimate its value function nonparametrically: the assumption of a large number of potential shipowners implies that the resale price equals the value of the ship. Our methodology is a “mirror image” of previous work on the estimation of dynamic settings, such as Rust (1987) on single agent dynamics and Bajari, Benkard and Levin (2008), Pakes, Ostrovsky and Berry (2008), Jofre-Bonet and Pesendorfer (2003) and others, on dynamic games. In those papers, estimated per period payoffs are used to recover value functions, which in turn reveal exit and entry costs. In this paper, we estimate value functions directly from the data and use them to retrieve nonparametrically entry and exit costs, as well as per period payoffs. This methodology

can be of use in other settings with a large number of homogeneous agents.

We test the model and empirical strategy against two time series not used in the estimation. In particular, we compare the estimated profits to an earnings index calculated by shipbrokers, and the estimated entry costs to the average shipbuilding price. We find that the model and estimation perform surprisingly well. In addition, we use the estimated profits of the model to discern the nature of competition in shipping services: we fit a Cournot game for cargo transport and detect that firms essentially act as price takers.

Kydland and Prescott (1982) argue that time to build is crucial in explaining aggregate economic fluctuations. A large theoretical literature has explored the impact of uncertainty on investment, such as Caballero and Pindyck (1996), Dixit (1989), Pindyck (1993) and many others. Empirical work on uncertainty and investment includes Bloom (2009) and Collard-Wexler (2008). Little empirical work on the impact of time to build can be found in the literature; Rosen, Murphy and Scheinkman (1994) explore time to build and periodicity in breeding cattle. The contribution of this paper to the literature is to measure the impact of both uncertainty and time to build in the context of a particular industry, that of bulk shipping. At the same time, this industry, representing 70% of world seaborne trade (in tons), is important in its own right: both freight rates and inventory costs caused by shipping delivery lags can significantly affect trade flows, as documented by Alessandria, Kaboski and Midrigan (2009).

We use the estimated model to address these issues. We compute the equilibrium of our model in a counterfactual world without shipbuilding capacity constraints, where time to build is a constant number of periods corresponding to pure construction lags, as well as a counterfactual world with no time to build. Time to build becomes a critical constraint at positive demand shocks. Simulating an unexpected positive demand shock we find that moving from time-varying, to constant, to no time to build reduces prices, while increasing the fleet and the surplus of industry participants. In particular, following a positive demand shock, in a world with time to build the supply response is limited, as entry is both lower and slower. Indeed, new ships arrive a number of periods later, while in the meantime supply is limited to the existing capacity constrained fleet. In addition, the incentives to enter are dampened since the shock fades due to a mean-reverting demand process and ships that are delivered late will not be able to take advantage of the increased demand for shipping services. As a result, time to build leads to high profits and prices and low consumer (i.e. shipper) surplus. Preliminary results from longrun simulations of the industry suggest that in the absence of time to build prices are 20% lower and 20% less volatile, while the fleet is larger, younger and more volatile. We also take a first

step towards evaluating the impact of shipbuilding subsidies on the shipping industry. Subsidies increase shipbuilding capacity and decrease the delivery lag. We ask what would happen if China had not launched its “Long and Medium Term Plan for the Shipbuilding Industry 2006 – 2015”. To do so, we adjust the dependence of time to build on the backlog and find that entry is lower.

The remainder of the paper is organized as follows: Section 2 provides a short description of the industry. Section 3 presents the model. Section 4 describes the data used. Section 5 presents the empirical strategy employed and the estimation results. Section 6 provides the counterfactual experiments and Section 7 concludes.

2 Description of the Industry¹

Bulk shipping concerns vessels designed to carry a homogeneous unpacked dry or liquid cargo, for individual shippers on non-scheduled routes. The entire cargo usually belongs to one shipper (owner of the cargo). Bulk carriers operate like taxi drivers, instead of buses: they carry a specific cargo individually in a specific ship for a certain price. In practice there are a few different types of cargo transport agreements that involve different methods of payment (e.g. price per ton, price per day) and duration. Bulk ships, nevertheless, do not operate on scheduled itineraries like container ships, but only via individual contracts. Dry bulk shipping involves mostly raw materials, such as iron ore, steel, coal, bauxite, phosphates, but also grain, sugar, wood chips and chemicals. There are four different categories of bulk carriers based on size: Handysize (10,000 – 40,000 DWT), Handymax (40,000 – 60,000 DWT), Panamax (60,000 – 100,000 DWT) and Capesize (larger than 100,000 DWT). Vessels in different categories in general carry different products, take different routes and approach different ports. Practitioners treat them as different markets. The industry consists of a large number of shipowning firms. For instance, Figure 3 shows the firm size distribution of Handysize bulk carriers (on which the empirical analysis will focus) in terms of the number of ships owned by each firm, as well as its market share (captured by its fleet share). Even though there is some scope for differentiation (based on the age of the ship, the shipyard where it was built at, the reputation of the shipowner, etc.) shipping services are perceived as homogeneous.

Demand for shipping services is driven by world seaborne trade and is thus subject to world economy fluctuations. It is also rather inelastic. In recent years, the growth and infrastructure building at several countries led to increased imports of raw materials,

¹This section draws from Stopford (2009), as well as Beenstock and Vergottis (1989).

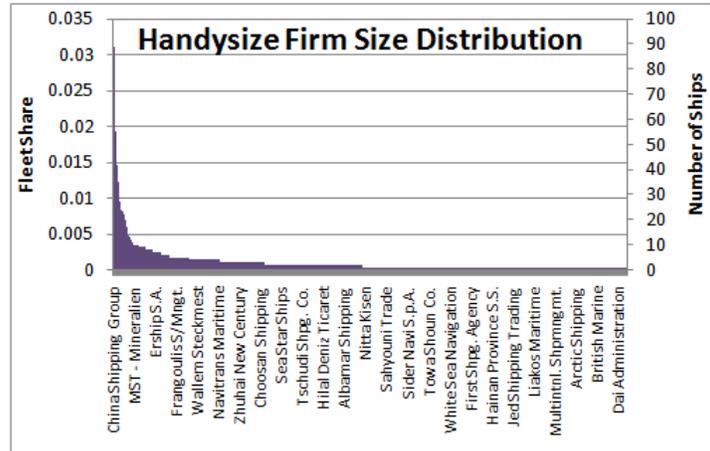


Figure 3: Fleet and fleet shares of shipowning firms.

significantly boosting demand for bulk transport. For instance, Chinese imports grew by 15% per year between 2004 and 2007 (UNCTAD (2009)), while Asian imports represent more than 50% of world imports; Figure 4 shows China's imports of iron ore and coal.

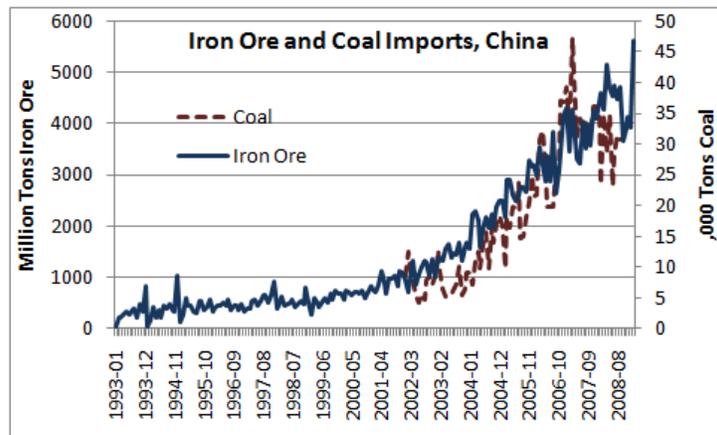


Figure 4: China's iron ore and coal imports.

Supply of shipping services is determined in the short run by the number of voyages carried out by shipowners. The number of voyages is proportional to the speed of the vessel: each period, shipowners choose their navigation speed which determines the number of voyages they supply. Voyage costs include fuel, port/canal dues and cargo handling. They are convex in the chosen speed: there are only so many trips that a ship can execute in a given time period. In addition, voyage costs increase with the ship's age, as its fuel efficiency and overall operation deteriorates over time. Ships also face fixed costs

that include maintenance and insurance. Finally, in the long run, the supply of cargo transportation is restricted by the size of the world fleet and adjusts via the building and scrapping of ships.

Exit in the industry occurs when shipowners scrap their ships. During the scrapping procedure, the interiors of the ship (machinery, furniture, oils, etc.) are removed from the vessel and sold. The ship is then taken to a scrapyards, where it is dismantled and the steel it is made of is recycled.²

Entry in the industry occurs when shipowners buy new ships from world shipyards. The building of new ships is characterized by significant construction lags. In addition to the actual construction time, shipyards often face binding capacity constraints due to their limited number of assembly docks. In that case queues of ship orders are formed, that may considerably increase the time to build. As indicated in Figure 5, the average

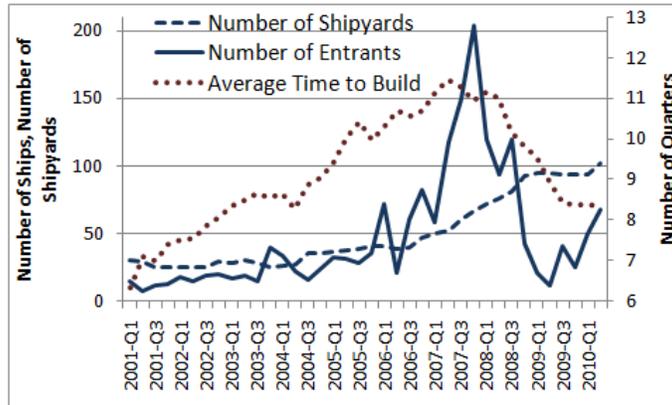


Figure 5: Entry, time to build and number of shipyards with outstanding Handysize orders.

time to build almost doubled between 2001 and 2008 due to the increase in ship orders, despite the threefold increase in shipbuilding. A large fraction of the new shipbuilding capacity was concentrated in China, as shown in Figure 6, as part of China’s “Long and Medium Term Plan for the Shipbuilding Industry 2006-2015”, launched in 2006. More generally, 85% of the shipbuilding activity is located in South Korea, China, and Japan and is heavily subsidized by governments. This industry has been historically supported by governments, as shown in Figure 7, which documents the history of shipbuilding.

²The shipowner also has the option of selling the ship in the second-hand market. In this case the ship continues to operate and does not exit the industry.

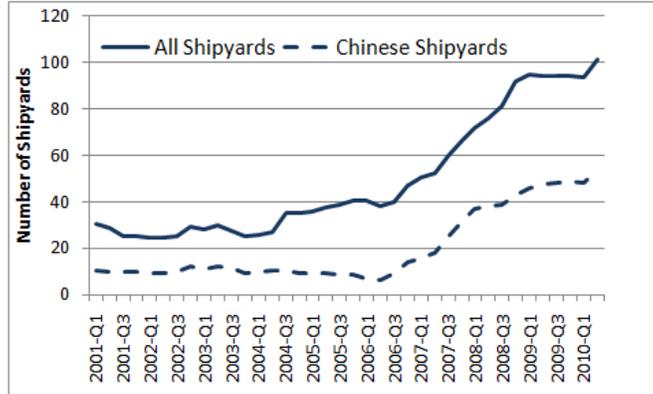


Figure 6: Number of world and Chinese shipyards with outstanding Handysize orders.

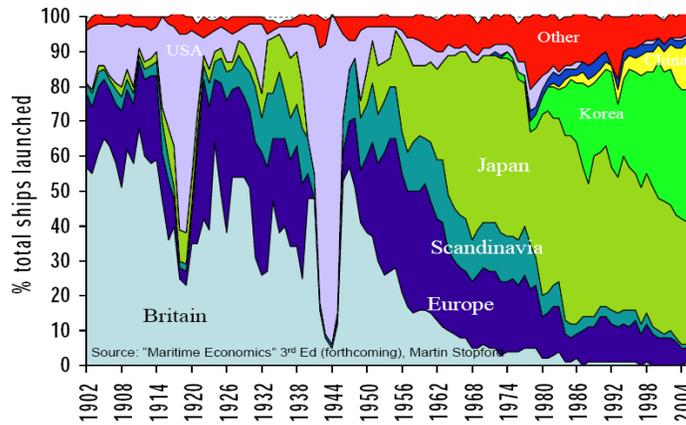


Figure 7: Shipbuilding History. Source: Stopford 2007.

3 Model

In this section, we present a dynamic model for the world bulk shipping industry. Our model lies within the general class of dynamic games studied in Ericson and Pakes (1995) and is closest to the model of entry and exit in Pakes, Ostrovsky and Berry (2007). There are two distinct differences from Pakes, Ostrovsky and Berry (2007): our model includes firm heterogeneity (due to age) and time to build.

3.1 Environment

Time is discrete and the horizon is infinite. A ship is a firm and the terms firm and ship will be used interchangeably. Constant returns to scale with respect to the fleet seems a reasonable assumption for this industry, while the firm size distribution shown in Figure

3 suggests the majority of firms consist indeed of one ship. There are two types of agents: incumbent firms and a large number of identical potential entrants. There is time to build, so that a firm begins its operation a number of periods after its entry decision. The state variable of a ship consists of:

1. its age, $j \in \{0, \dots, A\}$ ³
2. the age distribution of the fleet, $s_t \in \mathbb{R}^A$, where $s_t = [s_t^0 \ s_t^1 \ \dots \ s_t^A]$ and s_t^i is the number of ships of age i
3. the backlog $b_t \in \mathbb{R}^{\bar{T}}$, where $b_t = [b_t^1 \ b_t^2 \ \dots \ b_t^{\bar{T}}]$ and b_t^i , $i = 1, \dots, \bar{T}$, is the number of ships scheduled to be delivered at period $t + i$, while \bar{T} is the maximum possible construction lag⁴
4. the aggregate demand for shipping services, $d_t \in \mathbb{R}$, captured by the intercept of the linear inverse demand curve for freight transport.

Per period payoffs of a ship of age j are given by $\pi_j(s_t, d_t)$. Firms face a linear inverse demand curve,

$$P_t = d_t - \eta Q_t \tag{1}$$

where P_t is the price per voyage and Q_t the total number of voyages offered in period t . $\pi_j(s_t, d_t)$ are the profits resulting from the combination of the demand curve (1) and an assumption on the nature of competition (e.g. Cournot, perfect competition, etc.). We do not specify the nature of competition, allowing for flexibility. The demand state variable d_t , follows an exogenous first order Markov process. Note that the backlog b_t does not affect per period payoffs; it only affects the transition of the state and thus the timing of different profit levels.

Every period, incumbents privately observe a scrap value, ϕ drawn from a distribution F_ϕ and decide whether they will exit the market and obtain ϕ or continue operating.

There is free entry into the market, subject to an entry cost $\kappa(s_t, b_t, d_t)$. All entrants of period t receive the same time to build, T_t . This number of periods T_t changes with the state of the industry (s_t, b_t, d_t) . As the number of orders under construction increases,

³The maximum age A is determined endogenously by the model. Due to data limitations, in the empirical model below we assume that A is 20 years and that it thus represents the age group of "older than 20 years old". In other words, once the ship reaches age 20 it stays in that age group until it (endogenously) exits. See the next section and Appendix A for more details.

⁴We assume that $T_t = T(b_t) \leq \bar{T}$ for all b_t , i.e. ships can be delivered in finite time. This assumption should hold for any equilibrium b_t , as shipowners would not order ships delivered in an infinite number of periods.

shipbuilding capacity constraints bind at different stages of the ship production process. Waiting time is added to the construction time and the time to build increases. Time to build is a deterministic function of the backlog, $T : \mathbb{R}^{A+\bar{T}+1} \rightarrow \mathbb{R}$, so that

$$T_t = T(s_t, b_t, d_t)$$

Ideally, time to build should also be a function of shipbuilding capacities. In that case, shipbuilding capacity would become a state variable in the firms' dynamic optimization problem. It is straightforward to accommodate for this state variable in the theoretical model, as long as its transition is specified. However, introducing capacities in the empirical application is cumbersome, as the already large state space would be further augmented. Specifying the time to build as a function of the entire state (rather than only b_t for instance) captures queues due to period t 's entry, which is a function of the industry state (i.e. time to build depends on N_t). Moreover, it may accommodate for changes in capacities that occur in response to the industry state. Finally, the assumption that the shipbuilding industry is exogenous in the model does not seem unreasonable in light of Figure 7, which shows the history of this industry: governmental support has been crucial for shipyard functioning. Taking shipyards as exogenous, therefore, does not seem an unreasonable assumption.

The timing in each period is as follows: incumbents and potential entrants first observe the state of the industry (s_t, b_t, d_t) . Incumbents then receive an exogenous shock with probability δ , which forces them to sell the ship in the second-hand market. This shock is independent to the value of the ship (a liquidity or retirement shock). In this model, since all shipowners share the value of a ship, the price of a ship in the second hand market equals its value. In addition, as a firm is a ship and the identity of a firm does not matter, sales can be ignored: sales are neither entry -the ship was already in operation- or exit -the ship will continue operating. At a sale the ship just changes ownership. We revisit sales in the empirical part of the paper, where we treat them as observations on the value function. After incumbents observe the state and trade in the second hand market, they privately observe their scrap value and decide whether they will exit the market or continue operating, while potential entrants make their entry decisions simultaneously. Incumbents receive their operating profits, $\pi_j(s_t, d_t)$. Finally, the entry and exit decisions are implemented and the state evolves.

3.2 Firm Behavior

An incumbent firm of age j maximizes its expected discounted stream of profits, deciding only on exit. Its value function before privately observing the scrap value, ϕ , drawn from F_ϕ , is:

$$V_j(s_t, b_t, d_t) = \pi_j(s_t, d_t) + \beta E_\phi \max\{\phi, VC_j(s_t, b_t, d_t)\} \quad (2)$$

where its continuation value is given by,

$$VC_j(s_t, b_t, d_t) \equiv E[V_{j+1}(s_{t+1}, b_{t+1}, d_{t+1}) | s_t, b_t, d_t] \quad (3)$$

The expectation in the expression in (3) is over the number of firms that will enter, N_t and exit, Z_t , as well as the demand for shipping services, d_{t+1} . Randomness in this model results from these three variables, as the remainder of the state evolves deterministically conditional on the current state. In particular, the transition of the age fleet distribution, s_t is as follows:

$$\begin{aligned} s_{t+1}^0 &= b_t^1 \\ s_{t+1}^j &= s_t^{j-1} - Z_t^j, \quad j = 1, \dots, A \end{aligned}$$

The number of age 0 ships in period $t + 1$ equals the first element of the backlog. The remaining elements of s_t age one period and the number of firms that exit at time t , Z_t^j , is subtracted. The transition of the backlog is as follows:

$$\begin{aligned} b_{t+1}^i &= b_t^{i+1}, \quad i \neq T_t, \bar{T} \\ b_{t+1}^{T_t} &= b_t^{T_t+1} + N_t \end{aligned}$$

(unless $T_t = \bar{T}$ in which case, $b_{t+1}^{\bar{T}} = b_t^{\bar{T}} + N_t$ as above, we must set $b_{t+1}^{\bar{T}} = 0$ to maintain the size of $b_t \in \mathbb{R}^{\bar{T}}$). The first element of the backlog enters the market and the remaining elements move one period closer to delivery. Period t entrants, N_t , enter at position $T_t = T(b_t)$. If this position is not the maximum construction lag \bar{T} , zero ships are added at the end of the backlog so that the backlog vector b_t retains its size. Finally, as mentioned above, the demand state variable d_t , follows an exogenous first order Markov process.

An incumbent firm exits if the scrap value drawn, ϕ , is higher than its continuation

value, $VC_j(s_t, b_t, d_t)$. This event occurs with probability:

$$\Pr(\phi > VC_j(s_t, b_t, d_t)) = 1 - F_\phi(VC_j(s_t, b_t, d_t)) \quad (4)$$

where F_ϕ is the distribution of the scrap values. The number of firms that exit at age j , Z_t^j , thus follows the Binomial distribution. We approximate the exit rate by a Poisson process. This approximation is strongly supported by the data, as shown in the following sections and formally justified if the number of firms at age j is large, while the exit probability is small.

We next turn to potential entrants. The value of entry for a ship is:

$$VE(s_t, b_t, d_t) \equiv \beta^{T_t} E[V_0(s_{t+T_t}, b_{t+T_t}, d_{t+T_t}) | s_t, b_t, d_t]$$

Note that due to time to build, a firm that decides to enter in period t will begin operating $T_t = T(s_t, b_t, d_t)$ periods later. The firm, thus, needs to compute the time to build T_t and then forecast the state that will prevail T_t periods in the future. It therefore forms expectations over the number of entrants until delivery, $\{N_t, N_{t+1}, \dots, N_{t+T_t-1}\}$, the number of firms that will exit $\{Z_t^j, Z_{t+1}^j, \dots, Z_{t+T_t-1}^j\}$, all j , as well as demand d_{t+T_t} . Finally, the ship is of age 0 when it is delivered.

A free entry condition determines the entry rate:

$$VE(s_t, b_t, d_t) = \kappa(s_t, b_t, d_t) \quad (5)$$

Note that the entry cost κ varies over time, with the state of the industry, in contrast to the empirical literature on dynamic games, where entry costs are *iid* draws from a distribution. Potential entrants play a symmetric mixed entry strategy. In that case, the number of actual entrants, N_t , follows the Poisson distribution with a state-dependent mean, $\lambda(s_t, b_t, d_t)$.⁵

The equilibrium concept is that of a Markov Perfect Equilibrium. An equilibrium to our model consists of a pair of entry $\lambda(s, b, d)$ and exit $\mu(s, b, d)$ that satisfy conditions (2) to (5). In equilibrium, agents' expectations in VC and VE are formed using the equilibrium strategies. Existence of equilibrium follows by the analysis of Doraszelski and Satterthwaite (2010), Weintraub, Benkard and Van Roy (2008) and Ericson and Pakes (1995).

⁵For a proof see Weintraub, Benkard and Van Roy (2008).

4 Data

The data employed in this paper are taken from the publications of Clarksons, a leading shipbroking firm based in the UK. Four different datasets are utilized. The first consists of world second-hand ship sale transactions. It reports the date of the transaction, the name, age and size of the ship sold, the seller and buyer and the price in million US dollars. The dataset includes sales that occurred between August 1998 and June 2010. We use data on Handysize bulk carriers, the size sector with most second hand sales. We are left with 1838 observations for 48 quarters.

The second dataset consists of shipping voyage contracts and includes the date of transaction, the name and size of the ship and the price per ton. The data are used to estimate the inverse demand curve for shipping services. The dataset includes a subset of voyage contracts between January 2001 and June 2010. There are two limitations to these data. First, as Handysize vessels represent the smallest type of bulk carriers, they carry out a large variety of routes making the exhaustive tracking of vessel movement impossible. To correct for the fact that the contracts observed are, therefore, a strict subset of contracts realized, we compute and use the fleet utilization rate of Capesize and Panamax bulk carriers (the two largest categories whose routes are limited) where the full sample of contracts is observed. Second, the number of per period contracts only partly captures transportation offered by Handysize vessels, as trips vary in both time and distance; ideally, we would want the ton-miles realized each quarter by Handysize ships. Our use of instruments (see below) may correct for this measurement error. The above limitations, as well as the multitude of different types of shipping contracts (and the missing ton-miles realized by age of ship) render the estimation of per period payoffs extremely difficult.

The third dataset consists of quarterly time-series for the orders of new ships (i.e. entrants), deliveries, demolitions (i.e. exitors), fleet and total backlog. It is used to construct the age fleet distribution, $s_t \in \mathbb{R}^A$. Due to data limitations, the final element of s_t , s_t^A , represents the age group of “older than 80 quarters”: once a ship becomes 20 years old, it enters this absorbing age and remains there until it (endogenously) decides to exit. In addition, we never observe a ship younger than 20 years old exiting and we therefore assume that a firm can decide to exit only once it enters the age group of “older than 20 years old”. Therefore, in the empirical model $s_t = [s_t^0, s_t^1, \dots, s_t^A] \in \mathbb{R}^A$ where $A = 80$ and

its transition is:

$$\begin{aligned} s_{t+1}^0 &= b_t^1 \\ s_{t+1}^j &= s_t^{j-1}, \quad j = 1, \dots, A-1 \\ s_{t+1}^A &= s_t^A + s_t^{A-1} - Z_t \end{aligned}$$

Appendix A provides more detailed information on how the age fleet distribution and the backlog are constructed, as well as other data details.

Finally, the fourth dataset is the ship orderbook. It lists all ships under construction each month between 2001 and 2010. The list includes the size, owner, shipyard and delivery date of each ship under construction. We use this dataset to construct the backlog state variable, b_t and estimate the time to build function.

5 Empirical Strategy and Estimation Results

The difficulty in the estimation of dynamic games results from the need to compute continuation values. Recently, a number of papers (e.g. Bajari, Benkard and Levin (2007), Pakes, Ostrovsky and Berry (2007) and Jofre-Bonet and Pesendorfer (2003))⁶ have proposed estimation procedures for dynamic games in which per period payoffs are used to compute continuation values, which in turn allow the estimation of entry and exit costs. Following one of these procedures for the bulk shipping industry using our data would be very difficult: estimating per period payoffs requires detailed and reliable information on cargo shipping contracts which we do not have. In this paper, we use the second-hand ship sale transactions to recover value functions nonparametrically directly from the data. In particular, since in our model there is a large number of identical potential shipowners, a second-hand transaction price must equal the value of the ship. Value functions, along with estimated state transitions provide the necessary information to recover the primitives of the model: the entry costs, the scrap value distribution and the per period payoffs. These primitives are then used to conduct counterfactual experiments.

The empirical strategy is summarized in the following steps:

1. Estimate a demand curve for shipping and create the state variable d_t .
2. Estimate the value function from the second-hand sales

⁶Applications of these methodologies have occurred in several contexts, e.g. Collard-Wexler (2008), Ryan (2009), Dunne, Klimek, Roberts and Xu (2009), Besanko and Doraszelski (2004).

3. Estimate the state transition matrix
4. Estimate nonparametrically the scrap value distribution, F_ϕ
5. Calculate profits from the Bellman equation (2)
6. Estimate the time to build function
7. Calculate entry costs from the free entry condition

We next describe each step, along with the estimation results. Appendix B provides further details on the implementation of these steps. The horizontal axis in all figures in this section consists of the quarters of our sample, $Q3 - 1998$ to $Q2 - 2010$. Each of these quarters corresponds to a state (s, b, d) . We therefore show the results only on this subset of the states that estimation is performed on. This is done for expositional simplicity (the estimation is performed on a larger grid, as described in Appendix B).

5.1 Demand for Shipping Services

We estimate an inverse demand for shipping services via instrumental variables regression. The empirical analogue of (1) is:

$$P_t = \eta_0 + \eta_1 H_t + \eta Q_t + \varepsilon_t \quad (6)$$

where P_t is the average price per voyage in a quarter, H_t includes demand shifters, while Q_t is the total number of voyage contracts realized in a quarter. H_t includes the index of world industrial production (taken from the OECD), China's steel production, the fleet of Handymax ships (bigger size category of ships), the index of food prices, agricultural raw material prices and minerals prices (taken from UNCTAD). We correct for endogeneity and measurement error in a first stage regression of Q_t on the demand shifters and instruments that include supply shifters: the total fleet, the mean and the standard deviation of the age distribution of the fleet.

Results are given in Table 1. The world index of industrial production and China's steel production positively affects prices, while the number of voyages - corrected for endogeneity and measurement error by instruments- has a negative sign. It is not clear what the appropriate sign of commodity prices is, as these capture shifts in both the demand and the supply of commodities and may affect shipping prices either way. The same is true for the Handymax fleet, which may act as a substitute (as suggested by the negative sign), but it may also capture higher overall demand for shipping services.

constant	-11703619 (3009670)*
world industrial production index	137693 (29294)*
agr raw material price	9753.4 (7178.4)
minerals price	-11056.4 (3007.2)*
food price	7241.8 (4210.4)
China's steel production	9143.9 (6603.7)*
Handymax fleet	-2943.5 (737.5)*
\widehat{Q}_t	-644.4 (564.9)

Table 1: Inverse demand curve for freight transport: IV regression results. Standard errors reported in parentheses. Star indicates significance at the 0.05 level.

The demand state variable d_t is taken to be the intercept of the inverse demand curve, i.e.

$$d_t = \log(\widehat{\eta}_0 + \widehat{\eta}_1 H_t + \widehat{\varepsilon}_t)$$

and is plotted in Figure 8. The residual of the IV regression (6) $\widehat{\varepsilon}_t$ is included in d_t , as it captures omitted demand shifters. We consider the log of the intercept as the state to the dynamic game, to guarantee positive prices. Otherwise, the autoregressive scheme specified for d_t below may lead to negative prices.

5.2 Value Function

Under the assumption of a large number of identical potential shipowners, the second hand sale price of a ship must equal its value. Indeed, since shipowners share the value for the ship a seller is willing to sell at price p^{SH} only if $p^{SH} \geq V_j(s, b, d)$, while a buyer is willing to buy only if $p^{SH} \leq V_j(s, b, d)$. As mentioned in the previous section, sales in this model occur due to random shocks independent to the value of the ship that hit shipowners. We therefore use the resale dataset to estimate the value function of a ship as a function of its age j and the remaining state variables (s, b, d) . The resale price is thus given by:

$$p^{SH} = V_j(s, b, d) + \varepsilon$$

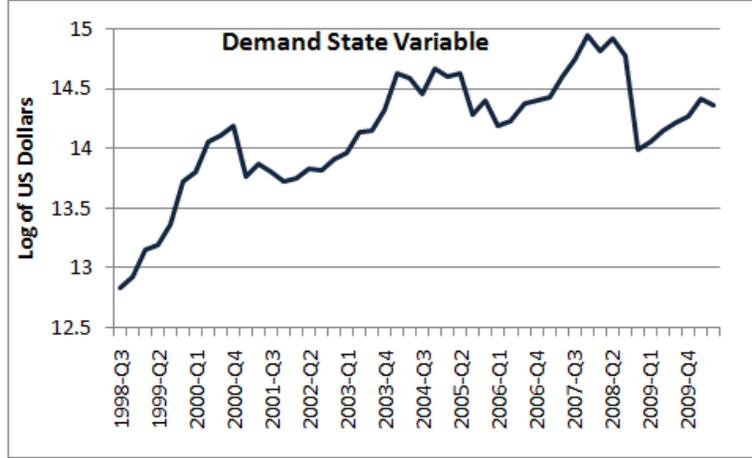


Figure 8: Demand state variable.

where ε captures measurement error. Note that this model does not allow for unobservable state variables. Such a state should be an argument of the value function and not in ε . The value function at a grid point (j, s, b, d) is given by the conditional expectation of the second hand price, p^{SH} on (j, s, b, d) :

$$V_j(s, b, d) = E [p^{SH} | j, s, b, d]$$

We estimate the value function by local linear regression, i.e. we perform linear regression at every grid point $x = (j, s, b, d)$, including a weighting scheme that downweighs the contributions of data points away from $x = (j, s, b, d)$. The construction of the grid on which value functions are estimated is described in the next section and Appendix B. We therefore solve the following local weighted least squares problem, for all grid points x :⁷

$$\min_{\beta_0, \beta_j, \beta_s, \beta_B, \beta_d} \sum_i \left\{ \begin{array}{l} p_i^{SH} - \beta_0(x) - \beta_j(x)(j_i - j) - \beta_s(x)(s_i - s) \\ - \beta_B(x)(b_i - b) - \beta_d(x)(d_i - d) \end{array} \right\}^2 K_h(x_i - x)$$

where $\{p_i^{SH}, j_i, s_i, b_i, d_i\}_{i=1}^{1838}$ represent the data of second hand sale transactions, while K_h is a normal kernel with bandwidth h , so that $K_h(x_i - x) = \frac{1}{2h\sqrt{\pi}} \exp(-((x_i - x)^2 / 2h))$, where $x_i = [j_i, s_i, b_i, d_i]$, $x = [j, s, b, d]$. The bandwidth is different for each state variable based on its observed range. We have experimented with k -nearest neighbor bandwidths,

⁷For more information on local linear regression, see Fan and Gijbels, (1996) or Pagan and Ullah (1999).

cross-validation, as well as non-normal kernels.⁸

Figure 9 shows the estimated value functions. Each quarter is represented by its corresponding state (s, b, d) which is the same for ships of all ages. Different curves represent different ages and the order is preserved so that older ships have a lower value function for all (s, b, d) . A subset of all ages (corresponding to 1 year old, 2 years old and up to 20+ years old) is depicted in the figure, for expositional simplicity.

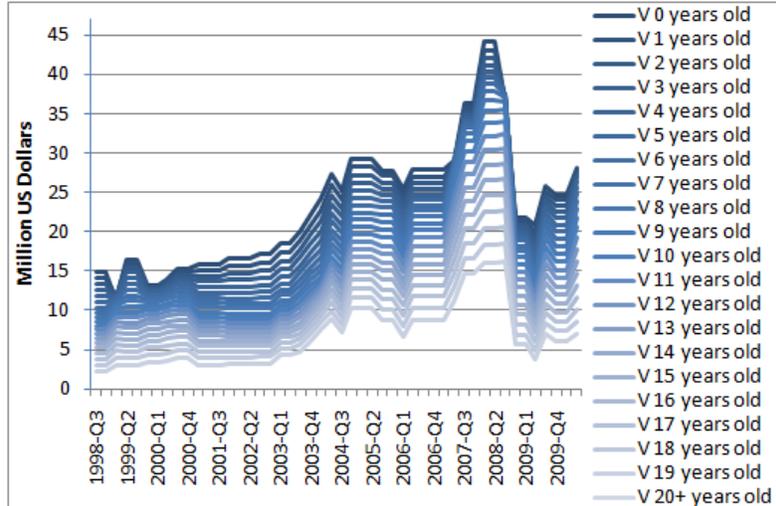


Figure 9: Plots of value functions for different ages over observed state path.

Pointwise confidence intervals can be computed using the asymptotic variance of the local linear regression estimator (see Fan and Gijbels (1996)). Figure 10 shows the value function of a 0 year old ship and that of a 20 year old ship with the computed 95% confidence intervals. As more observations are available for old ships, confidence intervals are tighter.

Estimating value functions using resale prices is central in our empirical strategy. One may worry about both adverse selection, as well as sample selection. Adverse selection would arise in the model if different shipowning firms were allowed to privately observe the quality of the ship. Practitioners assure that this is not a concern: once a ship enters the resale market its history of maintenance and accidents is public information. In addition, potential buyers thoroughly inspect the ship; quoting Stopford (2009): “Full details of the ship are drawn up, including the specification of the hull, machinery, equipment, class, survey status and general equipment. ... (Inspections) will generally include a physical

⁸Adland and Koekebakker (2007) also use second hand ship sale prices. They aim to uncover non-parametrically how a ship’s value in the second hand market is affected by its age, its size and freight rates.

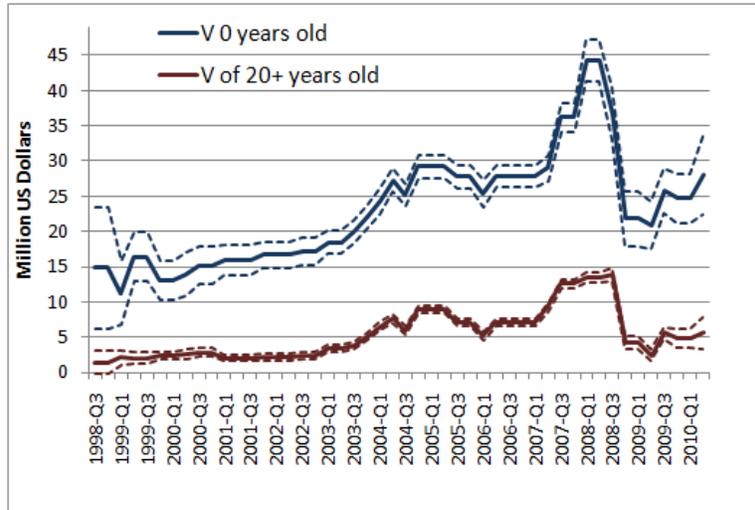


Figure 10: Value Functions and Confidence Intervals.

inspection of the ship, possibly with a dry docking or an underwater inspection by divers ... The buyer will also inspect the classification society records for information about the mechanical and structural history of the ship”. Selection on unobservables seems a more worrisome concern. If only low quality ships are traded in the market, even if this is common knowledge, it will lead to a bias in the estimated value functions. The second-hand market for ships is thick: about 10% of the fleet is traded annually. In addition, price variation conditional on a state (j, s, b, d) is very small: usually significantly lower than 10%. Unfortunately, we do not observe the distribution of observable characteristics for ships not in our second-hand sample. But our concerns for econometric selection are put to rest when we turn to the external validation of our model and estimation, presented at the end of this section. In particular, Figure 18 compares the estimated entry costs to the average shipbuilding price. The graph is surprisingly good suggesting that selection can not be a severe issue. Finally, as the assumption of a large number of identical potential shipowners is crucial, let us also discuss the possibility of search costs: based on Stopford (2009) search costs are not a severe concern either; “The shipowner ... offers the vessel through several broking companies. On receipt of the instruction the broker will telephone any client ... (and) call up other brokers in order to market the ship”.

In this paper, second-hand prices are the tool we use to estimate our dynamic model of ship entry and exit. Following the discussion above, the bulk shipping industry provides an ideal setup for our empirical methodology. A large literature examines resales in other, more concentrated industries, such as Gavazza (2010) who examines the case of

commercial aircraft.

5.3 Transition Matrix

As pointed out in Section 3, the state variable consists of the age distribution of the fleet s_t of dimension A , the backlog b_t of size \bar{T} and the demand d_t (as well as the ship's age which evolves deterministically and we thus omit from this discussion). The overall state has dimension $A + \bar{T} + 1$ and evolves according to the law of motion described in Section 3. In the empirical model, we specify the evolution of the demand state variable d_t as a first order autoregressive process with non-normal disturbances,

$$d_t = c + \rho d_{t-1} + \sigma \varepsilon_t \quad (7)$$

where ε_t follows the Student t distribution with ν degrees of freedom, so that

$$f_\varepsilon(\varepsilon) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\nu\pi}\Gamma\left(\frac{\nu}{2}\right)} \left(1 + \frac{\varepsilon^2}{\nu}\right)^{-\left(\frac{\nu+1}{2}\right)}$$

The Student- t innovations allow for a fat tail distribution (compared to normal) of the demand state variable. Extreme shocks are thus more likely under this specification.

Consider the state

$$x_t = \left[s_t^0, \dots, s_t^A, b_t^1, \dots, b_t^{\bar{T}}, d_t \right]$$

Then,

$$\begin{aligned} x_{t+1} &= \left[s_{t+1}^0, \dots, s_{t+1}^A, b_{t+1}^1, \dots, b_{t+1}^{\bar{T}}, d_{t+1} \right] \\ &= \left[b_t^1, s_t^0, \dots, s_t^{A-2}, s_t^A + s_t^{A-1} - Z_t, b_t^2, \dots, b_t^{T(b_t)} + N_t, \dots, b_t^{\bar{T}-1}, 0, d_{t+1} \right] \end{aligned}$$

where N_t indicates the number of entrants (ship orders) at time t when the state is x_t and Z_t represents the number of ships destined for scrap at state x_t . In the stochastic setting we are interested in the random variables ε_{t+1} , N_t and Z_t . More precisely, ε_{t+1} is i.i.d Student- t distributed and summarizes the uncertainties inherent in demand for shipping services. Entrants N_t and exitors Z_t are distributed Poisson with rates $\lambda(x_t)$ and $\mu(x_t)$ respectively, introducing nonlinearity in the above state transitions. Under the assumption that N_t , Z_t and ε_{t+1} are conditionally independent given the state x_t , the

transition probability from x_t to x_{t+1} is given by:

$$Q(x_t, x_{t+1}) f_N(N_t = n | x_t) f_Z(Z_t = z | x_t) f_\varepsilon(\varepsilon_{t+1}) \quad (8)$$

where

$$Q(x_t, x_{t+1}) = \delta\{s_{t+1}^0 = b_t^1\} \prod_{i=1}^{A-1} \delta\{s_{t+1}^i = s_t^{i-1}\} \prod_{i \neq T(b_t)} \delta\{b_{t+1}^i = b_t^{i+1}\} \quad (9)$$

$n = b_{t+1}^{T(b_t)} - b_t^{T(b_t)+1}$, $z = s_t^A + s_t^{A-1} - s_{t+1}^A$, $e = d_{t+1} - c - \rho d_t$, while $\delta\{\cdot\}$ is an indicator function, f_N and f_Z are the Poisson densities with parameter $\lambda(x_t)$ and $\mu(x_t)$ respectively and f_ε is Student- t with ν degrees of freedom. In words, the first element of the age distribution s_t equals the first element of the backlog (scheduled deliveries for this period), while all remaining elements of s_t age one period (by shifting one position). All elements of the backlog move one period closer to delivery. The estimated time to build function $T(s_t, b_t, d_t)$ (see below) is used to determine the position of entrants. Finally, exit is subtracted from s^A and entry is added in the last position of the backlog.

To construct the transition matrix, we first estimate the Poisson processes of entry and exit and the demand transition (7). Second, we use clustering techniques to maintain a tractable state space. The state transitions described in (8) and (9) pose significant computational challenges that arise from the high state dimension and the large number of integer values the state components may take. To develop manageable finite state space approximations we employ aggregation and clustering techniques. Aggregation serves to drive down the dimension of the state space. We consider three age groups, $S_t = (S_t^1, S_t^2, S_t^3)$, where $S_t^1 = \sum_{i=0}^{A_1} s_t^i$, $S_t^2 = \sum_{i=A_1+1}^{A_2} s_t^i$ and $S_t^3 = \sum_{i=A_2+1}^{A_3} s_t^i$,⁹ the total backlog $B_t = \sum_{i=1}^T b_t^i$ and demand d_t . Aggregation compresses linearly the state into a much smaller vector. Clustering techniques are employed to reduce the number of admissible values for these aggregated states, by constructing representative states via an iterative algorithm (see Appendix B for details.).

Once a finite state space is constructed we compute the transition matrix in two steps. First, we simulate thousands of states $x_t \in \mathbb{R}^{A+T+1}$ using the estimated Poisson processes for N_t and Z_t and the $AR(1)$ process for d_t . Second, we replace each simulated state by its corresponding aggregate state $[S_t, B_t, d_t] \in \mathbb{R}^5$ and we compute a frequency transition matrix, as detailed in Appendix B.¹⁰

⁹In particular, the first age group includes the number of ships between 0 and 10 years old, so that $S_t^1 = \sum_{i=0}^{40} s_t^i$, the second group includes the number of ships between 10 and 20 years old so that $S_t^2 = \sum_{i=41}^{79} s_t^i$ and the third group is the number of ships older than 20 years old $S_t^3 = s_t^{80}$.

¹⁰Another possible approach is to not reduce the size of the finite state. This approach entails integrat-

We next turn to the results. We assume that the Poisson entry rate equals:

$$\lambda(S_t, B_t, d_t) = e^{(\gamma_0 + \gamma_S S_t + \gamma_{B1} B_t + \gamma_{B2} B_t^2 + \gamma_d d_t)}$$

and that the Poisson exit rate equals

$$\mu(S_t, B_t, d_t) = e^{(\delta_0 + \delta_S S_t + \delta_{B1} B_t + \delta_{B2} B_t^2 + \delta_d d_t)}$$

Since λ and μ are greater or equal than zero, they can be expressed as $\exp(f(S_t, B_t, d_t))$. We can therefore think of the chosen specification as a Taylor approximation to $f(\cdot)$ (first order for S_t and d_t and second order for B_t).

We estimate the parameters γ and δ via OLS using the time series of entrants and exitors. It is possible to estimate $\lambda(S_t, B_t, d_t)$ and $\mu(S_t, B_t, d_t)$ nonparametrically as a function of the state via local linear regression, but due to the short time series and the large number of states, we have chosen the parametric approach. Results from the entry regression are reported in Table 2. As expected, entry is increasing in demand d : as the

	constant	S^1	S^2	S^3	B	B^2	d
Param	3.57	-0.003	-0.0011	-0.0016	0.011	$-8.25e - 006$	0.213
s.e.	(4.51)	(0.002)	(0.0011)	(0.001)	(0.0018)*	($1.47e - 006$)*	(0.135)*

Table 2: Entry regression estimates. Standard errors reported in parentheses. Star indicates significance at the 0.5 level.

demand curve for freight shifts outward, more entry is attracted to the market. Moreover, entry is decreasing in the number of competitors, $S = (S^1, S^2, S^3)$, with the number of young competitors S_1 affecting entry more strongly. Figure 11 compares the estimated mean entry rate $\lambda(S_t, B_t, d_t)$ to the observed number of entrants, N_t . The Poisson entry process fits the data well. Examining the observed number of entrants N_t one notices that the variance seems to follow the mean, as in the Poisson density. Last, the observed series N_t provide another testament of the large investment swings in bulk shipping: note the spike in new ships orders during 2006 – 2008.

Similarly, we estimate the parameters of the exit Poisson process. Table 5 presents the results. As expected, exit is decreasing in demand d_t : as the demand curve for

ing the value function at all possible states in the next period in order to compute continuation values. It is computationally burdensome. Moreover, integrating the value function T times in the future in order to recover the entry costs introduces noisyness, as local linear regression does not perform as well at such distant states (especially at the end of the sample).

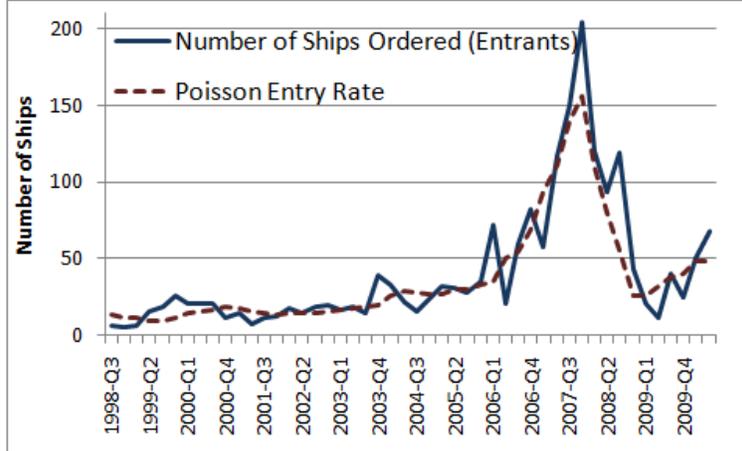


Figure 11: Observed Entrants and Poisson entry process.

	constant	S^1	S^2	S^3	B	B^2	d
Param	-1.547	0.0064	0.0029	0.0029	-0.012	$1.055e - 005$	-0.34
s.e.	(7.09)	(0.0038)	(0.0017)	(0.0023)	(0.0029)*	$(2.31e - 006)^*$	(0.21)

Table 3: Exit regression estimates. Standard errors reported in parentheses. Star indicates significance at the 0.5 level.

freight shifts outward, exit becomes a less attractive option. Exit is increasing in the number of competitors $S = (S^1, S^2, S^3)$. Figure 12 compares the estimated mean exit rate $\mu(S_t, B_t, d_t)$ to the actual number of firms that exit, Z_t . The Poisson exit rate fits the data relatively well. Our model implies that the number of firms that exit each period follows the Binomial density. The observed number of ships older than 20 years old, s^A , ranges in our sample from 1206 to 1718 ships. As shown in Figure 12, the number of exitors is very small compared to the magnitude of s^A , thus justifying the Poisson approximation. Finally, note the practically zero exit rate between 2005 and 2007 when demand for shipping services was high, followed by a spike in exit at the crisis of 2008.

Next, we estimate the transition of d_t which is assumed to follow an $AR(1)$ with Student- t with v degrees of freedom innovations, given in (7). We estimate (c, ρ, σ^2, v) via Maximum Likelihood, as in Hamilton (2005). Results are shown in Table 4. The presence of non-normal innovations significantly increases the likelihood, supporting the presence of fat tails in the distribution. Note that the demand process exhibits high persistence, with $\rho = 0.91$ at a quarterly level. It is stationary though, as $\rho < 1$.

Once we obtain the transition matrix, $P \in \mathbb{R}^{L \times L}$ -where L is the sized of the reduced

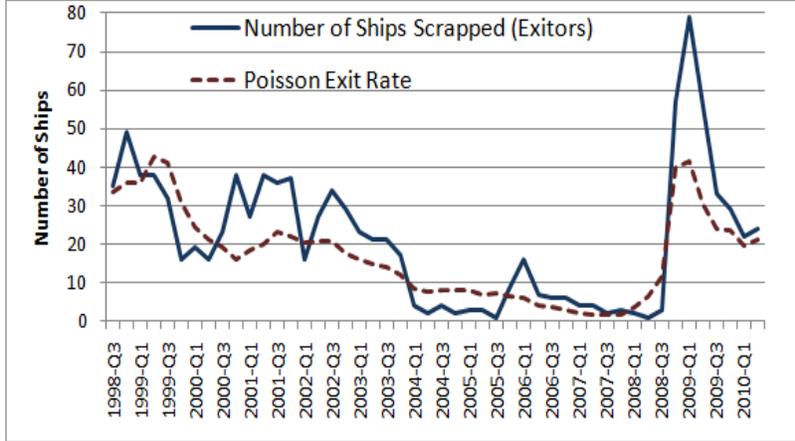


Figure 12: Observed Exitors and Poisson exit process.

	c	ρ	σ^2	v
Param	1.36	0.91	0.096	2.39
s.e.	(0.475)	(0.034)	(0.0036)	(1.09)

Table 4: Demand transition parameters. Standard errors reported in parentheses.

state space- and the value function $V \in \mathbb{R}^{L \times A}$, the continuation value of an age $j < A$ ship is simply $VC_j = PV_{j+1} \in \mathbb{R}^L$, while $VC_A = PV_A$. We use the continuation value of age A ships, VC_A to estimate the scrap value distribution, since exit is only allowed from age group A .

5.4 Scrap Value Distribution

The estimated value function of old ships V_A provides information directly on the scrap values. To obtain the density of scrap values, we use the estimated continuation value and the observed exit combined with a simple exit equation.

We estimate the scrap value distribution nonparametrically, via local linear regression of the observed exit probabilities $\frac{Z_t}{s_t^A}$ on the continuation value of old ships, $VC_A(x_t)$. Indeed, the probability of exit at continuation value v equals $[1 - F_\phi(v)]$, or

$$\frac{Z}{s^A} = [1 - F_\phi(v)]$$

We run the following weighted regression at all points of interest v ,

$$\min_{\beta_0, \beta_{VC}} \sum_t \left\{ \frac{Z_t}{s_t^A} - \beta_0(v) - \beta_{VC}(v)(VC_A(x_t) - v) \right\}^2 K_h(VC_A(x_t) - v)$$

where K_h is a normal kernel with bandwidth h . We therefore get that $[1 - F_\phi(v)] = \beta_0(v)$. Note also, that the local linear regression estimator implies that β_{VC} is equal to the derivative of $\beta_0(v)$ at v . Therefore, $\beta_{VC}(v)$ is the density of the scrap values at v .

The density of scrap values is shown in Figure 13, while the distribution is shown in Figure 14. Our data provide information directly on the value function of an old ship, which in turn is translated in the quartiles of the scrap value density. The horizontal axis in Figure 13 is essentially derived from the range of continuation values.¹¹

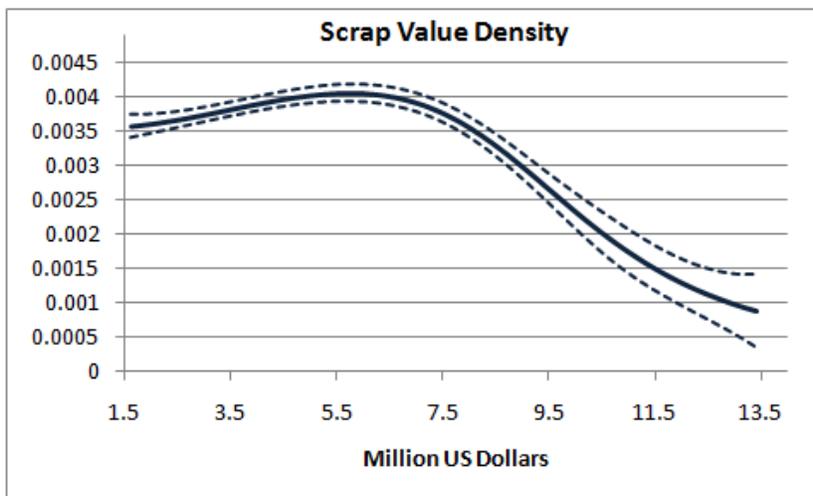


Figure 13: Density of scrap values. Estimated via local linear regression.

5.5 Profits

We compute profits $\pi \in \mathbb{R}^{L \times A}$ from the Bellman equations (2), which we repeat here for convenience:

$$V_j = \pi_j + \beta E_\phi \max \{ \phi, VC_j \}$$

¹¹Since profits can be negative due to fixed costs, value functions can be negative. The scrap value density can be defined on negative numbers as well. Negative scrap values amount to incurring costs in demolishing the ship (e.g. scrapyards are capacity constrained and the ship needs to be navigated far away, it is hard to sell its interiors, etc.). The value function in the observed range of data, however, is positive.

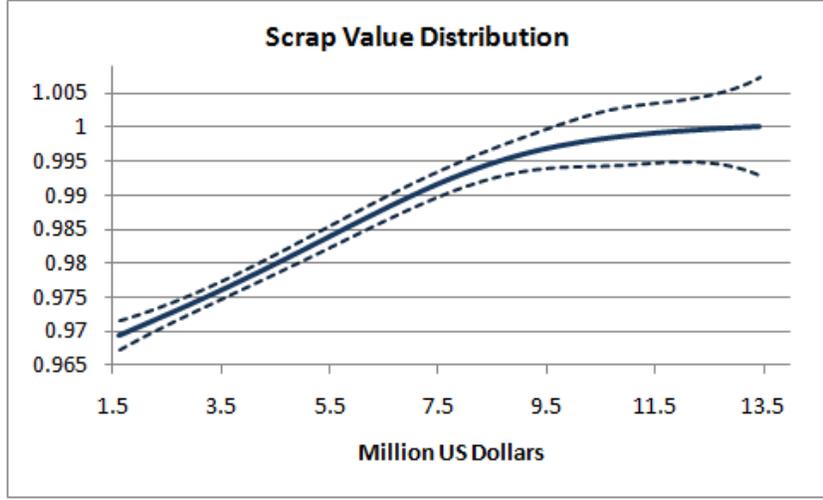


Figure 14: Distribution of scrap values. Estimated via local linear regression.

Profits then can be estimated from:

$$\pi_j = V_j - \beta VC_{j+1} \quad (10)$$

for $j = 0, 1, \dots, A - 1$, while

$$\pi_A = V_A - \beta E_\phi \max \{ \phi, VC_A \}$$

or

$$\pi_A = V_A - \beta (1 - p^X) VC_A - \beta p^X E [\phi | \phi > VC_A] \quad (11)$$

where $VC_j = PV_{j+1}$ and $p^X \in \mathbb{R}^L$ is the exit probability, and $A = 80$ so that

$$p^X = \frac{\mu(S, B, d)}{s^A}$$

Note that V_j , all j , as well as the transition matrix, P , were estimated in the previous steps of the estimation. Moreover, the exit Poisson rate $\mu(S, B, d)$ was estimated towards the construction of the transition matrix.

Figure 15 presents the estimated profits. We show the profits of older than 20 years old ships, as well as the average profits of ships 10 – 20 years old. Young ships are traded in the second-hand market less than older ones, leading to few observations on younger ships in our second-hand ship sale dataset. As a result, the estimates of value functions and profits are more noisy for young ships.

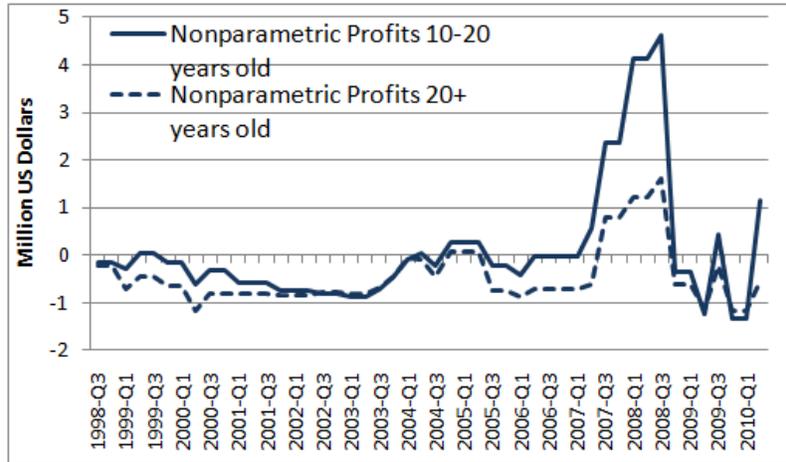


Figure 15: Estimated profits. Average across age 10-20 years and 20+ years old.

In order to assess the validity of our model and estimation results, we use an earnings index constructed by shipping brokers. This index is not used in the estimation. The scale of this index is uninformative (not only is it an index, but also earnings calculated by brokers do not include fixed costs, opportunity costs, etc.) and we only use it to assess the evolution of profits over time. Our estimated profits track broker calculated profits well, as shown in Figure 16. Note that profits fluctuate close to zero between 1998 and 2006, spike during the boom of 2007 – 2008 and collapse at the crisis.

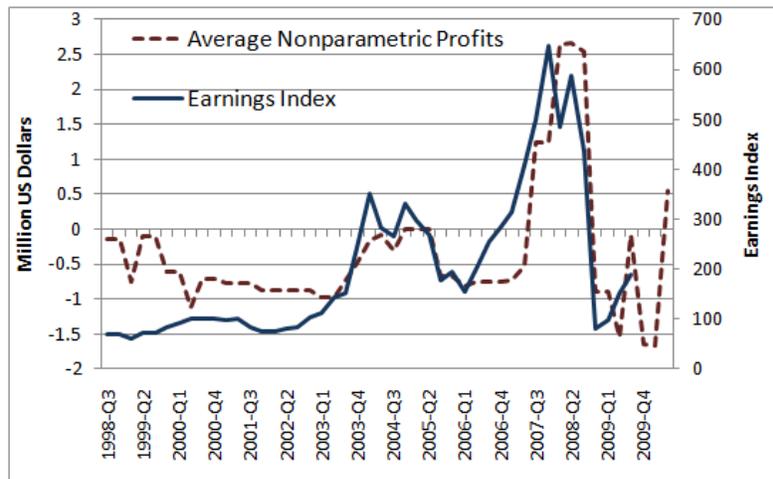


Figure 16: Comparison of the weighted average (across ages) nonparametric profits and the earnings index constructed by brokers. The index is based on a 45000DWT Handymax vessel and takes into account freight rates and fuel/port costs for several representative routes.

5.6 Time to Build Function

We use the ship orderbook data to construct the average time to build corresponding to each period. Consistent with our state space reduction, we assume that time to build is a function of the aggregate state, (S_t, B_t, d_t) . We estimate the following time to build function $T(S_t, B_t, d_t)$:

$$\log(T_t) = a_0^{TTB} + a_S^{TTB} S_t + a_B^{TTB} B_t + a_d^{TTB} d_t$$

Figure 17 presents the actual and predicted time to build, while Table provides the estimated parameters. Consistent with conversations with practitioners, the minimum

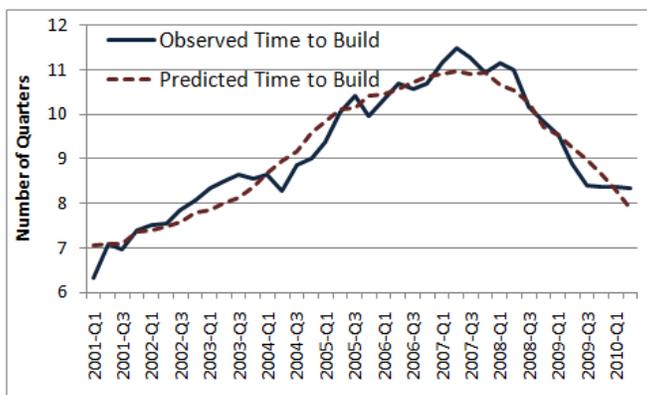


Figure 17: Average observed time to build and Estimated time to build.

time to build is 4-5 quarters under ideal conditions (i.e. a high quality shipyard with experience in the same vessel).

	constant	S^1	S^2	S^3	B	d
Param	2.25	-0.0008	-0.0006	0.0001	$1.4e - 005$	0.051
s.e.	(0.8)*	(0.0003)*	(0.0002)*	(0.0002)	($5.4e - 005$)	(0.03)

Table 5: Exit regression estimates. Standard errors reported in parentheses. Star indicates significance at the 0.5 level.

It is possible to recover $T(\cdot)$ nonparametrically via local linear regression, but the parametric form allows for a counterfactual exploring the impact of shipbuilding subsidies, as explained in the Counterfactual section.

5.7 Entry Costs

The free entry condition states that at each aggregate state $x = (S, B, d)$, the value of entry equals the entry cost, $\kappa(x)$. The value of entry is the expectation of the value function of age 0, $T(x)$ periods ahead. To construct the value of entry at x we first compute the $T(x)^{th}$ power of the transition matrix. We then multiply the row of that matrix corresponding to state x with $V_0 \in \mathbb{R}^L$ as well as $\beta^{T(S,B,d)}$, i.e.,

$$\beta^{T(x)} P_x^{T(x)} V_0 = \kappa(x) \quad (12)$$

where $\kappa \in \mathbb{R}^L$ are the entry costs, $V_0 \in \mathbb{R}^L$ is the value of an age 0 ship and $P_x^{T(x)}$ refers to the row of $P^{T(x)}$ that corresponds to state x . Having computed the transition matrix P , the time to build function $T(x)$ and the value function V_0 we can directly estimate entry costs κ from the free entry condition.¹²

In practice, the largest portion of entry costs is the price charged by the shipyard. Additional costs include legal fees, negotiation costs with the shipyard and engineering supervision during construction. We use the average shipbuilding price to evaluate the power of our model and estimation results. Like the earnings index in Figure 16, the shipbuilding price is not used in the estimation and only acts as an “additional degree of freedom”. Moreover, as mentioned above, our estimate of the value function of an age 0 ship is noisy due to shortage of observations on young ships. We use this noisy estimate, as well as the $T(x)^{th}$ power of our transition matrix to generate entry costs. Figure 18 compares entry costs to the shipbuilding price. Considering that the noisiest estimate is used and combined with the free entry condition implied by our model and the estimated transition matrix, the fit is surprisingly good.

¹²The free entry condition holds with equality only when entry is strictly positive. If entry is zero, then the value of entry provides only a lower bound for the entry costs.

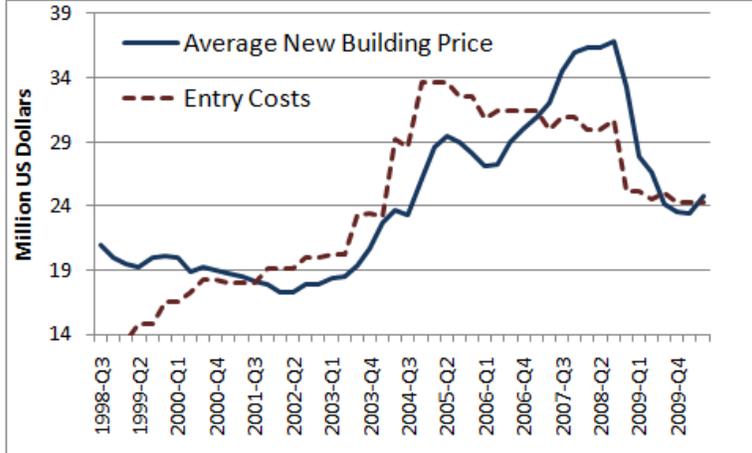


Figure 18: Estimated entry costs and newbuilding price index.

6 Counterfactuals

In this section, we quantify the impact of demand uncertainty and time to build on freight rates, investment in new ships and surplus. In order to analyze the behavior of shipping prices and surplus, we need to take a stand on the nature of per period competition in the shipping industry. Even though counterfactuals concerning the investment in new ships can be performed employing the nonparametric estimates of the model, counterfactuals concerning prices and surplus require structure on the profit function. We therefore assume that firms engage in competition in quantities (Cournot game) in the market for freight transport and use the estimated nonparametric profits to estimate the cost parameters.

6.1 A Cournot Game for Profits

We assume that firms engage in competition in quantities (Cournot game), facing a common demand function and convex operating costs that increase in the ship's age. In particular, every period, a firm of age j chooses the number of voyages q in order to maximize profits:

$$\max_q \{Pq - c_j q^3 - F_j\} = \max_q \{(d - \eta Q)q - c_j q^3 - F_j\}$$

where $P = d - \eta Q$ is the inverse demand function (as estimated in step 1), $Q = \sum q$, and $(c_j q^3 + F_j)$ are total costs that include fixed costs F_j . Based on Stopford (2009), fuel consumption is proportional to the square or cube of the vessel's speed, depending on its

engines. We further assume that the number of trips is proportional to speed and that total costs are proportional to fuel costs. Assuming that $c_j \in \{c_{0-10}, c_{10-20}, c_{20+}\}$ and $F_j \in \{F_{0-10}, F_{10-20}, F_{20+}\}$ so that costs change when the ship's age crosses a new decade, we can solve for profits $\{\pi_{0-10}, \pi_{10-20}, \pi_{20+}\}$ as a function of the state $[S^1, S^2, S^3, d]$. Price is also expressed as a function of the state $[S^1, S^2, S^3, d]$.

We use the estimated profits from equations (10) and (11)- in particular, we use the mean recovered profit for ages 10 – 20 for π_{10-20} and π_{20+} , as well as observed prices and recover the cost parameters via GMM (matching the mean and standard deviation of each equation). We impose a linear dependence of these parameters on age, so that $c_j = \gamma_0 + \gamma_1 j$ and $F_j = \delta_0 + \delta_1 j$ and recover $\{\gamma_0, \gamma_1, \delta_0, \delta_1\}$. The cubic cost function does not allow closed-form expressions for prices and profits and we therefore solve the game during the optimization. The fit of the Cournot game is shown in Appendix D.

We choose to model per period competition as a Cournot game because the assumptions of homogeneous firms and a market clearing price seem reasonable in our framework. In addition, Cournot competition is consistent with perfect competition when the number of firms is large. In fact, our empirical findings suggest that firms do act as price takers, as shown below.

Comparing the predicted Cournot price to the estimated marginal costs, $\{3c_j q_j^2\}$ we find that the price is very close to marginal costs. Indeed, the mean difference between the price and $3c_1 q_1^2$ is 249 dollars with standard deviation 126 dollars. The mean difference between price and $3c_2 q_2^2$ is 229 dollars with standard deviation 116 dollars. Finally, the mean difference between the price and $3c_3 q_3^2$ is 211 dollars with standard deviation 107 dollars. Note, however, that firms can still make positive profits, due to non-constant marginal costs.

6.2 Time to Build

We now quantify the impact of demand uncertainty and time to build on freight rates, investment in new ships and surplus. We do so by computing the equilibrium of the model under different counterfactual scenarios. In particular, we compute the entry policy $\lambda \in \mathbb{R}^L$, the exit policy $\mu \in \mathbb{R}^L$ and value function $V \in \mathbb{R}^{L \times A}$, using the estimated primitives, i.e. the profits $\pi \in \mathbb{R}^{L \times A}$, the scrap value distribution F_ϕ and the entry costs $\kappa \in \mathbb{R}^L$.¹³

We compute the equilibrium of the game under two counterfactual scenarios: constant

¹³See Appendix C for details.

time to build and no time to build. We compare these two worlds to the true world of endogenous time to build, i.e. $T_t = T(S_t, B_t, d_t)$. The case of constant time to build captures a situation without shipbuilding capacity constraints, allowing for only a time invariant period of construction. Regardless of how many ship orders are placed, construction begins immediately and they are all delivered a fixed number of periods later. We choose this constant time to build to equal six quarters, a quarter less the minimum time to build we observe. In a world without time to build, firms that decide to enter at period t , begin operating at period $t + 1$.

We first explore the impact of a positive shock, as this is the case where time to build becomes a crucial constraint. Following a positive demand shock, supply is restricted because existing ships face capacity constraints and new ships require time to build. This leads to high profits and prices and low shipper surplus. If time to build were reduced or removed, supply would be more elastic and the size of fleet would increase in order to satisfy the increase in demand, reducing prices and profits. We simulate the model 500 times: after 7 quarters, we hit the economy with the positive shock that has probability of occurrence 0.12. We continue the simulations for another 25 quarters.

As shown in Figure 19 when the positive shock hits the industry, entry spikes. In the case of no time to build as well as constant time to build this spike is large. In contrast, in the case of the endogenous time to build, the entry spike is considerably smaller. As the demand process is mean-reverting the shock fades and ships that are delivered late can't take advantage of the increased demand for shipping services. In addition, entrants face the risk of raising the time to build too much. Finally, more ships enter in the case of no time to build compared to the case of constant time to build, benefiting from the positive shock.

Figure 20 demonstrates the response of exit to the positive demand shock. Time to build induces lower exit activity. As it takes time for new ship orders to be delivered, existing old ships prefer to remain in the industry and reap the benefits of increased demand, rather than exit. In the case of no time to build exit is even higher than in the case of constant time to build as new more efficient ships can immediately take the place of older ships.

Freight rates are lower as time to build decreases, as shown in Figure 21. Prices increase when the shock hits, as in all three cases entrants need to wait at least one period before operating. The decline, however, is steeper moving from endogenous to constant to no time to build. Moreover, during the first six quarters (the duration of the constant time to build) price does not differ in a world with constant versus endogenous

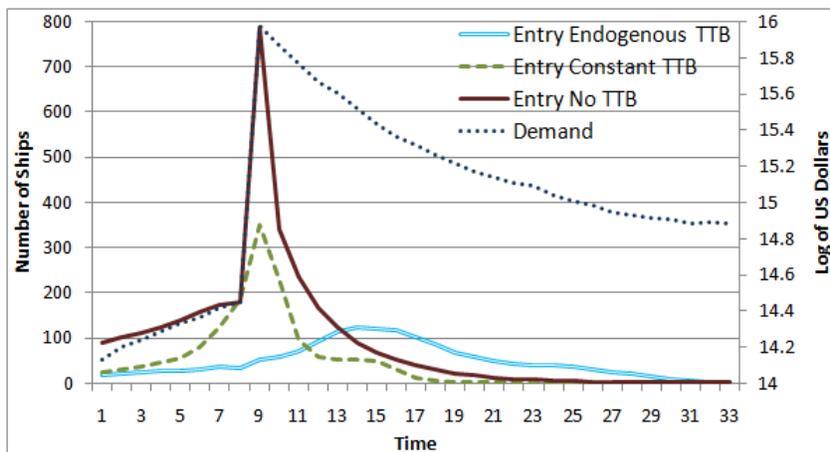


Figure 19: Entry under endogenous, constant and no time to build (TTB). Demand depicted on right axis.

time to build. As a result of the above, consumer surplus is also higher as time to build decreases, as shown in Figure 22. In this case, consumer surplus refers to the surplus of the customers of shipping companies, i.e. the firms that require transportation services (shippers), rather than the consumers of the final goods.

Figures 23 and 24 show total and per ship profits. Total profits are lower as time to build increases. In contrast, profits per ship are higher as time to build increases. The difference between the two figures, captures the difference between the extensive and the intensive margin. Moving from endogenous to constant to no time to build, the fleet increases and so total profits increase. In the presence of time to build and a positive demand shock, however, operating firms reap high profits, as new firms can't enter immediately. Therefore, moving from no to constant to endogenous time to build, profits per ship increase.

Figure 25 shows the sum of consumer surplus and total profits, which increases as time to build decreases when a positive shock hits the industry.

To assess the impact of time to build on fleet evolution, as well as on the volatility of investment and prices, we perform long-run (240 quarters) simulations of the model. Preliminary results suggest that the size of the fleet is 30% larger in a world without time to build: entry is higher and even though exit is also higher, the first effect dominates. The fleet is also younger on average as time to build declines. The fleet exhibits oscillatory behavior in all cases (endogenous, constant, no time to build). Fleet cycles under endogenous time to build lag those under constant time to build. The covariance of the fleet with demand increases as time to build moves from endogenous to constant

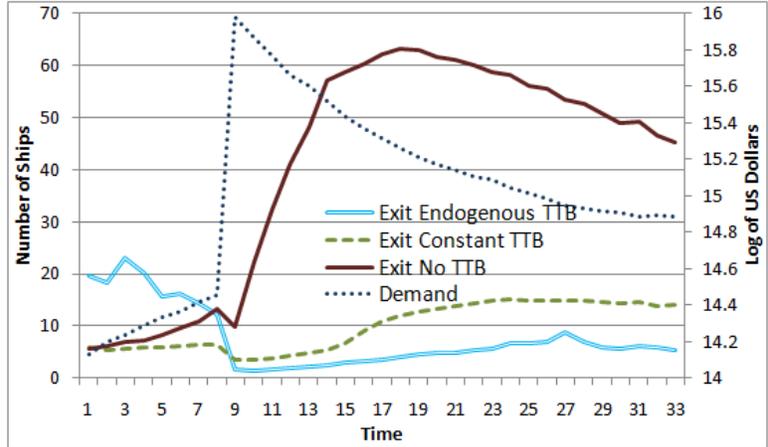


Figure 20: Exit under endogenous, constant and no time to build (TTB). Demand depicted on right axis.

to absent. In addition, the fleet is more volatile in the absence of time to build, as entry and exit respond instantly to demand changes, while this response is larger.

As a result of the larger fleet, shipping prices are 20% lower in the absence of time to build. More voyages are travelled, as we move from endogenous to constant to no time to build. The number of voyages per ship, however, is smaller: when there is time to build, the intensive margin takes over and the fewer operating ships offer more transportation. Price volatility is lower in the absence of time to build by 20%. In addition, shipping prices are less responsive to the demand process, so that the covariance of prices and demand is 15% lower. Price volatility persists in a world without time to build mostly because of the presence of entry costs: under a mean reverting demand process, even in the presence of high demand shocks, entry is limited by irreversible costs of entry. Ships are long-lived capital which depreciates over time; even though in the absence of time to build entrants can immediately take advantage of high demand, they still need to compare the stream of profits during the ship’s lifetime to the cost of entry. Therefore, in the presence of irreversible entry costs, depreciation, demand mean reversion (as well as some demand elasticity), time to build alone cannot account for the total of price volatility.

We also take a first step towards evaluating the impact of shipbuilding subsidies on the shipping industry. Subsidies increase shipbuilding capacity and decrease the construction lag. We ask what would happen if China had not launched its “Long and Medium Term Plan for the Shipbuilding Industry 2006 – 2015”. To do so, we modify the dependence of time to build on the backlog in an effort to capture a decrease in China’s subsidization program. More specifically, we assume that the level of capacity affects the constant

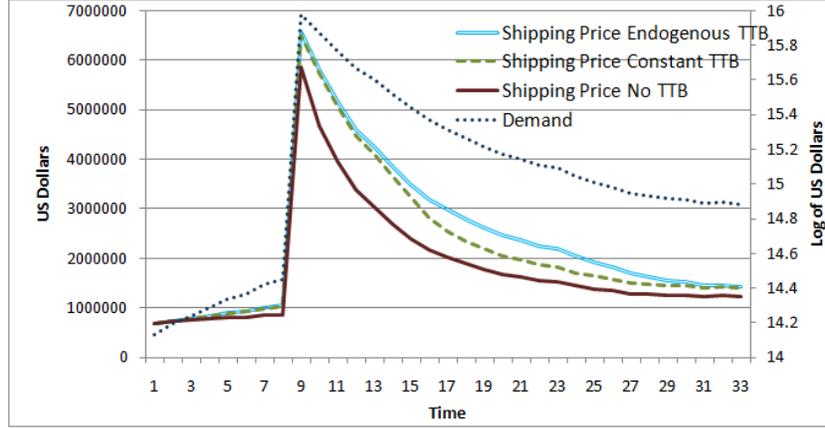


Figure 21: Price under endogenous, constant and no time to build (TTB). Demand depicted on right axis.

term in the time to build function. In particular, we assume that $a_0^{TTB} = \gamma/\bar{c}_t$ where \bar{c}_t is the mean number of operating shipyards in the sample. In this counterfactual, we replace a_0^{TTB} by $a_0^{TTB'}$ where $a_0^{TTB'} = \gamma/\bar{c}_t'$ where \bar{c}_t' is the mean number of non-Chinese operating shipyards and we solve the equilibrium of the game with this time to build function. Results of this counterfactual- which represents an increase in time to build- follow the intuition of the previous experiments. Entry in a world without China's subsidies is 8% lower compared to entry in the baseline model, in response to a positive shock in demand.

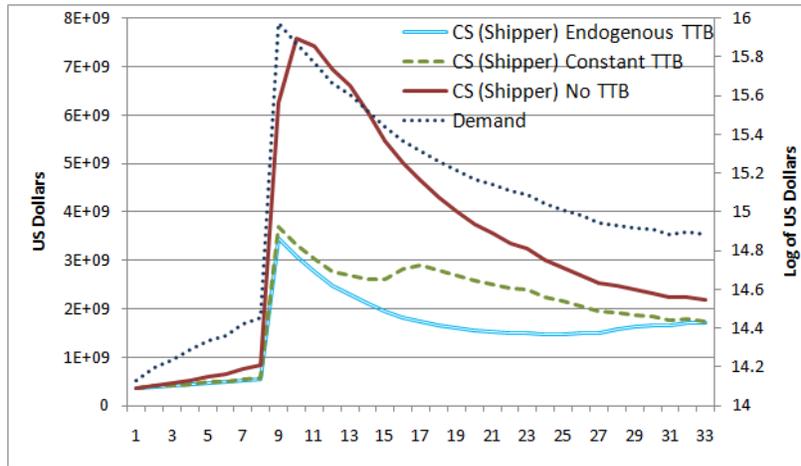


Figure 22: Consumer surplus (i.e. surplus of freight customers) under endogenous, constant and no time to build (TTB). Demand depicted on right axis.

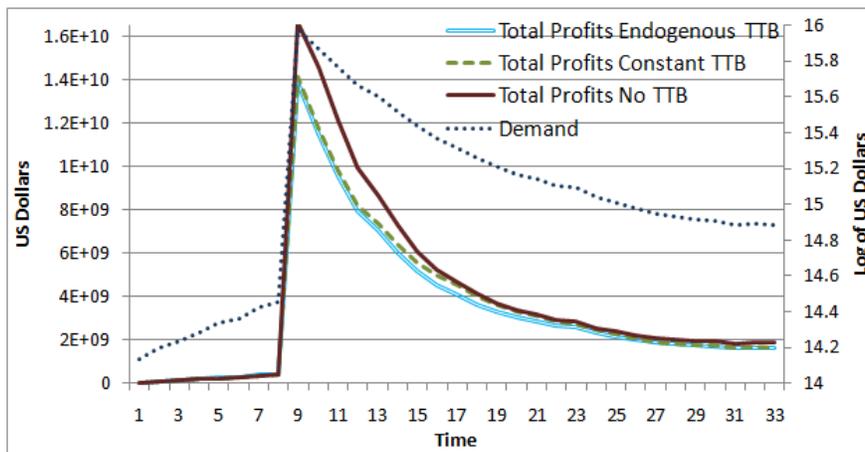


Figure 23: Total (sum of) profits under endogenous, constant and no time to build (TTB). Demand depicted on right axis.

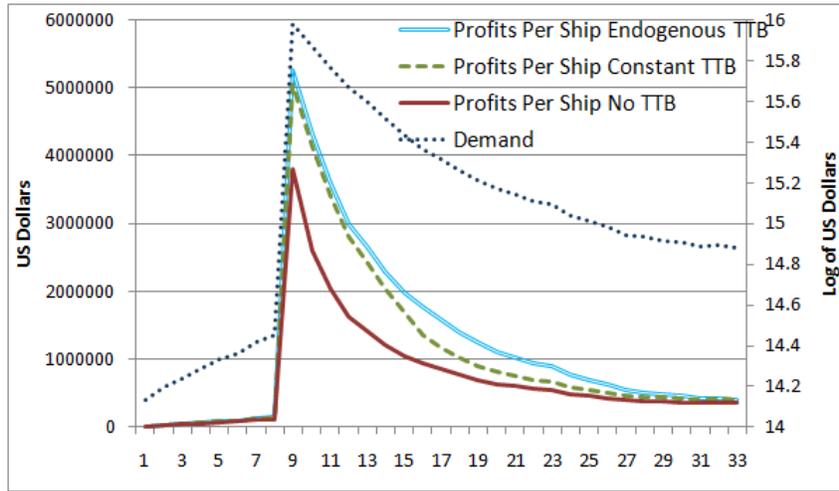


Figure 24: Profits per ship under endogenous, constant and no time to build (TTB). Demand depicted on right axis.

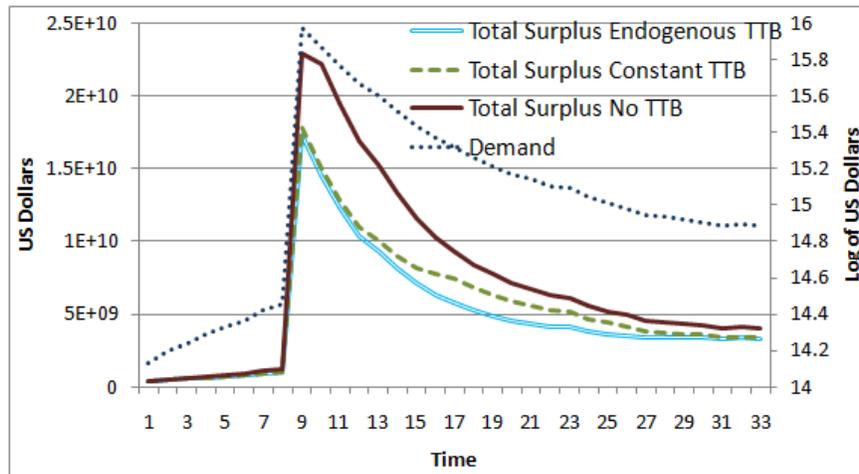


Figure 25: Total surplus of participants in the shipping industry, under endogenous, constant and no time to build (TTB). Demand depicted on right axis.

7 Conclusion

Demand uncertainty and long construction lags in production are observed in multiple industries. This paper argues that these two factors can have a significant impact on firm behavior. Entry and exit decisions subsume different profiles, which in turn affect prices and surplus. Dynamic stochastic games provide an appropriate framework for addressing these issues. The overall analysis requires three steps: estimation of the game primitives, computation of equilibrium policies and counterfactual analysis. In this paper, this methodology is applied to the world bulk shipping industry. A dynamic model of ship entry and exit is estimated nonparametrically and used to quantify the impact of uncertainty and time to build on prices, investment and industry surplus. We are able to estimate a ship's value function directly and use it to recover entry costs, scrap values and per period payoffs. We find that moving from time-varying to constant to no time to build reduces prices, while increasing entry and the surplus of industry participants.

APPENDIX A: DATA

The index of world industrial production used in the IV regression performed to construct the demand state variable d_t , is the weighted average formed by the industrial production index of all OECD countries, as well as Brazil, India, Russia, and South Africa. The indices of individual countries are taken from SourceOECD and the base is Q1 – 1999. The weights are based on 2007 annual GDP levels, taken from Penn World Tables. China doesn’t report a world industrial production index (only a year-on-year percentage change). We therefore add China’s steel production. Commodity prices (for food, minerals and agricultural raw materials) are taken from UNCTAD.

To construct the age distribution of the fleet, s_t , we use quarterly time series of the total fleet, demolitions, deliveries, entrants, losses and other additions and removals (such as conversions from other ship types). These series date back to 1970, while the second hand sales data begin in 1998. Moreover, the demolitions time series provides no information on the age of ships scrapped. In a small sample of individual scrap contracts, no ships younger than 20 years old are observed. We construct the age fleet distribution only up to 20 years of age or older. The last element in s_t represents the absorbing age group of “older than 20 years (80 quarters) old”.

The second-hand sales dataset includes the year of build for each ship sold. Since a quarter is the unit of time, age in years must be converted into age in quarters. We construct the distribution of deliveries across quarters of each year in our sample. Then, based on that distribution, we draw an age in quarters for each ship in the second-hand sale sample, from the 4 possible quarters within its year of built.

To construct the backlog vector, we take the following steps. First, we estimate the time to build function using the total backlog (from our timeseries data) and the average time to build. In particular, we use the fourth dataset which lists the orderbook to identify entrants and compute the average time to build offered. Second, we use the timeseries of entrants and total backlog, as well as the estimated time to build function and create the backlog vector for each quarter of our sample, $b_t \in \mathbb{R}^T$.

All prices are deflated using the CPI taken from the US Bureau of Labor Statistics. The basis used is Q1 – 2009 US dollars.

APPENDIX B: ESTIMATION

In this Appendix, we provide more information on the construction of the state transition matrix. As outlined in Section 3, we use clustering techniques in order to reduce

the number of possible values that the aggregated state $[S_t, B_t, d_t]$ can take. In particular, we use vector quantization, a method that divides a large set of vectors in a set of groups. Each group is represented by its centroid (as in the k-means algorithm). The vector quantization algorithm calculates iteratively the group representatives (the codebook values) until the mean square error between the codebook and the provided data is minimized. We use vector quantization for $[S_t, B_t]$ and scalar quantization (using the Lloyd algorithm which is the same as the vector quantizer for the case of a scalar variable) for d_t . We then take all possible combinations of $[S_t, B_t]$ groups and d_t groups. The vector quantizer guarantees that the observed data is represented well by the groups, while including unobserved combinations of d_t with the remaining aggregate state enriches the state space.

In order to construct the transition matrix, we count transition frequencies between the groups. We simulate thousands of states $x_t = [s_t^0, \dots, s_t^A, b_t^1, \dots, b_t^{\bar{t}}, d_t] \in \mathbb{R}^{A+T+1}$ using the estimated Poisson processes for N_t and Z_t and the $AR(1)$ process for d_t . We then encode the simulated states to their corresponding aggregate state and count frequencies. Finally, we smooth the frequency transition matrix as follows: each element (i, j) of the transition matrix is the inner product of the j -th column and a row of weights reflecting the closeness of the i -th state to the rest of the states. The weights are based on a multivariate normal kernel.

APPENDIX C: COUNTERFACTUALS

In this Appendix we outline the computation of the equilibrium of our model.

For computational reasons we replace the nonparametric scrap value distribution with a normal distribution. Its mean and standard deviation (μ_ϕ, σ_ϕ) are estimated via Maximum Likelihood:

$$\max_{\mu_\phi, \sigma_\phi} \sum_t \log \left[F_\phi (VC_A(x_t); \mu_\phi, \sigma_\phi)^{s_t^A - Z_t} (1 - F_\phi (VC_A(x_t); \mu_\phi, \sigma_\phi))^{Z_t} \right]$$

where the expression above results from the fact that Z_t follows the Binomial distribution. We find $\mu_\phi = -13.53$ and $\sigma_\phi = 8.58$.

In addition, to make sure that the primitives (entry costs, scrap value density and profits) are internally consistent, we use the Cournot profits to compute entry costs, by recovering the value function corresponding to the Cournot profits. These value functions are close to the estimated ones but smoother, as suggested by the fit of the Cournot game.

The goal is to recover the entry policy $\lambda \in \mathbb{R}^L$, the exit policy $\mu \in \mathbb{R}^L$ and value function $V \in \mathbb{R}^{L \times A}$, using the estimated primitives, i.e. the Cournot profits $\pi \in \mathbb{R}^{L \times A}$, the scrap value distribution $N(\mu_\phi, \sigma_\phi)$ and the entry costs $\kappa \in \mathbb{R}^L$ under different counterfactual scenarios. Below we describe the fixed point algorithm used to compute the equilibrium of the game, which is inspired by Weintraub, Benkard and Van Roy (2009). At each iteration n , we use the policies $(\lambda_{n-1}, \mu_{n-1})$ and value functions V_{n-1} from the previous iteration to update to (λ_n, μ_n, V_n) . We repeat the iterations until (λ_n, μ_n) don't change in n as follows:

1. Update exit $\mu^* = s^{80} (1 - F_\phi [P_{\lambda_{n-1}, \mu_{n-1}} V_{80, n-1}])$
2. Compute the new transition matrix P_{λ_{n-1}, μ^*} as described in Appendix B
3. Update the value functions:

$$\begin{aligned} V_{80, n} &= \pi_{20+} + \beta (1 - \mu^*) P_{\lambda_{n-1}, \mu^*} V_{80, n-1} + \beta \mu^* E [\phi | \phi > P_{\lambda_{n-1}, \mu^*} V_{80, n-1}] \\ V_{79, n} &= \pi_{20+} + \beta P_{\lambda_{n-1}, \mu^*} V_{80, n} \\ &\vdots \\ V_{0, n} &= \pi_{0-10} + \beta P_{\lambda_{n-1}, \mu^*} V_{1, n} \end{aligned}$$

4. Update entry $\lambda^* = \lambda_{n-1} \frac{\beta^T (P_{\lambda_{n-1}, \mu^*})^T V_{n-1}}{\kappa}$
5. Set $\lambda_n = \alpha_n \lambda^* + (1 - \alpha_n) \lambda_{n-1}$ and $\mu_n = \gamma_n \mu^* + (1 - \gamma_n) \mu_{n-1}$ where α_n and γ_n go to zero as n increases.

Note that in order to perform step 4 we need to compute $E [\phi | \phi > P_{\lambda^*, \mu^*} V_{80, n-1}]$, via:

$$E [\phi | \phi > A] = \frac{\int_A^\infty \phi f_\phi(\phi) d\phi}{1 - F_\phi(A)}$$

where F_ϕ is the normal distribution.

Once the optimal policies are recovered, we regress them on the state.

APPENDIX D: COURNOT FIT

Figures 26 and 27 show the fit of the Cournot game in terms of the prices and profits. The final profits for the three age groups implied by the Cournot game are shown in Figure 28.

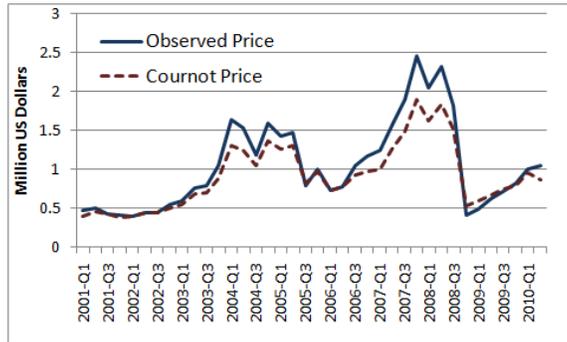


Figure 26: Fit of Cournot price.

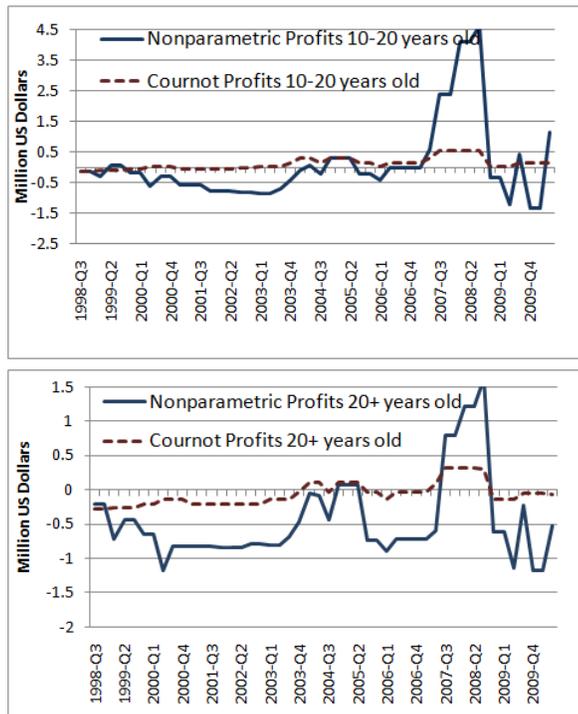


Figure 27: Fit of Cournot profits.

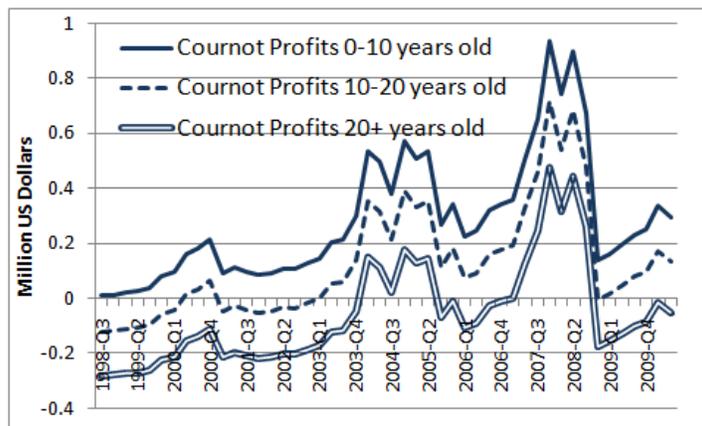


Figure 28: Cournot Profits.

References

- [1] Adland R., and S. Koekebakker (2007), “Ship Valuation Using Cross-Sectional Sales Data: A Multivariate Non-parametric Approach”, *Maritime Economics and Logistics*, 9, 105 – 118.
- [2] Alessandria G., Kaboski J. and V. Midrigan (2009), “Inventories, Lumpy Trade and Large Devaluations”, Mimeo, New York University.
- [3] Bajari P., L. Benkard and J. Levin (2007), “Estimating Dynamic Models of Imperfect Competition”, *Econometrica*, 75 (5), 1331 – 1370.
- [4] Beenstock M., and A. Vergottis (1989), “An Econometric Model of the World Market for Dry Cargo Freight and Shipping”, *Applied Economics*, 21 (3), 339 – 356.
- [5] Benkard L., (2004), “A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft”, *Review of Economic Studies*, 71, 581 – 611.
- [6] Besanko D. and U. Doraszelski (2004), “Capacity Dynamics and Endogenous Asymmetries in Firm Size”, *Rand Journal of Economics*, 35 (1), 23 – 49.
- [7] Bloom N., (2009), “The Impact of Uncertainty Shocks”, *Econometrica*, 77 (3), 623 – 685.
- [8] Caballero R.J. and R.S. Pindyck (1996), “Uncertainty, Investment and Industry Evolution”, *International Economic Review*, 37 (3), 641 – 662.
- [9] Collard-Wexler A. (2010), “Demand Fluctuations in the Ready-Mix Concrete Industry”, Mimeo NYU.
- [10] Dixit A.K. (1989), “Entry and Exit Decisions Under Uncertainty”, *The Journal of Political Economy*, 97 (3), 620 – 638.
- [11] Dixit A.K. and R.S. Pindyck (1994). “Investment Under Uncertainty”, Princeton University Press.
- [12] Doraszelski U. and M. Satterthwaite (2009), “Computable Markov-perfect Industry Dynamics”, *Rand Journal of Economics*, 41 (2), 215 – 243.
- [13] Dunne T., Klimmek S.D., Roberts M.J. and D.Y. Xu (2009), “Entry, Exit and the Determinants of Market Structure”, Mimeo.

- [14] Ericson R. and A. Pakes (1995), “Markov-Perfect Industry Dynamics: A Framework for Empirical Work”, *The Review of Economic Studies*, 62(1), 53 – 82.
- [15] Fan J, and I. Gijbels (1996), “Local Polynomial Modelling and Its Applications”, Chapman & Hall/CRC.
- [16] Gavazza A. (2010), “Leasing and Secondary Markets: Theory and Evidence from Commercial Aircraft”, *The Journal of Political Economy*, forthcoming.
- [17] Gavazza A. (2010), “Asset Liquidity and Financial Contracts: Evidence from Aircraft Leases”, *Journal of Financial Economics*, 95 (1), 62 – 84.
- [18] Gavazza A. (2010), “The Role of Trading Frictions in Real Asset Markets”, *American Economic Review*, forthcoming.
- [19] Hamilton J.D. (2005), “What’s Real About the Business Cycle?”, Federal Reserve Bank of Saint Lewis Review, 87 (4), 435 – 452.
- [20] Jofre-Bonet M. and M. Pesendorfer (2003), “Estimation of a Dynamic Auction Game”, *Econometrica*, 71 (5), 1443 – 1489.
- [21] Koeva P. (2000), “The Facts About Time to Build”, IMF Working Paper.
- [22] Kydland E. and E. Prescott (1982), “Time to Build and Aggregate Fluctuations”, *Econometrica*, 50(6), 1345 – 1370.
- [23] OECD (2008), “The Shipbuilding Industry in China”.
- [24] Pagan A. and A. Ullah (1999), “Nonparametric Econometrics”, Cambridge University Press.
- [25] Pakes A., M. Ostrovsky and S. Berry (2007), “Simple Estimators for the Parameters of Discrete Dynamics Games”, *The RAND Journal of Economics*, 38 (2), 373 – 399.
- [26] Rosen S., K. Murphy and J. Scheinkman (1994), “Cattle Cycles”, *The Journal of Political Economy*, 102 (3), 468 – 492.
- [27] Rust J. (1987), “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher”, *Econometrica*, 55 (5), 999 – 1033.
- [28] Ryan S.P. (2009), “The Costs of Environmental Regulation in a Concentrated Industry”, Mimeo, MIT

- [29] Stopford, M. (1999), “Maritime Economics”, Routledge New York.
- [30] UNCTAD (2009), “Review of Maritime Transport”
- [31] Weintraub G., L. Benkard and B. Van Roy (2008), “Markov Perfect Industry Dynamics with Many Firms”, *Econometrica*, 76 (6), 1375 – 1411.
- [32] Weintraub G., L. Benkard and B. Van Roy (2009), “Computational Methods for Oblivious Equilibrium”, Mimeo, Yale University.