American Economic Review 2014, 104(2): 564–608
http://dx.doi.org/10.1257/aer.104.2.564

Time to Build and Fluctuations in Bulk Shipping†

By Myrto Kalouptsidi*

This paper explores the nature of fluctuations in world bulk shipping by quantifying the impact of time to build and demand uncertainty on investment and prices. We examine the impact of both construction lags and their lengthening in periods of high investment activity, by constructing a dynamic model of ship entry and exit. A rich dataset of secondhand ship sales allows for a new estimation strategy: resale prices provide direct information on value functions and allow their nonparametric estimation. We find that moving from time-varying to constant to no time to build reduces prices, while significantly increasing both the level and volatility of investment. (*JEL* G31, L11, L62, L92)

Adjustment costs and irreversibilities figure prominently in both the modern theory of firm-level investment under uncertainty as well as in aggregate business cycle models.[1] Yet there is little empirical evidence of their quantitative importance in specific settings. It is the goal of this paper to fill this gap by focusing on a particular industry, that of oceanic bulk shipping. This industry lends itself to this analysis,[2] as it provides an outstanding example of the interaction between uncertainty and adjustment costs, which take the form of time to build: shipping firms—conducting 70 percent of world seaborne trade (in tons)—face long lags between the order and delivery of a new vessel, while the uncertain demand for sea transport may substantially alter conditions during this wait. For instance, the recent growth of raw material imports, particularly in China, led to sustained increases in freight rates and a sevenfold surge in the new ship backlog between 2003 and 2008. The crisis of 2008 led to an idling of the existing fleet, at the same time that another 70 percent of that fleet was still scheduled for delivery by 2012.

To study this issue we construct a dynamic model of ship entry and exit. We structurally estimate this model employing a novel estimation strategy based on the observation that resale prices reflect value functions. In counterfactual experiments

* Princeton University, Department of Economics, 315 Fisher Hall, Princeton, NJ 08544-1021 (e-mail: myrto@princeton.edu). I am extremely thankful to my advisors, Steve Berry, Phil Haile, and Lanier Benkard for their numerous useful comments and suggestions. I am also grateful to Jan De Loecker and Richard Rogerson for many helpful remarks. I have gained a lot from discussions and comments from Costas Arkolakis, Nick Bloom, Chris Conlon, Eduardo Engel, Kirill Evdokimov, Manolis Galenianos, Rick Mansfield, and Theodore Papageorgiou. Finally, I'd like to thank the anonymous referees.

† Go to http://dx.doi.org/10.1257/aer.104.2.564 to visit the article page for additional materials and author disclosure statement(s).

[1] In the context of firm-level investment, see, for example, Abel (1983); Pindyck (1991); Dixit (1992); Caballero (1999). In the context of business cycle models, see, for example, Kydland and Prescott (1982) who emphasized time to build as a key element in matching the properties of aggregate investment, or the recent business cycle models such as Christiano, Eichenbaum, and Evans (2005); Smets and Wouters (2007).

[2] Since the very beginnings of mathematical economics, Tinbergen (1931) viewed shipping cycles as interesting illustrations of business cycles more generally.

we examine the dynamic evolution of the industry, as well the industry's response to shocks and find that time to build has a substantial impact on the level and volatility of investment; in particular, in the absence of time to build investment in new ships is higher, but also significantly more volatile.

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. Our time to build generalizes the conventional fixed construction lag (e.g., Kydland and Prescott 1982) by exhibiting cyclical variation, as delivery lags are lengthened during periods of high investment activity due to capacity constraints and queueing at shipyards. 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 next estimate the primitives of this dynamic model. Our estimation strategy is based on a rich dataset of secondhand ship sales which provides direct information on key dynamic objects of the model and allows for their nonparametric estimation. In particular, we treat ships' resale prices as observations of their value functions, under the assumption of homogeneous firms and a large number of potential entrants. Our methodology is a "mirror image" of previous work on the estimation of dynamic games (e.g., Bajari, Benkard, and Levin 2007; Pakes, Ostrovsky, and Berry 2007), where estimated per period payoffs and policy functions are used to recover value functions, which in turn reveal exit and entry costs. In this article, we estimate value functions directly from the data and use them to retrieve nonparametrically per period payoffs, as well as entry and exit costs. Data limitations, as well as the complicated operational environment of ships, do not allow the use of the estimation procedures suggested by the aforementioned work. Furthermore, our estimation procedure allows for a rich set of competition modes in the industry spot market without altering the estimation burden.[3] Finally, our methodology can be of use in other settings with homogeneous agents. In particular, we show how observed transactions of a good, whose valuation is the unknown solution of an underlying dynamic framework, can be used to uncover the dynamic parameters of interest, thereby reducing the computational and economic restrictions.

Our estimation results suggest that profits exhibit patterns of hysteresis: shipowners delay exit decisions and suffer losses in anticipation of better times, consistent with anecdotal evidence. Hysteresis results from investment irreversibility and time to build. On one hand, the exit decision is irreversible, as scrapped ships can't be "unscrapped," while in addition scrap values are considerably lower than entry costs (approximately 15 percent on average, but ranging from 6 percent to 32 percent). On the other hand, entry is slow, making good times even more profitable for existing firms. Hysteresis can also lead to chronic excess capacity (Pindyck 1991), a common feature of bulk shipping.[4]

As a test of the estimated model, we compare the retrieved primitives against two time series not included in the estimation. In particular, we compare the estimated

[3] In our case, postestimation we are able to detect that firms act as price takers.

[4] Indeed, during the crisis of 2008 the press featured several pictures of armies of idle ships waiting for cargo.

profits to an earnings index calculated by shipbrokers, and the estimated entry costs to the average shipbuilding price of a new ship. We find that they match surprisingly well and interpret this as support for the model and empirical strategy.[5] We then proceed to quantifying the impact of both pure construction lags, as well as the endogenously variable time to build on investment, shipping prices and surplus. Exploring the comparative statics of time to build has the descriptive value of measuring the magnitude of investment adjustment costs per se, as well as their variability brought by industry evolution. In addition, time to build can indeed change via governmental subsidies to the shipbuilding industry, which are prevalent in East Asia. We next provide a summary of our counterfactuals' findings.

We begin by examining the industry's response to a positive demand shock, when time to build becomes a critical constraint. We find that moving from time-varying, to constant, to no time to build reduces prices, while increasing the fleet and shipper (consumer) surplus. In particular, following a positive demand shock, in a world with time to build the supply response is limited, as entry is both slower and lower, while existing supply is limited to the 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 temporarily increased demand for shipping services. As a result, time to build leads to high profits and prices and low shipper (consumer) surplus.

We next explore the dynamic evolution of the industry via long-run simulations of the model. We find that time to build has a dramatic impact on the volatility of investment in terms of both flow (entry) and stock (fleet): the fleet is 45 percent more volatile under constant time to build and twice more volatile under no time to build, while entry is twice more volatile under constant time to build and seven times more volatile in the absence of time to build. Consistent with hysteresis, higher adjustment costs render firms less likely to respond to demand shocks, leading to a smoother investment process. At first this may seem to contradict Kydland and Prescott (1982) whose work associated time to build with volatility. But Kydland and Prescott (1982) do not assess the sign of this relationship; rather they document that time to build is crucial in matching the properties of investment. In addition, time to build has a sizeable impact on the level of investment: in the absence of time to build the fleet is about 15 percent larger. Finally, prices are lower and less volatile as time to build declines, by 14 percent in the case of no time to build.

Finally, we explore the oscillatory and steady-state behavior of investment. We find that time to build leads to investment oscillations that fade extremely slowly. The industry exhibits strong "echo effects" (a periodic-like fleet movement that follows after an entry impulse) that vanish considerably more slowly under time to build. This pattern is consistent with the higher responsiveness of the fleet to demand changes, as well as the higher turnover of ships in the no time to build case, both of which force any "echo effects" to vanish relatively fast.

This article contributes to two strands of literature. First, a voluminous literature has studied the impact of irreversibility or adjustment costs and uncertainty on firm investment decisions (e.g., Abel 1983; Caballero 1991; Pindyck 1991; Dixit 1992;

[5] A series of robustness exercises also reveals that results are surprisingly robust.

Pindyck 1993a, b; Abel and Eberly 1994; Caballero and Pindyck 1996; Caballero and Engel 1999). Time to build was underlined by Kydland and Prescott (1982),[6] while more recent business cycle models often include lags and adjustment costs (see references above). The second strand of literature we contribute to is that of industry dynamics, consisting of a long list of theoretical (e.g., Lambson 1991; Hopenhayn 1992; Ericson and Pakes 1995), as well as empirical (e.g., Benkard 2004; Aguirregabiria and Mira 2007; Bajari, Benkard, and Levin 2007; Collard-Wexler 2013; Pakes, Ostrovsky, and Berry 2007; Pesendorfer and Schmidt-Dengler 2008; Xu 2008; Ryan 2012) work. The present article focuses on industry dynamics in an environment of investment irreversibilities, time to build and demand uncertainty, which alter firm entry and exit behavior.

The remainder of the article is organized as follows: Section I provides a description of the industry. Section II presents the model. Section III describes the data used. Section IV presents the empirical strategy employed and the estimation results. Section V provides the counterfactual experiments, and Section VI concludes.

I. Description of the Industry

Bulk shipping concerns vessels designed to carry a homogeneous unpacked dry or liquid cargo, for individual shippers on nonscheduled routes.[7] 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. Bulk ships 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, and wood chips. 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 can carry different products, take different routes, and approach different ports. Practitioners treat them as different markets. Each such market consists of a large number of small shipowning firms. For instance, Figure 1 shows the firm size distribution of Handysize bulk carriers (where the empirical analysis will focus): note that the largest fleet share is 3 percent, while there is a long tail of small shipowners. Even though there is some scope for differentiation (based on the age of the ship, the shipyard where it was built, the reputation of the shipowner, etc.) shipping services are largely perceived as homogeneous.

Demand for shipping services is driven by world seaborne trade and is, thus, subject to world economy fluctuations. In recent years, the growth and infrastructure building at several countries led to increased imports of raw materials, significantly boosting demand for bulk transport. For instance, Chinese imports grew by 15 percent every year between 2004 and 2007 (UNCTAD 2009), while Asian imports

[6] Bar-Ilan and Strange (1996) and Majd and Pindyck (1987) consider time to build within the setup of the investment under uncertainty literature. A key conceptual difference from our setup is that investment can be abandoned during the lead time. Empirical work on uncertainty and investment includes Bloom, Bond, and Van Reenen (2007) and Bloom (2009), while Rosen, Murphy, and Scheinkman (1994) explore time to build and periodicity in breeding cattle.

[7] This section draws from Stopford (1997), as well as Beenstock and Vergottis (1989).

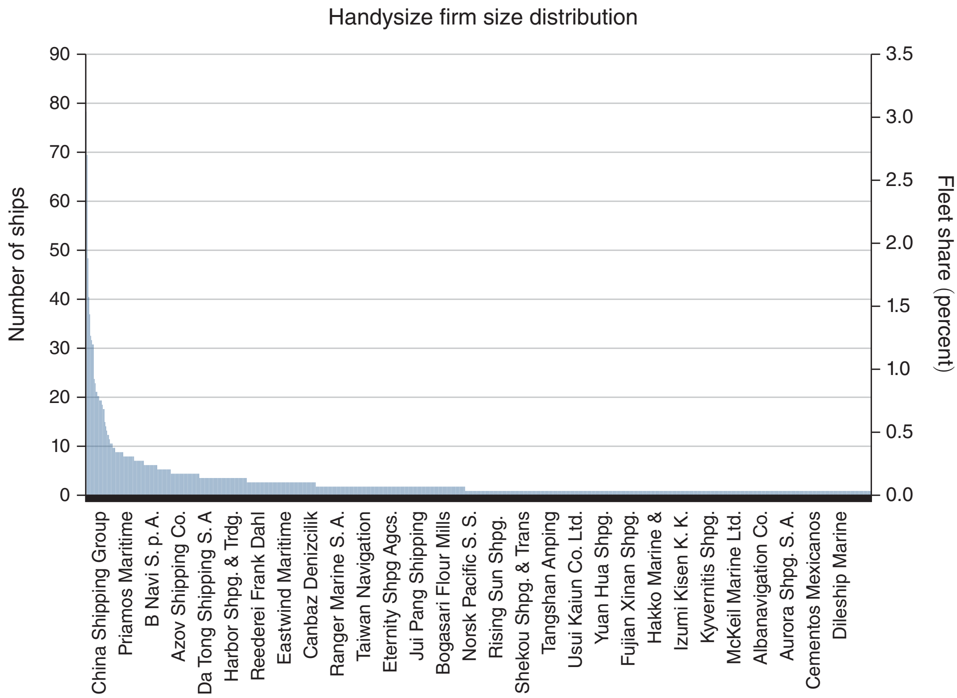

Figure 1. Fleet and Fleet Shares of Shipowning Firms

represent more than 50 percent of world imports; Figure 2 shows China's imports of iron ore and coal.

In the short run, the supply of shipping services is determined by the number of voyages carried out by shipowners. As Stopford (1997) notes, "even though the physical capacity of ships is fixed at a point in time, the available transportation capacity is flexible," as shipowners can adjust the ton-miles they offer by adjusting their speed of sail. Even so, short-run supply is rather inelastic as voyage costs (which include fuel, port/canal dues and cargo handling) are convex in speed. They also increase with the ship's age, as its fuel efficiency and overall operation deteriorates over time. The fluctuating shipping demand combined with the inelastic supply leads to volatile shipping prices, shown in Figure 3.

In the long run, the supply of cargo transportation 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 scrapyard, where it is dismantled and its steel hull is recycled. 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 increase the time to build. For instance, the average time to build almost doubled between 2001 and 2008 (documented in Figure 8). Finally, it is worth noting that canceling a shipyard order

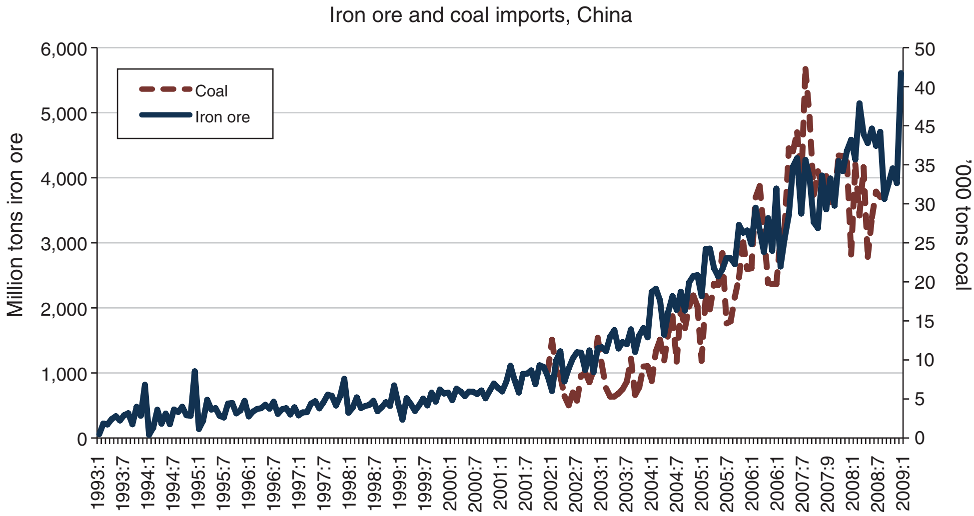

FIGURE 2. CHINA'S IRON ORE AND COAL IMPORTS

FIGURE 3. THE BALTIC DRY INDEX

Notes: Daily index based on weighted average of rates on 20 representative bulk routes. Compiled by the Baltic Exchange. 1/11/1999 = 1,334.

is difficult; even during the 2008–2009 crisis, only about 10 percent of the entire orderbook (in all ship types) was canceled (UNCTAD 2009).

II. Model

In this section, we present a dynamic model of the world bulk shipping industry, which lies within the general class of dynamic games studied in Ericson and Pakes (1995), as modified by Doraszelski and Satterthwaite (2010).[8]

[8] A large number of models are built on this framework, such as Benkard (2004) and Besanko and Doraszelski (2004).

A. *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 is a reasonable assumption for this industry, while the firm size distribution shown in Figure 1 suggests that many 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; in other words, a firm begins its operation a number of periods after its entry decision. The state variable of a ship consists of:

(i) its age, $j \in \{0, \ldots, A\}$

(ii) the age distribution of the fleet, $\mathbf{s}_t \in S \subset \mathbb{R}^{A+1}$, where $\mathbf{s}_t = \left[s_t^0\ s_t^1 \ldots s_t^A\right]$ and s_t^i is the number of ships of age i

(iii) the backlog $\mathbf{b}_t \in B \subset \mathbb{R}^{\bar{T}}$, where $\mathbf{b}_t = \left[b_t^1\ b_t^2 \ldots b_t^{\bar{T}}\right]$ and b_t^i, $i = 1, \ldots, \bar{T}$, is the number of ships scheduled to be delivered at period $t + i$

(iv) the aggregate demand for shipping services $d_t \in D \subset \mathbb{R}$, capturing shifts in the inverse demand curve for freight transport.

Firms face the inverse demand curve,

$$P_t = P(d_t, Q_t), \tag{1}$$

where P_t is the price per voyage and Q_t is the total number of voyages offered in period t. The demand state variable, d_t, follows an exogenous first-order Markov process. Per period payoffs of a ship of age j are given by $\pi_j(\mathbf{s}_t, d_t)$, which is bounded and nonincreasing in j for all $(\mathbf{s}_t, d_t)$. $\pi_j(\mathbf{s}_t, d_t)$ is a reduced form reflecting the equilibrium of the industry spot market and results from the combination of the demand curve (1) and the nature of competition (e.g., perfect competition or Cournot). We do not specify the nature of competition, allowing for flexibility, but as in Ericson and Pakes (1995) and Weintraub, Benkard, and Van Roy (2008) we require that it generates a complete preorder $\succeq$ over s, which defines the strength of competition, so that profits are nonincreasing in the sense of $\succeq$, for all (j, d_t). Finally, note that the backlog $\mathbf{b}_t$ does not affect per period payoffs.

Every period, incumbents privately observe a scrap value, ϕ, drawn i.i.d. across time and firms from a distribution F_ϕ which has a well defined inverse function.[9] They decide whether to exit the market and obtain ϕ or continue operating. In practice, part of the scrap value of a ship consists of the steel value of its hull, which might be expected to be common across firms.[10] Stopford (1997), however, notes: "the (scrap) price also varies from ship to ship, depending on its suitability for

[9] The introduction of private information over exit is needed to establish existence of equilibrium, as shown in Doraszelski and Satterthwaite (2010).

[10] Data limitations and computational constraints make difficult the inclusion of steel price as a state in the empirical exercise.

scrapping" and includes the interiors of a ship (engines, furniture, etc.) which are sold separately. These are likely to be private and may be changing in an i.i.d. fashion over time depending on demand for these parts of the ship.

There is free entry into the market, subject to an entry cost $\kappa(\mathbf{s}_t, \mathbf{b}_t, d_t)$. All entrants of period t receive the same time to build, T_t, which is a deterministic function of the industry state, $T : (\mathcal{S} \times \mathcal{B} \times \mathcal{D}) \rightarrow [1, \overline{T}]$, so that

$$T_t = T(\mathbf{s}_t, \mathbf{b}_t, d_t)\,. \tag{2}$$

Time to build changes over time, capturing shipbuilding capacity constraints which add waiting time to construction time when new-building orders increase. Specifying time to build as a function of the entire state (rather than, say, only $\mathbf{b}_t$) captures queues due to period t's entry, which is a function of the industry state. In addition, allowing for both time to build and entry costs to depend on the industry state may accommodate for changes in shipbuilding capacity that occur in response to it.[11] Finally, it is worth noting this is the first empirical paper to our knowledge to allow for non-i.i.d., state-dependent entry costs.

The timing in each period is as follows: incumbents and potential entrants first observe the state of the industry $(\mathbf{s}_t, \mathbf{b}_t, d_t)$. Incumbents then privately observe their scrap value and decide whether they will exit the market or continue operating, while simultaneously potential entrants make their entry decisions. Incumbents receive their operating profits, $\pi_j(\mathbf{s}_t, d_t)$. Finally, the entry and exit decisions are implemented and the state evolves.

To sum up, this is a dynamic model where firms enter, age, and endogenously exit. In practice, shipowners also have the option of selling their ship in the market for used ships. In this model, since all shipowners share the value of a ship, the price of a ship in the secondhand 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: they are neither entry—the ship was already in operation—nor exit—the ship will continue operating. At a sale the ship just changes ownership, and shipowners are always indifferent between selling a ship and operating it themselves. We revisit sales in the empirical part of the article, where we treat them as observations on the value function.[12]

B. *Firm Behavior*

An incumbent firm of age j maximizes its expected discounted stream of profits, deciding only on exit. Let N_t be the number of entrants in period t and Z_t^j the number of firms of age j that exit. The incumbent's value function before privately observing the scrap value, ϕ, drawn from F_ϕ, is

[11] Ideally time to build and entry costs should also be a function of shipbuilding capacity, which would then become a state variable. Even though it is straightforward to accommodate it in the model, introducing it in the empirical application is difficult, due to lack of detailed data and the increase of an already large state space.

[12] We can think of sales as occurring due to some exogenous shock; in particular, we can add a step in the timing of each period after incumbents and potential entrants observe the state of the industry, in which incumbents receive an exogenous shock with some probability, which forces them to sell the ship in the secondhand market. This shock is independent of the value of the ship (a liquidity or retirement shock). After incumbents observe the state and trade in the secondhand market, they privately observe their scrap value, and the remainder of the timing is left unchanged. Finally, the value functions defined in the following section would not change.

$$V_j(\mathbf{s}_t, \mathbf{b}_t, d_t) = \pi_j(\mathbf{s}_t, d_t) + \beta E_\phi \max\{\phi, VC_j(\mathbf{s}_t, \mathbf{b}_t, d_t)\}. \tag{3}$$

where its continuation value is given by

$$VC_j(\mathbf{s}_t, \mathbf{b}_t, d_t) \equiv E[V_{j+1}(\mathbf{s}_{t+1}, \mathbf{b}_{t+1}, d_{t+1}) \,|\, \mathbf{s}_t, \mathbf{b}_t, d_t]. \tag{4}$$

The expectation in (4) is over N_t and Z_t^j for all j, as well as the demand for shipping services, d_{t+1}. Randomness in this model results from these variables, as the remainder of the state evolves deterministically conditional on the current state. In particular, the transition of the fleet age distribution, $\mathbf{s}_t$, is as follows:

$$s_{t+1}^0 = b_t^1$$

$$s_{t+1}^j = s_t^{j-1} - Z_t^j, \qquad j = 1,\ldots,A.$$

The number of age zero ships in period $t+1$ equals the first element of the backlog. The remaining elements of $\mathbf{s}_t$ 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:

$$b_{t+1}^i = b_t^{i+1}, \qquad i \neq T_t, \overline{T}$$

$$b_{t+1}^{T_t} = b_t^{T_t+1} + N_t, b_{t+1}^{\overline{T}} = 0, \qquad \text{if } T_t < \overline{T}$$

$$b_{t+1}^{\overline{T}} = b_t^{\overline{T}} + N_t, \qquad \text{if } T_t = \overline{T}.$$

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 $T_t = T(\mathbf{s}_t, \mathbf{b}_t, d_t)$. If this position is not the maximum time to build $\overline{T}$, zero ships are added at the end of the backlog. 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, $VC_j(\mathbf{s}_t, \mathbf{b}_t, d_t)$. This event occurs with probability

$$\zeta_j(\mathbf{s}_t, \mathbf{b}_t, d_t) \equiv \Pr(\phi > VC_j(\mathbf{s}_t, \mathbf{b}_t, d_t)) = 1 - F_\phi(VC_j(\mathbf{s}_t, \mathbf{b}_t, d_t)). \tag{5}$$

The number of firms that exit at age j, Z_t^j, thus follows a binomial distribution.

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

$$VE(\mathbf{s}_t, \mathbf{b}_t, d_t) \equiv \beta^{T_t} E[V_0(\mathbf{s}_{t+T_t}, \mathbf{b}_{t+T_t}, d_{t+T_t}) \,|\, \mathbf{s}_t, \mathbf{b}_t, d_t].$$

Note that due to time to build, a firm that decides to enter in period t will begin operating $T_t = T(\mathbf{s}_t, \mathbf{b}_t, d_t)$ 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, $\{N_t, N_{t+1}, \ldots, N_{t+T_t-1}\}$, the number of firms that will exit $\{Z_t^j, Z_{t+1}^j, \ldots, Z_{t+T_t-1}^j\}$, all j, as well as demand d_{t+T_t}. Finally, the ship is age zero upon delivery.

We assume that a large pool of potential entrants play a symmetric mixed entry strategy, so that the number of actual entrants, N_t, is well approximated by the Poisson distribution with a state-dependent mean, $\lambda(\mathbf{s}_t, \mathbf{b}_t, d_t)$.[13] This is endogenously determined by satisfying a zero expected discounted profits condition.

We end this section by defining the model's equilibrium. The equilibrium concept is that of a Markov Perfect Equilibrium (MPE) in the sense of Maskin and Tirole (1988a, b). We also assume that equilibrium is symmetric, such that all firms use a common stationary exit strategy that takes the form of a cut-off rule: the pure exit strategy prescribes exit if $\phi > VC_j(\mathbf{s}_t, \mathbf{b}_t, d_t)$. As discussed in Doraszelski and Satterthwaite (2010), this strategy is equivalently described by the exit probability $\zeta : \{0, \ldots, A\} \times (\mathcal{S} \times \mathcal{B} \times \mathcal{D}) \rightarrow [0, 1]$. Similarly, define the entry rate $\lambda : (\mathcal{S} \times \mathcal{B} \times \mathcal{D}) \rightarrow [0, \overline{\lambda}]$. We rewrite the incumbent's value function (3) to demonstrate its dependence on strategies; let $\mathbf{x} = (\mathbf{s}_t, \mathbf{b}_t, d_t)$ and $V_j(\mathbf{x}; \zeta', \zeta, \lambda)$ be the expected net present value for a firm at state $(j, \mathbf{x})$ given that its competitors each follow a common strategy ζ, the entry rate function is λ, and the firm itself follows strategy ζ':

$$V_j(\mathbf{x}; \zeta', \zeta, \lambda) = \pi_j(\mathbf{s}_t, d_t) + \zeta'\beta E\big(\phi \,|\, \phi > VC_j(\mathbf{x}; \zeta, \lambda)\big)$$
$$+ (1 - \zeta')\,\beta VC_j(\mathbf{x}; \zeta, \lambda),$$

where the continuation value $VC_j(\cdot)$ assumes that the firm does not exit and depends on the entry rate and other firms' exit strategies ζ through state transitions.

An equilibrium consists of an exit strategy and entry rate function that satisfy the following two conditions:

(i) Incumbent firm strategies represent a MPE:

$$\sup_{\zeta_j'(\mathbf{x}) \in [0,1]} V_j(\mathbf{x}; \zeta', \zeta, \lambda) = V_j(\mathbf{x}; \zeta, \zeta, \lambda),$$
$$\text{all } j \in \{0, 1, \ldots, A\}, \mathbf{x} \in (\mathcal{S} \times \mathcal{B} \times \mathcal{D}).$$

(ii) At each state, either entrants have zero expected discounted profits or the entry rate is zero (or both):

$$VE(\mathbf{x}) \leq \kappa(\mathbf{x}), \qquad \text{all } \mathbf{x} \in (\mathcal{S} \times \mathcal{B} \times \mathcal{D}) \tag{6}$$

with equality if $\lambda(\mathbf{x}) > 0$.

Under the assumptions made in this section regarding profits, the scrap value density, and finiteness of the state space (i.e., S, B, D finite), a symmetric equilibrium exists as proven by Doraszelski and Satterthwaite (2010) (continuity of the state transition in strategies is also required and is satisfied due to (10) and (11)). Our setup differs from the general model of Doraszelski and Satterthwaite (2010) in

[13] For a proof, see Weintraub, Benkard, and Van Roy (2008).

Table 1—Secondhand Sales Summary Statistics

	Mean	SD
Price in million US dollars	9.97	8.54
Age in years	18.7	7.6
Size in DWT	27,876	6,718
Sales in a quarter	41.85	15.3
Observations	1,838	

three respects: we do not have firm investment, we allow for an aggregate state d_t, and we adopt the entry process of Weintraub, Benkard, and Van Roy (2008) rather than private i.i.d. entry costs; the same steps of Doraszelski and Satterthwaite (2010) can be followed to show existence in this case as well.

The model presented in this section presumes that firms behave strategically; yet Figure 1 illustrates that the industry consists of a large number of small firms. Our model nests perfect competition as a special case: for instance, if firms are Cournot competitors in the freight market as the number of firms grows large, the market becomes approximately perfectly competitive. Our empirical strategy does not require us to take a stand on the nature of competition in the freight market.[14] Indeed, we are able to recover profits nonparametrically and then examine what equilibrium assumption fits best (see Section VA).

III. Data

The data employed in this article are taken from Clarksons Research, a shipbroking firm based in the United Kingdom. Four different datasets are utilized. The first consists of the world secondhand 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 secondhand sales. We are left with 1,838 observations for 48 quarters. Table 1 provides summary statistics of this dataset. Prices exhibit high variance unconditional on the state (in contrast to the variance conditional on the state). Moreover, old ships are sold more often.

The second dataset consists of shipping voyage contracts and includes the date of transaction, the name and size of the ship, and the ship's price per trip. The dataset includes a subset of all voyage contracts between January 2001 and June 2010. There are two important 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 where

[14] The empirical strategy of this paper is based on the assumption of homogeneous agents and a large number of potential shipowners but no further assumptions regarding the nature of competition. In fact, nothing would change in the empirical exercise if we had defined a perfectly competitive equilibrium rather than an MPE.

the full sample of contracts is observed. Second and most important, the number of per period contracts only partly captures transportation offered by Handysize vessels, as trips vary in both time and distance—indeed, the ideal measure of quantity would be in ton-miles.

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, $\mathbf{s}_t \in \mathbb{R}^{A+1}$. A limitation of these series is that they date back to 1970, while the secondhand 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 up to 20 years of age or older, so that the last element in $\mathbf{s}_t$, s_t^A, represents the absorbing age group of "older than 20 years (80 quarters) old." We also 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 $\mathbf{s}_t = \left[s_t^0, s_t^1, \ldots, s_t^A\right] \in S \subset \mathbb{R}^{A+1}$, where $A = 80$ and its transition is

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

Figure 4 shows the number of ships that enter and exit, as well as the total fleet. The entrant series provides another testament of the large investment swings in bulk shipping: note the spike in new ships orders during 2006–2008. Note also 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.

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 and delivery date of each ship under construction. We use this dataset to construct the backlog state variable, $\mathbf{b}_t$ (see online Appendix), as well as estimate the time to build function.

IV. 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., Aguirregabiria and Mira 2007; Bajari, Benkard, and Levin 2007; Pakes, Ostrovsky, and Berry 2007; Pesendorfer and Schmidt-Dengler 2008; and Jofre-Bonet and Pesendorfer 2003)[15] have proposed estimation procedures for dynamic games in which choice probabilities (and in some of these methods also per period payoffs) are used to compute continuation values, which in turn allow the estimation of entry and exit costs. Data limitations, as well as the complicated operational environment of ships, do not allow the use of these estimation procedures. In this article, we apply a new estimation strategy which uses the secondhand ship sale transactions to recover value functions

[15] Applications of these methodologies have occurred in several contexts, e.g., Snider (2008); Xu (2008); Ryan (2013); Dunne et al. (2009).

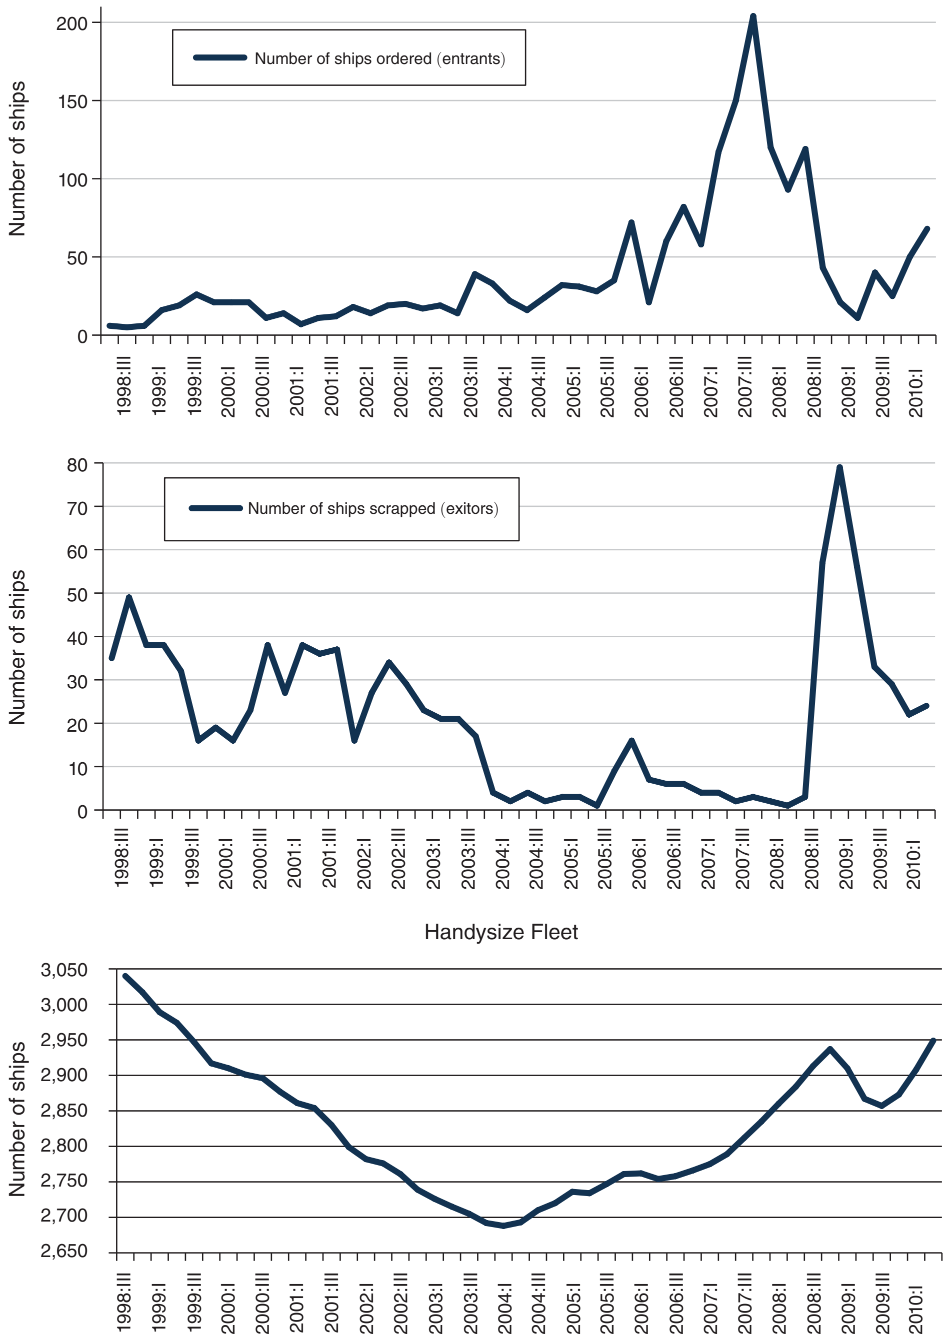

Figure 4. Observed Entry, Exit, and Fleet

nonparametrically directly from the data. In particular, since in our model agents are homogeneous and there is a large number of potential firms, a secondhand transaction price must equal the value of the ship. A ship sale in the secondhand market is "outside" of the dynamic model: owners are always indifferent between selling the ship and operating it themselves. As a result, under these assumptions, resale prices can be treated as observations of the value function. 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. Our methodology can be of use in other settings with homogeneous agents, as we show how observed transactions of a good, whose valuation is the unknown solution of an underlying dynamic framework, can be used to uncover the dynamic parameters of interest, thereby reducing the computational and economic restrictions.

The empirical strategy consists of two steps. In the first step we estimate the equilibrium objects necessary to obtain the model primitives, i.e., the value functions and the transition matrix. In this step we also estimate some exogenous objects that are not determined by the dynamic equilibrium: the demand curve for shipping given in equation (1), the transition of the demand state variable d_t, and the time to build function given in (2).

In the second step, we estimate the model primitives, i.e., the profits, the scrap value density, and the entry costs. Profits are calculated from the Bellman equations (3). The scrap value distribution is estimated nonparametrically using the estimated continuation values and observed exit behavior. Finally, the entry costs are obtained from the free entry condition.

Time-series variation is key in our empirical strategy (the global nature of this industry does not provide cross-sectional variation). In fact, an appealing feature of this industry is the exogenous, strong variability of demand due to international trade and business cycle movements, which helps identify supply parameters of interest. That said, the relatively short time series available (12 years) creates challenges, for instance in the creation of the demand state variable, its transition (see Section IVA), and the entry and exit policy functions (see Section IVA). It is important, however, that we observe both a boom and a bust in demand in our sample period.

A. *First Step: Estimate Exogenous and Equilibrium Objects*

We begin with the estimation of the inverse demand curve for shipping services. We then proceed to estimate the value functions nonparametrically on a discretized state space. The transition rule for d_t, the time to build function, and the estimated entry and exit policy functions are combined to create the transition matrix. Each estimation task is described below and followed by the results.

Demand for Shipping Services.—We estimate the inverse demand for shipping services via instrumental variables regression. Our short time series does not allow for a nonparametric demand curve, and we choose the fairly standard assumption of log-linearity. In particular, we assume that the demand curve in (1) takes the form

$$P_t = D_t Q_t^{\eta},$$

where $d_t = \log D_t$. The empirical analog of this demand curve is

$$\log (P_t) = \eta_0 + \eta_1 \log(\mathbf{H}_t) + \eta \log(Q_t) + \varepsilon_t, \tag{7}$$

where P_t is the average price per voyage in a quarter, $\mathbf{H}_t$ includes demand shifters, while Q_t is the total number of voyage contracts realized in a quarter. $\mathbf{H}_t$ includes the index of world industrial production (taken from the OECD), China's steel production, the fleet of Handymax ships, the index of food prices, agricultural raw material prices, and minerals prices (taken from UNCTAD).

In the first-stage regression, we project Q_t on the demand shifters and instruments. The instruments we admit are the following supply shifters: the total fleet, as well as the mean and the standard deviation of the fleet age distribution. The fleet is a key determinant of industry supply capacity, considering the convexity of ship operating costs. The age of the fleet determines industry supply, as younger ships face lower costs and can, thus, perform more voyages. These instruments correct both for endogeneity and for measurement error (as mentioned above, we observe only the number of trips realized, rather than ton-miles).

Table 5 in Appendix A reports the first- and second-stage results of the demand curve estimation. In the second stage, the index of world industrial production (WIP) and China's steel production positively affect 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. The impact of all shifters $\mathbf{H}_t$, though, will be lumped into the state variable d_t as described below. Measurement error, functional form, as well as the quality of the instruments are likely to influence the elasticity, which is estimated at -6.17 (implying a change of \$10,000 in price leads to a change in quantity of about 60 trips). Measurement error in the number of trips is inherent in the data. Regarding functional form, we find that a linear demand curve would lead to a mean elasticity of -1.6 (while its range on observed data is $[-4.2, -0.58]$, (see Table 3 in the online Appendix).[16] Finally, we explore whether adding more instruments can improve the estimated elasticity by augmenting them with fuel prices: the elasticity is pushed down to -4.4, while the demand state variable d_t witnesses an upward shift that is practically constant. We repeat our estimation with the new demand state variable and find that the nonparametrically estimated primitives are identical to our baseline specification results.[17]

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

$$d_t = \widehat{\eta_0} + \widehat{\eta_1} \log(\mathbf{H}_t) + \widehat{\varepsilon}_t,$$

and is plotted in Figure 5. The residual of the IV regression (7) $\widehat{\varepsilon}_t$ is included in d_t, as it captures omitted demand shifters.

[16] In the online Appendix, we provide a robustness exercise with respect to the demand specification: we repeat the entire estimation procedure under a linear instead of log-linear demand curve. We find that the estimated primitives are robust. The drawback of the linear specification is that it yields a mediocre fit in the parametric profits (see Section VA).

[17] To also assess the impact of the change in the elasticity we repeat our counterfactual experiments with this demand state variable and elasticity of -4.4 and conclude that our findings are robust.

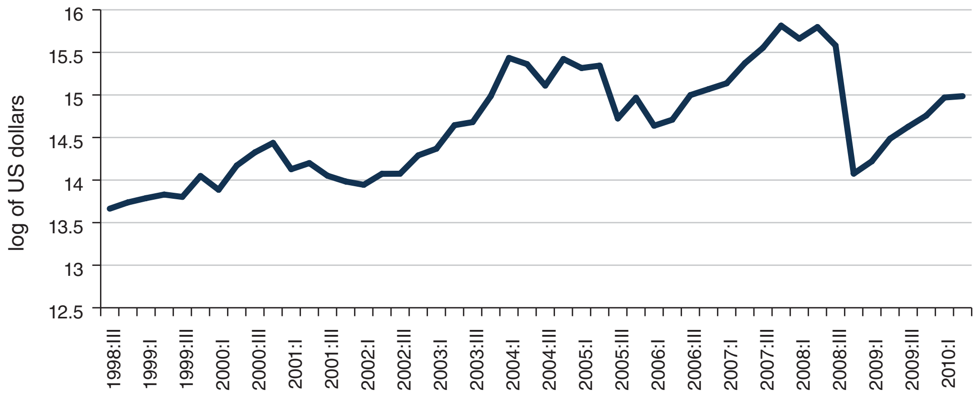

Figure 5. Demand State Variable

The benefit of using the demand estimation toward the construction of a demand state variable is twofold; on one hand, it provides a nice way of combining several observed indices affecting global transportation (as opposed to using, say, a directly observed proxy, like the index of world industrial production). In addition, it is consistent with the behavior of firms in the freight market and allows us to predict shipping prices in the counterfactual exercises.

Demand State Transition.—We specify the evolution of the demand state variable d_t as a first-order autoregressive process with nonnormal disturbances; departing for normality is important as it allows for large shocks and is discussed in detail below:

$$d_t = c + \rho d_{t-1} + \varepsilon_t. \tag{8}$$

The error ε_t follows a continuous mixture of normal distributions: if its variance σ were known, ε_t is normally distributed, $\varepsilon_t \sim N(0, \sigma^2)$. The precision parameter (inverse variance, $\frac{1}{\sigma^2}$), however, is not known but is a positive random variable that is assumed to follow the Gamma distribution. The probability density of the resulting model is given in closed form (see Hogg, Craig, and McKean 2004) by

$$f_\varepsilon(\varepsilon) = \frac{\Gamma\left(\alpha + \frac{1}{2}\right)}{\gamma^\alpha \sqrt{2\pi}\,\Gamma(\alpha)} \left(\frac{2\gamma}{2 + \gamma\varepsilon^2}\right)^{\alpha + \frac{1}{2}}, \tag{9}$$

where $\Gamma(\cdot)$ is the Gamma function and $\alpha > 0$, $\gamma > 0$ are the parameters of the Gamma density. If (α, γ) are restricted to satisfy $4\alpha\gamma = 1$, then f_ε becomes a Student t density with 2α degrees of freedom. The continuous mixture distribution is a generalization of the normal distribution, featuring a fat tail.

We estimate the parameters associated with (8) and the continuous mixture, $(c, \rho, \alpha, \gamma)$, via maximum likelihood. Results are shown in Table 2. The demand

Table 2—Demand Transition Parameters

	c	ρ	a	γ
Parameter	0.7483	0.9542	1.575	25.356
Standard error	(0.698)	(0.048)	(0.044)	(7.784)

Note: Standard errors based on 500 bootstrap samples.

process exhibits high persistence, with $\rho = 0.954$ at the quarterly level. It is stationary though, as $\rho < 1$. The reported standard errors are based on 500 bootstrap samples (the resampling is done on the error term), to account for the fact that d_t is constructed from the inverse demand curve.

The adoption of a distribution with fat tails for the innovations is important because the process allows for larger shocks compared to the normal. Indeed, the normal distribution is strongly rejected with a likelihood ratio test. Moreover, the presence of nonnormal innovations significantly increases the likelihood (from -8.16 to 35.74). We have experimented with various distributional assumptions for the demand innovations, such as the Student t distribution (as in Hamilton 2005) and the mixture of two normals, and have found the estimation results to be robust, while the likelihood was highest under the chosen specification. Finally, we can compute the probability of observing a shock as large as the crisis, i.e., $\Pr[\min(\widehat{\varepsilon}_t)] = \Pr[\min(d_t - \hat{c} - \hat{\rho} d_{t-1})]$ under each estimated specification. Normally distributed innovations assign zero probability to such shocks, even though these are observed during our sample—in contrast, this probability is significantly raised under the continuous mixture of normals.

One may worry that demand for freight services is, in fact, nonstationary, possibly trending upwards. As is well known, it is impossible to distinguish a unit root from a highly persistent but stationary $AR(1)$ process in finite samples (see, for instance, Hamilton 1994). Adopting a nonstationary model requires the use of time-varying dynamics in the firm's problem, significantly complicating the analysis and data requirements; in particular, time becomes a state variable and finite horizon techniques are used to solve the dynamic optimization of firms. As ships are long lived, it is difficult to adopt a finite horizon, especially given our data sample.[18]

Value Function.—Under the assumption of homogeneous shipowners and a large number of potential entrants, the secondhand sale price of a ship must equal its value. As shipowners have the same valuation for a ship, a seller is willing to sell at price p^{SH} only if $p^{SH} \geq V_j(\mathbf{s}, \mathbf{b}, d)$, while a buyer is willing to buy only if $p^{SH} \leq V_j(\mathbf{s}, \mathbf{b}, d)$. The resale price of a ship at age j and industry state $(\mathbf{s}, \mathbf{b}, d)$ is, thus, given by

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

[18] For robustness purposes, we experiment with Markov switching regimes. We estimate the demand process of such a model (as in Hamilton 1994) and find that d_t is always in the same regime.

where ε captures measurement error. Note that this model does not allow for unobservable state variables, as they would be an argument of the value function and not in ε.

The value function at a state $(j, \mathbf{s}, \mathbf{b}, d)$ is given by the conditional expectation of the secondhand price, p^{SH} on $(j, \mathbf{s}, b, d)$:

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

We estimate the value function, $V_j(\mathbf{s}, \mathbf{b}, d)$, nonparametrically on a grid of states, via local linear regression. We first describe how the grid of states is obtained and then proceed to the estimation on this grid, along with the estimation results. We then provide a discussion of the assumptions underlying our estimation procedure. Finally, we summarize some results from the robustness analysis presented in the online Appendix.[19]

The state variable consists of the age of the ship, $j \in \{0, 1, \ldots, A\}$, the age distribution of the fleet $\mathbf{s}_t$ of dimension $A + 1$, the backlog $\mathbf{b}_t$ of size $\overline{T}$, and the scalar demand d_t. The overall state has dimension $A + \overline{T} + 2$. Define:

$$\mathbf{x}_t \equiv [s_t^0, \ldots, s_t^A, b_t^1, \ldots, b_t^{\overline{T}}, d_t] \in (\mathcal{S} \times \mathcal{B} \times \mathcal{D}) \subset \mathbb{R}^{A+\overline{T}+2}.$$

In our case of quarterly data, $A = 80$ and $\overline{T} = 11$ (the maximum time to build observed), so that the state $[j, \mathbf{x}_t]$ has 94 components. Performing local linear regression on 94 variables is impossible. In addition to this dimensionality problem, each state $[j, \mathbf{x}_t]$ can take on a large number of values, as each of its 94 components can itself take a large number of values. The high state dimension and the large number of integer values the state components can take pose significant computational challenges.

To develop a manageable finite state space we employ aggregation and clustering techniques. Aggregation addresses the first problem of dimensionality, serving to drive down the dimension of the state space. We reduce the 93 dimensions of $\mathbf{x}_t$ to five by considering three age groups, $\mathbf{S}_t = (S_t^1, S_t^2, S_t^3)$, 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 $S_t^3 = \sum_{i=A_2+1}^{A_3} s_t^i$, the total backlog $B_t = \sum_{i=1}^{\overline{T}} b_t^i$ and demand d_t. In particular, the first age group includes the number of ships between zero and ten years old, so that $S_t^1 = \sum_{i=0}^{39} s_t^i$, the second group includes the number of ships between ten and 20 years old so that $S_t^2 = \sum_{i=40}^{79} s_t^i$, and the third group is the number of ships older than 20 years old $S_t^3 = s_t^{80}$. Aggregation linearly compresses $\mathbf{x}_t$ into the much smaller vector, $\mathbf{X}_t$:

$$\mathbf{X}_t = [S_t^1, S_t^2, S_t^3, B_t, d_t] \in \mathbb{R}^5.$$

Clustering techniques are employed to address the second problem of dimensionality and reduce the number of values these aggregated states $\mathbf{X}_t$ can take. Indeed, like the original state $[j, \mathbf{x}_t]$, the aggregated state $[j, \mathbf{X}_t]$ can also take a large number of values: S_t^i and B_t are positive integers whose upper bound is considerably larger

[19] Adland and Koekebakker (2007) use secondhand ship sale prices to uncover nonparametrically how a ship's value in the secondhand market is affected by its age, its size, and freight rates.

that the upper bound of their respective summands s_t^i and b_t^i. To reduce the admissible values of the state, we employ vector quantization, a well-known technique in computer science aimed at grouping data points into a given number of categories.[20] In particular, we represent $[\mathbf{S}_t, B_t]$ by groups $[\mathbf{S}, B]$ and d_t by groups d. We then take all possible combinations of the $[\mathbf{S}, B]$ groups and the d groups. We end up with L representative states $\mathbf{X} = [\mathbf{S}, B, d] \in \mathbb{R}^{L\times 5}$ where L is a few hundred (and combined with the 81 possible ages this leads to about 16,000 states). The vector quantization technique guarantees that the observed data are represented well by the groups, while including unobserved combinations of d with the remaining aggregate state enriches the state space.

We estimate the value function by local linear regression, i.e., we perform linear regression at every grid point $[j, \mathbf{X}] = [j, \mathbf{S}, B, d]$, including a weighting scheme that downweights the contributions of data points away from $[j, \mathbf{X}]$. We therefore solve the following local weighted least squares problem, for all grid points $[j, \mathbf{X}]$:[21]

$$\min_{\beta_0,\beta_j,\beta_s,\beta_B,\beta_d} \sum_i \{p_i^{SH} - \beta_0(j,\mathbf{X}) - \beta_j(j,\mathbf{X})(j_i - j) - \beta_s(j,\mathbf{X})(\mathbf{S}_i - \mathbf{S}) - \beta_B(j,\mathbf{X})(B_i - B) - \beta_d(j,\mathbf{X})(d_i - d)\}^2 K_h([j_i,\mathbf{X}_i] - [j,\mathbf{X}]),$$

where $\{p_i^{SH}, j_i, \mathbf{S}_i, B_i, d_i\}_{i=1}^{1838}$ represent the data of secondhand sale transactions (so that $\mathbf{S}_i = [S_i^1, S_i^2, S_i^3] = [\sum_{l=0}^{39} s_i^l, \sum_{l=40}^{79} s_i^l, s_i^{80}]$ and $B_i = \sum_{l=1}^{\bar{T}} b_i^l$), while K_h is a multivariate normal kernel with bandwidth $h \in \mathbb{R}^6$, so that $K_h(\mathbf{Y}_i - \mathbf{Y}) = \exp((\mathbf{Y}_i - \mathbf{Y})'H^{-1}(\mathbf{Y}_i - \mathbf{Y}))/\sqrt{2\pi}\prod_j h_j$, where $\mathbf{Y}_i = [j_i, \mathbf{S}_i, B_i, d_i]$ and H is a diagonal matrix with diagonal entries the bandwidths h_j. The unit of observation i corresponds to a quarter in our sample and identifies the common to all ships state $[\mathbf{S}_i, B_i, d_i]$. It also designates a specific ship sold identified by its age, j_i. The bandwidth is different for each state variable based on its observed range. We have experimented with $k-$nearest neighbor bandwidths and nonnormal kernels, and our results are robust.

Figure 6 shows the estimated value functions. Each quarter is represented by its closest state $\mathbf{X} = (\mathbf{S}, B, d)$. We find that older ships have a lower value function for all $\mathbf{X}$, as is expected.

Pointwise confidence intervals are computed via bootstrap samples; we use the standard bootstrap method for local linear regression described in Fan and Gijbels (1996). In particular we resample prices and states without replacement 500 times. Figure 7 shows the value function of a zero-year-old ship and that of a 20-year-old ship with the computed 95 percent confidence intervals. As more observations are available for old ships, confidence intervals are tighter.

[20] Vector quantization is a generalization of the K-means algorithm (performed on scalars rather than vectors) which partitions data into mutually exclusive clusters. Each cluster in the partition is defined by its member objects and by its centroid. In practice this is an iterative algorithm that minimizes the sum of distances from each object to its cluster centroid, over all clusters, and is provided by programming languages like MATLAB.

[21] For more information on local linear regression, see Fan and Gijbels (1996).

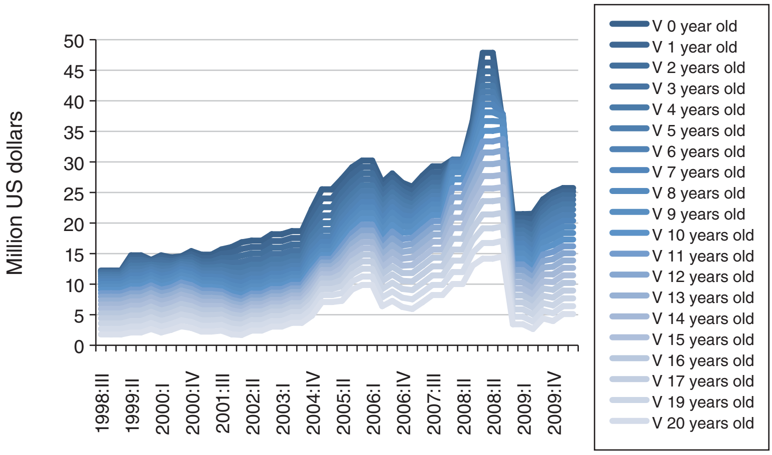

Figure 6. Value Functions of Yearly Ages

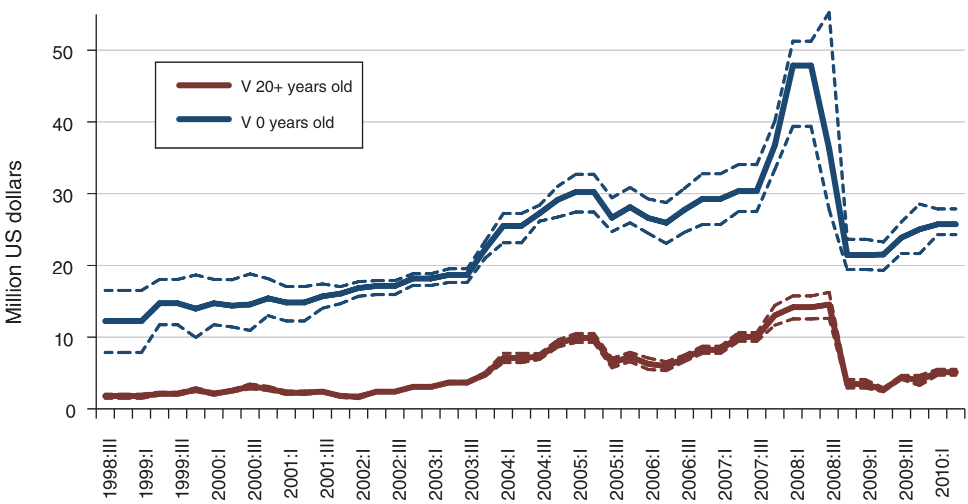

Figure 7. Value Functions and Confidence Intervals

Estimating value functions using resale prices is central in our empirical strategy. One may worry about both adverse selection and sample selection. Adverse selection would arise in the model if different shipowning firms 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 (1997): "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 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. Unfortunately, we do not observe the distribution of observable characteristics for ships not in our secondhand sample to test for selection further. There is, however, strong suggestive evidence against selection. The secondhand market for ships is thick: at least 10 percent of the fleet is traded annually (even during the crisis). In addition, price variation conditional on a state $(j, \mathbf{S}, B, d)$ is quite small. In particular, we calculate the coefficient of variation of prices at each state. The median (across states) coefficient of variation is 0.15. Moreover, as mentioned above, in the online Appendix we estimate value functions via LASSO; the post-LASSO r-square is 0.81. Part of the unexplained variation in prices arises from assigning the same age to all ships older than 20 years. The median coefficient of variation falls to 0.13 if old ships are ignored, while the post-LASSO r-square is 0.87 if the observed individual age of the ship in years is employed.[22] 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 the Section IV. In particular, Figure 12 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, let us also discuss the possibility of search costs: based on Stopford (1997), 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."[23]

We conclude this section with a discussion of the robustness exercises performed in the online Appendix. We first examine how robust our results are regarding the state aggregation by reestimating the model using a different aggregation. Indeed, instead of three age groups, we use the total fleet younger and older than 20 years old, as well as the mean and standard deviation of the age of the fleet (younger than 20 years old). We find that the estimated primitives (i.e., profits, entry cost, and scrap value distribution) are very similar to the baseline. We also examine robustness with respect to the local linear regression estimator, by employing sieve estimation via OLS as well as LASSO on the aggregate states $[j, \mathbf{X}_t]$. Value functions are surprisingly robust to the LASSO/OLS estimator for several order polynomials.

Time to Build Function.—We use the ship orderbook data to construct the average time to build corresponding to each quarter in our sample, T_t. Consistent with our state space reduction, we assume that time to build is a function of $\mathbf{X}_t = (\mathbf{S}_t, B_t, d_t)$. We estimate the following time to build function, and Figure 8 presents the actual and predicted time to build.

[22] Unfortunately, even though we know the age of all ships sold, we have to group those older than 20. The reason is twofold. First, we can't transform the age in years to age in quarters, as we don't know the distribution of deliveries prior to 1970—see the online Appendix. Second, we don't have data on the ages of ships demolished and therefore can't allow for age-specific exit rates. The adoption of the "20 plus group" may be biasing scrap values upward if most demolitions occur at ages significantly older than 20.

[23] In this article, secondhand prices are the tool we use to estimate our dynamic model of ship entry and exit. A large literature examines resales and their frictions in other industries, such as Gavazza (2011), who examines the case of commercial aircraft.

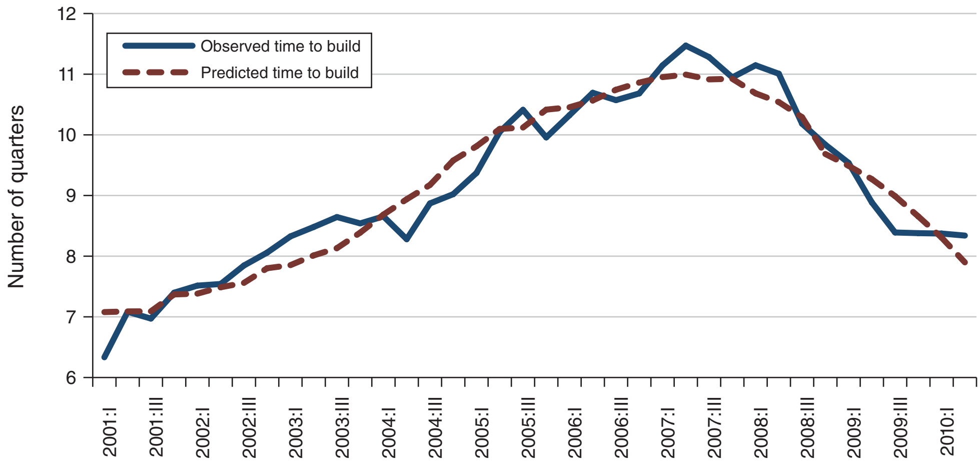

Figure 8. Average Observed and Estimated Time to Build

$$\log\left(T_t\right) = a_0^{TTB} + a_S^{TTB}\mathbf{S}_t + a_B^{TTB}B_t + a_d^{TTB}d_t.$$

Transition Matrix.—The evolution of the state is as follows:[24]

$$\mathbf{x}_t = \left[s_t^0, \ldots, s_t^A,\, b_t^1, \ldots, b_t^{\overline{T}},\, d_t\right]$$

and

$$\mathbf{x}_{t+1} = \left[s_{t+1}^0, \ldots, s_{t+1}^A,\, b_{t+1}^1, \ldots, b_{t+1}^{\overline{T}},\, d_{t+1}\right]$$

$$= \left[b_t^1, s_t^0, \ldots, s_t^{A-2}, s_t^A + s_t^{A-1} - Z_t, b_t^2 \ldots, b_t^{T(\mathbf{X}_t)} + N_t, \ldots, b_t^{\overline{T}-1}, 0, c + \rho d_t + \varepsilon_{t+1}\right].$$

The driving stochastic processes are, ε_{t+1}, N_t, and Z_t. More precisely, ε_{t+1} follows a continuous mixture of normals and summarizes the uncertainties inherent in demand for shipping services. The number of entrants N_t is distributed Poisson with rate $\lambda(\mathbf{x}_t)$. We approximate the binomial random variable Z_t by a Poisson random variable with mean $\mu(\mathbf{x}_t)$. This approximation is formally justified if the number of firms s_t^A is large, while the exit probability is small—both conditions hold in our data, as shown in Figure 4. Under the assumption that N_t, Z_t, and ε_{t+1} are conditionally independent given the state $\mathbf{x}_t$, the transition probability from $\mathbf{x}_t$ to $\mathbf{x}_{t+1}$ is given by

$$\Pr\left(\mathbf{x}_{t+1} \,|\, \mathbf{x}_t\right) = Q(\mathbf{x}_t, \mathbf{x}_{t+1}) f_N(N_t = n \,|\, \mathbf{x}_t) f_Z(Z_t = z \,|\, \mathbf{x}_t) f_\varepsilon(\varepsilon_{t+1}), \tag{10}$$

[24] The time to build function $T(\mathbf{X}_t)$ can be used to determine the position of entrants (for state $\mathbf{x}_t$ we compute the corresponding group state $X_t = (\mathbf{S}_t, B_t, d_t)$).

where

$$(11)\quad Q(\mathbf{x}_t, \mathbf{x}_{t+1}) = \delta\{s_{t+1}^0 = b_t^1\} \prod_{i=1}^{A-1} \delta\{s_{t+1}^i = s_t^{i-1}\} \prod_{i \neq T(\mathbf{X}_t)} \delta\{b_{t+1}^i = b_t^{i+1}\}$$

and

$$n = b_t^{T(\mathbf{X}_t)} - b_t^{T(\mathbf{X}_t)+1}$$

$$z = s_t^A + s_t^{A-1} - s_{t+1}^A$$

$$e = d_{t+1} - c - \rho d_t,$$

while $\delta\{\cdot\}$ is an indicator function, f_N and f_Z are the Poisson densities with parameter $\lambda(\mathbf{x}_t)$ and $\mu(\mathbf{x}_t)$, respectively, and f_ε is given in (9).

As explained in Section IVA, we base the estimation on the aggregated states $\mathbf{X}$, rather than $\mathbf{x}_t$ due to the large dimension of the latter. State aggregation creates complications in the computation of expectations: even though the transition from $\mathbf{x}_t$ to $\mathbf{x}_{t+1}$ is known (given in (10) and (11)), the transition from $\mathbf{X}_i$ to $\mathbf{X}_j$ is unknown. We propose a method for computing the transition matrix $\mathbf{P} \in \mathbb{R}^{L \times L}$ whose (i, j) element P_{ij} specifies the probability of transiting from state $\mathbf{X}_i$ to state $\mathbf{X}_j$. This method proceeds in two steps and employs the known state dynamics of $\mathbf{x}_t$.

First, note that (10) and (11) are known if the entry $\lambda(\mathbf{x}_t)$ and exit $\mu(\mathbf{x}_t)$ processes, as well as the demand parameters $(c, \rho, \alpha, \gamma)$, are known. The first step of the transition matrix computation consists, therefore, of recovering the Poisson processes of entry and exit (results presented below) and the demand transition (8) (results presented in Section IVA).

Second, once the transition of $\mathbf{x}_t$ is fully specified from the first step, we simulate thousands of states $\mathbf{x}_t \in \mathbb{R}^{A+T+2}$ using the estimated Poisson processes for N_t and Z_t and the $AR(1)$ process for d_t. For each simulated state $\mathbf{x}_t$ we find its closest group state $\mathbf{X} = [\mathbf{S}, B, d] \in \mathbb{R}^5$ (by computing its distance from each of the L group states and taking the minimum). We now have simulated paths of group states and can compute the frequency transition matrix (i.e., count the number of times $\mathbf{X}_i$ transits to $\mathbf{X}_j$ and divide by the total number of visits to $\mathbf{X}_i$).[25]

We next turn to the results.

We assume that the entry and exit rates are given by the following functional forms $\lambda(\mathbf{S}_t\ B_t, d_t) = \exp(\gamma_0 + \gamma_{S1} S_t^1 + \gamma_{S2} S_t^2 + \gamma_{S3} S_t^3 + \gamma_B B_t + \gamma_{B^2}(B_t)^2 + \gamma_d d_t)$ and $\mu(\mathbf{S}_t, B_t, d_t) = \exp(\delta_0 + \delta_{S_1} S_t^1 + \delta_{S_2} S_t^2 + \delta_{S3} S_t^3 + \delta_B B_t + \delta_{B^2}(B_t)^2 + \delta_d d_t)$.[26] We estimate the parameters γ and δ via OLS using the time series of entrants[27] and exitors. This step is similar to the first estimation step in Bajari, Benkard, and Levin

[25] We also smooth the frequency transition matrix as follows: each element (i, j) of the transition matrix is the inner product of the jth column and a row of weights reflecting the closeness of the ith state to the rest of the states. The weights are based on a multivariate normal kernel.

[26] Since λ and μ are greater or equal than zero, they can be expressed as $\exp(f(\mathbf{S}_t, B_t, d_t))$. We can therefore think of the chosen specifications as Taylor approximations to $f(\cdot)$ (first order in S_t^1, S_t^2, S_t^3 and d_t and second order in B_t).

[27] Instead of OLS we could estimate via local linear regression and recover $\lambda = E(N_t | \mathbf{S}_t, B_t, d_t)$. If we then perform OLS on the nonparametrically estimated $E(N_t | \mathbf{S}_t, B_t, d_t)$, we get the same results.

(2007) or Pakes, Ostrovsky, and Berry (2007) where policy functions are estimated "offline." Ideally, these policy functions would be recovered nonparametrically, but as is common in this literature, data limitations force us to use parametric assumptions. Table 7 in Appendix A reports the results. As expected, entry is increasing in demand: as the demand curve for freight shifts outward, more entry is attracted to the market. Similarly, exit is decreasing in demand: as the demand curve for freight shifts outward, exit becomes a less attractive option. Entry is decreasing in the number of competitors, (S^1, S^2, S^3), with the number of young competitors S_1 affecting entry more strongly. Exit is increasing in the number of competitors (S^1, S^2, S^3). Note, however, that entry is increasing in the backlog for low backlog levels and then decreasing, while exit is first decreasing and then increasing in the backlog. This problem is due to the "offline" estimation of policy functions found in most two-step estimation procedures for dynamic games which may suffer from endogeneity of competition variables. To avoid this issue, we choose more parsimonious specifications. We have experimented with various functional forms, as well as several estimation methods, and concluded to the following forms: $\lambda(\mathbf{S}_t, B_t, d_t) = \exp(\gamma_0 + \gamma_{S_1} S_t^1 + \gamma_{S_2} S_t^2 + \gamma_d\, d_t)$ and $\mu(\mathbf{S}_t, B_t, d_t) = \exp(\delta_0 + \delta_{S_1} S_t^1 + \gamma_{S_2} S_t^2 + \gamma_{S_3} S_t^3 + \delta_d\, d_t)$. Table 8 in Appendix A reports the results. As shown in Appendix A, only d_t appears to be significant when bootstrap standard errors are used (to account for the construction of d_t). The regressions, however, are strongly jointly significant at the 0.01 level, implying that the relevant variables are important in determining the dynamics of the industry. It is perhaps not surprising given the short time series and measurement error in d_t that the demand state variable turns out to be the only significant term individually, as in a sense it "swamps" everything else. Collectively, however, the states matter and, importantly, have the correct economic signs and are thus used in computing the transitions of the industry.

B. *Second Step: Estimate the Model Primitives*

In the second step of our estimation strategy we estimate the model primitives. The discount factor is set to 0.9877 which corresponds to 5 percent annual interest rate.

Scrap Value Distribution.—Once we obtain the transition matrix, $\mathbf{P} \in \mathbb{R}^{L \times L}$ and the value function $\mathbf{V} \in \mathbb{R}^{L \times (A+1)}$, the continuation value is simply

$$\mathbf{VC}_j = \mathbf{PV}_{j+1}, \qquad j = 0, 1, \dots, A - 1$$

$$\mathbf{VC}_A = \mathbf{PV}_A.$$

The estimated continuation value function of old ships $\mathbf{VC}_A$ provides direct information on scrap values and is employed to recover their density, in combination with observed exit frequencies and the model implied exit equation. Indeed, we estimate the scrap value distribution nonparametrically, via local linear regression of the observed exit probabilities Z_t/s_t^A on the continuation value of old ships, $VC_A(\mathbf{x}_t)$. The probability of exit at continuation value v equals $[1 - F_\phi(v)]$, or

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

We run local linear regression at all points of interest v:

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

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

The distribution and density of scrap values are shown in Figure 9. Our data provide information directly on the value function of an old ship, which in turn is translated in the quantiles of the scrap value density. The horizontal axis in Figure 9 is essentially derived from the range of continuation values.

Confidence intervals for the scrap value distribution are constructed as follows. We use the 500 bootstrap samples of the entry and exit regressions, as well as the autoregressive process for the demand state variable to create 500 bootstrap transition matrices. As the errors are independent, we multiply the 500 bootstrap transition matrices with the 500 bootstrap value functions and create 500 bootstrap continuation values for old (20 years old and over) ships. For each bootstrapped continuation value, we calculate the scrap value distribution and density by following the bootstrap method for local linear regression (i.e., we create a bootstrap sample for Z_t/s_t^A and the corresponding continuation value and perform local linear regression as above).

A common concern with two-step estimators for dynamic setups is that the first-stage exit policy, which is estimated "offline" and thus does not impose equilibrium behavior, needs to be close to the model predicted exit policy:

$$\tilde{\mu}(\mathbf{S}, B, d) = s^A\big(1 - F_\phi(\mathbf{VC}_A)\big).$$

The trade-offs associated with the level of flexibility for the first stage are discussed in Bajari, Benkard, and Levin (2007). Following Li and Racine (2007), we test whether the two exit policy functions are equal using bootstrap critical values and find that we cannot reject the null of equality (the p-value is 0.59).

Profits.—We compute profits $\boldsymbol{\pi} \in \mathbb{R}^{L\times(A+1)}$ from the Bellman equations (3), which we repeat here for convenience:

$$\mathbf{V}_j = \boldsymbol{\pi}_j + \beta E_\phi \max\{\phi, \mathbf{VC}_j\}.$$

Profits then can be estimated from

$$\boldsymbol{\pi}_j = \mathbf{V}_j - \beta\mathbf{VC}_j, \qquad j = 0, 1, \ldots, A - 1 \tag{12}$$

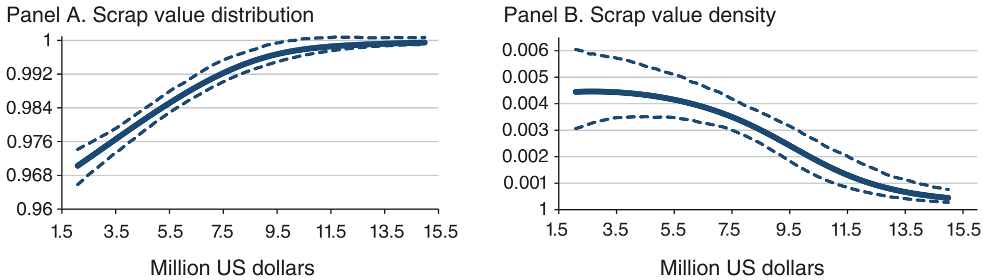

Figure 9. Distribution and Density of Scrap Values

Note: Dashed lines denote 95 percent bootstrap confidence intervals.

and

$$\boldsymbol{\pi}_A = \mathbf{V}_A - \beta E_\phi \max\{\phi, \mathbf{VC}_A\}$$

or

$$\boldsymbol{\pi}_A = \mathbf{V}_A - \beta(1 - \mathbf{p}^X)\,\mathbf{VC}_A - \beta \mathbf{p}^X E[\phi \,|\, \phi > \mathbf{VC}_A], \tag{13}$$

where $\mathbf{p}^X \in \mathbb{R}^L$ is the exit probability, so that $p^X(\mathbf{S}, B, d) = \tilde{\mu}(\mathbf{S}, B, d)/s^A$.

Figure 10 presents the estimated profits of ships older than 20 years old, as well as the average profits of ships ten to 20 years old. Confidence intervals are constructed by calculating the profits corresponding to each of the 500 bootstrap samples of value functions and continuation values (for the 20+ age category, the bootstrapped scrap value density and exit rate are also used).[28]

Note that profits fluctuate close to and below zero between 1998 and 2006, spike during the boom of 2007–2008, and collapse at the crisis. This pattern is consistent with hysteresis. Shipowners are willing to suffer losses and delay exit decisions in anticipation of better times, when demand is high and the supply response is limited in the short run due to time to build and convex operating costs of existing ships. Hysteresis results from investment irreversibility—on one hand, scrapped ships can't be "unscrapped"; on the other hand, scrap values are considerably lower than entry costs, a fact taken into consideration at the entry decision. To get some rough measure of this irreversibility, we compare the average scrap price of a Handysize ship during our sample period to the average price of a new Handysize vessel that prevailed 20 years in the past. We find that scrap prices are on average equal to 16 percent of new ship prices (of 20 years ago), while this percentage fluctuates violently from 5.8 percent to 32 percent. Hysteresis can also lead to chronic excess capacity, a not unusual phenomenon in oceanic bulk shipping. Indeed, during the crisis of 2009 the press featured several pictures of idle ships anchored. Stopford (2009) also provides anecdotal evidence of these patterns: "Each company faces the challenge of navigating its way through the succession of booms and depressions

[28] Value functions and profits for young ships are noisier as we have fewer observations on them. In counterfactuals, we will impose a parametric assumption on profits so that this won't be as important a concern; see Section VA.

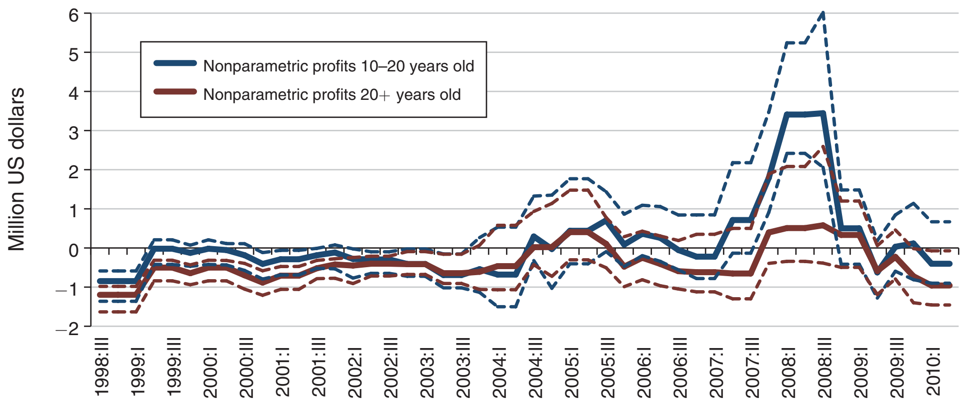

Figure 10. Estimated Profits

Notes: Average across age 10–20 years and 20+ years old. 0.95 bootstrap confidence intervals.

that characterize the shipping market," while giving a specific example, "For several years the company had accepted this drain on its cashflow, in the hope that the market would improve."

In order to assess the validity of our model and estimation results, we use an earnings index constructed by Clarksons. 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 use it only to assess the evolution of profits over time. Our estimated profits track broker calculated profits well, as shown in Figure 11.

Entry Costs.—The free entry condition states that at each state $\mathbf{X} = (\mathbf{S}, B, d)$, the value of entry equals the entry cost, $\kappa(\mathbf{X})$. The value of entry is the expectation of the value function of age zero, $T(\mathbf{X})$ periods ahead. To construct the value of entry at $\mathbf{X}$ we first compute the $T(\mathbf{X})$th power of the transition matrix. We then multiply the row of that matrix corresponding to state $\mathbf{X}$ with $\mathbf{V}_0 \in \mathbb{R}^L$ as well as $\beta^{T(\mathbf{X})}$, i.e.,

$$\beta^{T(\mathbf{X})}\, \mathbf{P}_X^{T(\mathbf{X})}\, \mathbf{V}_0 = \kappa(\mathbf{X}), \tag{14}$$

where $\kappa \in \mathbb{R}^L$ are the entry costs, $\mathbf{V}_0 \in \mathbb{R}^L$ is the value of an age zero ship, and $\mathbf{P}_X^{T(\mathbf{X})}$ refers to the row of $\mathbf{P}^{T(\mathbf{X})}$ that corresponds to state $\mathbf{X}$. Having computed the transition matrix $\mathbf{P}$, the time to build function $T(\mathbf{X})$, and the value function $\mathbf{V}_0$ we can directly estimate entry costs κ from the free entry condition.[29] Confidence intervals for the entry costs are constructed by computing the entry costs of each bootstrap sample of a zero-year-old ship and a bootstrap transition matrix.

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

[29] The free entry condition holds with equality only when entry is strictly positive. If mean entry is zero, then the value of entry provides only a lower bound for the entry costs. There is no zero entry at any of our states.

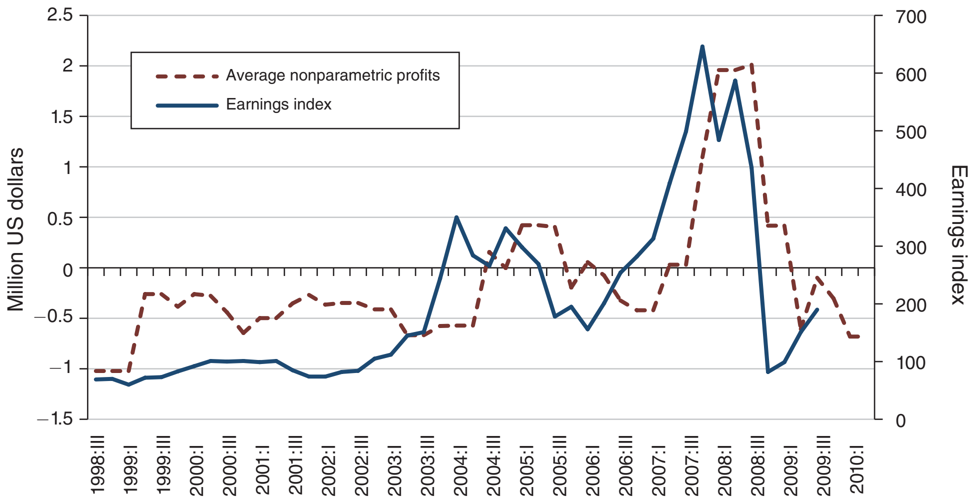

Figure 11. Weighted Average of Nonparametric Profits and Earnings Index

Note: Index based on 45000DWT Handymax, based on freight rates and fuel/port costs for representative routes.

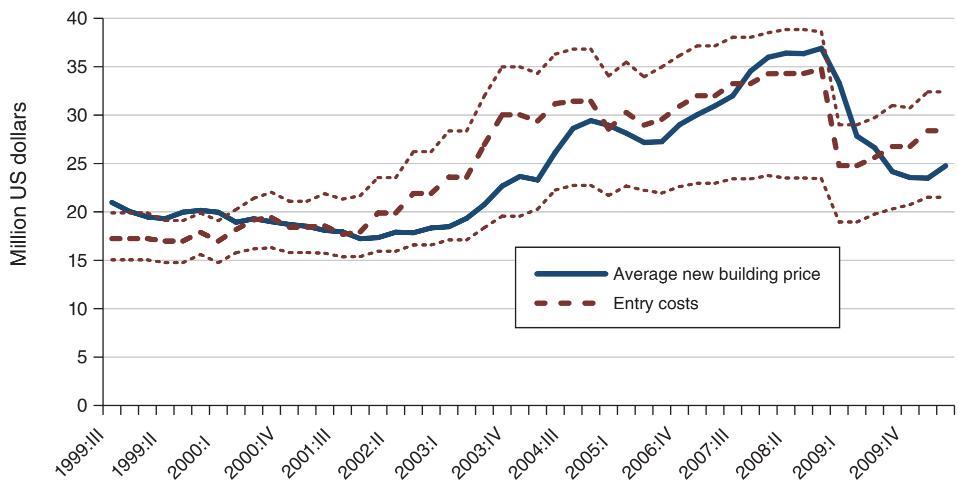

Figure 12. Estimated Entry Costs and Average Shipbuilding Price

Note: Dashed lines denote 95 percent bootstrap confidence intervals.

to evaluate the power of our model and estimation results. The average shipbuilding price is not used in the estimation and acts only as an "additional degree of freedom." Considering that entry costs result from the combination of the estimated value function of a young ship—which is our noisiest estimated value function (due to the shortage of observations on young ships)—with the estimated transition matrix to the $T(\mathbf{X})$th power and the free entry condition implied by our model, the fit is surprisingly good. Following Li and Racine (2007) and using bootstrap critical values, we find that we cannot reject the null that the entry costs are equal to the shipbuilding price (the p-value is 0.66).

V. Counterfactuals

In this section, we assess how quantitatively important time to build is in shaping the dynamic evolution of the industry, as well as the industry's response to shocks. In particular, we explore the impact of time to build on investment in new ships, freight rates, and surplus. First, we assume there is perfect competition in the market for freight transport and estimate a parametric form for profits. Next, we compute the equilibrium of our model in two counterfactual worlds of constant and no time to build. We explore the impact of demand impulses and perform long-run simulations.

A. *Perfectly Competitive Profits*

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. We assume there is perfect competition in the market for freight transport and use the estimated nonparametric profits to estimate the operating cost parameters. The perfectly competitive behavior ties well with the large number of firms. We also experimented with the assumption that firms engage in competition in quantities (Cournot game) but found that the quantities chosen lead to marginal cost pricing.

We assume that firms choose quantities, taking prices as given and subject to convex operating costs that increase in the ship's age. In particular, every period, a ship of age j chooses the number of voyages q in order to maximize profits:

$$\max_q \{Pq - c_j q^3 - F_j\}, \tag{15}$$

where $P = e^d Q^\eta$ is the inverse demand curve (as estimated in Section IVA) and $(c_j q^3 + F_j)$ are total costs that include fixed costs F_j. Based on Stopford (1997), fuel consumption is proportional to the cube of the vessel's speed. Furthermore, the number of trips is proportional to speed, and cost is proportional to fuel cost. Costs change when the ship's age crosses a new decade, so that

$$c_j \in \{c_{0-39}, c_{40-79}, c_{80+}\}$$

$$F_j \in \{F_{0-39}, F_{40-79}, F_{80+}\}.$$

Under these assumptions, price can be expressed as a function of the state $[S^1, S^2, S^3, d]$:

$$P(S^1, S^2, S^3, d) = (e^d)^{\frac{2}{2-\eta}} \left(\frac{S^1}{\sqrt{3c_{0-39}}} + \frac{S^2}{\sqrt{3c_{40-79}}} + \frac{S^3}{\sqrt{3c_{80+}}} \right)^{\frac{2\eta}{2-\eta}}. \tag{16}$$

Replacing the above in (15) we obtain the profits $\{\pi_{0-39}, \pi_{40-79}, \pi_{80+}\}$ as a function of the state $[S^1, S^2, S^3, d]$ and the parameters (c_j, F_j), given below:

$$\pi_j(S^1, S^2, S^3, d) = (e^d)^{\frac{3}{2-\eta}} \left(\frac{S^1}{\sqrt{3c_{0-39}}} + \frac{S^2}{\sqrt{3c_{40-79}}} + \frac{S^3}{\sqrt{3c_{80+}}} \right)^{\frac{3\eta}{2-\eta}} \frac{0.38}{\sqrt{c_j}} \tag{17}$$

Table 3—Profit Parameters

	c_{0-39}	c_{40-79}	c_{80+}	F_{0-39}	F_{40-79}	F_{80+}
Parameters	300,387	466,683	651,456	492,307	648,282	821,588
Standard errors	(103,871)	(156,016)	(216,789)	(347,982)	(347,977)	(347,971)

Note: Standard errors based on 500 bootstrap samples.

for $j \in \{0\text{–}39, 40\text{–}79, 80+\}$ representing the three age groups. Assuming that operating costs change by age decade ensures that profits depend on the specified state variables.

We use the estimated profits from equations (12) and (13), as well as observed prices, and recover the cost parameters via Nonlinear Least Squares. In order to aid the extrapolation to younger ships which are not used in the estimation, we impose a linear dependence of the cost parameters on age, so that $c_{0-39} = \gamma_0 + \gamma_1$, $c_{40-79} = \gamma_0 + 10\gamma_1$, $c_{80+} = \gamma_0 + 20\gamma_1$, and similarly $F_j = \delta_0 + \delta_1$, $F_{40-79} = \delta_0 + 10\delta_1$, $F_{80+} = \delta_0 + 20\delta_1$. We recover $\theta = \{\gamma_0, \gamma_1, \delta_0, \delta_1\}$ via

$$\min_{\theta}\left\{\sum_{i=40}^{79}\left(||\widehat{\pi}_i - \pi_i(\theta)||^2\right) + ||\widehat{\pi_{80+}} - \pi_{80+}(\theta)||^2 + ||P - P(\theta)||^2\right\},$$

where $\{\widehat{\pi}_i\}_{i=40}^{79}$ and $\widehat{\pi_{80+}}$ are the nonparametrically estimated profits from Section IVB, $\{\pi_i(\theta)\}_{i=40}^{79}$, $\pi_{80+}(\theta)$, $P(\theta)$ are the profits and price resulting from the perfectly competitive behavior, and P the observed prices from (17) and (16).[30] The estimated parameters are given in Table 3. Standard errors are calculated using the bootstrapped profits.

We compare our estimates to cost figures found in Stopford (1997) and find that the model and estimation perform well. Based on Stopford (1997), fixed costs include operating costs such as manning, stores, lubricants, repairs, insurance, administration, as well as periodic maintenance. The fixed costs of a ten-year-old Handymax vessel,[31] based on Stopford (1997), are \$586,000 quarterly, which is close to our estimate of F_{40-79} at \$648,282. Similarly, the costs of 20-year-old Handymax vessel are \$827,460, close to our value of \$821,588. Finally, Figure 13 shows the fit of the perfectly competitive market in terms of the prices and profits, as well as the final perfectly competitive profits for the three age groups. The implied perfectly competitive profits are less variable than the nonparametric profits, as is expected, especially in the ten-to-20-years-old age group.

B. *Counterfactual Results*

We now quantify the impact of time to build and demand uncertainty on investment in new ships, freight rates, and surplus. We compute the model equilibrium

[30] We don't include profits from young ships due to their not as precise nonparametric estimates (parameters, however, are rather robust to their inclusion).

[31] Unfortunately, the costs for Handysize vessels, as well as younger vessels, are not reported.

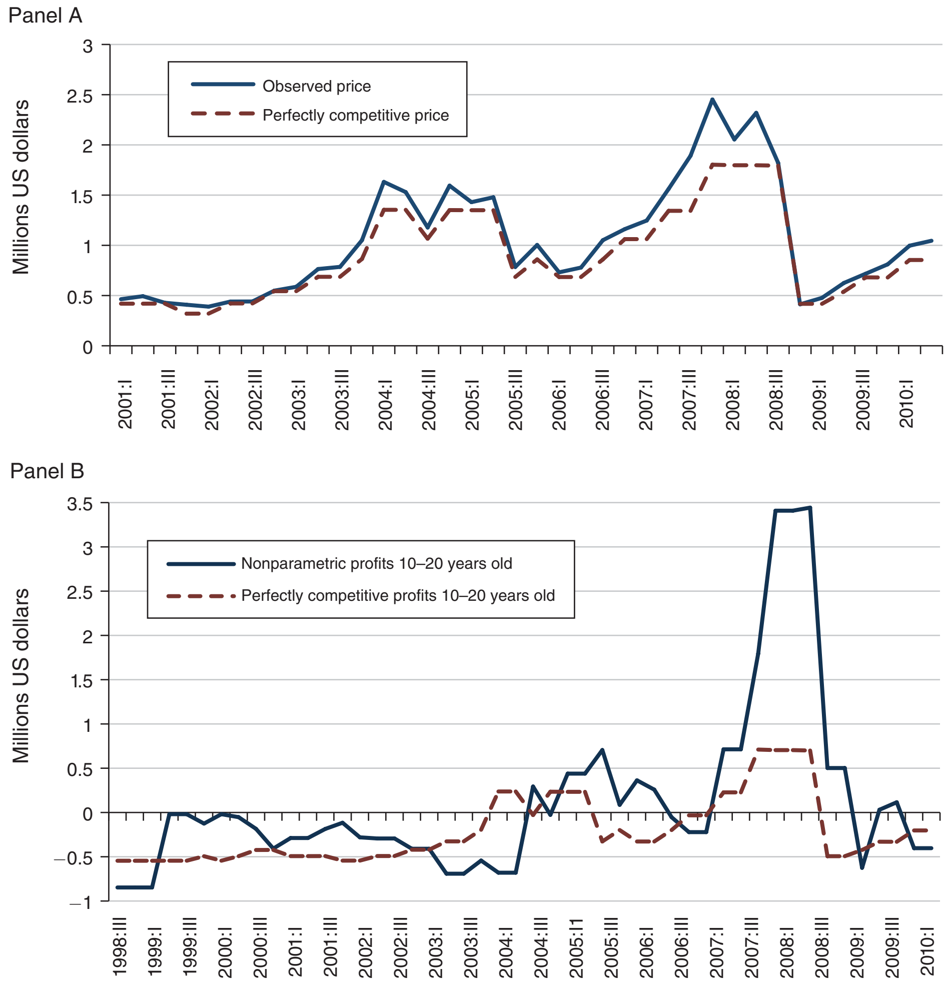

Figure 13. Results of Perfectly Competitive Model

under two counterfactual scenarios: constant time to build and no time to build. We compare these two worlds to the true world of endogenous time to build, i.e., $T_t = T(\mathbf{S}_t, B_t, d_t)$. In the extreme case of no time to build, firms that decide to enter at period t begin operating at period $t + 1$. The case of constant time to build captures a situation without shipbuilding capacity constraints, allowing for only a time invariant construction period. We choose this constant time to build to equal six quarters, the minimum time to build we observe. To construct the counterfactual worlds, we compute the equilibrium of the model, i.e., we solve for the entry policy $\boldsymbol{\lambda} \in \mathbb{R}^L$, the exit policy $\boldsymbol{\mu} \in \mathbb{R}^L$, and value function $\mathbf{V} \in \mathbb{R}^{L\times(A+1)}$, using the estimated primitives, i.e., the perfectly competitive profits $\boldsymbol{\pi} \in \mathbb{R}^{L\times(A+1)}$, the scrap value distribution F_ϕ, and the entry costs $\boldsymbol{\kappa} \in \mathbb{R}^L$.[32]

[32] See the online Appendix for details.

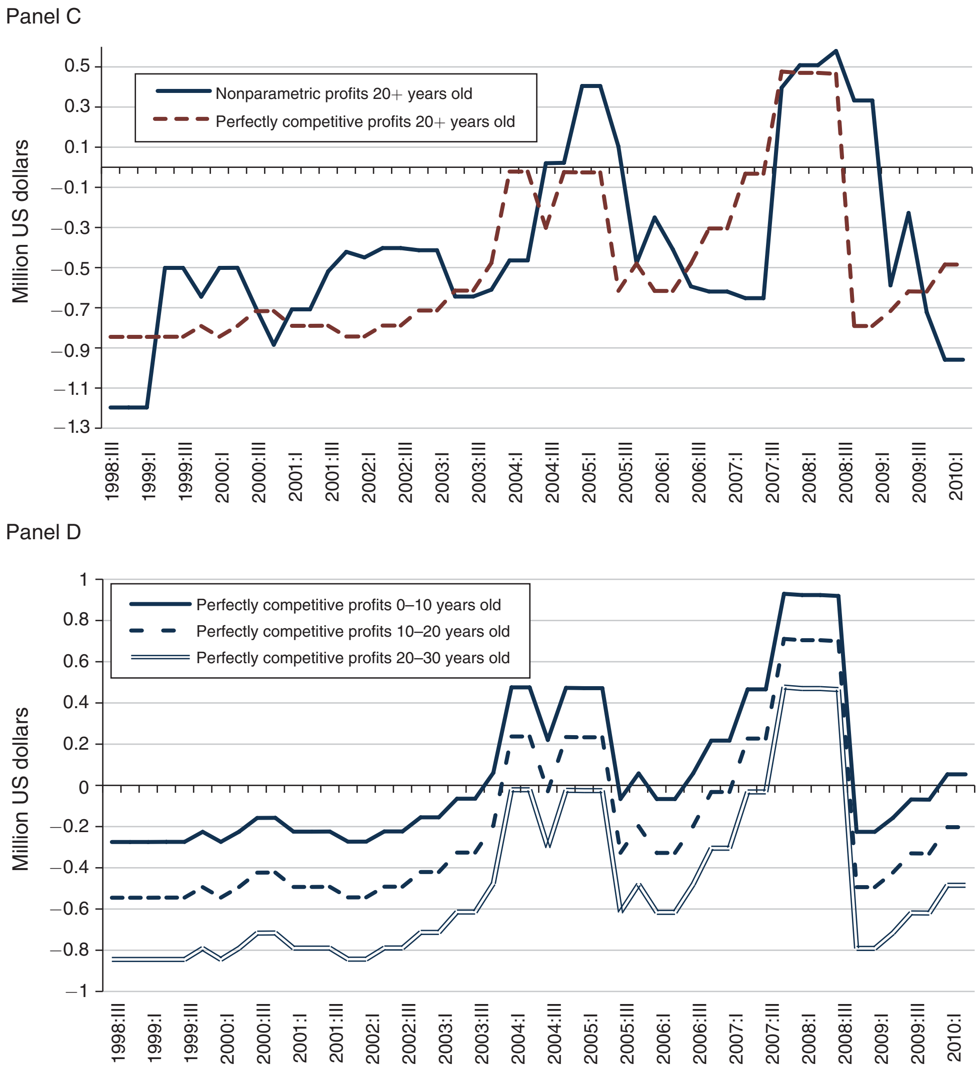

Figure 13. Results of Perfectly Competitive Model (*Continued*)

Impulse Response.—We first explore the impact of a positive shock, as this is the case where time to build becomes a critical 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, as during the boom of this decade. If time to build were reduced or removed, supply would be more elastic, and the size of the fleet would increase in order to satisfy the increase in demand, reducing prices and profits. In the case of a negative demand shock, time to build is not a critical constraint, as the fleet can immediately adjust through ship exit. We simulate the model 1,000 times for 20 quarters each: after two quarters, we hit the economy with the positive shock.

As shown in Figure 14 when the positive shock hits the industry, entry spikes. In a world with time to build, entry is both lower and slower. Indeed, new ships arrive a

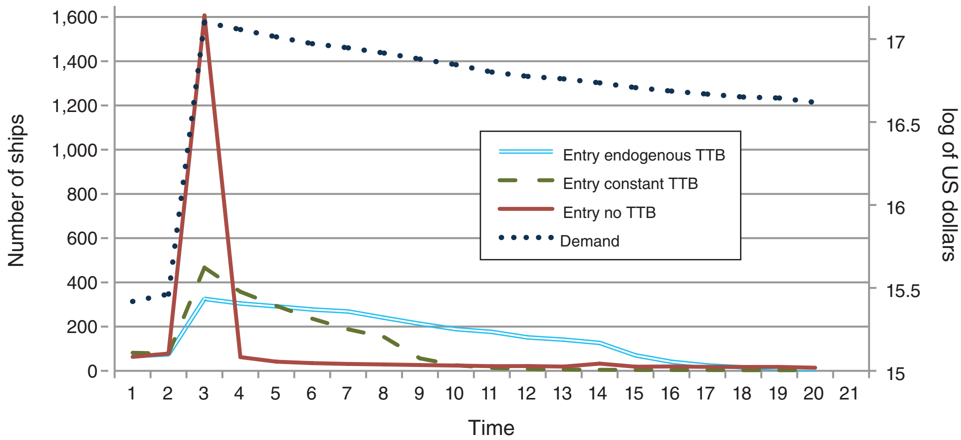

Figure 14. Entry Under Endogenous, Constant, and No Time to Build (TTB)

Note: Demand depicted on right axis.

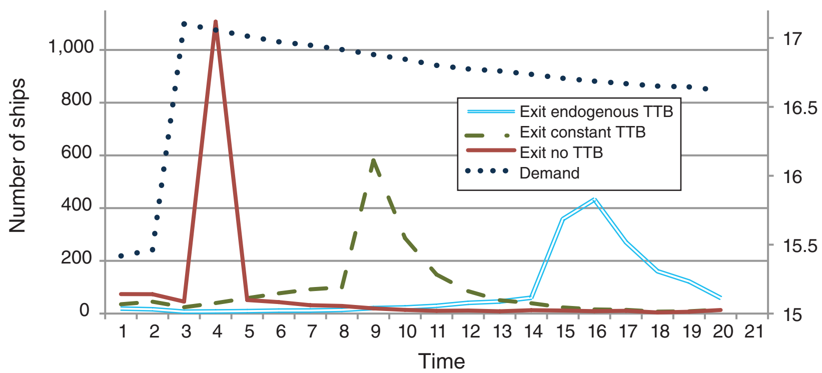

Figure 15. Exit under Endogenous, Constant, and No Time to Build (TTB)

Note: Demand depicted on right axis.

number of periods later, while the incentives to enter are dampened since the shock fades due to the mean-reverting demand process, and ships that are delivered late will not be able to take advantage of the increased demand. As a result, the response of entry is significantly larger in the absence of time to build.

Figure 15 demonstrates the response of exit to the positive demand shock. Exit occurs in "waves" as we move from no to constant to endogenous time to build. In the absence of time to build, old ships exit immediately, as a large number of young, more efficient entrants floods the market. As a result, past a positive demand shock, in a world without time to build the turnover of ships is higher compared to a world with time to build, with young ships quickly replacing old ships. In contrast, time to build induces lower and slower exit activity, as existing old ships prefer to remain in the industry, reap the benefits of increased demand, and not exit (Figure 4 documents that old ships indeed didn't exit during the boom) until new ships are

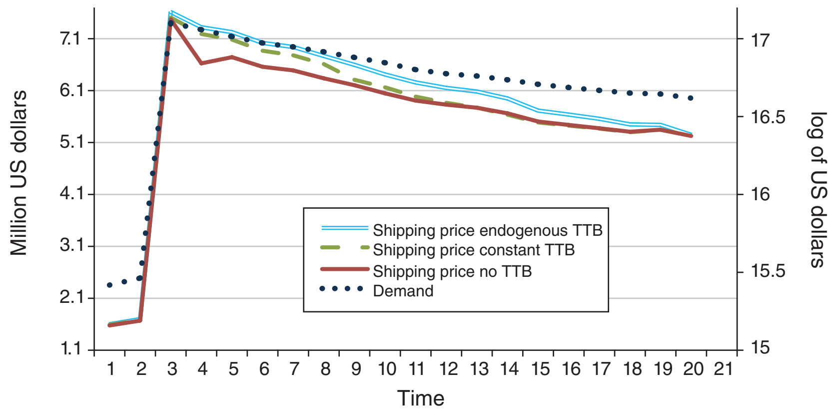

Figure 16. Price under Endogenous, Constant, and No Time to Build (TTB)

Note: Demand depicted on right axis.

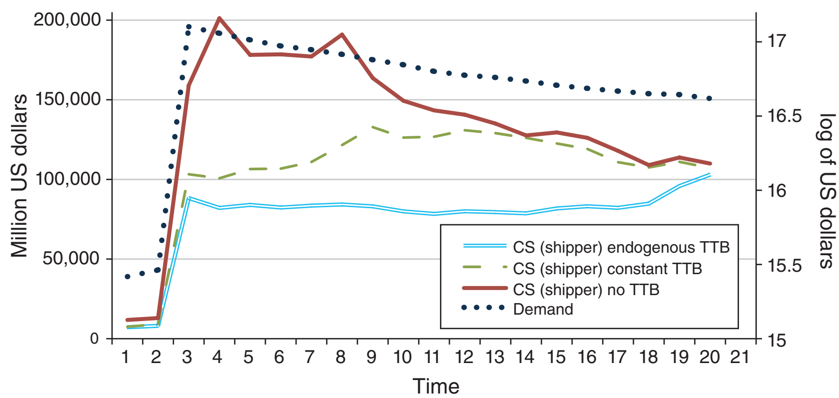

Figure 17. Consumer Surplus (*i.e., surplus of freight customers*) under Endogenous, Constant, and No Time to Build (TTB)

Note: Demand depicted on right axis.

delivered from the shipyards. The exit spike is lower when there is time to build, as entry was lower.

As a result of the entry and exit behavior, freight rates are lower as time to build decreases—shown in Figure 16—and consumer surplus (calculated using the model in Section VA) is higher—shown in Figure 17.

Figure 18 shows total and per ship profits by age group. Total profits are lower as time to build increases, simply because the fleet increases. In contrast, profits per ship are higher as time to build increases for all age groups, because past a positive demand shock operating firms reap high profits, as new firms can't enter immediately, while fewer firms enter as well.

Long-Run Simulations.—In this section we explore the dynamic evolution of the dry bulk shipping industry by performing long-run simulations of the model. Our

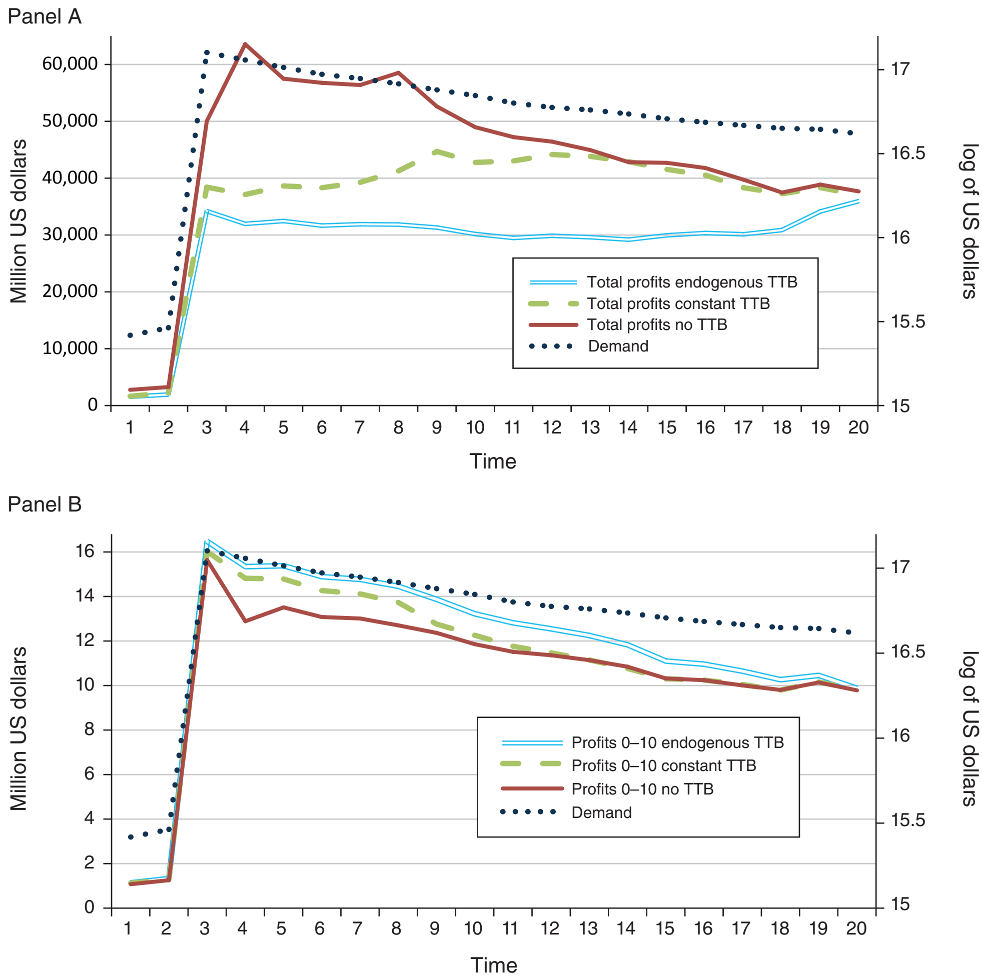

Figure 18. Total (*sum of*) and per Ship Profits for All Age Groups under Endogenous, Constant, and No Time to Build (TTB)

Note: Demand depicted on right axis.

goal is to understand and quantify the impact of time to build on the fluctuations of investment and prices. We perform 1,000 simulations of our model, each 80 years (320 quarters) long. We use these long-run simulations to explore the properties of our model during both specific and average demand paths.

Table 4 compares the ratio of constant to endogenous time to build (CT/ET), as well as the ratio of no to endogenous time to build (NT/ET) in terms of the level and variance of the total fleet, entry and price, as well as their covariance with demand. We can think of entry as the flow investment and fleet as the stock.

We find that on average, across time and simulations, both entry and the total fleet are larger by about 16 percent in the absence of time to build and by about 1 percent in the case of pure construction lags: shutting down time to build has a sizable impact on the level of investment, while shutting down its endogenous variation has

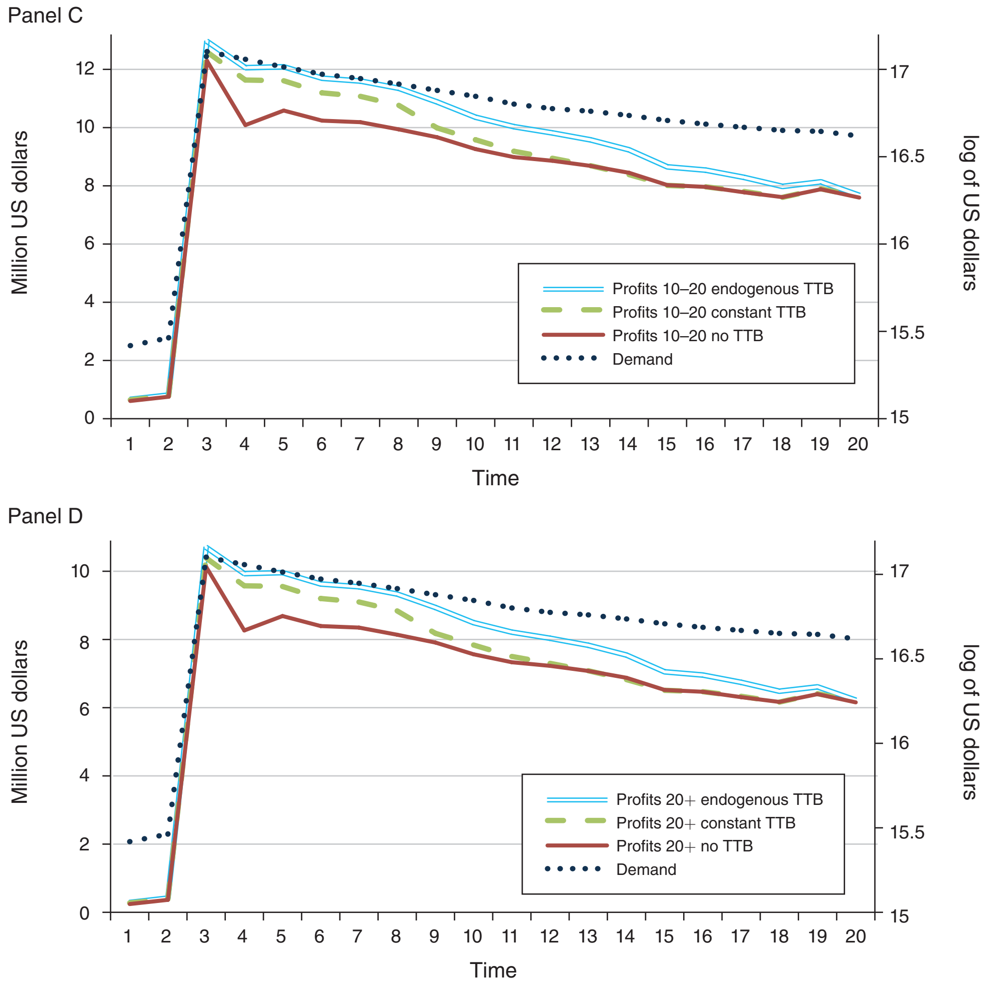

Figure 18. Total (*sum of*) and per Ship Profits for All Age Groups under Endogenous, Constant, and No Time to Build (TTB) (*Continued*)

Note: Demand depicted on right axis.

only a small effect on mean investment. This is reminiscent of work studying investment under uncertainty and irreversibility, which finds that incentives to invest are weakened. As a result of the larger fleet, shipping prices are lower in the absence of time to build. The fleet is also younger on average as time to build declines.

One of our main findings is that time to build has a dramatic impact on the volatility of investment. In the absence of time to build, entry—and therefore the fleet as well—is more responsive to demand and, as a result, prices are less responsive to demand. Indeed, we find that investment volatility is significantly higher as time to build declines: the fleet is 45 percent more volatile under constant time to build and twice more volatile under no time to build, while entry is twice more volatile under constant time to build and seven times more volatile in the absence of time to build. That entry is more volatile than the fleet is also consistent with macro magnitudes.

Table 4—Comparison of Ratio of Constant to Endogenous Time to Build (CT/ET)

	CT/ET	NT/ET
Fleet level	1.008	1.16
Price level	0.999	0.97
Entry level	1.016	1.165
Fleet variance	1.44	2.06
Price variance	0.98	0.86
Entry variance	2.11	7.64
Fleet and demand covariance	1.19	2.42
Price and demand covariance	0.995	0.945
Entry and demand covariance	1.19	1.763

Notes: Average across 1,000 simulations and 320 quarters. First 50 quarters removed.

In contrast, prices are less volatile as time to build declines, by 14 percent in the no time to build case.[33] Consistent with this finding, the covariance of the fleet and the demand process is larger in the absence of time to build, while the covariance of prices and demand is lower. At first the increased investment volatility in the absence of time to build may seem to contradict Kydland and Prescott (1982) whose work associated time to build with volatility. But Kydland and Prescott (1982) do not assess the sign of this relationship; rather, they document that time to build is critical in matching the properties of investment.

As suggested by the analysis of the impulse response function above, in the absence of time to build entry and exit respond instantly to positive demand shocks, while their response is also larger: entry is higher as firms can take advantage of the positive demand shock. Even though exit is also higher in the absence of time to build, due to the young firms coming in immediately, the entry effect dominates.[34] Consistent with hysteresis, higher adjustment costs render firms less likely to respond to demand shocks, leading to a smoother investment process.

Figure 19 illustrates the increased investment volatility in the absence of time to build by showing one specific path of the industry (this is a single typical simulation of the model). This simulation corroborates the above analysis. The fleet always follows the particular demand path; its response under time to build, however, is much smoother.

Exploring the comparative statics of time to build has the descriptive value of measuring the magnitude of investment adjustment costs per se, as well as their variability brought by industry evolution. In addition, time to build can indeed change

[33] Price volatility persists in a world without time to build mostly because of the presence of entry costs which limit entry under a mean reverting demand process. 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, if a shock is temporary, the fleet may not adjust fully, and the price increases. In addition, as revealed by the simulations, more voyages are traveled, as we move from endogenous to constant to no time to build. The number of voyages per ship, however, is smaller: under time to build the fewer operating ships offer more transportation. Therefore, in the presence of irreversible entry costs, depreciation, and demand mean reversion, time to build alone cannot account fully for price volatility.

[34] In contrast, the industry response to negative demand shocks depends on the backlog level when the shock hits. If the backlog is low enough, entry is higher in a world with time to build compared to one without, as demand will likely be higher upon delivery of the vessel. At the same time, exit may be lower under time to build due to the option value to waiting for the higher demand, when time to build restrains the new entrants.

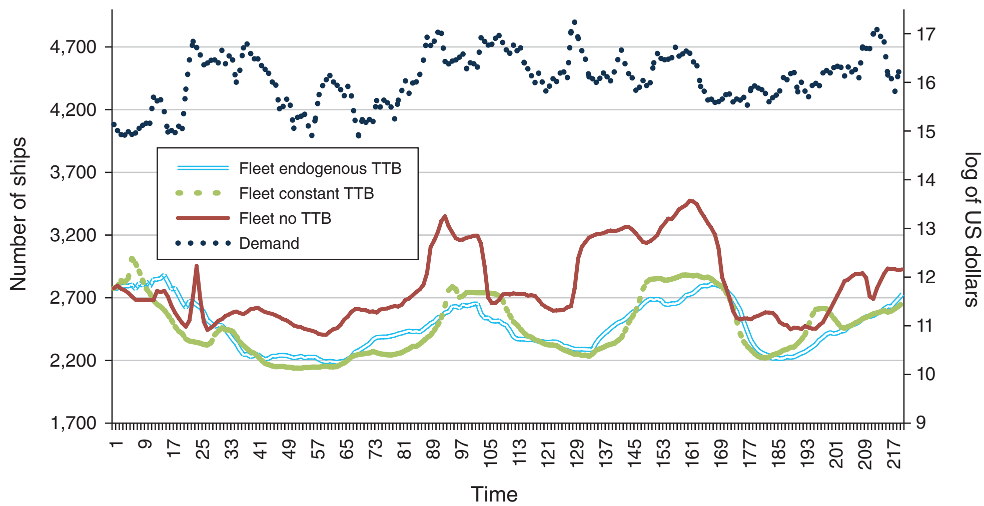

FIGURE 19. A SAMPLE INDUSTRY PATH

Note: Demand depicted on right axis.

via governmental subsidies to the shipbuilding industry, which are prevalent in East Asia and aimed at building more shipbuilding docks. Our results suggest that such subsidies impact not only first moments of investment and prices, but also second moments. The literature on irreversible investment under uncertainty usually considers comparative statics in demand uncertainty, in order to study the impact of government policies regarding uncertainty on aggregate investment. In our setup, there is no obvious way that institutions can alter uncertainty, as uncertainty stems from international trade flows and business cycle movement. Time to build, however, is of interest as it depends on firm behavior and can be manipulated in part by government subsidies. In addition, as the above mentioned literature stresses, uncertainty and adjustment costs are tightly linked and are both necessary to affect investment behavior.

We next examine the adjustment of the industry to its steady state (by averaging across simulations). As shown in Figure 20 demand has settled (after an initial adjustment period) to its long-run average level. In the absence of time to build, the total fleet approaches a constant value under small fading oscillations. In contrast, time to build leads to stronger oscillatory investment behavior which fades, albeit extremely slowly, with the fleet cycles in the constant time to build case leading those of the endogenous time to build case. A way to interpret these oscillations is via "echo effects," which refer to the following phenomenon: if a large number of ships enter at time t (say, because of an abnormally good shock) then it is likely that a large number of ships will enter again after a time period approximately equal to a ship's lifetime. This pattern of strong oscillations in the time to build worlds is consistent with the higher turnover of ships in the no time to build case. The higher exit rates under no time to build imply that the average life duration of a ship is lower. As mentioned above, the fleet is more responsive to demand changes in the absence of time to build, and, thus, any "echo effects" vanish relatively fast. Figure 21, which

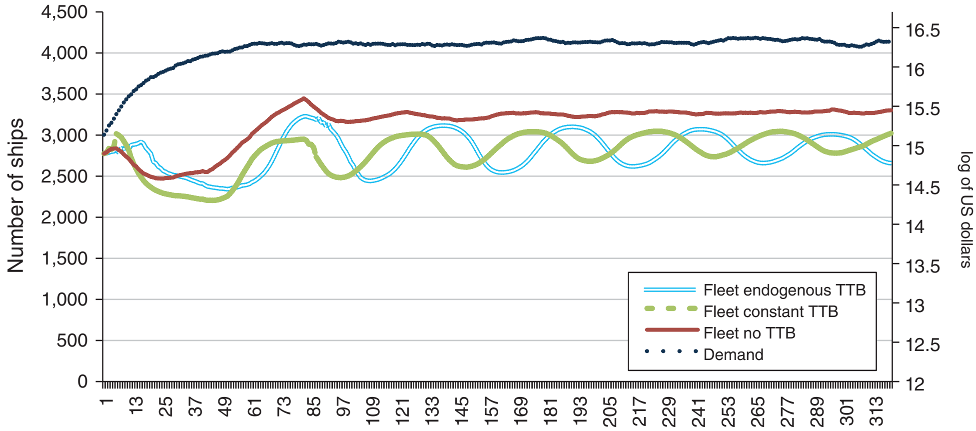

FIGURE 20. AVERAGE FLEET FROM LONG-RUN SIMULATIONS

Note: Demand depicted on right axis.

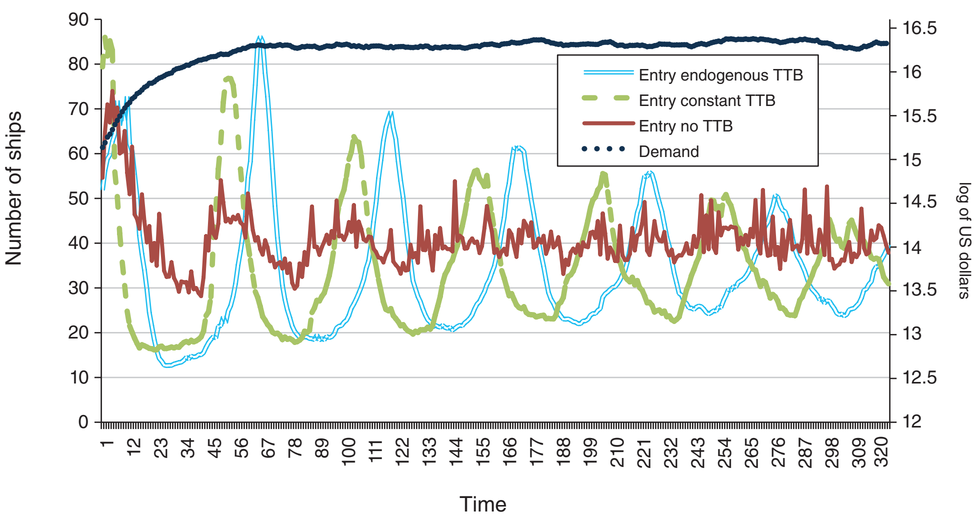

FIGURE 21. AVERAGE ENTRY FROM LONG-RUN SIMULATIONS

Note: Demand depicted on right axis.

depicts average entry, provides further evidence of this phenomenon: in the no time to build case entry happens in small and constant adjustments (note that the Poisson distribution creates further deviations). In contrast, under time to build entry is lumpy and happens in waves, thus leading to more pronounced "echo effects."

Appendix B further motivates and explains these results. We analyze the dynamics of the mean and covariance of the state vector. The analysis and findings are reminiscent of Rosen, Murphy, and Scheinkman (1994). We find that the matrix governing the dynamics of the mean state has eigenvalues whose norm is always smaller than one. Several eigenvalues are complex, indicating slowly fading oscillatory behavior.

VI. Conclusion

Adjustment costs and irreversibilities are key ingredients of leading macro models, as well as the theory of firm-level investment under uncertainty. This article empirically assesses their quantitative impact on an industry's dynamic evolution, in the context of oceanic bulk shipping. We find that investment volatility is significantly higher as time to build declines: the fleet is 45 percent more volatile under constant time to build and twice more volatile under no time to build, while entry is twice more volatile under constant time to build and seven times more volatile in the absence of time to build. We also find that the fleet is larger by about 15 percent in the absence of time to build.

A dynamic model of entry and exit was constructed and estimated using a rich dataset of secondhand ship sales, matched with a new estimation strategy. Our main estimating assumption is that the price of a ship in the secondhand market is equal to its value function. From there, we are able to uncover the primitives that govern the dynamic industry equilibrium. The proposed empirical strategy may be of use in other setups but is applied here on an industry that is characterized by severe demand variability and investment fluctuations. At the same time, this industry 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 (2010).

Appendix A

This Appendix presents the results from the first step of the proposed estimation method (see Section IVA).

Table 5—Inverse Demand Curve for Freight Transport: IV Regression Results

	First stage, dep. variable Q_t		Second stage, dep. variable P_t	
	Parameter	SE	Parameter	SE
const	2.01	(20)	−7.601	(23.8)
WIP	−5.05	(3.4)*	9.501	(4.51)**
agr raw mat P	1.291	(0.97)*	2.969	(1.32)**
mineral P	0.394	(0.57)	−1.658	(0.565)**
food P	−0.548	(0.715)	−0.346	(0.702)
China steel	0.365	(0.716)	1.534	(0.592)**
Handymax	−2.03	(2.12)	−4.705	(1.324)**
fleet	0.0013	(0.0014)		(0.597)
mean age fl	0.287	(0.150)**		
std age fl	0.5823	(0.335)**		
$\widehat{Q}_t$			−0.162	

**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table 6—Time to Build Regression Estimates

	Constant	S^1	S^2	S^3	B	d
Parameters	2.536	−0.00082	−0.00063	0.00011	$1.93e-005$	0.0303
Standard errors	(1.266)	(0.00058)	(0.00036)	(0.00036)	$(8.3e-005)$	(0.019)

Notes: Standard errors based on 500 bootstrap samples. Coefficients joint significant at the 0.01 level.

Table 7—Entry and Exit Regression Estimates

	Constant	S^1	S^2	S^3	B	B^2	d
Entry							
Parameters	−3.68	−0.00134	−0.00019	−0.00014	0.0069	$-5.37e-006$	0.486
Standard errors	(7.99)	(0.0049)	(0.0019)	(0.0018)	(0.0024)*	$(2.12e-006)$*	(0.22)*
Exit							
Parameters	22.68	0.0040	0.000178	0.00050	−0.0042	$4.141e-006$	−1.585
Standard errors	(10.15)*	(0.0065)	(0.0023)	(0.0023)	(0.0036)	$(3.09e-006)$	(0.34)**

Note: Standard errors based on 500 bootstrap samples.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table 8—Entry and Exit Regression Estimates

	Constant	S^1	S^2	S^3	d
Entry					
Parameters	−8.425	−0.0024	−0.00045		0.934
Standard errors	(4.90)	(0.0025)	(0.00075)		(0.244)**
Exit					
Parameters	22.728	0.0073	0.00093	0.00104	−1.859
Standard errors	(4.89)**	(0.0016)**	(0.00092)	(0.0008)	(0.242)**

Note: Standard errors based on 500 bootstrap samples.
**Significant at the 5 percent level.

Appendix B

In this Appendix, we analyze the dynamics of the mean of the state vector, in order to explore the oscillatory and steady-state behavior of the industry. The covariance has also been considered but is omitted due to space limitations.

Consider the state $\mathbf{x}_t \equiv \left[s_t^0, \dots, s_t^A,\, b_t^1, \dots, b_t^{\bar{T}},\, d_t\right] \in \mathbb{R}^{A+\bar{T}+2}$ and its transition as described in Section IVA, which we rewrite here as follows:

$$\mathbf{x}_{t+1} = \mathbf{V} + \mathbf{F}\mathbf{x}_t + \mathbf{G}\begin{bmatrix} N_t \\ Z_t \end{bmatrix} + \mathbf{h}\mathbf{e}_{t+1}. \tag{18}$$

The precise form of (18) depends on the time to build assumption (endogenous, constant, no). We focus on constant time to build of T periods; the no time to build case results when $T = 1$. The endogenous time to build case is harder to analyze because the matrices **F** and **G** vary with the state. Under the assumption of constant time to build, (18) takes the form

$$\mathbf{F} = \begin{pmatrix} A & 0 & B \\ 0 & C & 0 \\ 0 & 0 & A \\ 0 & 0 & \rho \end{pmatrix} \in \mathbb{R}^{(81+T+1)\times(81+T+1)},$$

where

$$\mathbf{A} = \begin{pmatrix} 0_{1\times 39} & 0_{39\times 1} \\ I_{39} & 0 \end{pmatrix} \qquad \mathbf{C} = \begin{pmatrix} 0_{1\times 39} & 0_{39\times 1} \\ I_{39} & 1 \end{pmatrix}$$

and $\mathbf{B} \in \mathbb{R}^{T\times T}$ is a zero matrix whose $(1, T)$ element equals 1, $\mathbf{v} = [0\ 0\ \cdots\ 0\ c]' \in \mathbb{R}^{81+T+1}$, $\mathbf{h} = e_1 \in \mathbb{R}^{81+T+1}$, $\mathbf{G} = [e_2\ \ -e_{81+T+1}] \in \mathbb{R}^{(81+T+1)\times 2}$ and $\mathbf{e}_i$ a vector which has all entries zero except for the ith entry, which equals one. **F** has multiple eigenvalues at 0, 1, ρ. Let $\mathbf{m}_t = E[\mathbf{x}_t \,|\, \mathbf{x}_0]$ for some initial condition $\mathbf{x}_0$. Then, after some manipulations we get that $m(0) = \mathbf{x}_0$ and

$$\mathbf{m}_{t+1} = \mathbf{v} + \mathbf{F}\mathbf{m}_t + \mathbf{G}\begin{bmatrix} E[\lambda(\mathbf{x}_t)\,|\,\mathbf{x}_0] \\ E[\mu(\mathbf{x}_t)\,|\,\mathbf{x}_0] \end{bmatrix}. \tag{19}$$

The stationary points of (19), $\mathbf{m}$, satisfy (19) for $\mathbf{m}_{t+1} = \mathbf{m}_t = \mathbf{m}$ and are of the form

$$\mathbf{m} = \left[a,\, a, \dots, a,\, a^{80},\, a, \dots, a, \frac{c}{1-\rho}\right],$$

where $a = E[\lambda(\mathbf{x}_t)\,|\,\mathbf{x}_0] = E[\mu(\mathbf{x}_t)\,|\,\mathbf{x}_0]$ and α^{80} solves $a = E[\mu(\mathbf{x}_t)\,|\,\mathbf{x}_0]$. For the counterfactual performed in Section V with $T = 6$ we find that the steady-state fleet which is equal to $80a + a^{80}$ is close to the mean fleet value obtained by the long-run simulations.

The behavior of (small) deviations from stationary points is explored by linearizing (19), using the functional forms $\lambda(\mathbf{x}) = e^{\gamma_0 + \gamma' \mathbf{x}}$ and $\mu(\mathbf{x}) = e^{\delta_0 + \delta' \mathbf{x}}$. The matrix that governs the dynamics of the linearized system is

$$\mathbf{F} + a\mathbf{GK}, \text{ where } \mathbf{K} = \begin{bmatrix} \gamma' \\ \delta' \end{bmatrix}.$$

The eigenvalues of $\mathbf{F} + a\mathbf{GK}$ determine the return of the system to the stationary point, the rate of convergence, and the oscillatory features. They also influence the behavior of variables that are a function of the state, such as the total fleet and prices. We compute the eigenvalues of $\mathbf{F} + a\mathbf{GK}$ using the estimated c, ρ, and the (γ, δ) found from the constant time to build counterfactual. $\mathbf{F} + a\mathbf{GK}$ has multiple complex eigenvalues leading to oscillatory behavior. The largest eigenvalue (in terms of its norm) is 0.99, which means that the system returns to steady state slowly, with fluctuations that fade very slowly. This finding is consistent with Figure 20. As the total fleet is a linear function of the state, we are able to isolate the eigenvalues of $\mathbf{F} + a\mathbf{GK}$ that determine its behavior; these are also complex, with the largest one equal to 0.98. Price is a nonlinear function of the state, and the dynamics of its mean and covariance are more complicated, unless first-order approximations are made.

REFERENCES

Abel, Andrew B. 1983. "Optimal Investment under Uncertainty." *American Economic Review* 73 (1): 228–33.

Abel, Andrew B., and Janice C. Eberly. 1994. "A Unified Model of Investment under Uncertainty." *American Economic Review* 84 (5): 1369–84.

Adland, Roar, and Steen Koekebakker. 2007. "Ship Valuation Using Cross-Sectional Sales Data: A Multivariate Non-parametric Approach." *Maritime Economics and Logistics* 9 (2): 105–18.

Aguirregabiria, Victor, and Pedro Mira. 2007. "Sequential Estimation of Dynamic Discrete Games." *Econometrica* 75 (1): 1–53.

Alessandria, George, Joseph P. Kaboski, and Virgiliu Midrigan. 2010. "Inventories, Lumpy Trade, and Large Devaluations." *American Economic Review* 100 (5): 2304–39.

Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin. 2007. "Estimating Dynamic Models of Imperfect Competition." *Econometrica* 75 (5): 1331–70.

Bar-Ilan, Avner, and William C. Strange. 1996. "Investment Lags." *American Economic Review* 86 (3): 610–22.

Beenstock, Michael, and Andreas Vergottis. 1989. "An Econometric Model of the World Market for Dry Cargo Freight and Shipping." *Applied Economics* 21 (3): 339–56.

Benkard, C. Lanier. 2004. "A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft." *Review of Economic Studies* 71 (3): 581–611.

Besanko, David, and Ulrich Doraszelski. 2004. "Capacity Dynamics and Endogenous Asymmetries in Firm Size." *RAND Journal of Economics* 35 (1): 23–49.

Bloom, Nicholas. 2009. "The Impact of Uncertainty Shocks." *Econometrica* 77 (3): 623–85.

Bloom, Nick, Stephen Bond, and John van Reenen. 2007. "Uncertainty and Investment Dynamics." *Review of Economic Studies* 74 (2): 391–415.

Caballero, Ricardo J. 1991. "On the Sign of the Investment-Uncertainty Relationship." *American Economic Review* 81 (1): 279–88.

Caballero, Ricardo J. 1999. "Aggregate Investment." In *Handbook of Macroeconomics*. Vol. 1B, edited by John B. Taylor and Michael Woodford, 813–62. Amsterdam: Elsevier, North-Holland.

Caballero, Ricardo J., and Eduardo M. R. A. Engel. 1999. "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S, s) Approach." *Econometrica* 67 (4): 783–826.

Caballero, Ricardo J., and Robert S. Pindyck. 1996. "Uncertainty, Investment, and Industry Evolution." *International Economic Review* 37 (3): 641–62.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113 (1): 1–45.

Clarksons Research. 1998–2010. Shipping Intelligence Network. http://www.clarksons.net/sin2010/default.aspx (accessed July 2010).

Collard-Wexler, Allan. 2013. " Demand Fluctuations in the Ready-Mix Concrete Industry." *Econometrica* 81 (3): 1003–37.

Dixit, Avinash. 1992. "Investment and Hysteresis." *Journal of Economic Perspectives* 6 (1): 107–32.

Doraszelski, Ulrich, and Mark Satterthwaite. 2010. "Computable Markov-Perfect Industry Dynamics." *RAND Journal of Economics* 41 (2): 215–43.

Dunne, Timothy, Mark Roberts, Shawn Klimek, and Yi Xu. 2009. "The Dynamics of Market Structure and Market Size in Two Health Services Industries." In *Producer Dynamics: New Evidence from Micro Data*, edited by Timothy Dunne, J. Bradford Jensen, and Mark J. Roberts, 303–28. Chicago: University of Chicago Press.

Ericson, Richard, and Ariel Pakes. 1995. "Markov-Perfect Industry Dynamics: A Framework for Empirical Work." *Review of Economic Studies* 62 (1): 53–82.

Fan, Jianqing, and Irène Gijbels. 1996. *Local Polynomial Modelling and Its Applications*. Boca Raton, FL: Chapman & Hall/CRC.

Gavazza, Alessandro. 2011. "Leasing and Secondary Markets: Theory and Evidence from Commercial Aircraft." *Journal of Political Economy* 119 (2): 325–77.

Hamilton, James D. 1994. *Time Series Analysis*. Princeton, NJ: Princeton University Press.

Hamilton, James D. 2005. "What's Real about the Business Cycle?" *Federal Reserve Bank of St. Louis Review* 87 (4): 435–52.

Hogg, Robert V., Allen T. Craig, and Joseph McKean. 2004. *Introduction to Mathematical Statistics*. Upper Saddle River, NJ: Pearson Education, Prentice Hall.

Hopenhayn, Hugo A. 1992. "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica* 60 (5): 1127–50.

Jofre-Bonet, Mireia, and Martin Pesendorfer. 2003. "Estimation of a Dynamic Auction Game." *Econometrica* 71 (5): 1443–89.

Kalouptsidi, Myrto. 2014. "Time to Build and Fluctuations in Bulk Shipping: Dataset." *American Economic Review*. http://dx.doi.org/10.1257/aer.104.2.564.

Kydland, Finn E., and Edward C. Prescott. 1982. "Time to Build and Aggregate Fluctuations." *Econometrica* 50 (6): 1345–70.

Lambson, Val Eugene. 1991. "Industry Evolution with Sunk Costs and Uncertain Market Conditions." *International Journal of Industrial Organization* 9 (2): 171–96.

Li, Qi, and Jeffrey Scott Racine. 2007. Nonparametric Econometrics: Theory and Practice. Princeton, NJ: Princeton University Press.

Majd, Saman, and Robert S. Pindyck. 1987. "Time to Build, Option Value, and Investment Decisions." *Journal of Financial Economics* 18 (1): 7–27.

Maskin, Eric, and Jean Tirole. 1988a. "A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs." *Econometrica* 56 (3): 549–69.

Maskin, Eric, and Jean Tirole. 1988b. "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles." *Econometrica* 56 (3): 571–99.

Pakes, Ariel, Michael Ostrovsky, and Steven Berry. 2007. "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples)." *RAND Journal of Economics* 38 (2): 373–99.

Pesendorfer, Martin, and Philipp Schmidt-Dengler. 2008. "Asymptotic Least Squares Estimators for Dynamic Games." *Review of Economic Studies* 75 (3): 901–28.

Pindyck, Robert S. 1991. "Irreversibility, Uncertainty, and Investment." National Bureau of Economic Research Working Paper 3307.

Pindyck, Robert S. 1993a. "Investments of Uncertain Cost." *Journal of Financial Economics* 34 (1): 53–76.

Pindyck, Robert S. 1993b. "A Note on Competitive Investment under Uncertainty." *American Economic Review* 83 (1): 273–77.

Rosen, Sherwin, Kevin M. Murphy, and Jose A. Scheinkman. 1994. "Cattle Cycles." *Journal of Political Economy* 102 (3): 468–92.

Ryan, Stephen P. 2012. "The Costs of Environmental Regulation in a Concentrated Industry." *Econometrica* 80 (3): 1019–61.

Smets, Frank, and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97 (3): 586–606.

Snider, Connan. 2008. "Predatory Incentives and Predation Policy: The American Airlines Case." Unpublished.

Stopford, Martin. 1997. *Maritime Economics*. 2nd ed. New York: Routledge.

Tinbergen, Jan. 1931. "Ein Schiffbauzyklus?" *Weltwirtschaftliches Archiv* 34: 152–64.

United Nations Conference on Trade and Development (**UNCTAD**)**.** 2009. *Review of Maritime Transport*. New York: United Nations Publications.

Weintraub, Gabriel Y., C. Lanier Benkard, and Benjamin Van Roy. 2008. "Markov Perfect Industry Dynamics with Many Firms." *Econometrica* 76 (6): 1375–411.

Weintraub, Gabriel Y., C. Lanier Benkard, and Benjamin Van Roy. 2009. "Computational Methods for Oblivious Equilibrium." Unpublished

Xu, Daniel Yi. 2008. "A Structural Empirical Model of R&D, Firm Heterogeneity and Industry Evolution." Unpublished.