

Waves in Ship Prices and Investment^{*}

Robin Greenwood and Samuel G. Hanson

Harvard Business School

First draft: February 2013

Abstract

We study the returns to owning dry bulk cargo ships. Ship earnings exhibit a high degree of mean reversion, driven by industry participants' competitive investment responses to increases in demand. This mean reversion is not fully reflected in ship prices. We show that high current ship earnings are associated with high secondhand ship prices and heightened industry investment, but forecast low future returns. We suggest and estimate a behavioral model that can account for the evidence. In our model, individual firms overestimate their ability to respond to common shocks and underestimate the ability of the competition, leading to excessive industry investment during booms and low subsequent returns on capital. Our model nests both rational expectations at one extreme and Kaldor's (1938) cobweb theory at the other, in which producers naively set current production quantities based on lagged prices. Formal estimation of our model suggests significant competition neglect in the shipping industry.

^{*} We thank Andreas Beroutsos, Ben Esty, Fritz Foley, Fotis Giannakoulis, Charlie Norz, David Scharfstein, Andrei Shleifer, Erik Stafford, Jeremy Stein, Martin Stopford, Larry Summers, Adi Sunderam and seminar participants at Harvard for helpful suggestions. We are grateful to Jared Dourdeville for excellent research assistance and to Clarkson Research for providing access to data.

I. Introduction

Dry bulk shipping is a highly volatile and cyclical industry in which earnings, investment, and returns on capital appear in waves. In 2001, a 5-year old “Panamax” ship commanded daily lease rates of \$5,325 and could be purchased for \$14 million. By December 2007, daily lease rates had grown more than tenfold to \$61,000, and purchase prices had risen more than fivefold to \$89 million. By 2011, lease rates and secondhand prices had nearly returned to their 2001 levels. This boom-bust cycle occurred alongside enormous fluctuations in industry investment. In December 2001, outstanding orders for new ships amounted to less than 10% of the active fleet. However, by August 2008, outstanding orders for new ships exceeded 75% of the active fleet.

We study how cycles in investment, lease rates, and secondhand prices are connected to predictable variation in the returns to ship owners. Using monthly data on secondhand ship prices and ship lease rates between 1975 and 2011, we measure payoffs to an investor who purchased a dry bulk cargo ship, operated it for a period of time, earning a dividend stream in the form of market-based lease rates, and later sold the ship in the second-hand market. The annual returns to owning and operating a ship vary massively over time, from a low of -68% between December 2007 and December 2008 to a high of +87% between June 1978 and June 1979.

We show that returns to owning and operating a ship are predictable and closely related to industry-wide investment in capacity. High current ship earnings are associated with higher ship prices and higher industry investment, but predict low future returns on capital. Conversely, high levels of industry scrapping—a measure of disinvestment—forecast high returns. The economic magnitude of return predictability is enormous: our baseline regressions suggest *expected* two year forward excess returns range from -68% to +40% between 1976 and 2011.

One should not be surprised that ship charter rates—the payment that a ship owner earns for leasing out his ship—fluctuate significantly over time. The supply of bulk carriers is essentially fixed in the short-run, because building and delivering a new ship takes between one and three years. Coupled with inelastic demand for shipping services, this time-to-build problem means that temporary imbalances between global demand for shipping services and the size of the fleet can lead to large changes in ship charter rates.

At the same time, in competitive purchase and sale markets, fluctuations in short-term charter rates do not imply anything about the expected returns to owning and operating a ship.

Consider the natural benchmark in which the returns *required* by ship owners are constant over time. What is the competitive response to an unexpected spike in shipping demand in such a setting? Charter rates would temporarily spike, raising contemporaneous realized returns. But ship investors would respond to the jump in charter rates by building additional ships. Over time, charter rates would fall back to their steady-state level as additional ships were added to the fleet. The prices of used ships—which are long-lived capital assets—would initially jump modestly in anticipation of heightened near-term cash flows, before gradually returning to their steady-state level. In equilibrium, enough ships would be brought online each period to bring the expected return to investing in new ships down to the required return. In short, forward-looking rational behavior on the part of ship owners and competitive industry dynamics would ensure that earnings exhibit a high degree of mean reversion but that expected returns were constant. Rosen, Murphy, and Scheinkman (1994) explore dynamics of this sort in their model of cattle cycles.

Although this rational expectations benchmark is appealing, the data suggest a more complex story. In a boom, the subsequent glut of shipping supply pushes charter rates and second-hand ship prices below the rational-expectations level, resulting in low realized returns to ship owners. Thus, a simple calculation based on constant discount rates suggests that at peaks, market participants may have overpaid for ships by more than 50 percent. Conversely, in downturns, there is excess capacity, many ships are scrapped, and future ship returns are high.

What can explain these results? A first possibility is that the return investors *require* for owning and operating ships varies significantly time, perhaps because ships are priced by diversified investors whose overall risk exposure or attitudes toward risk vary over the macroeconomic cycle. In this case, investment is high during booms, but prices are fair because investors require far lower returns going forward. Owners fully expect charter rates to fall as much as they do and fully expect holding period returns to be low. While one can never fully rule out this type of explanation, we show that the returns to ship owners are not very correlated with the returns on US stocks or bonds. Nor do our predictors of ship returns correlate strongly with traditional predictors of the returns on the stock market. Thus, it seems difficult to reconcile the predictability we document in ship returns with straightforward macro-finance explanations in which the risk-premia on risky assets asset vary over the business cycle.

We explore a simple behavioral explanation in which there is too much investment in capacity during booms because ship owners incorrectly believe that investments will continue to reap high returns. Motivated by extensive evidence from psychology and accounts from industry experts, we suggest that industry participants exhibit a form of overconfidence known as “competition neglect” (Camerer and Lovo (1999) and Kahneman (2011)). According to laboratory experiments, subjects in competitive settings repeatedly overestimate their own skill and speed in responding to common observable shocks, and underestimate the skill and responsiveness of their competitors. This leads to over-entry following common shocks that boost profitability. As Kahneman (2011) notes, competition neglect can be particularly strong when investors receive delayed feedback about the consequences of their decisions, as one would expect in industries such as shipping that are characterized by a significant time to build.

We formalize the logic of competition neglect by adapting a simple q -theory model of investment dynamics in a competitive industry. We first explore the rational expectations benchmark, in which equilibrium implicitly requires each firm to perfectly anticipate the investment response of other firms. In the rational benchmark, prices and investment evolve precisely as described above.

We introduce competition neglect into the model by assuming each firm underestimates the investment response of its rivals by a constant proportion. This implies that positive shocks to shipping demand generate excessive investment responses. This investment predictably depresses future charter rates and ship prices, leading prices to overshoot their rational-expectations levels. In this model, even though the required return of shipping investors is constant, investors’ tendency to underestimate the competition generates variation in expected returns. We show that the dynamics generated by this model are qualitatively similar to a model in which investors simply over-extrapolate demand. However, as we explain, both biases may be quantitatively useful for explaining competitive industry dynamics.

Our model is closely related to Kaldor’s (1934, 1938) cobweb theory of industry cycles, which turns out to be a special case of our model. According to the cobweb theory, producers naïvely set current quantities based on lagged prices, generating oscillations in price and quantity that—under some conditions—converge to a steady state. In our model, this naïve cobweb model obtains when producers *completely* neglecting the competition. However, even with moderate

competition neglect, we show that our model can generate significant excess volatility in ship earnings, secondhand prices, and investment along with the attendant return predictability.

The model can fit many of the facts we establish about return predictability and investment in the dry bulk shipping industry. We estimate the model using the simulated method of moments (McFadden (1989)). The main parameter of interest is the degree of “competition awareness” among market participants, which we estimate to be about 50%. This means that ship owners do not fully anticipate the investment response of their peers when reacting to demand shocks. At the same time, competition neglect is not complete, meaning that market participants are somewhat forward looking. Thus, our results suggest that industry dynamics lie roughly halfway between the naïve cobweb model and the perfect foresight rational expectations model.

Our paper is closely related to the cobweb model of industry cycles, first outlined by Kaldor (1934, 1938) but explored empirically by many others (Freeman (1975), and Rosen, Murphy, and Scheinkman (1994))¹. Our main contribution to this literature is to show how price and investment dynamics are connected to return predictability. Our findings are also related to an extensive literature in asset pricing that documents return predictability in stock and bond markets (Campbell and Shiller (1989), Fama and French (2006), Cochrane (2010)). In contrast with this literature, we study the investment returns on physical capital directly as opposed to the returns on financial claims on corporate cash flows. Our work is also related to recent papers on industrial structure and stock returns (Hou and Robinson (2006) and Hoberg and Philips (2010)), especially Hoberg and Philips (2010) who study the role of industry competition in driving the value premium among US stocks. Finally, we draw on Stopford (2009) and Kalouptside’s (2012) useful accounts of the dry bulk shipping industry.

The next section provides historical background on the dry bulk carrier industry and summarizes our data. Section III summarizes the correlations between earnings, secondhand prices, investment, and future returns. These correlations motivate a model of industry investment and price cycles, which we develop in Section IV. Section V estimates the model using the simulated method of moments. Section VI concludes.

¹ See also Ezekiel (1938), Muth (1961), and Nerlove (1961).

II. Dry Bulk Carriers: Earnings, Prices, and Investment

Solid commodities such as grain and iron ore are transported globally in large cargo ships known as bulk carriers. In 2011, bulk carriers made up approximately 40% of the world's shipping fleet (tankers and container ships made up most of the rest), and had a combined cargo capacity of 609 million deadweight tonnes (DWT) across 8,868 ships. These ships had a combined market value of roughly \$180 billion, having peaked at \$549 billion in June 2008. The market for shipping cargo is highly competitive with hundreds of firms operating bulker ships, and no single firm owning more than a few percent of the fleet.² In 2011, more than half of bulk carriers were held by Greek, Japanese, or Chinese firms.

We obtain time-series data on the dry bulk shipping market from Clarkson Research's Shipping Intelligence Network, the leading provider of shipping data to market participants. Starting in 1976, Clarkson provides monthly estimates of short-term trip charter rates by route, time charter rates, and second hand prices for various types and vintages of ship. In addition, we obtain monthly information on the composition of the fleet, as well as data breaking out new additions to the fleet and scrapping decisions. Beginning in 1996, Clarkson provides the complete order book, which includes the number of ships on order, deliveries, and cancellations.

In 2011, bulker fleet capacity was split between smaller ships (Handymax and Handysize) with total capacity of 125 million DWT, medium-sized ships (Panamax) with a fleet-wide capacity of 154 million DWT, and larger ships (Capesize) with a total capacity of 244 million DWT.³ As shown in Figure 1, the share of bulk cargo carried by larger ships (Panamax and Capesize) has grown since the 1970s. However, beyond secular shifts in composition, investment across different ship types has been highly synchronized over time. Panel B of Figure 1 shows that if we define investment as the 2-year percentage change in fleet capacity, there is a high correlation between investment across different ship types.⁴

² In 2006, the top 19 owners excluding COSCO (which was owned by the Chinese government) controlled about 20% of total industry capacity (Bornozis (2006)).

³ DWT is measured in metric tonnes. A metric tonne is 1,000 kilograms or roughly 2,205 pounds. Panamax carriers are so called because they are small enough to transverse the Panama Canal. Capesize ships are too large to traverse the Panama Canal and must round Cape Horn to travel between the Pacific and Atlantic oceans.

⁴ This is especially true after 1990: between 1990 and 2011, the time-series correlation between the percentage changes in Panamax and Capesize fleet sizes was 62% (or 73% between Panamax and Handysize fleet sizes).

In our analysis below, we measure investment fleet-wide, so as to abstract away from changes in the composition of the global cargo fleet (and thus arbitrage between different ship types). Thus, we implicitly assume that different ship types are very close substitutes in the service they provide—i.e., two 35,000 DWT ships can satisfy the same demand for shipping services as one 70,000 DWT ship. Under this assumption, total fleet capacity is the relevant quantity to relate to ship earnings and prices.⁵

A. Earnings, prices, and value

Ship owners earn income either by transporting cargo for hire or by leasing out their vessels for a defined period of time in the “time charter” market. In this market, which is organized by a large network of brokers, a charterer pays the ship owner a daily hire rate for the entire length of the contract, which is typically 12 months. The owner furnishes the charterer with the ship and must pay the costs of the crew and periodic capital expenditures,⁶ but the remaining costs, including fuel, are borne by the charterer. In computing earnings and holding period returns, we assume that the owner leases the ship out rather than operating it directly.

Panel A shows our time-series data on earnings, defined as the (real 2011) dollars expected to be earned net of costs by leasing the ship over the next year. Clarkson provides us with monthly estimates of the one-year charter rate for many different ships, based on their polling of brokers in the market as well as recent transactions.⁷ We focus on the 76,000 DWT Panamax carrier. We use this ship because it is a fairly representative bulk carrier—neither the smallest nor the largest vessel—and because we can construct consistent time series on both earnings and secondhand prices for ships of this specific size over our full 1976-2011 sample.

⁵ While surely an approximation, carriers are seen as providing a highly homogenous service. And consistent with this assumption, the Internet Appendix shows that earnings and secondhand prices are nearly perfectly positively correlated across different ships in the time-series. For example, the real price of a 5-year old Capesize vessel (a larger boat) is 97% correlated with the real price of Panamax carriers that we consider.

⁶ The main cost associated with maintenance is the opportunity cost of not using the ship. A 5-year old ship will be docked approximately 8-days per year, with older ships requiring longer.

⁷ To verify the reliability of the Clarkson data, we obtained micro data on time charter rates and secondhand prices for a limited sample of transactions between 2009 and 2012. Monthly averages of time charter rates were 98.2% correlated with the time charter time series we use from Clarkson. Individual 5-year old Panamax sales for December 2009 to November 2012 were 99.8% correlated with the price index values.

For a 5-year old ship, the owner earns the charter rate for an average of 357 days per year; the boat is docked for maintenance for the remaining 8 days per year. Although the lessor pays fuel and insurance, the ship owner must provide a crew at a daily cost that we estimate to be approximately \$6,000 per day, adjusted for inflation (using the US CPI)⁸. Thus annual earnings in year t are given by

$$\pi_t \equiv 357 \cdot \text{DailyCharterRate}_t - 365 \cdot \text{DailyCrewCost}_t, \quad (1)$$

where both *DailyCharterRate* and *DailyCrewCost* are expressed in real 2011 dollars. Equation (1) is an approximation, but it is confirmed in case studies of the shipping industry by Stafford, Chao, and Luchs (2002) and Esty and Sheen (2012), as well as in due diligence we conducted with industry participants. The earnings data are summarized in Table I and shown in Figure 2.

As can be seen in Panel A of Figure 2, ship earnings are highly volatile. Recall that our earnings measure is based on 12-month charter rates, so we are already examining a smoothed, forward-looking version of shorter-term spot hire rates.⁹ Before 2002, annual real earnings had a monthly standard deviation of \$2.5 million, compared to a mean of \$3.7 million. Starting in 2002, volatility increased substantially to a monthly standard deviation of \$6.1 million. These estimates are consistent with longer historical analysis of ship earnings in Stopford (2009).

In addition to illustrating the high volatility of ship earnings, Panel A of Figure 2 also shows the high degree of mean reversion in earnings. Earnings are 96% correlated with earnings in the previous month, but only 61% correlated with earnings 6 months earlier, 24% correlated with earnings 12 months earlier, and uncorrelated after 16 months. This high degree of estimated mean reversion is not sensitive to the time-period in question.¹⁰

Panel A of Figure 2 also shows data on secondhand prices for 5-year old vessels. As for earnings, we rescale the price series by the US CPI so that it is expressed in real 2011 US dollars. As can be seen, the secondhand price tracks earnings extremely closely throughout the 1976-2011

⁸ This estimate is based on Stopford (2005) and conversations with ship owners. CPI potentially overstates the growth of crew costs because globalization has allowed ships to source crews from lower wage countries over time.

⁹ We do not have spot hire rates for the 76,000 DWT Panamax carrier. In principle, one could complete the exercise using spot rates, but, because rates are determined in a broker-intermediated market, we would have to make additional assumptions regarding how many days of the year a ship was unmatched to a potential lessor.

¹⁰ For example, pre-1995, earnings in month t are uncorrelated with earnings in month $t+20$. Post-1995, the mean reversion appears to be somewhat faster, with earnings in month t being uncorrelated with earnings in month $t+15$.

period; the correlation is 89%. Panel B shows that although earnings and prices are correlated, the ratio of earnings to price is far from constant. When earnings are high, ship prices are also high but prices do not rise quite as quickly, leading to lower price earnings multiples. This is what one would expect if firms understand that the level of earnings is mean reverting.¹¹

It is difficult to evaluate the apparent volatility in ship pricing without first considering a benchmark in which discount rates are constant. In the spirit of Shiller (1981), Figure 3 plots the actual time series of market ship prices versus a simple model-implied present value of a ship earnings based on a constant 11% real discount rate.¹² To calculate the present value, we assume that the buyer of a 5-year old ship receives the current charter rate for 12 months following purchase, and then signs a new charter for another 12 months thereafter. We estimate the rate on this new charter based on the full time series autocorrelation of charter rates (0.24 at a one-year horizon, from Table 1). This initial two year period, we assume the buyer receives the average real earnings calculated over the full time series. We make a proportional adjustment once the ship is 15 years old, reducing charter rates by 15%, because older ships tend to lease at lower prices (Stopford (2005)). Finally, we assume that ships have a useful economic life of 25 years, so the 5-year old ship will be scrapped in 20 years and the owner will receive a scrap value. The complete details of this calculation are provided in the Internet Appendix.

This present value calculation is meant to capture the mean reversion in earnings, and based on an 11% discount rate, the average model-implied present value of a ship (\$32.9 million) approximately matches the time-series average market price of \$32.7 million. Figure 3 shows that the model-implied PV of the cash flows from a ship are considerably less volatile than actual ship prices. Consistent with Shiller (1981) and subsequent work on the excess volatility of asset prices (e.g., Campbell (1991)), the standard deviation of model-implied present values is \$2.4 million, compared with a standard deviation of \$15.28 million for used ship market prices. This discrepancy is driven by the fact that the present value calculation is not very responsive to changes in current earnings, which are expected to be almost completely reverted away one year later. In contrast,

¹¹ To illustrate, the growing perpetuity formula says that $\pi/P = r - g$ where r is the required return and g is the expected earnings growth rate. Thus, if expected earnings growth is low when the level of earnings is high, we would expect π/P to also be high at these times.

¹² We use an 11% discount rate so our average model NPV estimate matches the average used ship price in sample.

actual market prices are extremely responsive to current ship earnings. Taken together, this suggests that investors value ships as if they anticipate considerably less mean reversion in earnings than there has been in the actual data.

B. Returns

Using earnings and prices, we can construct holding period returns. The one-year holding period return on a ship is the one-year change in the secondhand price, plus the earnings accruing to an owner who signed a one-year lease immediately after purchasing the ship, scaled by the period t secondhand price:

$$R_{t+1} = (P_{t+1} - P_t + \pi_t) / P_t, \quad (2)$$

where P is the secondhand ships price and π is defined according to equation (1). We use secondhand prices instead of new prices because a buyer of a secondhand ship has immediate access to the ship and thus rental income.¹³ Following the convention in modern asset pricing, we use log excess returns as our main dependent variable. Log excess returns are defined as $rx_{t+1} = \log(1 + R_{t+1}) - \log(1 + R_{f,t+1})$.

Multi-year returns can be computed according to (2), except that we must have a multi-year estimate of π . Here, we assume that after each year, the ship owner signs a new one-year time charter at the prevailing rates. Thus, we can compute 2- and 3- year cumulative log excess returns by simply summing 1-year log excess returns. Table 1 shows that holding period returns are incredibly volatile: average one-year returns are 10%, with a standard deviation of 30%.

C. Investment plans: The order book

At the firm-level, investment occurs when a secondhand ship is bought from another owner or a new ship is ordered from a shipyard. At the industry-level, net investment occurs only when a new ship is ordered. Beginning in 1996, Clarkson provides monthly data on the dry bulk order book, which is the ledger of ships that have been ordered at shipyards around the world. The order book evolves according to

$$Orders_t = Orders_{t-1} + Contracting_{t-1} - Deliveries_{t-1} - Cancel_{t-1}. \quad (4)$$

¹³ A buyer of a new ship typically must wait 18-36 months for delivery, depending on workflow at shipyards. Thus, a buyer could be justified in paying a higher price for a used than a new ship when current charter rates are high. Such a dynamic occurred in 2007-2008.

Thus, the change in the order book equals new order in each month (*Contracting*), minus ships delivered in that month (*Deliveries*), minus previous orders that were cancelled (*Cancel*). All items in equation (4) are measured in DWT so as to capture changes in total industry-wide fleet capacity, ignoring shifts in the composition of ships of different sizes.

Based on equation (4), we construct three measures of investment plans, all scaled by current fleet size: contracting between months $t-12$ and t , which is a measure of gross new orders ($Contracting/Fleet_t$); the net change in orders in the past 12-months ($\Delta_{12}Orders_t/Fleet_t$); and the size of the current order book ($Orders_t/Fleet_t$). Contracting is 99% correlated with the net change in the order book, and 67% correlated with the size of the order book. The average size of the current order book is 28% of the fleet during the 1996-2011 period for which we have order book data.

D. Investment realizations: changes in fleet capacity

The fleet size evolves according to

$$Fleet_t = Fleet_{t-1} + Deliveries_{t-1} - Demolitions_{t-1} + Conversions_{t-1} - Losses_{t-1}, \quad (5)$$

with all terms expressed in units of DWT. Changes in the bulker fleet are primarily driven by deliveries (i.e., when new ships come online and fall out of the order book) and demolitions (i.e., when old ships are scrapped). *Conversions* and *Losses* capture rare incidents in which ships are repurposed from one use to another (i.e., a dry bulk carrier is converted to a tanker or vice versa) or ships are lost in accidents.

The *Deliveries* term in equation (5) represents the realization of past investment plans. Once ordered, a ship typically takes two to three years to be built and delivered (Kalouptsidei 2012). *Demolitions* are driven by the aging of the fleet—as ships become older, they become more costly to maintain and eventually they are no longer safe to use and must be scrapped. However, the demolition of an old ship may be postponed when current lease rates are high and accelerates when lease rates are low. Thus, aggregate *Demolitions* partially reflect a series of active disinvestment decisions made by shipping firms.

Panel A of Figure 4 shows fleet changes since 1974. The dashed line shows that the fleet has grown considerably over time from a capacity of 100 million DWT to just over 600 million DWT in December 2011, reflecting strong global economic growth and increased international trade. On average during our sample, the fleet increases by approximately 5% per year. The bars on the figure denote fleet additions and deletions. Not surprisingly given the growth of fleet capacity over time,

deliveries tend to outpace demolitions on average. However, there is significant cyclical-ity in both series, and an unprecedented surge in additions after 2006. Panel B shows one measure of investment, which is the difference between fleet additions and fleet deletions over the past two years, scaled by the current fleet size.

III. Predictability of Shipping Returns

We now investigate the relationship between current earnings, prices, and investment and the subsequent returns to ship owners. We adopt the standard asset pricing approach of using time-series variables to forecast returns. The idea is that, if (i) required returns are constant over time and (ii) ships are always fairly priced, then expected returns would equal required returns at each date and, hence, returns would be unpredictable. If instead returns are predictable, this must either be because ship owners have time-varying required returns that are linked to the forecasting variable, or because the forecasting variable is linked to temporary mispricing. By adopting this approach, we avoid having to construct a model-implied notion of fair value, as we did earlier when constructing Figure 3. After establishing in this section that returns are indeed predictable—i.e., that expected returns vary significant over time—we then ask whether this reflects time variation in investors’ required returns or whether it is more consistent with the correction of temporary mispricing.

We organize our empirical investigation around forecasting regressions of the form

$$rx_{t+k} = a + b \cdot X_t + u_{t+k}, \quad (6)$$

where rx_{t+k} denotes the k -year cumulative log excess return between t and $t+k$ and X_t denotes earnings, prices, or investment at time t . Recall that the k -year cumulative excess return is the total return (in excess of the risk-free rate) received by an investor who buys a ship in the secondhand market, collects earnings for k years, and then sells the ship again in k years time.

A. Earnings, prices, and earnings yields

We start by studying the relationship between current earnings and future returns. This relationship is illustrated in Panel A of Figure 5 where we plot current earnings versus future 2-year excess returns. The figure shows that when current real earnings are high, future returns are low. The figure shows real earnings and returns to ship investors over the next two years, but Table III reveals that this pattern holds over both shorter and longer holding periods. For 1-year returns shown in Panel A, the regression coefficient is -0.022. This means that a one standard deviation

increase in real earnings (approximately \$3.9 million, according to Table 1) leads to 8.58 percentage points decline in expected returns over the following year. The results for two-year returns are approximately twice that magnitude: a one standard deviation increase in earnings leads to 16.38 percentage point decline in expected returns over the next two years. The economic magnitudes are impressive given the mean and standard deviation of shipping returns (e.g., 1-year excess returns have a mean of 10.0% and a standard deviation of 30.4%).

The middle columns of Table III show that secondhand ship prices also negatively forecast future returns. This is perhaps not surprising given the strong positive correlation between prices and earnings shown in Table II. However, the economic magnitude of these results is stunning. A one standard deviation increase in the price (approximately \$15 million) is associated with a 17.9 percentage point reduction in future 1-year returns. At the peak price of \$91 million (in real terms) in August 2008, the regression implies that the expected excess return over the following two years was -40% (the realized 2-year excess return was quite close to that, at -50.5%).

Since both high prices and high earnings forecast low future returns, it is an open question as to whether the earnings yield, π/P , forecasts returns. In most asset pricing contexts, high current yields forecast high returns (e.g., Kojien, Moskowitz, Pedersen, and Vrugt (2012)). As explained by Campbell (1991), this stylized fact is a manifestation of basic present value logic in a dynamic setting. Loosely speaking, the growing perpetuity formula says that $\pi_t/P_t \approx r_{t+K} - g_{t+K}$ where r_{t+K} is the realized return and g_{t+K} is the realized earnings growth rate over some long horizon K . So a high earnings-price yield must either forecast high future returns or low future earnings growth. In most asset pricing settings, variation in current yields has little ability to forecast cash flow growth and, thus, is associated with higher future returns.

The right columns of Table III show that, if anything, the opposite result obtains for ships. When ships have high earnings relative to prices, this forecasts low future returns, albeit with weak statistical significance. While this result may seem surprising, we note that unlike the stock market where dividend payouts are highly persistent, ship earnings mean revert quickly. As noted by Cochrane (2008), this means that a high earnings yield must forecast low future earnings growth.¹⁴

¹⁴ Specifically, we must have $1 = Cov(r_{t+K}, \pi_t/P_t)/Var(\pi_t/P_t) - Cov(g_{t+K}, \pi_t/P_t)/Var(\pi_t/P_t)$. Since the first term, $Cov(r_{t+K}, \pi_t/P_t)/Var(\pi_t/P_t)$, is negative the second term, $Cov(g_{t+K}, \pi_t/P_t)/Var(\pi_t/P_t)$, must be highly negative in order for the identity to hold.

The bottom panel of Table II shows the same specifications from Panel A, except that we now include a time trend in the regression. There is no strong theoretical justification for including a time trend in a forecasting exercise of this type, but we do it here to check that our results are not because we mismeasure the trends in operating costs. Including a trend has little impact on the results. We have also repeated these regressions excluding the 2006-2011 “super-cycle” period (not shown). The regression coefficients are of comparable magnitude and statistical significance.¹⁵

B. Investment

We now ask how investment decisions are related to future returns. Recall that we have measures of investment plans (gross new orders, net new orders, and the full order book) available starting in 1996. Panel B of Figure 5 shows plots the time series of net orders of new ships, expressed as a percentage of the current fleet, together with the future 2-year excess returns on ships. The figure shows a negative correlation ($\rho = -0.34$). The corresponding regressions are shown in Panel A of Table III, where we run specifications of the form

$$rx_{t+k} = a + b \cdot X_t + c \cdot t + u_{t+k}. \quad (7)$$

Panel A shows the univariate results and Panel B controls for a potential time-trend as in (7). As can be seen in Panel A, whether we measure investment as gross new orders or the outstanding order book, industry investment negatively forecasts shipping returns in the subsequent years. For instance, consider the specifications shown in the first and second columns of Table III. The coefficient of -1.06 implies that a one standard deviation increase in gross new orders is associated with a 10.7 percentage point decline in returns over the next year, and a 14.3 percentage point decline over the next two years combined. The results are economically and statistically stronger when we include a time trend to account for the secular growth of the order book over time.

The biggest limitation of these regressions is that data on the order book only becomes available in 1996. For a much longer sample (1975-2011), however, we have monthly measures of *realized* changes in the total fleet capacity. These changes are directly driven by past investment decisions. Under the assumption that orders in the past 12 months translate approximately into

¹⁵ In the earnings regressions, the statistical significance is slightly weaker, although the coefficient estimates are slightly larger, while in the price regressions the results are stronger across the board. The coefficients in the earnings yield regressions are similar but with reduced statistical significance. We also obtain qualitatively similar results with comparable significance if instead we use log earnings, log prices, and the log earnings-to-price ratio to forecast returns.

deliveries in the next year, we can measure past investment decisions as being reflected in current deliveries. The advantage of the deliveries data, vis-à-vis the orders regressions already shown, is that we can analyze a much longer time-series. The disadvantage is that our measure suffers from some look-ahead bias. This look-ahead is only a problem insofar as the precise timing of deliveries may depend on demand realizations. From this perspective, industry scrappage, which measures disinvestment, is a cleaner measure. This is because a ship can be scrapped immediately once the decision has been made.

In any case, the forecasting regressions using deliveries and scrappage are shown in the right-columns of Table III. High current deliveries—our proxy for past investment decisions—are associated with low future returns and, conversely, high lagged scrappage is associated with high future returns. These results hold with and without the time-trend control.

C. Summary and discussion

We have shown that when charter rates are high, ships sell at high prices in the secondhand market, new ships are ordered at a higher rate, and the future returns to ship owners are low. These results are summarized by Figure 6, where we plot realized 2-year excess returns alongside the fitted values from our forecasting regressions in Table II and Table III. Using a seemingly-unrelated-regression approach, the joint statistical significance of these forecasting regressions at horizons of one- and two-years yields a p -value 0.006.¹⁶

Can we interpret the return predictability as evidence of collective mistakes? Direct evidence that industry participants are acting based on mistaken beliefs is not possible without polling market participants, and these polls are subject to their own issues (Greenwood and Shleifer (2012)). However, there is one piece of evidence to support the idea that industry participants make investment errors: cancellations of orders placed with shipyards are predictable in much the same way as returns. Specifically, we estimate regressions of the form

$$Cancellations_{[t+18,t+30]} = a + b \cdot X_t + c \cdot t + u_{t+18}. \quad (8)$$

¹⁶ We run six time-series regressions of $rx_{t+k} = a + b \cdot X_t + u_{t+k}$ for $k=12$ months, 24 months and for $X=\pi, P$, and *Deliveries*. We test the joint hypothesis that $b = 0$ in all regressions. Formally, we take a system OLS approach to these six regressions and estimate the joint variance-covariance matrix using a Newey-West-style estimator that allows residuals to be correlated both within and across moment equations at up to 36 months.

The dependent variable is ship order cancelations between 18- and 30-months following date t , alternately scaled by the current stock of ships outstanding, or by total orders over the past 12 months. These results are shown in Table IV. When earnings, used ship prices, or investment are high, we can predict cancelations of orders 18 to 30 months later. One interpretation of these results is that investment “regret” is predictable. That is, firms appear to be systematically surprised, and can be expected to later reverse investment decisions made during boom times.¹⁷

Another explanation for predictable returns, suggested by a large literature in asset pricing and macro-finance, is that the variation in expected returns to dry bulk ships is driven by changes in investors’ *required* returns. Under such an explanation, what appears to be excessive investment in capacity during booms would reflect ship owners’ willingness to invest at lower than usual expected returns. That is, owners would expect charter rates to fall as much as they do during the subsequent bust, and would expect their future holding period returns to be low. Although our goal is not to evaluate these explanations in detail, we are somewhat skeptical of such explanations in this context for a few reasons.

According to these explanations, the expected excess return on ships at time t is

$$E_t[rx_{t+1}] = Corr_t[rx_{t+1}, -m_{t+1}]\sigma_t[rx_{t+1}]\sigma_t[m_{t+1}], \quad (9)$$

where the real stochastic discount factor (SDF), m_{t+1} , depends on the marginal utility of well-diversified investors. Equation (9) shows that time variation in required returns must either be driven by a time-varying correlation between shipping returns and investor well-being ($Corr_t[rx_{t+1}, -m_{t+1}]$), time-variation in the risk of investing in ships ($\sigma_t[rx_{t+1}]$), or time-variation in the economy-wide price of risk ($\sigma_t[m_{t+1}]$). Before discussing these possibilities in turn it is worth noting that the variation in expected returns we document is very large from an economic point of view—from as low as -68% to as high as +40% over a two-year holding period. This suggests that risk-based explanations for our findings would need to invoke enormous time-variation in one or more of the terms on the right-hand-side of (9).

¹⁷ However, predictable order cancelations do not necessarily imply investment regret. Depending on the contract terms, ordering a ship might correspond to a rational forward-looking purchase of a real investment option. The fact that this option is not subsequently exercised (i.e., that the order was cancelled) or time-variation in overall exercise rates need not imply that the initial option purchase was mistaken or that the firm experienced *ex post* regret.

First, there is little reason to suspect that $Corr_t[rx_{t+1}, -m_{t+1}]$ varies significantly over time—i.e., that ships have time-varying hedge value for diversified investors. And, there is even less reason to believe this correlation varies in a way that could explain our results—e.g., there is little reason to believe that $Corr_t[rx_{t+1}, -m_{t+1}]$ is low when ship earnings, prices, and investment are high.

Turning to the second term in equation (9), a far more natural alternative to consider is that time-variation in $\sigma_t[rx_{t+1}]$ explains our results—e.g., future shipping risk might be low during shipping booms when earnings, prices, and investment are high. This hypothesis fails resoundingly in the data, because current earnings, used prices, and investment strongly forecast future *increases* in the risk of ship earnings and prices. Specifically, we estimate regressions of the form $\sigma_{t \rightarrow t+k} = a + b \cdot X_t + u_{t+k}$, where the dependent variable is the standard deviation of one-month earnings or returns, computed over the next year (untabulated).¹⁸ For example, when we use current real earnings for X_t (as in Table II, where we use earnings to forecast returns), $b = 0.27$ with a t -statistic of 4.78 and an R^2 of 0.45. In contrast, earnings, prices and investment forecast *lower* returns, suggesting that the relationship between risk and return is reversed in the time-series.

Finally, turning to the third term in equation (9), modern asset pricing theories suggest that the economy-wide risk premia investors require to hold risky assets (i.e., $\sigma_t[m_{t+1}]$) may fluctuate over the business cycle due to changes in either the aggregate quantity of risk or in investors' willingness to bear risk. We take a simple approach to assess whether these theories might plausibly explain our findings. Specifically, we ask whether expected and realized returns on ships are correlated with traditional risk premia measures and risk factors from the equity market. By doing so, we are effectively asking whether the time-variation in expected shipping returns documented above can be naturally explained by an omitted economy-wide factor.

We start, in Panel A of Table V, by showing our main forecasting regressions (e.g., equation (6)), but now include *ex ante* proxies for equity risk premia as control variables. Specifically, we include the dividend price ratio, the smoothed earnings yield, and the risk-free rate, and adopt the interpretation that these variables measure *ex ante* required returns. The first three columns show that these variables do not by themselves forecast the returns to owning ships, except for the risk-

¹⁸ Computing the volatility of earnings is straightforward. Computing the volatility of returns is more complicated, because our measure of returns assumes that a ship-owner signs a one-year time-charter at the start of each year.

free rate, which negatively forecasts returns. The remaining columns of Table V show that our results are not much affected—if anything they are stronger—when we include these controls.

In Panel B of Table V, we perform similar horse races, except that we now include *ex-post* realizations of equity risk factors, including the excess return on the market (*MKTRF*), the realized return on high book-to-market stocks over low *B/M* stocks (*HML*), the realized return on small stocks over big stocks (*SMB*), and the return to the momentum factor *MOM*. In other words, Panel B of Table V asks whether we can forecast the CAPM and 4-factor “alphas” from investing in ships. And, if our results were simply driven by time-variation in the economy-wide price of risk (i.e., $\sigma_t[m_{t+1}]$), we should have no ability to forecast alpha. Panel B of Table V shows that the realized returns on ships are not strongly tied to the returns on these traditional asset pricing factors. For instance, the 24-month excess stock returns on the US stock market are 9% correlated with our returns. The right six columns in Panel B suggest that controlling for contemporaneous returns on equity risk factors if anything tends to strengthen our forecast results. Thus, based on these traditional approaches to risk-adjustment, our regressions suggest that there is predictable time-series variation in the alpha from owning ships.

In summary, while it is difficult to decisively prove that industry participants made *ex ante* mistakes, the results in this subsection suggest that at a minimum, waves in ship prices, investment, and returns are driven by factors that are unique to the shipping market.

IV. A Cobweb Model of Competition Neglect

In behavioral theories of time-varying expected returns, investor misperceptions of fundamental value may cause them to overpay for assets, even if their required returns are constant. In a competitive industry such as shipping, there are two plausible candidate forces that may drive such misperceptions. First, investors may underestimate the rate of mean reversion of underlying shifts in demand. Second, investors may underestimate the effect that competition will have in returning cash flows back to their steady-state levels. Camerer and Lovallo (1999) call the latter “competition neglect.”

The main idea of competition neglect is that when confronted with some change in market conditions, participants in a competitive industry *should* ask themselves, “How should I respond given how I expect all of my competitors to respond?” This is a subtle and complex question.

Instead of answering it, participants answer the simpler question of how they should respond, assuming that no one else reacts. In this way, they are substituting a simple partial equilibrium question for the far more difficult general equilibrium question. We consider a model in which both mistakes are permitted. For clarity, our exposition emphasizes the role of competition neglect, but in the Internet Appendix we work out a model in which both investor errors play a role.

There is ample experimental evidence supporting the existence of competition neglect. According to Camerer and Lovo (1999), agents appear to overestimate their own skill and speed in responding to common observable shocks and underestimate the skill and speed of their competitors. In the laboratory, this phenomenon can be dramatic when entry involves significant time-to-build—because participants only receive delayed feedback about the consequences of their entry and investment decisions.

We formalize the logic of competition neglect using a simple q -theory model of investment dynamics in a competitive industry. At one extreme, the model nests the rational expectations benchmark in which managers perfectly forecast the competition. As we show, at the other extreme, the model reduces to a dynamic version of Kaldor’s (1938) cobweb model where firms *completely* neglect the competition.

A. Setup

The aggregate supply of ships is fixed in the short-term at Q_t . The inverse demand curve for shipping services at time t is $H(A_t, Q_t) = A_t - BQ_t$, where H_t denotes the rental or “hire” rate for a 1-period shipping charter. A higher value of B is associated with a more inelastic demand curve. We assume that the aggregate demand parameter, A_t , follows an $AR(1)$ process

$$A_{t+1} = \bar{A} + \rho(A_t - \bar{A}) + \varepsilon_{t+1}, \quad (10)$$

with $\rho \in (0, 1)$ and $Var[\varepsilon_{t+1}] = \sigma_\varepsilon^2$, so high values of A_t signify high demand for ship charters.

There is a unit measure of identical shipping firms that make investment decisions in discrete time. These firms act as price takers in the spot rental market for shipping services and in the secondary market for ships.

We consider the capital budgeting problem of a representative firm in the industry. The fleet size maintained by the representative firm, denoted q_t , evolves according to

$$q_{t+1} = (1 - \delta)q_t + i_t, \quad (11)$$

where $\delta \in (0,1)$ is the depreciation rate and i_t is firm investment at time t . Analogously, the aggregate fleet size, denoted Q_t , evolves according to

$$Q_{t+1} = (1 - \delta)Q_t + I_t, \quad (12)$$

where I_t is aggregate investment at time t .

A firm must choose q_{t+1} at time t *before* it learns the realization of the aggregate demand shock A_{t+1} which it learns at $t+1$. In this way, the model captures the time-to-build delays that are a critical aspect of dry bulk shipping and many other industries. Since the resulting aggregate supply of ships is fixed in the short-term at Q_{t+1} , ship hire rates can fluctuate significantly over time in response to temporary demand imbalances in the spot charter market. We are interested in understanding how these temporary imbalances are capitalized into the price of ships.

We model competition neglect by assuming that each firm believes that $I_t = \theta i_t$ where $\theta \in [0,1]$ measures *competition awareness*. Thus, each firm directionally anticipates how its competitors will respond to common shocks, but if $\theta < 1$, underestimate the magnitude of the response. If $\theta = 1$, firms have fully rational expectations about how competitors will respond.¹⁹

Because all firms are the same, competition neglect leads investment to overreact to common shocks that affect firm profitability. We use $E_t[\cdot]$ to denote the subjective expectations of individual firms, which believe that $I_t = \theta i_t$ where $\theta \in [0,1]$. By contrast, we use $E_t^0[\cdot]$ to denote the unbiased expectations of an econometrician who knows that $I_t = i_t$.

We assume the profits of the representative shipping firm in period t are given by

$$\Pi(q_t, i_t, A_t, Q_t) = q_t \pi_t - p_r i_t - \frac{k}{2} (i_t)^2 = q_t (A_t - BQ_t - C) - p_r i_t - \frac{k}{2} (i_t)^2. \quad (13)$$

The firm's fleet size is q_t . The rental price of a ship is $H(A_t, Q_t) = A_t - BQ_t$ and operating costs are C , so the firm earns a net profit of $\pi_t = A_t - BQ_t - C$ on each unit of its current fleet (i.e., on each unit of installed capital). The price of raw ship materials—i.e., the replacement cost—is p_r . In

¹⁹ Although we refer to competition neglect as a “mistake,” we do not rule out the possibility that the relevant decision makers are behaving rationally based on some set of distorted incentives—i.e., that competition neglect stems from agency considerations. From a corporate finance perspective, we also take no stand on whether these mistakes (or distortions) are made by the managers of shipping firms, the outside investors who provide financing to them, or both.

addition, we assume there are convex adjustment costs, $k(i_t)^2 / 2$. The adjustment cost parameter k is inversely related to the elasticity of firm supply.²⁰

One should think of the firms in our model as vertically-integrated firms that build and operate ships. The model abstracts from the fact that ships tends to be manufactured by one set of firms and sold to another set of firms that own and operate them. Thus, one can interpret the adjustment costs in our model as the combined costs of adjusting the scale of ship-building capacity as well as shipping operations.

B. The capital budgeting problem of the representative firm

Each firm chooses its current investment to maximize the expected net present value of earnings. Standard dynamic programming arguments (see the Internet Appendix) show that firm investment is given by the familiar q -theory investment equation

$$i_t^* = \frac{P(A_t, Q_t) - p_r}{k}, \quad (14)$$

where p_r is the replacement cost of a ship and the market price of the ship,

$$\begin{aligned} P(A_t, Q_t) &= \frac{E_t[\pi_{t+j} + (1-\delta)P(A_{t+1}, Q_{t+1}) | A_t, Q_t]}{1+r} \\ &= \frac{1}{1+r} \sum_{j=1}^{\infty} \left(\frac{1-\delta}{1+r}\right)^{j-1} E_t[\pi_{t+j} | A_t, Q_t], \\ &= \frac{1}{1+r} \sum_{j=1}^{\infty} \left(\frac{1-\delta}{1+r}\right)^{j-1} E_t[A_{t+j} - BQ_{t+j} - C | A_t, Q_t], \end{aligned} \quad (15)$$

is just the expected present value of earnings. Thus, as in any q -theory setting, firms invest when the price of new ships exceeds their replacement cost. Conversely, firms disinvest, scrapping some portion of their existing fleet, when the replacement cost exceeds the price of a ship.

C. Equilibrium investment and ship prices

²⁰ Convex adjustment costs ensure that each firm's optimization problem is concave and facilitates micro-founded aggregation. One can think of these costs as arising from technological constraints which lead to convex production costs or from frictional costs of adjusting firm scale. Alternately, one can also interpret increases in k as reflecting more severe time-to-build delays: investment responds more gradually to shifts in demand for higher values of k .

To solve for equilibrium investment and ship prices, we rewrite the future aggregate fleet size as a function of the current fleet size and future aggregate investment. Iterating on equation (11), we obtain

$$Q_{t+j} = (1-\delta)^j Q_t + \sum_{s=0}^{j-1} (1-\delta)^{j-1-s} I_{t+s}. \quad (16)$$

The aggregate fleet size (i.e., the aggregate capital stock) at time $t+j$ is the sum of the depreciated initial fleet size plus appropriately depreciated future investment. Using (15) and (16), we have

$$P(A_t, Q_t) = \frac{1}{1+r} \sum_{j=1}^{\infty} \left(\frac{1-\delta}{1+r} \right)^{j-1} E_t \left[A_{t+j} - B(1-\delta)^j Q_t - B \sum_{s=0}^{j-1} (1-\delta)^{j-1-s} I_{t+s} - C \mid A_t, Q_t \right]. \quad (17)$$

Equation (17) shows that ship prices and hence optimal firm investment depend on firms' expectations about current and future industry-wide investment. Naturally, all else equal, $P(A_t, Q_t)$ and i_t^* are increasing in current aggregate demand, decreasing in current industry fleet size, and decreasing in expected future industry-wide investment.

We now solve for the equilibrium investment policy of the representative shipping firm. We conjecture that investment is linear in the two state variables

$$i_t = x_i + y_i A_t + z_i Q_t. \quad (19)$$

We need to solve for the unknown coefficients, namely, x_i , y_i , and z_i . To do so, we make use of our assumption regarding competition neglect

$$\begin{aligned} E_t[I_{t+j} \mid A_t, Q_t] &= \theta E_t[i_{t+j} \mid A_t, Q_t] \\ &= \frac{\theta}{k} \left(E_t[P(A_{t+j}, Q_{t+j}) \mid A_t, Q_t] - p_r \right) \\ &= \theta \left(x_i + y_i E_t[A_{t+j} \mid A_t, Q_t] + z_i E_t[Q_{t+j} \mid A_t, Q_t] \right). \end{aligned} \quad (20)$$

If $\theta < 1$, individual firms underestimate the extent to which industry-wide investment reacts to aggregate demand and fleet size. Equation (20) also shows that our approach to modeling of competition neglect is equivalent to assuming that all firms have adjustment costs k , but that each believes that its competitors have costs k/θ . In other words, firm overconfidence takes the form of assuming that one is able to adjust to common shocks more nimbly than one's competitors.

In equilibrium, the representative firm optimally chooses its investment given a conjecture about industry-wide investment. When $\theta=1$, the solution corresponds to a recursive rational expectations equilibrium in which the firm's conjecture about industry-wide investment is the

precisely same as the actual level of industry investment (see e.g., Ljungqvist and Sargent (2004)). When $\theta < 1$, the solution is a *biased expectations equilibrium* in which the representative firm's conjecture about industry-wide investment equals θ times the actual level of industry investment.

Solving for equilibrium investment and prices leads to our first result.

Proposition 1 (Equilibrium investment and prices): *There exists a unique equilibrium such that the investment of the representative firm is $I_t^* = x_i^* + y_i^* A_t + z_i^* Q_t$ and equilibrium ship prices are $P_t^* = p_r + kx_i^* + ky_i^* A_t + kz_i^* Q_t$. Firm investment and ship prices are increasing in current shipping demand (i.e., $y_i^* > 0$) and decreasing in current aggregate fleet size (i.e., $z_i^* < 0$). The slope coefficients (i.e., y_i^* and z_i^*) are functions of six exogenous parameters: k , r , δ , ρ , θ , and B . In addition to these six parameters, the intercept term (i.e., x_i^*) also depends on \bar{A} , C , and p_r .*

In addition, we have:

- (i) *A demand shock raises prices and investment, but prices and investment react more aggressively when competition neglect is more severe (i.e., $\partial y_i^* / \partial \theta < 0$).*
- (ii) *A drop in fleet size raises prices and investment, but prices and investment react more aggressively when competition neglect is more severe (i.e., $\partial z_i^* / \partial \theta > 0$).*

Proof: See Internet Appendix.

Properties (i) and (ii) of Proposition 1 follow from the logic of competition neglect. When firms underestimate the speed with which their competitors will adjust, they tend to overreact to elevated ship charter rates, whether these are due to high current demand for ships or a low industry supply of ships.

The Internet Appendix further characterizes the equilibrium. In addition to the results described in Proposition 1, the model generates a number of intuitive comparative statics. Both investment and prices react more aggressively to the current fleet size and level of demand when the demand curve is more inelastic. Investment and prices react less aggressively to current fleet size and the level of demand when required returns are higher. When investment adjustment costs are higher, investment reacts less aggressively to the fleet size and current demand, but prices react more aggressively. Investment and prices reacts more aggressively to demand shocks when demand

is more persistent, but the response to aggregate supply is independent of ρ . Finally, investment and prices react more aggressively to current aggregate supply when capital is longer lived.

D. Cobweb dynamics and capital asset returns

To illustrate the key intuitions of the model, we study the special case where capital is infinitely lived and shifts in demand are permanent and deterministic, i.e., where $\delta = 0$, $\rho = 1$, and $\sigma_\varepsilon = 0$. For simplicity, we also assume $C = 0$. As we will see, the model takes on a particularly simple form in this case.

Given an initial demand A_0 , the steady-state fleet size is $Q^*(A_0) = (A_0 - rp_r) / B$, the steady-state rental rate is $H^* = (A_0 - BQ^*(A_0)) = rp_r$, and the steady-state ship price equals replacement cost $P^* = p_r$. Thus, the steady state rental rate enables capital to just earn its required return, so economic profits are zero in the steady state.

Suppose there is a one-time unexpected shock at $t = 0$ that permanently raises demand to $A_0 + \varepsilon$. The new steady state fleet size is $Q^*(A_0 + \varepsilon) = (A_0 + \varepsilon - rp_r) / B$ and the steady-state rental rate and ship price are unchanged. Letting H_t denote the hire rate in period t , the system evolves according to

$$Q_{t+1} = Q_t + \overbrace{(-z_t^* / B)(H_t - rp_r)}^{I_t} \text{ and } H_{t+1} = A_0 + \varepsilon - BQ_{t+1}, \quad (21)$$

where z_t^* is the negative root of $0 = -k\theta z_t^2 + (kr + B\theta)z_t + B$ or

$$z_t^* = \frac{kr + B\theta}{2k\theta} - \sqrt{\left(\frac{kr + B\theta}{2k\theta}\right)^2 + \frac{B}{k\theta}} < 0 \text{ and } \lim_{\theta \rightarrow 0} z_t^* = -B / (kr). \quad (22)$$

Since $z_t^* < 0$, investment is positive (negative) when hire rates are above (below) their steady-state.

Iterating on (21), it is easy to see that

$$Q_t = Q^*(Q_0) - \varepsilon(z_t^* / B) \left[\sum_{j=0}^{t-1} (1 + z_t^*)^j \right] \text{ and } H_t = rp_r + \varepsilon \left[1 + z_t^* \sum_{j=0}^{t-1} (1 + z_t^*)^j \right]. \quad (23)$$

Thus, if $|1 + z_t^*| < 1$, $\lim_{t \rightarrow \infty} Q_t = Q^*(Q_0 + \varepsilon)$ and $\lim_{t \rightarrow \infty} H_t = rp_r$ and the system converges to its new steady-state following the shock.

Figure 7 uses this simplified version of the model to contrast industry dynamics in three simple cases: the cobweb model ($\theta = 0$), rational expectations ($\theta = 1$), and partial competition

neglect ($0 < \theta < 1$). In Figure 7, the sequence of equilibrium (Q_t, H_t) pairs are marked with dots. Vertical movements in the figure show the determination of spot rental rates, H_t , given the current fleet size, Q_t . These movements are dictated by the demand curve. Lateral movements in the figure depict firms' investment response to current hire rates. The lateral movements show the prices that each firm *expects* to prevail next period and, given those expectations, the quantity that each firm supplies. When firms suffer from competition neglect, actual hire rates differ from expected hire rates because actual industry investment differs from the industry investment firms had expected.

Panel A illustrates Kaldor's (1938) cobweb model in which firms choose the quantity to supply in period $t+1$ under the naïve assumption that prices will always be the same as they were in period, t . When $\theta = 0$, we have

$$Q_{t+1} = Q_t + \overbrace{k^{-1}(H_t / r - p_r)}^{I_t}, \quad (24)$$

—i.e., firms invest under the assumption that there will be *zero* competitive supply response, so hire rates will be constant in perpetuity. Thus, the lateral movements in Panel A are perfectly horizontal and the adjustment process traces out the cobweb-like pattern in price vs. quantity space.

Panel A of Figure 7 contrasts these cobweb dynamics with the rational expectations equilibration process that obtains when $\theta \rightarrow 1$. With positive adjustment costs, hire rates and ship prices must remain above their steady state levels for several periods to induce firms to invest, thus bringing fleet size up to its new steady state. However, since expected prices are the same as actual prices under rational expectations, the adjustment process simply traces out a straight line.

Panel B of Figure 7 illustrates the case where competition neglect is severe but θ is greater than zero. In this case, firms understand that prices in period $t+1$ will differ from prices in period t , but underestimate the supply response of their competitors. While firms correctly anticipate the direction of price adjustment, they seriously underestimate the *degree* of price adjustment following a demand shock. This generates a dampened cobweb-like pattern, depicted in Panel B of Figure 7. Specifically, in Panel B, the lateral movements are not perfectly horizontal. This reflects the fact that firms partially anticipate the future competitive response and its impact on hire rates.

As θ rises from 0 to 1, individual firms increasingly recognize how competitors are likely to respond, so industry investment becomes less sensitive to deviations of earnings from their steady state (i.e., $-\partial z_i^* / \partial \theta < 0$). For instance, comparing the green line (partial competition neglect) to the

blue line (complete neglect) in Panel B, we see that industry investment reacts less strongly to deviations of earnings from the state. For small enough values of θ , the dynamics can be oscillatory—i.e., $1 + z_i^* < 0$ so prices and quantities will swing back and forth as they converge to the new steady state. As θ rises toward 1, $1 + z_i^*$ rises above 0 so the dynamics become non-oscillatory—i.e., industry fleet size steadily rises to the new steady state and hire rates steadily fall back to rp_r . This is shown in Panel C which compares the dynamics under moderate competition neglect with those under rational expectations. Since competition neglect is only moderate, the lateral movements in Panel C are far from horizontal. Finally, Panel D illustrates the dynamics with moderate competition neglect when adjustment costs (k) are small. In this case, $1 + z_i^* < 0$, so we again obtain an oscillatory cobweb-like pattern.

The realized return from owning and operating a ship between time t and $t+1$ along the equilibrium path following the initial shock is

$$1 + R_{t+1} = \frac{H_{t+1} + P_{t+1}}{P_t} = \frac{(1+r)p_r + (1+z_i^*)((B - kz_i^*)/B)(H_t - rp_r)}{p_r - ((kz_i^*)/B)(H_t - rp_r)} \quad (25)$$

Since this is a purely deterministic example, realized returns and the unbiased expectations of an econometrician are one and the same. In the rational expectations case, realized returns R_{t+1} are the same as required returns r , irrespective of the value of H_t . In contrast, when firms suffer from competition neglect, next period returns are less than required returns when current rental rates are above their steady state level—i.e., if $\theta \in (0,1]$, then $R_{t+1} < r$ when $H_t > rp_r$ and $R_{t+1} > r$ when $H_t < rp_r$. The intuition for these results is intuitive. Along the rational expectations equilibrium path, firms expect any earnings in excess of the steady-state level to be precisely offset by capital losses from holding ships. When firms suffer from competition neglect, their investment overreacts to deviations of rental rates from the steady state. Because firms are surprised by the industry supply response which depresses future rental rates, subsequent realized returns are below required returns.

E. Equilibrium expected returns

We now explore this return predictability result more generally. The realized return from owning and operating a ship between time t and $t+1$ is

$$\begin{aligned}
1 + R_{t+1} &= \frac{\pi_{t+1} + (1 - \delta)P(A_{t+1}, Q_{t+1})}{P(A_t, Q_t)} \\
&= \frac{\left[\overbrace{(A_{t+1} - BQ_{t+1})}^{\text{Ship charter rate}} - \overbrace{\tilde{C}}^{\text{Operating cost}} - \overbrace{\delta P(A_{t+1}, Q_{t+1})}^{\text{Maintenance cost}} \right] + P(A_{t+1}, Q_{t+1})}{P(A_t, Q_t)}
\end{aligned} \tag{26}$$

By construction (see equation (15)),²¹ individual firms *expect* that the return on ships will equal the required return

$$E_t[1 + R_{t+1} | A_t, Q_t] = \frac{E_t[(A_{t+1} - BQ_{t+1} - C) + (1 - \delta)P(A_{t+1}, Q_{t+1})]}{P(A_t, Q_t)} = 1 + r. \tag{27}$$

However, the expected returns perceived by the econometrician, $E_t^0[1 + R_{t+1} | A_t, Q_t]$, may differ from firms' required returns when $\theta < 1$. In this case, we have

$$E_t^0[1 + R_{t+1} | A_t, Q_t] = \frac{x_r^* + y_r^* A_t + z_r^* Q_t}{p_r + kx_i^* + ky_i^* A_t + kz_i^* Q_t}, \tag{29}$$

where x_r^* , y_r^* , and z_r^* are constants given in the Internet Appendix. By differentiating expected returns, the model delivers the relationships between charter rates, secondhand prices, investment, and future returns that we saw empirically in Section III. These relationships are summarized by Proposition 2 below.

Proposition 2 (Comparative statics for expected returns): *When firms exhibit competition neglect (i.e., when $\theta < 1$),*

$$(i) \quad \partial E_t^0[R_{t+1} | A_t, Q_t] / \partial A_t < 0 \text{ and}$$

$$(ii) \quad \partial E_t^0[R_{t+1} | A_t, Q_t] / \partial Q_t > 0.$$

Under mild regularity conditions given in the Internet Appendix, in the steady state distribution of the model with $0 \leq \theta < 1$ prices, earnings, and investment will each negatively predict returns in a univariate forecasting regression.

The intuition for Proposition 2 is natural. When demand is high or when the aggregate fleet size is small, ships hire rates and prices will be high and firms will want to invest. However,

²¹ Our current empirical approach to constructing ship returns ignores the depreciation term. Although it is a small detail, we will add this at some later stage. Specifically, we could define holding period returns as $R_{t+1} = (\pi_{t+1} + 0.95 \cdot P_{t+1} - P_t) / P_t$ corresponding to an annual depreciation rate of 5%.

because each firm underestimates the response of other competitors, firms will be surprised by the resulting level of industry investment which will push hire rates below what firms had expected, resulting in low future returns. Since ship earnings, ship prices, and investment are each increasing in current demand and decreasing in industry fleet size, Proposition 2 explains why each of these variables is associated with low future expected returns.

Further comparative statics about expected returns are given by Proposition 3.

Proposition 3 (The role of competition neglect, inelastic demand, and elastic supply on expected returns): *The comparative statics results in Proposition 2 become more pronounced when competition neglect is more severe: $\partial^2 E_t^0 [R_{t+1} | A_t, Q_t] / \partial A_t \partial \theta > 0$ and $\partial^2 E_t^0 [1 + R_{t+1} | A_t, Q_t] / \partial Q_t \partial \theta < 0$. Furthermore, the comparative statics are more pronounced when shipping demand is more inelastic (i.e., when B is large) or when shipping supply is more elastic (i.e., when k is small): $\partial^2 E_t^0 [R_{t+1} | A_t, Q_t] / \partial A_t \partial B < 0$, $\partial^2 E_t^0 [R_{t+1} | A_t, Q_t] / \partial Q_t \partial B > 0$, $\partial^2 E_t^0 [R_{t+1} | A_t, Q_t] / \partial A_t \partial k > 0$, and $\partial^2 E_t^0 [R_{t+1} | A_t, Q_t] / \partial Q_t \partial k < 0$.*

The first part of Proposition 3 is intuitive and follows from the logic of competition neglect. The intuition for the second half of Proposition 3 follows directly from the logic of Kaldor's (1938) cobweb theorem. Specifically, the cobweb-style oscillations about the steady state become larger in amplitude when demand becomes more inelastic or as supply becomes more elastic. Thus, a novel prediction of our model is that a given amount of competition neglect should lead to more pronounced return predictability in industries that (i) face inelastic customer demand or (ii) are characterized by elastic firm supply.

F. Impulse response functions

We now use the model to draw impulse response functions following a shock to demand. These diagrams illustrate much the same logic as the cobweb depictions in Figure 7, but allow us to simultaneously examine realizations of earnings, investment, prices, and returns.

Before proceeding, we define the steady-state fleet size Q^* as total quantity of ships such that industry investment just equals industry depreciation²² when shipping demand is average:

$$i_t^*(\bar{A}, Q^*) = x_i^* + y_i^* \bar{A} + z_i^* Q^* = \delta Q^* \Rightarrow Q^* = \frac{x_i^* + y_i^* \bar{A}}{\delta - z_i^*}. \quad (29)$$

We draw impulse response functions assuming that the industry is initially in this steady state at $t = 0$ (i.e., we assume that $A_0 = \bar{A}$ and that $Q_0 = Q^*$). We then assume that the industry is hit by a demand shock at $t = 1$. Specifically, using equation (12) which gives the dynamics of shipping demand, we assume that $\varepsilon_1 = \sigma_A$ and that $\varepsilon_t = 0$ for all $t > 1$. We then trace out the impulse response of various equilibrium quantities under different assumptions.²³

We first compare the impulse response functions under competition neglect with those under rational expectations. These pictures are shown in Figure 8 and underscore the cobweb intuition for our model. Under competition neglect earnings, investment, ship prices, and ship holding period returns all exhibit a oscillatory pattern, jumping at $t = 1$ before oscillating thereafter. Competition neglect leads aggregate investment to overreact to the initial demand shock. Thereafter, however, competition neglect again leads firms to react to the resulting glut of supply and all firms disinvest too aggressively under the assumption that their competitors will be slow to respond.

The final panel of Figure 8 shows 1-period realized returns on capital following this shock. These returns are equal to the expected returns from the econometrician's perspective with the exception of the shock at $t = 0$. There is a large positive return at $t = 1$ since the demand shock is unanticipated. However, firms that exhibit competition neglect over-react to this common shock at $t = 1$, thus, driving ship earnings, prices, and hence expected returns down at $t = 2$. Firms overreact

²² When $\delta > 0$, firms suffering from competition neglect perceive a different steady state than the econometrician. Firms believe the steady-state is $Q^{**} = (x_i^* + y_i^* \bar{A}) / (\delta / \theta - z_i^*) > Q^*$. Letting $\pi^{**} = \bar{A} - BQ^{**} - C$, denote perceived steady-state level of earnings, the steady-state price of ships is just $P^{**} = \pi^{**} / (r + \delta)$ —i.e., by the formula for a declining perpetuity whose cash flows shrink at rate δ .

²³ As shown in the Internet Appendix, the dynamics of the system are governed by $1 - \delta + z_i^* < 1$. The condition for the system to have convergent dynamics is $|1 - \delta + z_i^*| < 1$; we have non-oscillatory dynamics if $1 - \delta + z_i^* > 0$ and oscillatory dynamics if $1 - \delta + z_i^* < 0$. In general, we have $0 < 1 - \delta + \theta z_i^* < 1$ and thus $1 - \delta + z_i^* < 1$. Thus, in the rational expectation case ($\theta = 1$) we always have non-oscillatory and convergent dynamics. However, when $0 \leq \theta < 1$, we can have convergent, non-oscillatory dynamics (i.e., $0 < 1 - \delta + z_i^* < 1$), convergent, oscillatory dynamics (i.e., $-1 < 1 - \delta + z_i^* < 0$), or divergent, oscillatory dynamics (i.e., $1 - \delta + z_i^* < -1$). These explosive dynamics obtain if (i) θ is sufficient close to 0 and (ii) B is sufficiently large or k is sufficiently small.

to the resulting glut of supply, disinvesting aggressively at $t = 2$, leading to positive returns at $t = 3$. This oscillatory pattern continues as returns gradually converge to their steady-state level of r .

Figures 9 and 10 depict some of the comparative statics described in Proposition 3. Figure 9 shows the impulse response of investment and ship returns for five different levels of θ : rational expectations ($\theta = 1$), minor competition neglect ($\theta = 0.75$), moderate competition neglect ($\theta = 0.50$), severe competition neglect ($\theta = 0.05$), and complete competition neglect ($\theta = 0$). As θ gets closer to zero, the dynamics become larger in magnitude, turning from non-oscillatory to oscillatory. Indeed, for some parameter values the system becomes explosive when θ is sufficiently small. Figure 10 shows the impulse response of returns for different demand and supply elasticities. Panel A shows that the effects are larger in magnitude when demand is more inelastic (i.e., when B is larger). Conversely, Panel B shows the effects are large when supply is more elastic (i.e., when k is small).

G. Extension: Competition neglect versus over-extrapolation

For clarity, our exposition has emphasized the implications of competition neglect for prices, investment, and returns. But we can also solve the model allowing for a more traditional form of investor extrapolation, or allow both types of errors to be present.

Specifically, we have solved our model under the assumption that the true persistence ρ_0 of the demand shocks is less than the persistence expected by firms, ρ . Specifically, we suppose that the true law of motion for demand shocks is given by $A_{t+1} = \bar{A} + \rho_0(A_t - \bar{A}) + \varepsilon_{t+1}$, but firms believe the law of motion is given by equation (10) where $\rho > \rho_0$. As a result, the coefficients x_i^* and y_i^* in $I_t^* = x_i^* + y_i^*A_t + z_i^*Q_t$ depend on the perceived persistence, $\rho > \rho_0$.

Not surprisingly, even when there is no competition neglect (i.e. $\theta = 1$), if firms over-extrapolate demand shocks ($\rho > \rho_0$), they are predictably surprised following positive shocks since demand reverts to its steady state faster than they expect. Figure 11 shows impulse response functions in the case where $\theta = 1$ and $\rho > \rho_0$. In contrast with the pictures for competition neglect, the dynamics here are no longer oscillatory. Investment over-reacts to the initial demand shock. Expected returns remain below required returns for several periods, but do not oscillate. Investors gradually recognize that demand will not be as high as they had foreseen, leading to persistent underperformance. One can show that, in general, if investors over-estimate the persistence of

demand shocks but do not suffer from competition neglect then the dynamics will always be non-oscillatory as in Figure 11. Specifically, since $\theta = 1$ in this case, the dynamics of the system are governed by $1 - \delta + z_i^* = 1 - \delta + \theta z_i^* > 0$.

Because it is straightforward, although analytically less elegant, to include both types of errors in the analysis, when we estimate the model below we allow both exist.

V. Estimating the model in the dry bulk shipping industry

We now estimate the parameters our structural model using a Simulated Method of Moments (SMM) procedure (McFadden (1989)). Estimating the structural model enables us to assess whether one needs to posit severe competition neglect (i.e., θ near 0) to make sense of our findings. Alternately, it may be the case that moderate levels of competition neglect can rationalize the return predictability we uncover in the data. The SMM estimation exercise also allows us to run a horse race between competition neglect and more traditional forms of over-extrapolation—i.e., where firms overestimate the persistence of the underlying demand shocks.

A. SMM Estimation Procedure

The SMM procedure is explained in detail in the Internet Appendix. The basic intuition behind the estimation procedure relies on the analogy principle. We are interested in estimating an L -dimensional vector of unknown parameters of the structural model. We choose $M \geq L$ time-series moments that are jointly informative about these parameters. Specifically, these moments include means and variances of various time series (e.g., ship earnings, prices, returns, and industry investment) as well as the regression coefficients from a variety of time-series regressions described in Section III (e.g., the coefficient from a regression of returns on prior prices, earnings, and investment). For a given set of model parameters, we simulate a very long time series of data in the model and then compute the analogous moments in the simulated data. The resulting simulated moments are non-linear functions of the underlying parameters of the structural model.²⁴ To estimate the model parameters, we search for the parameter value such that the moments estimated

²⁴ Specifically, we simulate 100,000 years of data for each set of model parameters. Simulating an extremely long time-series enables us to treat the simulated moments as a deterministic function of the model parameters. This means that we can ignore any “simulation noise” when computing standard errors for our parameter estimates (see e.g., Pakes and Pollard (1986) for more on simulation noise).

using the simulated model data are as close as possible to the moments we observe in the actual data. We obtain the steady-state, stationary distribution induced by the structural model for a given set of model parameters by starting the simulation at the steady state defined in equation (29).

Asymptotic standard errors for these parameter estimates are obtained in the usual fashion. Intuitively, standard errors for the structural model parameters are smaller when the empirical moments are estimated more precisely and when the simulated moments are more sensitive to changes in the structural model parameters. Since we are trying to match $M \geq L$ moments, we must minimize a weighted distance between the empirical moments and the simulated moments. We weight each moment inversely to its estimated variance.²⁵

We hardwire C , p_r , B , and r in our estimation procedure. That is, we do not estimate these parameters using SMM, but simply assume calibrated values for these parameters in the estimation procedure. However, the resulting estimates of θ are not sensitive to the chosen values of these parameters. We calibrate the value of these parameters as follows:

- **Ship operating costs, C :** We use the estimate of real operating costs, C , discussed in Section II. Specifically, we assume a per annum operating cost of $C = 365 \times (6/1000) = \2.19 million for the 76,000 DWT Panamax tanker.
- **Ship replacement costs, p_r :** We assume that the replacement cost of a 76,000 DWT Panamax tanker is \$20 million in real current dollars.
- **The required return on ships, r :** The average return on ships is not precisely estimated in the data, so we assume that $r = 15\%$. This is consistent with the average excess return of 11% and an average real rate of 4% is reasonable for our 1976-2011 sample.
- **The slope of the shipping demand curve, B :** We assume that $B = 0.15$. In the absence of an exogenous supply shifter—i.e., an instrument for shipping supply, we cannot separately estimate k which governs the slope of firm investment supply and B which governs the slope of demand for shipping services. Since we are primarily interested in firms' investment respond, we hardwire B and choose to estimate k .²⁶

This leaves us with six parameters to estimate: ρ , \bar{A} , σ_ε , δ , k , and θ . These are, respectively, the persistence of demand, the average level of demand, the volatility of demand shocks, the annual rate of depreciation, the investment adjust cost parameter, and the degree of competition neglect.

²⁵ That is, we use a diagonal weighting matrix that weights each moment inversely to its estimated variance. As note in the Internet Appendix, this procedure differs slightly from efficient SMM, but has superior finite-sample behavior.

²⁶ This parallels much empirical research in industrial organization on demand systems where researchers are interested in estimating demand elasticities and, thus, use calibrated or hardwired supply elasticities.

We choose model parameters to match the following 20 empirical moments. We provide some intuition about the parameter identification that comes from each of the moments.

1. **$E[\pi_t]$, the average level of real earning:** This moment is informative about \bar{A} . $E[\pi_t]$ is also informative about k , δ , and θ since the steady-state level of earnings depend on these parameters (see equation (A24) in the Internet Appendix).
2. **$Var[\pi_t]$, the variance of real earnings:** This moment is informative about σ_ε . In addition, since $Var[\pi_t]$ is proportional to $Var[A_t] = \sigma_\varepsilon^2 / (1 - \rho^2)$, this moment is informative about ρ and about k since, for reasonable parameter values, $Var[\pi_t]$ is larger when firm supply is less elastic—i.e., when k is larger.
3. **$E[R_{t+1}]$, the average return:** This is informative about ρ , k , δ , and θ .
4. **$Var[R_{t+1}]$, the variance of returns:** $Var[R_{t+1}]$ is informative about σ_ε , ρ , and k .
5. **$E[P_t]$, the average real price of ships:** This moment is informative about \bar{A} . In addition, $E[P_t]$ is informative about k , δ , and θ since the steady-state level of prices depend on these parameters (see equation (A22) in the Internet Appendix).
6. **$Var[P_t]$, the variance of real ship prices:** $Var[P_t]$ is informative about σ_ε , ρ , and k .
7. **$E[I_t/Q_t]$, the average level of investment:** This moment is highly informative about δ . Indeed in the model's steady-state we have $I_t/Q_t = \delta$ by construction.
8. **$Var[I_t/Q_t]$, the variance of investment:** This is informative about σ_ε and k .
9. **$Corr[\pi_t, \pi_{t+1/12}]$, 1-month autocorrelation of earnings:** We assume there is no supply response at a monthly horizon. Since demand shocks follow an AR(1), we match $Corr[\pi_t, \pi_{t+1/12}]$ in the data with $\rho^{1/12}$ in our simulation. Thus, by construction, this moment is only informative about ρ .
10. **$Corr[\pi_t, \pi_{t+1}]$, 12-month autocorrelation of earnings:** This is informative about ρ . $Corr[\pi_t, \pi_{t+1}]$ is also increasing in k and, thus, helps to pin down adjustment costs. Intuitively, when $k > 0$, earnings will be positively auto-correlated, but not be as auto-correlated as the underlying demand shocks due to the gradual supply response.
11. **$Corr[\pi_t, \pi_{t+2}]$, 24-month autocorrelation of earnings:** Contains information similar to that in $Corr[\pi_t, \pi_{t+1}]$.
12. **$\beta(R_{t+1}, \pi_t)$, slope from a forecasting regression of R_{t+1} on π_t :** This moment is primarily about θ . This moment also contains information about σ_ε , ρ , and k .
13. **$\beta(R_{t+1}R_{t+2}, \pi_t)$, slope from a forecasting regression of $R_{t+1}R_{t+2}$ on π_t :** Contains similar information to that in $\beta(R_{t+1}, \pi_t)$.
14. **$\beta(R_{t+1}, P_t)$, slope from a forecasting regression of R_{t+1} on P_t :** This moment is primarily about θ . This moment also contains information about σ_ε , ρ , and k .
15. **$\beta(R_{t+1}R_{t+2}, P_t)$, slope from a forecasting regression of $R_{t+1}R_{t+2}$ on P_t :** Contains similar information to that in $\beta(R_{t+1}, P_t)$.

16. $\beta(R_{t+1}, H_t/P_t)$, slope from a forecasting regression of R_{t+1} on H_t/P_t : This moment is primarily about θ . This moment also contains information about σ_ε , ρ , and k .
17. $\beta(R_{t+1}R_{t+2}, H_t/P_t)$, slope from a forecasting regression of $R_{t+1}R_{t+2}$ on H_t/P_t : Contains similar information to that in $\beta(R_{t+1}, H_t/P_t)$.
18. $\beta(R_{t+1}, I_t/Q_t)$, slope from a forecasting regression of R_{t+1} on I_t/Q_t : This moment is primarily about θ . This moment also contains information about σ_ε , ρ , and k .
19. $\beta(R_{t+1}R_{t+2}, I_t/Q_t)$, slope from a forecasting regression of $R_{t+1}R_{t+2}$ on I_t/Q_t : Contains similar information to that in $\beta(R_{t+1}, I_t/Q_t)$.
20. $\beta(I_t/Q_t, P_t)$, slope coefficient from a regression of I_t/Q_t on P_t : Informative about k since the coefficient from this regression is roughly $(kQ^*)^{-1}$ where Q^* is as in equation (29). Because our model features a 1-year time to build whereas in the real world the time to build is roughly 2 years, we match this to the coefficient from a regression of $Deliveries_{[t,t+12]}/Fleet_t$ on P_{t-12} in the data.

B. Estimation results

Estimation results are shown in Table VI. We first estimate the model allowing for competition neglect (i.e., allowing $\theta < 1$), but under the assumption that there is not primitive demand over-extrapolation (i.e., we impose $\rho = \rho_0$).

Table VI lists the empirical moments from the actual data and the associated standard errors and t -statistics. To deal with serial correlation, t -statistics for these empirical moments are based on Newey-West (1987) standard errors allowing for serial correlation at up to 36 months.²⁷ The table then shows the value of these moments in the simulated data using the estimated parameters of our structural model. The SMM parameter estimates along with standard errors are shown below. We also report t -statistics for several relevant hypothesis tests.

In Table VI, we estimate that $\hat{\theta} = 0.29$ with a standard error of 0.12. Thus, we have considerable power against both the perfectly rational expectations hypothesis (i.e., the t -statistic for the hypothesis that $\theta = 1$ is $t = 6.10$) and the hypothesis that the naïve cobweb model holds (i.e., the t -statistic for the hypothesis that $\theta = 0$ is $t = 2.45$). The other five model parameters in Table VI are also precisely estimated and the point estimates seem sensible.

C. Additional estimation results

²⁷ Formally, we take a system OLS approach to our vector of our empirical moments and estimate the joint variance-covariance matrix of using a Newey-West-style estimator that allows residuals to be correlated both within and across moment equations at up to 36 months.

In Table VII, we also allow for both competition neglect (i.e., $\theta < 1$) and over-extrapolation of demand shocks (i.e., $\rho_0 < \rho$) in our SMM estimation procedure. This horserace enables us to assess the extent to which each of these biases is useful in explaining industry dynamics and returns.

Our SMM estimates suggest that both competition neglect and demand over-extrapolation are useful for explaining our empirical results. Specifically, we now obtain $\hat{\theta} = 0.47$ and $\hat{\rho} = 0.96 > 0.71 = \hat{\rho}_0$ and we can reject (i) the hypothesis that $\theta = 0$ ($t = 2.29$), (ii) the hypothesis that $\theta = 1$ ($t = 2.51$), as well as (iii) the hypothesis that $\rho = \rho_0$ ($t = 3.47$). Furthermore, comparing the simulated and empirical moments between Tables VI and VII, we can see that allowing for both competition neglect and pure over-extrapolation in Table VII allows us to do a better job of matching the data than in Table VI where we only allow for competition neglect.

VI. Conclusion

Economists have long recognized that time-to-build problems may play an important role in explaining economic fluctuations. While early theories such as Kaldor's (1938) cobweb model suggested that these fluctuations could arise from rules of thumb, these theories fell into disrepute because such rule-of-thumb behavior is not fully rational or forward looking (Muth (1961)).

We develop a model of industry investment, earnings, and asset prices in which industry capacity must be set in advance, and market participants are boundedly rational. In our model, positive shocks to demand generate excessive investment responses. We show that this investment predictably depresses future earnings and the price of capital, leading prices to overshoot their rational-expectations levels. Our model nests both the naïve view implicit in the cobweb model and the fully rational view espoused in later research. We show that modest errors by market participants—i.e., modest competition neglect—can result in dramatic predictability in prices and investment, particularly in settings in which demand is inelastic.

We estimate the model using data on earnings, secondhand prices, and investment in the dry bulk shipping industry. We can safely reject the rational expectations null hypothesis of perfect foresight. Instead, our evidence points to substantial amounts of competition neglect. More broadly, our paper suggests that pro-cyclical investor expectations, driven by a form of competition neglect, may dramatically amplify economic fluctuations in competitive industries.

References

- Baker, Malcolm, Jeremy C. Stein, and Jeffrey Wurgler, 2003. "When Does The Market Matter? Stock Prices And The Investment Of Equity-Dependent Firms," *Quarterly Journal of Economics* 118, 969-1005.
- Bansal, Ravi and Amir Yaron, 2004, "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance* 59 (4): 1481-1509.
- Ben-David, Itzhak, John R. Graham and Campbell R. Harvey, 2010. "Managerial Miscalibration," NBER Working Paper 16215.
- Bornozis, Nicolas, 2006, "Dry Bulk Shipping: The Engine of Global Trade," white paper.
- Camerer, Colin F. and Lovallo, Dan, 1999, "Overconfidence and Excess Entry: An Experimental Approach," *American Economic Review* 89.
- Campbell, John Y., and John H. Cochrane, 1999, "By Force of Habit: A Consumption-based Explanation of Aggregate Stock Market Behavior", *Journal of Political Economy* 107, 205–251.
- Campbell, John Y., and Samuel B. Thompson, 2008. "Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?," *Review of Financial Studies* 21, 1509-1531.
- Campbell, John Y. and Albert S. Kyle, 1993, "Smart Money, Noise Trading and Stock Price Behaviour," *Review of Economic Studies* 60, 1-34.
- Campbell, John Y. and Robert J. Shiller, 1988, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies* 1, 195-228.
- Campbell, John Y. and Robert J. Shiller, 1988, "Stock Prices, Earnings, and Expected Dividends," *Journal of Finance* 43, 661-676.
- Campbell, John Y. & Yogo, Motohiro, 2006. "Efficient tests of stock return predictability," *Journal of Financial Economics* 81, 27-60.
- Case, Karl E., Robert J. Shiller, and Anne Thompson, 2012, "What Have They Been Thinking? Home Buyer Behavior in Hot and Cold Markets." *Brookings Papers on Economic Activity*.
- Cochrane, John, 2011, "Discount Rates," *Journal of Finance* 66, 1047-1108.
- Esty, Benjamin C., and Albert Sheen, 2011, "Vereinigung Hamburger Schiffsmakler und Schiffsagenten e.V. (VHSS): Valuing Ships." Harvard Business School Case 210-058, April 2011.
- Fama, Eugene F., 1970, "Efficient Capital Markets: A Review of Theory and Empirical Work", *Journal of Finance* 25, 383-417.
- Fama, Eugene F., and Kenneth R. French, 1988, "Dividend Yields and Expected Stock Returns", *Journal of Financial Economics* 22, 3-25.
- Gennaioli, Nicola, and Andrei Shleifer, 2010, "What Comes to Mind," *Quarterly Journal of Economics* 125, 1399-1433.
- Hayashi, Fumio, 1982, "Tobin's Marginal Q and Average Q: A Neoclassical Interpretation," *Econometrica* 50, 213–224.

- Kahneman, Daniel, 2011, *Thinking Fast and Slow*, Macmillan, New York.
- Kaldor, Nicholas, 1934, "A Classificatory Note on the Determination of Equilibrium," *Review of Economic Studies* 1, 122-136.
- Kaldor, Nicholas, 1938, "The Cobweb Theorem", *Quarterly Journal of Economics* 52, 255–280.
- Kalouptsi, Myrto, 2010, "Time to Build and Shipping Prices", Working paper.
- Kewei Hou & David T. Robinson, 2006. "Industry Concentration and Average Stock Returns," *Journal of Finance* 61, 1927-1956.
- Kydland E. and E. Prescott, 1982, "Time to Build and Aggregate Fluctuations," *Econometrica* 50, 1345-1370.
- Lamont, Owen, 2000, "Investment Plans and Stock Returns", *Journal of Finance* 55, 2719-2745.
- Malmendier, Ulrike, and Geoff Tate, 2005, CEO Overconfidence and Corporate Investment. *Journal of Finance* 60, 2661-2700.
- Moskowitz, Tobias, Lasse H. Pedersen, Ralph Koijen, and Evert B. Vrugt, 2012, "Carry", Working paper.
- Muth, John, 1961, "Rational Expectations and the Theory of Price Movements", *Econometrica* 29, 315–335.
- Nerlove, Mark 1958, "Adaptive Expectations and Cobweb Phenomena", *Quarterly Journal of Economics* 72, 227-40.
- Pakes, Ariel and David Pollard, 1989, "Simulation and the Asymptotics of Optimization Estimators," *Econometrica* 57, 1027-1057.
- Review of Maritime Transport, 2009, Report by the UNCTAD secretariat.
- Rosen, Sherwin, Kevin Murphy and Jose A. Scheinkman, 1994, "Cattle Cycles", *Journal of Political Economy* 102, 468-492.
- Shiller, Robert J., 1981, "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?" *American Economic Review* 71, 421–436.
- Shleifer, Andrei, 1986, "Implementation Cycles", *Journal of Political Economy* 94, 1163-1190.
- Shleifer, Andrei, and Robert W Vishny, 1992, Liquidation Values and Debt Capacity: A Market Equilibrium Approach, *Journal of Finance* 47, 1343-66.
- Stafford, Erik, Angela Chao, and Kathleen Luchs, 2002, "Ocean Carriers." Harvard Business School Case 202-027, April 2002.
- Stein, Jeremy C., 1996, "Rational Capital Budgeting in an Irrational World", *Journal of Business* 69, 429-455.
- Stopford, Martin, 2009, *Maritime Economics*, 3rd Edition. Routledge, New York.
- Tversky, Amos., and Daniel Kahneman, 1974. "Judgment Under Uncertainty: Heuristics and Biases", *Science* 185, 1124-1131.
- Welch, Ivo, and Amit Goyal, 2008, "A Comprehensive Look at The Empirical Performance of Equity Premium Prediction", *Review of Financial Studies* 21, 1455-1508.