



City Research Online

City, University of London Institutional Repository

Citation: Moutzouris, I. (2017). Asset valuation in dry bulk shipping. (Unpublished Doctoral thesis, City, University of London)

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/18410/>

Link to published version:

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

ASSET VALUATION IN DRY BULK SHIPPING

Ioannis C. Moutzouris

Cass Business School, City University of London

June 2017

A thesis submitted for the degree of

Doctor of Philosophy

Supervisor: Professor Nikos K. Nomikos

Second Supervisor: Dr Nikos C. Papapostolou

Table of Contents

Abstract	v
Chapter 1: Introduction.....	1
1.I. Description of the Dry Bulk Shipping Industry	1
1.I.A. Shipping Demand	1
1.I.B. Shipping Supply.....	3
1.I.C. The Shipping Freight Rate Mechanism	7
1.I.D. The Forward Freight Agreements Market	7
1.II. Contribution	9
1.II.A. Summary of Chapter 2.....	10
1.II.B. Summary of Chapter 3.....	11
1.II.C. Summary of Chapter 4.....	13
1.II.D. Future Research	16
Chapter 2: The Earnings Yield and Predictability of Earnings in the Dry Bulk Shipping Industry.....	18
2.I. Introduction	18
2.II. Data and Variables of Interest	23
2.III. Predictability of Net Earnings in Shipping	36
2.III.A. Predictive Regressions.....	36
2.III.B. Variance Decomposition.....	42
2.III.C. An Extension of the Campbell and Shiller (1988a) VAR Framework to Shipping	44
2.IV. Economic Interpretation and Discussion	49
2.IV.A. The Shipping Supply and Demand Mechanism.....	52
2.IV.B. Comparison to Other Markets	60
2.V. Conclusion.....	64
2.V.A. Connection to Chapter 3.....	65
2.V.B. Connection to Chapter 4.....	66
A.2. Appendix 2	68
A.2.A. Exact Present Value Relation	68
A.2.B. Linearisation of the Present Value Relation	69
A.2.C. Variance Decomposition.....	73
A.2.D. Extension of the Campbell and Shiller (1988a) VAR Framework to Shipping.....	73
Chapter 3: Heterogeneous Expectations and the Second-Hand Market for Dry Bulk Ships.....	79
3.I. Introduction	79
3.II. Environment and Model Solution	83

3.III.	Empirical Estimation of the Model in the Dry Bulk Shipping Industry	92
3.III.A.	Data on Net Earnings, Prices, and Trading Activity	92
3.III.B.	Simulation Methodology and Results	94
3.III.C.	Economic Interpretation.....	101
3.III.D.	Sensitivity Analysis.....	103
3.III.E.	Impulse Response Functions.....	106
3.III.F.	Expectations of Returns and Realised Returns	110
3.IV.	Robustness	113
3.V.	Conclusion	116
3.V.A.	Connection to Chapter 4	117
A.3.	Appendix 3	120
A.3.A.	Derivation of the Demand Functions for the Age-Specific Vessels	120
A.2.B.	Proposition and Corollaries	123
A.3.C.	Omitted Figures.....	132
A.2.D.	Changes in Investor Wealth	134
	Chapter 4: The Formation of FFA Rates in Dry Bulk Shipping: Spot Rates, Risk Premia, and Heterogeneous Expectations	136
4.I.	Introduction	136
4.II.	Data and Variables of Interest	142
4.III.	Predictable Variation in the FFA Market	150
4.III.A.	Predictability of Future Market Conditions and Risk Premia from the Basis	151
4.III.A.i.	A Simple Variance Decomposition Framework.....	151
4.III.A.ii.	Interpretation of the Results and Comparison to Other Markets.....	155
4.III.B.	Predictability of Risk Premia from Lagged Risk Premia and Spot Market Indicators	160
4.III.C.	Predictability of Risk Premia from Economic Variables	166
4.III.D.	Trading Activity, Spot Market Conditions, and Risk Premia.....	169
4.III.E.	Unbiased Expectations Hypothesis	175
4.III.E.i.	The Wald Test Approach.....	176
4.III.E.ii.	The Johansen Cointegration Approach.....	177
4.IV.	A Heterogeneous Expectations Model for the FFA Market	180
4.IV.A.	Economic Environment and Model Solution.....	181
4.IV.B.	Simulation of the Model	190
4.IV.B.i.	Scenario 1: Rational Benchmark without Hedging Pressure	190
4.IV.B.ii.	Scenario 2: Rational Benchmark with Constant Hedging Pressure	191

4.IV.B.iii. Scenario 3: Rational Benchmark with Time-Varying Hedging Pressure	192
4.IV.B.iv. Scenario 4: Distorted Expectations and Constant Hedging Pressure	194
4.IV.B.v. Scenario 5: Distorted Expectations and Time-Varying Hedging Pressure	197
4.V. Conclusion	198
A.4. Appendix 4	201
A.4.A. Complementary Results for the 3- and 4-Month Maturity Contracts	201
Chapter 5: Conclusion	209
References	211

Abstract

This thesis examines the dry bulk sector of the shipping industry. We begin by analysing the relation between second-hand vessel prices, net earnings, and holding period returns. Specifically, we provide strong statistical evidence that almost the entire volatility of shipping earnings yields can be attributed to variation in expected net earnings growth; almost none to expected returns variation and almost none to varying expectations about the terminal earnings yield. According to our results, earnings yields are negatively and significantly related to future net earnings growth. Furthermore, we find no consistent, strong statistical evidence supporting the existence of time-varying risk premia in the valuation of dry bulk vessels. Accordingly, we integrate the examination of the second-hand market by incorporating in the analysis the trading activity related to dry bulk vessels. For this purpose, we develop a heterogeneous expectations asset pricing model that can account for the actual behaviour of vessel prices and the positive correlation between net earnings, vessel prices, and second-hand vessel transactions. The proposed economy consists of two agent types who form heterogeneous expectations about future net earnings and at the same time under(over)estimate the future demand responses of their competitors. Formal estimation of the model suggests that the average investor expectations in the second-hand market for ships must be “near-rational”. In particular, the investor population must consist of a very large fraction of agents with totally – or very close to – rational beliefs while the remaining ones must hold highly extrapolative beliefs; thus, there must exist significant heterogeneity of beliefs in the market. Having concluded the analysis of the second-hand physical shipping market we turn to the derivative market for Forward Freight Agreements (FFAs) related to the dry bulk shipping sector. Accordingly, we illustrate formally that the bulk of volatility in the FFA basis can be attributed to expectations about future physical market conditions rather than expectations about future risk premia. Despite this finding, though, we document the existence of a bias in the FFA rates in the form of “contango” but also of both a strong momentum effect and significant predictability of risk premia by price-based signals and economic variables reflecting physical market conditions. The evidence of bias is further supported by the results of three econometric tests which suggest rejection of the unbiased expectations hypothesis. Finally, to justify these findings, we develop a dynamic asset pricing framework that can incorporate both the “hedging pressure” feature and a heterogeneous-beliefs explanation.

Chapter 1: Introduction

1.1. Description of the Dry Bulk Shipping Industry

The shipping industry plays a substantial role in the global economy since approximately 90% of the world trade is carried through vessels (UNCTAD, 2015). Each commodity has bespoke characteristics and requires a specific type of vessel to be transported. As a result, there is a large market for overseas transportation and, subsequently, many shipping firms – i.e., providers of the shipping service.

This thesis focuses on the dry bulk sector of the shipping industry mainly because it represents by far the largest segment in terms of both cargo carrying capacity and quantity transported (Alizadeh and Nomikos, 2010). Namely, in 2014 dry bulk vessels carried out approximately 42.9% of the world seaborne trade. Furthermore, the nature of competition and, especially, the distinct supply and demand mechanism that characterise this market give us the opportunity to interpret our empirical estimation results using straightforward microeconomic principles and rationale. Finally, investigating the dry bulk shipping market, as opposed to the tanker and container ones, allows us to employ significantly larger data regarding both the time dimension and the number and variety of incorporated variables.


Dry bulk shipping refers to the transportation of homogeneous unpacked dry cargoes – that is, raw materials in the form of solid, bulk commodities such as iron ore and grains – on non-scheduled routes, mainly on a “one ship-one cargo” basis (Alizadeh and Nomikos, 2010; Kalouptsi, 2014). Dry bulk carriers transport a wide variety of solid cargoes, ranging from the so-called major bulks (i.e., iron ore, coal, grains, bauxite/alumina, and phosphate rock) to the majority of minor bulk cargoes (e.g., steel products and chemical parcels).

1.1.A. Shipping Demand


Demand for dry bulk shipping services translates into demand for dry bulk seaborne trade which, in turn, is driven by five main factors. Undoubtedly, the most important one is the world economy: as Stopford (2009) documents, seaborne trade is highly correlated with world GDP cycles. In addition, seaborne trade is highly affected by the prevailing conditions in the related commodity trades – that is, the dry bulk commodity trends and prices. Note that commodity markets affect the demand for shipping in both the short- and long-term. Regarding the former, the observed short-term fluctuations in shipping market conditions are mainly caused by the seasonal character of some

trades (e.g., other hand, fluctuations can attributed to –

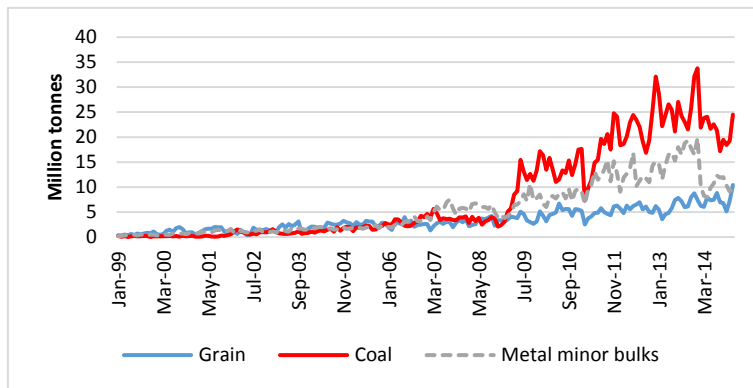
grains). On the long-term be mainly changes in – the



Panel A: Major bulk and total dry bulk seaborne trade from 1983 to 2014.



Panel B: World steel production from 1/1995 to 12/2014.



Panel C: China's coal, grain, and metal minor bulk imports from 1/1999 to 12/2014.

Figure 1.1: Demand for Dry Bulk Shipping Services.

Panel A illustrates the evolution of major bulk (measured in trillion tonne miles) and total dry bulk (measured in billion tonnes) seaborne trade. The dataset used is in an annual frequency. Panel B shows the world steel production (measured in million tonnes) in a monthly frequency. Finally, Panel C demonstrates the evolution of China's coal, grain, and metal minor bulk imports (measured in million) tonnes in a monthly frequency.

economic characteristics of the markets that import and export the corresponding commodities.

Despite these two factors, which are exogenous to the shipping industry, demand is also affected – however, at a significantly lower degree – by “the average haul of the trade” (measured in tonne miles)¹ and the costs of transportation (Stopford, 2009). Importantly, while these four variables can be in general predicted – of course, up to a certain level – and, therefore, accounted for by market participants, it is the existence of substantial random shocks that perturb the shipping equilibrium and result in the well-known shipping boom-bust cycles or, equivalently, generate the extraordinary volatility that characterises the industry. These unique and unpredictable shocks in shipping demand can be caused by either economic disturbances superimposed on business cycles – such as the two oil crash shocks in 1973 and 1979 and the recent financial crisis – or political events – such as wars, revolutions, and strikes (Stopford, 2009).

Consequently, demand is considered as rather inelastic and exogenous to the shipping industry. Panel A of Figure 1.1 presents the evolution of dry bulk seaborne trade for the period 1983 to 2014, measured in both tonnes and tonne miles.² Evidently, the aggregate demand variable follows an upward sloping trend. Specifically, the total increase over the period 1983-2014 is equal to 348.1%, corresponding to an annual average compound growth rate of 4.1%. However, as we observe in Panels B and C, commodity-specific and country-specific demand fluctuate significantly around this upward trend. In line with Stopford (2009), Panel C of Figure 1.2 illustrates that annual demand changes of around 10% are not an unusual phenomenon in this industry. In conclusion, we can characterise the evolution of dry bulk demand as a mean-reverting process around a substantial upward drift.³

1.1.B. Shipping Supply

The supply component of the shipping mechanism corresponds to the cargo carrying capacity of the dry bulk fleet. Depending on the size of the vessel, the dry bulk fleet can be subdivided into four main sectors which researchers and industry participants treat as different markets (Kalouptsidi, 2014); namely, the Capesize, Panamax, Handymax, and Handysize sectors. At the largest end of the range, Capesize carriers have a cargo carrying capacity that exceeds 100,000 dwt and heavily depend on the trades of iron ore and coal.⁴ Panamax carriers (60,000-99,000 dwt) serve mainly the coal, grain, bauxite, and the larger minor bulk trades. At the lower end of the range are the Handymax

¹ Tonne miles are defined as the product of the tonnage of shipped cargo times the transportation distance (Stopford, 2009).

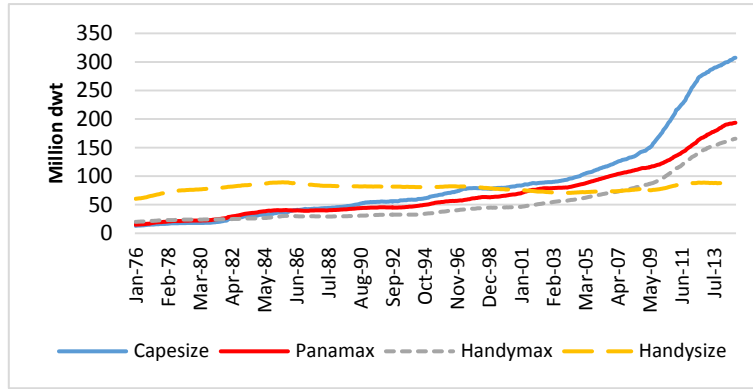
² Data are obtained from Clarksons Shipping Intelligence Network.

³ The assumption of a simple mean-reverting process for demand has been imposed in the literature by Kalouptsidi (2014) and Greenwood and Hanson (2015).


⁴ The abbreviation dwt stands for deadweight tonnage and measures the cargo carrying capacity of a vessel.

(40,000-59,000 dwt) carriers. Handyise mainly geared

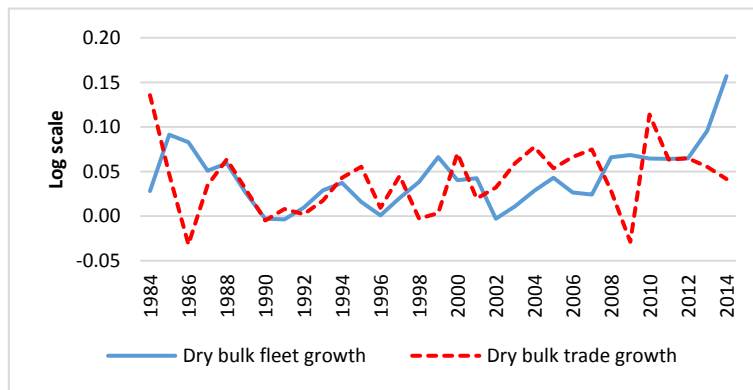
dwt) and (10,000-39,000 dwt) carriers. These ships are and serve as



Panel A: Dry bulk sectors fleet development from 1/1976 to 1/2014.



Panel B: Dry bulk fleet and trade development from 1983 to 2014.



Panel C: Dry bulk fleet and trade growth from 1983 to 2014.

Figure 1.2: Dry Bulk Shipping Supply and Correlation with Demand.

Panel A illustrates the fleet development for each dry bulk sector (measured in million dwt). The dataset employed is in a monthly frequency. Panel B provides a comparison between the total dry bulk fleet development (measured in million dwt) and the evolution of the total dry bulk trade (measured in billion tonnes). Finally, Panel C compares the evolutions of total dry bulk fleet and total dry bulk trade growth. The data corresponding to Panels B and C are in an annual frequency.

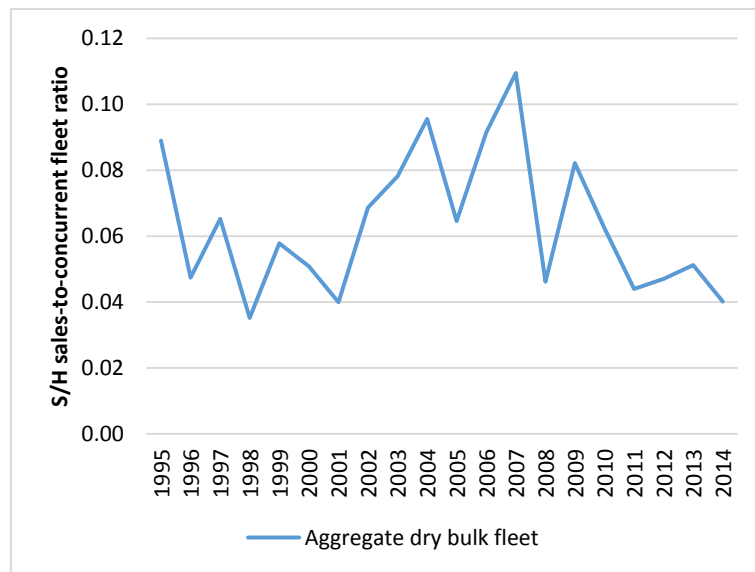
versatile workhorses in trades where parcel size and dimensional restrictions require smaller vessels. Usually, they carry minor bulks and smaller quantities of major bulks (Stopford, 2009). In December 2014, the Capesize, Panamax, Handymax, and Handysize dry bulk sectors consisted of 1,635, 2,442, 3,112, and 3,128 vessels, respectively. Equivalently, the total cargo carrying capacity amounted to approximately 756 million dwt.

However, each sector – and the dry bulk industry as a whole – consists of a substantial number of ship owning corporations that essentially act as price-takers. Therefore, from an economic point of view, dry bulk shipping is considered as a highly competitive industry (Kalouptsidi, 2014; Greenwood and Hanson, 2015). Panel A of Figure 1.2 illustrates the evolution of fleet capacity for each of the four sectors while Panel B depicts the development of the aggregate dry bulk fleet (all in terms of dwt). Noticeably, the evolutions of the sector-specific and aggregate supply variables are very similar to the one of aggregate demand. However, the aggregate dry bulk supply has realised even more significant increase compared to demand; namely, the total growth rate of aggregate vessel capacity over the period 1983-2014 equals 420.3% which is equivalent to a 4.7% average annual increase.


In contrast to demand, shipping supply is solely determined by the investment decisions of market agents; therefore, it is endogenous to the dry bulk industry. In particular, it can be increased through the ordering of newbuilding vessels and decreased through the demolition of existing ones. Consequently, supply is highly elastic in the long run. To quantify this inherent feature of the shipping industry, consider the following stylised fact. Following the market peak of 2008, the order book in 2009 was approximately equal to 77% of the corresponding fleet (in terms of million dwt).⁵ As a result, the net increase in the fleet between 2008 and 2014 – that is, after accounting for scrapping activity – was equal to 85% (Panels B and C of Figure 1.2). There are not many real asset industries where we can observe comparable fluctuations in the supply side in such a limited period. For example, in developed real asset markets – due to zoning and regulatory restrictions – the increase in the supply of premises is significantly bounded.

While the scrapping of a vessel can occur immediately, the delivery of a newbuilding order requires a time-to-build which can vary from 18 to 60 months, heavily depending on the prevailing market conditions (Kalouptsidi, 2014). Therefore, shipping supply is significantly inelastic in the short horizon. What is more, due to this time-to-build characteristic, supply adjusts sluggishly to demand (Greenwood and Hanson, 2015). Consequently, as Panels B and C of Figure 1.2 demonstrate, while the aggregate shipping supply and demand variables exhibit a high degree of co-movement in terms of levels (the estimated correlation coefficient is 0.97), their respective growth rates are extremely less

⁵ The order book measures the number – and the cargo carrying capacity – of vessels under construction or awaiting construction (Papapostolou *et al*, 2014).



Panel A: Aggregate S/H sales-to-concurrent fleet ratio from 1995 to 2014.



Panel B: Sector-specific S/H sales-to-concurrent fleet ratio from 1995 to 2014.

Figure 1.3: Trading Activity in the Second-Hand Market.

Panel A presents the evolution of the annual second-hand sales-to-concurrent fleet ratio for the aggregate dry bulk fleet while Panel B demonstrates the evolution of the sector-specific ratios. The corresponding period is from 1995 to 2014 and the sample is in an annual frequency.

correlated (the corresponding correlation coefficient is 0.31). The implications of this feature – as discussed thoroughly in Chapter 2 of this thesis – are very important in terms of shipping lease rates and, in turn, shipping net earnings⁶ but also of vessel prices.

Finally, from an industry perspective, activity in the second-hand market for vessels is not considered as (dis)investment since it solely affects the ownership distribution of the existing transport capacity. Therefore, in terms of industrywide investment and cargo carrying capacity, second-hand activity can be characterised as a zero-sum game. Note that – as analysed in Chapter 3 of this thesis – second-hand vessel markets are characterised by relatively low liquidity. Specifically, during the period 1995-2014, the average ratio of annual aggregate dry bulk second-hand vessel sales to the respective total dry bulk fleet was equal to 6.3% while it ranged from 3.5% to 11%.⁷ Figure 1.3 presents the evolution of the second-hand sales-to-concurrent fleet ratio for the aggregate dry bulk fleet and the sector-specific ones.

1.1.C. The Shipping Freight Rate Mechanism


Since the dry bulk shipping freight rate mechanism is explicitly analysed in Chapter 2 of this thesis, in this subsection we briefly outline how it operates. Namely, as it is well-documented in the literature (Stopford, 2009; Greenwood and Hanson, 2015), random shocks in demand drastically perturb the short-run shipping equilibrium – since supply is highly inelastic in the short run – and, consequently, the prevailing lease rates, that is, the shipping cash flows. In turn, changes in the prevailing lease rates have an indirect dramatic effect on future cash flows through the current investment decisions of shipping investors. Specifically, as analysed above, due to the time-to-build characteristic, changes in shipping supply will not be realised immediately (excluding the scrapping activity) but in future periods. This fact, accompanied by the mean-reverting (around an upward trend) character of the exogenous demand result in extremely volatile shipping cash flows. Consequently, shipping cash flows are not exogenously but partially endogenously determined by the investment decisions of shipping industry participants.

1.1.D. The Forward Freight Agreements Market

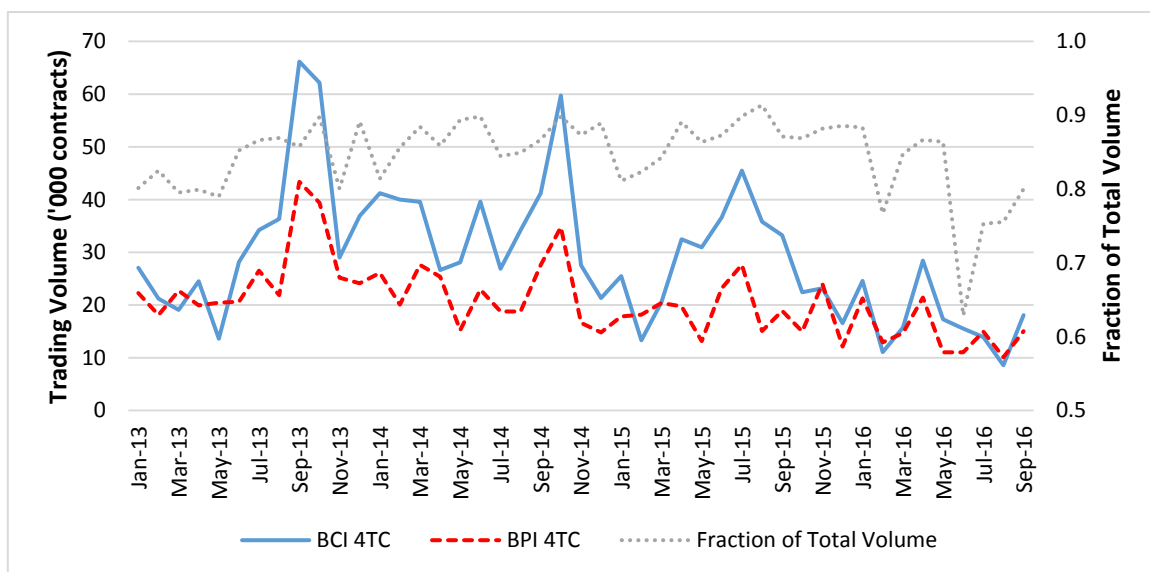
While in Chapters 2 and 3 of this thesis we focus on the physical shipping market for second-hand dry bulk vessels, Chapter 4 examines the derivative market for Forward Freight Agreements (FFAs)

⁶ Net earnings are defined as the operating profit for the owner of the vessel.

⁷ Regarding the Capesize sector, during the period 1995-2014, the average, maximum, and minimum values of the annual second-hand sales-to-concurrent fleet ratio were equal to 5.8%, 11.3%, and 3.4%, respectively. In the Panamax sector, the corresponding values were 7.4%, 13.2, and 2.6%. In the Handymax sector, those were 6.5%, 11.6%, and 2.8%, respectively. Finally, in the Handysize sector, they corresponded to 5.9%, 10%, and 3.1%.



Panel A: Trading Volume by Sector.



Panel B: Trading Volume by Contract Type.

Figure 1.4: FFA Trading Volume.

Panel A depicts FFA weekly trading volume related to all contracts in the Capesize and Panamax dry bulk sectors from July 2007 to September 2016. The grey dotted line plots the sum of Capesize and Panamax volumes as a fraction of the contemporaneous total dry bulk volume. Panel B depicts monthly trading volume

related to the BCI 4TC and BPI 4TC contracts from January 2013 to September 2016. The grey dotted line plots the sum of BCI 4TC and BPI 4TC volume as a fraction of the contemporaneous total dry bulk volume.

related to the dry bulk sector of the shipping industry. Specifically, the market for Forward Freight Agreement (FFA) contracts was established in 1992 as a hedging instrument for participants in the physical shipping market. Following Alizadeh and Nomikos (2009), an FFA contract is “an agreement between two counterparties to settle a freight rate or hire rate, for a specified quantity of cargo or type of vessel, for one or a basket of the major shipping routes in the dry-bulk or the tanker markets at a certain date in the future. The underlying asset of FFA contracts is a freight rate assessment for an underlying shipping route or basket of routes... FFAs are settled in cash on the difference between the contract price and an appropriate settlement price”.

In the context of this thesis, we focus on the Capesize and Panamax dry bulk FFA contracts since they constitute by far the most liquid instruments. In particular, trading volume in the Capesize and Panamax sectors accounts on average for approximately 46% and 42%, respectively, of the contemporaneous total volume in the FFA dry bulk contracts. Panel A of Figure 1.4 illustrates the evolution of trading volume in these two sectors (i.e. the summation of cleared and OTC contracts) over the period from July 2007 to September 2016, on a weekly basis. Notice that trading activity in these two sectors is significantly correlated; the correlation coefficient is 0.57. In addition, we plot the summation of Capesize and Panamax trading volume as a fraction of the contemporaneous total volume. Noticeably, the average value of this fraction is 0.88 while it is almost always above 0.7.

Regarding the specific FFA contracts, the bulk of trades is related to the BCI 4TC and BPI 4TC ones. These contracts correspond to the equally weighted average of the four trip-charters of the Baltic Capesize Index and the Baltic Panamax Index, respectively. Since market practitioners use these basket contracts to hedge their average monthly TC earnings, the corresponding settlement rate is estimated as the arithmetic average of the TC routes over all trading days of the month. During the period from January 2013 to September 2016 (for which we have data from the London Clearing House), monthly trading volume related to the BCI and BPI 4TC contracts accounted for approximately 49% and 36% of the total FFA dry bulk volume, respectively. In analogy to Panel A, Panel B of Figure 1.4 depicts the evolution of trading volume related to these two contracts over the respective period. As with the entire sectors, trading volume in these two contracts is highly correlated; the correlation coefficient is equal to 0.84. Panel B also plots the summation of BCI and BPI 4TC trading volume as a fraction of the corresponding total volume in the dry bulk market. Specifically, the average value of this fraction is 0.85 while it is almost constantly above 0.7.

1.II. Contribution

This thesis examines the dry bulk sector of the shipping industry. Specifically, we focus on the physical shipping market for second-hand vessels (that is, in Chapters 2 and 3) and the derivative market for Forward Freight Agreements (FFAs) related to the dry bulk sector of the shipping industry (that is, in Chapter 4).

1.II.A. Summary of Chapter 2

We begin by examining in Chapter 2 the relation between second-hand vessel prices, net earnings, and holding period returns in the Capesize, Panamax, Handymax, and Handysize dry bulk sectors. Namely, we analyse empirically the formation of the most frequently incorporated vessel valuation ratio – that is, the shipping earnings yield – through the Campbell-Shiller variance decomposition and vector autoregression (VAR) frameworks.

Our contribution to the literature is threefold. First, from a technical perspective, we extend the Campbell-Shiller variance decomposition and vector autoregression (VAR) frameworks (1988b and 1988a, respectively) to account for both “forward-looking” valuation ratios and economic depreciation in the value of the respective asset – that is, to be able to capture in a mathematically rigorous manner the case of real assets with limited economic lives. To the best of our knowledge, this is the first time that these features are explicitly incorporated in this asset pricing framework. Accordingly, the proposed methodologies can be used for the valuation of assets in other real asset economies with similar characteristics, such as the commercial real estate and airline industries.

Second, using the extended Campbell-Shiller (1988b) variance decomposition framework, we provide strong statistical evidence that the bulk of variation in net earnings yields reflects varying expectations about net earnings growth, not time-varying expected returns, and not varying expectations about the terminal earnings yield. In particular, shipping earnings yields are negatively and significantly related to future net earnings growth. Furthermore, there is no consistent, strong statistical evidence supporting the existence of time-varying risk premia in the formation of earnings yields. Equivalently, from a vessel valuation point of view, our results imply that dry bulk vessel prices vary mainly due to news related to expected net earnings, not due to expected returns, and not due to the terminal – scrap – price of the vessel. This latter argument is further reinforced using the modified Campbell-Shiller (1988a) VAR framework. Specifically, we illustrate formally that actual price-net earnings ratios can be replicated sufficiently well through a VAR model with constant required returns. To the best of our knowledge, these stylised facts had never been documented formally in the shipping literature before.

Third, since shipping is a capital-intensive industry with distinct, directly observable supply and demand determinants and mechanism, it provides an ideal environment to build a bridge between the incorporated empirical asset pricing framework and the – economic – characteristics of the

market under consideration. Subsequently, this reasoning can be extended to comparable real asset industries. Specifically, from an economic point of view, we argue that in order for valuation ratios to move due to expectations about future cash flows, the latter should be predictable by market agents using the current information set. Vice versa, if future cash flows are not predictable using current market information then they can neither be predicted by the earnings yield. Accordingly, we state that the major determinants of valuation ratios are the second-order effects (SOEs) that current cash flows have on current prices through the future cash flow stream. If there are no profound SOEs, then there is no reason for future cash flows to be predictable by the current information filtration. From a statistical perspective – and in line with recently obtained evidence (Chen *et al*, 2012; Rangvid *et al*, 2014) – we argue that the significant predictability of earnings growth by the earnings yield is driven by the extreme volatility of shipping net earnings.

In a cross-industry comparison, our results are diametrically opposed to the ones in the post-WWII U.S. equity markets and residential (housing) real estate markets but in line with the ones obtained from both the pre-WWII U.S. equity markets and the bulk of international equity markets as well as the majority of the commercial real estate industry – and the REIT index market. Therefore, this chapter provides strong evidence for further discussion regarding the economic principles that drive the forecasting properties of valuation ratios.

1.II.B. Summary of Chapter 3

In Chapter 3, we integrate and conclude the examination of the second-hand market by incorporating in the analysis the trading activity related to dry bulk vessels. Namely, we investigate the joint behaviour of vessel prices, net earnings, and second-hand trading activity. For this purpose, we develop and estimate empirically a heterogeneous expectations asset pricing model with microeconomic foundations that can account and, in turn, provide a plausible economic interpretation for numerous empirical findings related to this market. While the empirical analysis focuses on the Handysize sector, our results have been tested to the remaining dry bulk sectors and are both qualitatively and quantitatively robust; thus, our conclusions are representative of the entire dry bulk industry.

Specifically, the proposed partial equilibrium framework explains the observed behaviour of second-hand vessel prices; in particular, we are mainly interested in the actual price volatility, the autocorrelation of prices, and the high correlation between prices and prevailing net earnings. In addition, our model reproduces and justifies the stylised fact that trading activity is positively related to both market conditions and absolute changes in net earnings between two consecutive periods. In our sample, the two correlation coefficients are equal to 0.53 and 0.65, respectively, implying that

investors trade more aggressively during prosperous market conditions but also when net earnings have significantly changed compared to the previous period.

Moreover, our model implicitly captures the fact that second-hand markets for vessels are rather illiquid: as analysed in Subsection 1.1, during the period 1995-2014, the average annual sale and purchase turnover in the Handysize sector was approximately 5.8% of the corresponding fleet size. Finally, the proposed framework also accounts for the stylised features presented in Chapter 2 of this thesis; namely, for the finding that net earnings yields are highly positively correlated with the prevailing market conditions and, in turn, strongly negatively forecast future net earnings growth but also for the fact that the bulk of the earnings yield's volatility is attributed to expected cash flow variation and not to time-varying expected returns.

Our discrete-time economy consists of two agent types, conservatives and extrapolators, who form heterogeneous expectations about future net earnings and at the same time under (over) estimate the future demand responses of their competitors. Interestingly, formal estimation of the model suggests that, to simultaneously match the empirical regularities, the average investor expectations in the second-hand market for ships must be "near-rational". In particular, the investor population must consist of a very large proportion of agents (conservatives) with totally – or very close to – rational beliefs while the remaining fraction (extrapolators) must hold highly extrapolative beliefs; thus, there must exist significant heterogeneity of beliefs in the market.

From an economic perspective, this finding is in accordance with the nature of the shipping industry; namely, the large fraction of conservative investors corresponds to the large number of established shipping companies that operate in the industry. In some instances, ship owning families have been present in the market for more than a century (Stopford, 2009) and, consequently, have strong prior experience and expertise about the key supply and demand drivers of the industry. In turn, their superior knowledge translates into more accurate forecasts about future market conditions compared to relatively new investors.

Extrapolators, on the other hand, reflect new entrants such as diversified investors (e.g., private equity firms) with little or no previous experience of the market. It is well-documented that during prosperous periods, new entrants impressed by the high prevailing earnings and short-term returns are eager to buy vessels which, subsequently, are more than keen to sell as conditions deteriorate. In contrast, there are many cases where traditional owners have realised significant returns by selling vessels at the peak of the market and buying at the trough – a strategy known as "playing the cycles" (Stopford, 2009).

In conclusion, the contribution of this chapter to the literature can be summarised in the following.

First, this is the first time in the shipping literature that a structural economic model incorporates the coexistence of heterogeneous beliefs agents to explain the joint behaviour of observed vessel prices, net earnings, and second-hand vessel transactions. Regarding the existing shipping literature, Beenstock (1985), Beenstock and Vergottis (1989), and Kalouptsidei (2014) construct and estimate rational expectations general equilibrium models in a homogeneous agents' setting which, however, does not allow for the explanation of the second-hand market activity. Greenwood and Hanson (2015) develop a homogeneous beliefs model in which the behavioural mechanism is similar to the one proposed here, however, they focus on the newbuilding and demolition markets as opposed to the one for second-hand vessels as is the case in our context. Furthermore, in contrast to Greenwood and Hanson (2015), the introduction of two types of agents allows us to simultaneously capture the observed behaviour of prices, net earnings, and second-hand activity in the market.

Second, to the best of our knowledge, this is the first time in the asset pricing literature that a structural heterogeneous beliefs asset pricing model is applied to a real asset economy and, in particular, shipping. Therefore, our model looks at the main features of heterogeneous agents' models but also introduces important modifications which are required to capture the characteristics of the shipping industry. Namely, the fact that we examine an asset with finite life that is significantly affected by economic depreciation due to wear and tear provides different challenges in the economic modelling of the market compared to the case of an infinitely lived financial one (e.g., equity).

Moreover, in contrast to the bulk of the behavioural equity markets literature, in our model there is cash flow and not return extrapolation. The motivation for this is based on actual market practice and the economics of the industry. Namely, shipping industry participants characterise market conditions based on the prevailing – and forecasts of future – net earnings and not on realised returns. Thus, it is much more plausible for investors to form biased expectations regarding fundamentals rather than returns. In contrast, in equity markets, recent evidence from surveys (Greenwood and Shleifer, 2014) suggests that many investors extrapolate stock market returns. Accordingly, we provide a framework that can be incorporated and, accordingly, empirically evaluated in other markets with similar characteristics, such as the airplane and the commercial real estate industries. Finally, the fact that our model allows, in a straightforward manner, agents to hold distorted beliefs at different degrees, renders it easily expandable and applicable to other real asset markets characterised by – even alternative forms of – distorted investor behaviour.

1.II.C. Summary of Chapter 4

Having concluded the analysis of the physical shipping market for second-hand vessels, in Chapter 4, we examine the derivative market for Forward Freight Agreements (FFAs) and, particularly, the formation of dry bulk FFA rates. Namely, as analysed in Subsection 1.1, the empirical analysis concentrates upon the Capesize BCI 4TC and Panamax BPI 4TC monthly contracts. Our contribution to the literature is threefold. First, by applying a variance decomposition framework for the first time in the FFA market, we illustrate the significant forecasting power of FFA contracts regarding future market conditions. More importantly, we provide both an economic interpretation of this result and a comparison with the ones obtained from other industries. Second, for the first time in the literature, we document several noticeable empirical regularities related to FFA rates and risk premia: in particular, the existence of a bias in the dry bulk FFA market. Third, we develop a theoretical heterogeneous agents' behavioural asset pricing model that can account for the observed regularities.

We begin by analysing empirically the formation of the most frequently incorporated FFA valuation ratio, that is, the FFA basis. Accordingly, by applying a variance decomposition framework – for the first time to shipping derivative markets – we provide strong statistical evidence that the bulk of volatility in the FFA basis can be attributed to expectations about future physical market conditions rather than expectations about future risk premia as is commonly suggested in the commodity markets literature (Fama and French, 1987). Noticeably, our finding validates and extends the economic arguments presented in the seminal commodity market papers (Hazuka, 1984; French, 1986; Fama and French, 1987) that examine the forecasting power of derivative contracts. What is more, this result is perfectly aligned with our respective finding regarding the physical market for ships that the bulk of earnings yields' volatility can be attributed to variation in future market conditions rather than expected returns.

While, however, the bulk of FFA basis' volatility is attributed to future spot growth, we cannot exclude the existence of – time-varying – risk premia. Accordingly, for the first time in the shipping literature, we provide evidence of numerous stylised features that might be of interest to both academic researchers and market participants. First, in contrast to most futures and forwards commodity markets, there is no sign of “backwardation” in any type of contract or maturity in the dry bulk FFA market. More importantly, we find strong statistical evidence of “contango” in the 1-month contracts. Second, we document the existence of a momentum effect in the FFA market; namely, lagged risk premia positively forecast future risk premia in a strong statistical manner. Third, we provide further evidence that there exists – both economically and statistically – significant predictability of future risk premia in this derivative market. The documented predictability is more robust for the Panamax contracts but also for shorter maturities.

In particular, FFA risk premia can be forecasted by both price-based signals and economic indicators related to commodity trade and shipping demand. Regarding the former, there appears to be strong predictability using two lagged spot market indicators and the FFA basis. Regarding the latter, we illustrate that – changes in – economic variables such as commodity prices (e.g., iron ore) and trade indicators (such as the quantities of imported and exported dry bulk commodities) strongly negatively forecast future risk premia. In addition, we provide evidence that future risk premia can also be – negatively – forecasted by past trading activity in the sale and purchase market for second-hand vessels. Interestingly, note that trading activity has been used as an indicator of market liquidity in Chapter 3 of this thesis. Finally, we also test whether future market conditions and risk premia can be predicted by market activity variables that incorporate the FFA trading volume and open interest figures related to the corresponding contracts. While there appears to exist some sort of predictability, mainly in the Capesize sector, the results cannot yet be generalised given the small size of the employed dataset.

From an economic point of view, the documented stylised facts contradict the unbiased expectations hypothesis and, in turn, the efficiency of the FFA market. We further examine the validity of the hypothesis by performing three frequently incorporated econometric tests. Despite the sensitivity of these tests to the model specification, the obtained results unequivocally suggest that there exists a bias in the formation of the 1-month FFA rates in both contracts. Regarding the 2-month contracts, our findings point towards the existence of a bias, especially in the Panamax BPI 4TC case. Consequently, our empirical estimation results are robust and consistent. Therefore, we demonstrate formally, for the first in the shipping literature, the existence of a bias in the dry bulk FFA market.

Accordingly, in order to justify these findings, we develop a dynamic asset pricing framework that can incorporate both the familiar “hedging pressure” feature – the rational dimension – and a heterogeneous-beliefs explanation – the irrational dimension. The distinct feature of our framework is that, apart from having different objective functions, agents – that is, ship owners, charterers, and speculators – might also differ in the way they form expectations about future market conditions. Specifically, speculators are assumed to have distorted beliefs for two reasons: due to asymmetric and imperfect information but mainly due to a behavioural bias known as “the law of small numbers” or “gambler’s fallacy”.

From an economic perspective, the assumption of asymmetric and imperfect information can be justified by the fact that ship owners and charterers – who participate also in the physical market and, thus, have “inside” information regarding the actual future market conditions – are expected to

be able to form more accurate forecasts about future spot rates than speculators – who participate only in the FFA market.

Regarding the behavioural bias assumption, speculators in our model expect that a realised shock in current spot prices will be followed by one of the opposite sign in the next period and, as a result, they adopt a contrarian investment. It is well-documented (Grinblatt and Keloharju, 2000; Kaniel *et al*, 2008; Bloomfield *et al*, 2009) that, in practice, traders frequently follow contrarian strategies which can be influenced or motivated by behavioural biases such as the “gambler’s fallacy”. Specifically, there is market evidence that mainly uninformed and inexperienced investors usually adopt contrarian behaviour. Those findings are particularly related to our model since speculators correspond to financial investors who, as non-participants in the physical market, are assumed to be less sophisticated and informed regarding future shipping market conditions compared to traditional physical market agents.

At this point, recall that our empirical analysis in Chapter 3 concluded that the average investor expectations regarding future market conditions must be “near-rational”. In turn, note that the “average investor” of Chapter 3 corresponds to the “ship owner” agent type in Chapter 4. Furthermore, charterers can be plausibly assumed to form rational expectations since they participate in the physical market as well. Therefore, the average physical investor expectations in Chapter 4 can be plausibly assumed to be “near-rational” as well – for simplicity and without loss of generality, we assume that physical players are totally rational.⁸

Finally, since, there are no surveys regarding shipping industry participants’ beliefs and investment strategies as in the equity markets literature (Greenwood and Shleifer, 2014), we further test and justify our heterogeneous expectations explanation by contradiction, that is, using both theoretical predictions and numerical simulations of the proposed framework. Specifically, it is illustrated formally that, to simultaneously match the observed regularities, one must depart from the rational expectations benchmark of the model. While the predictions are not particularly sensitive to the degree of information asymmetry, this is not true for the behavioural bias feature; namely, a fraction of investors must suffer from the “gambler’s fallacy” and, in turn, follow a contrarian investment strategy.

1.II.D. Future Research

As analysed above, this thesis examines both the second-hand market for vessels and the FFA market of the dry bulk shipping industry from a partial equilibrium perspective. Accordingly, an idea

⁸ Note that it is straightforward to account for slightly extrapolative beliefs on behalf of ship owners in our framework. Even if we do so, however, the qualitative predictions and conclusions of our model are not affected.

for future research is to extend the heterogeneous agents economy to be able to account for the aggregate dry bulk industry from a general equilibrium point of view. Namely, we would like to develop a theoretical model to analyse how the newbuilding, scrapping, and sale and purchase investment decisions of shipping agents are jointly determined in equilibrium and, in turn, their interrelation with the respective spot and FFA rates. To the best of our knowledge, this has not been examined through a structural heterogeneous agents' model in the shipping literature before. Furthermore, the proposed methodologies in Chapters 2 and 3 can be incorporated for the valuation of assets – and, in turn, tested empirically – in other real asset economies such as the real estate and airline industries. Finally, since Chapter 4 illustrates the existence of a bias in FFA rates, it would be interesting (especially from an industry participant's perspective such as "shipping commodity hedge funds") to examine potential profitable trading strategies that incorporate the documented stylised facts.

Chapter 2: The Earnings Yield and Predictability of Earnings in the Dry Bulk Shipping Industry

Abstract. This chapter examines the relation between second-hand vessel prices, net earnings, and holding period returns in the shipping industry. Specifically, we concentrate on the Capesize, Panamax, Handymax, and Handysize dry bulk sectors. We demonstrate that the bulk of variation in shipping earnings yields reflects varying expectations about net earnings growth, not time-varying expected returns, and not varying expectations about the terminal earnings yield. Equivalently, dry bulk vessel prices – mainly – move due to news about net earnings and not due to news about returns. Technically, we contribute to the literature by extending the Campbell-Shiller framework to real assets with limited economic lives and incorporating a forward-looking definition of the corresponding valuation ratio. Our results strongly indicate that shipping earnings yields negatively forecast future net earnings growth while there is no consistent, significant statistical evidence of time-varying risk premia in the second-hand dry bulk shipping industry. In addition – by examining a real, capital intensive industry with distinct supply and demand mechanism – we provide an economic interpretation for the obtained results. Accordingly, we argue that for significant cash flow predictability to exist, current cash flows must have a profound second-order effect on the current price of the asset through the future cash flow stream. Based on this argument, we explain the similarities and differences in the respective findings across different industries. In particular, our results are in sharp contrast to the empirical asset pricing literature corresponding to the post-WWII U.S. equity markets. Importantly, however, our findings agree with recent researches in the pre-WWII U.S. and the bulk of global equity and the U.S. real estate markets. From a statistical perspective – and in line with recently obtained evidence – we argue that the significant predictability of earnings growth by the earnings yield is driven by the extreme volatility of shipping net earnings. To the best of our knowledge, these stylised facts had never been documented formally in the shipping literature before.

Keywords: Asset Pricing, Vessel Valuation, Expected Earnings, Expected Returns, Earnings Yield, Variance Decomposition

2.1. Introduction

This chapter investigates the formation of vessel valuation ratios and, in turn, second-hand vessel prices in the shipping industry. In particular, our empirical analysis focuses on the Capesize, Panamax, Handymax, and Handysize dry bulk sectors. Our contribution to the literature is threefold. First, we extend mathematically the Campbell-Shiller variance decomposition and vector autoregression (VAR) frameworks (1988b and 1988a, respectively) to capture the case of real assets with limited economic lives. Second, using this extended framework, we illustrate formally for the

first time in the shipping literature that the bulk of variation in vessel valuation ratios – that is, shipping earnings yields – reflects varying expectations about net earnings growth, not time-varying expected returns, and not varying expectations about the terminal ratio. Namely, our results strongly indicate that the shipping earnings yield negatively forecasts future net earnings growth while there is no consistent, significant statistical evidence of time-varying risk premia in the dry bulk shipping industry. Equivalently, dry bulk vessel prices mainly move due to news about future net earnings and not due to news about future returns. Third, we provide an economic interpretation for the obtained results based on which we further explain, in a theoretical manner, the similarities and differences in the respective findings across different industries.

We begin by analysing empirically the formation of the most frequently incorporated vessel valuation ratio, that is, the shipping earnings yield,⁹ defined as the (log) ratio of the one-period net earnings – or, equivalently, operating profits – to the prevailing second-hand vessel price. As it is well-analysed in the empirical asset pricing literature, predictability of future returns and/or predictability of future cash flow growth constitute the rational benchmark for the interpretation of variation in assets' valuation ratios (Bansal and Yaron, 2004). Accordingly, to explain which of the sources is the major driving force behind the volatility of asset valuation ratios, researchers have been – mainly – applying the Campbell-Shiller (1988b) empirical framework (Chen *et al*, 2012). Specifically, this framework is based on and, in turn, answers a question of relative predictability which can be applied to a variety of asset classes, both financial and real ones.

It has been extensively demonstrated that in the aggregate post-WWII U.S. equity markets virtually all variation in dividend yields – defined as the (log) ratio of the one-period dividend to the prevailing stock price – is the result of time-varying expected future returns or, equivalently, time-varying discount rates. In particular, dividend yields positively forecast future returns while future dividend growth appears to be unpredictable (Cochrane, 2005). Consequently, the bulk of empirical asset pricing research has concentrated on time-varying discount-rate theories in order to explain the formation of asset prices (Cochrane, 2011).¹⁰

Regarding the international equity markets, recent evidence suggests that these patterns do not extend uniformly to a cross-country setting. Specifically, while in larger countries – in terms of market capitalisation – like France, Germany, the United Kingdom, and the U.S.A., dividend growth

⁹ The terms “earnings yield” and “net earnings-price ratio” are used interchangeably in the context of this research. Accordingly, for the “dividend yield” and “dividend-price ratio” in equity markets and “rent yield” and “rent-price ratio” for the real estate ones.

¹⁰ These developments contradict the prevailing belief during the first era of asset pricing per which returns are unpredictable and variation in dividend yields reflects variation in expected dividends. In line with this former belief, Chen (2009) presents evidence of strong dividend growth predictability by the dividend yield in the aggregate U.S. equity markets data during the pre-WWII period – specifically, during the period 1872-1945. Interestingly, though, this predictability entirely disappears in the post-WWII period.

rates are not predicted by dividend yields (Ang and Bekaert, 2007), this is not true for smaller ones such as Denmark and Sweden (Engsted and Pedersen, 2010). In addition, Rangvid *et al* (2014) broaden the empirical research by examining dividend growth and return predictability using a sample consisting of 50 countries. Their results suggest that in most of these countries dividend yields do strongly predict future dividend growth rates.

Apart from the wide literature that focuses on financial markets, the Campbell-Shiller (1988b) empirical framework has also been expanded into the real estate sector. The respective forecasting variable in this case is the rent yield which corresponds to the (log) ratio of the one-period rent (net of all operating and maintenance expenses) divided by the estates' current price (Hamilton and Schwab, 1985). However, the major problem with the real estate literature is that the empirical results and, in turn, the conclusions cannot be easily generalised due to the severe heterogeneity that characterises the market. This heterogeneity stems from the large variety in both geographical and physical characteristics of properties (Capozza *et al*, 2004). The most necessary distinction imposed by researchers is the one between residential and commercial properties due to the different dynamics, cash flow, and return properties related to the two types of premises (Geltner and Miller, 2006). Despite this heterogeneity, however, most of the recent U.S. real estate literature suggests that there is no strong statistical evidence of future returns predictability by the rent yield. Interestingly, though, in most cases future rent growth appears to be strongly predictable. In conclusion, as Ghysels *et al* (2012) illustrate, the bulk of volatility in the U.S. real estate industry can be attributed to variation in future rent growth as opposed to future returns.

Despite the expansion of this empirical framework to real estate markets, not much research has been conducted in other real assets' economies and, specifically, shipping. As mentioned above, similar to the dividend and rent yields, the respective valuation ratio in shipping is the earnings yield, defined as the (log) ratio of annual net earnings to the current price of the respective – 5-year old – vessel. Note that – in analogy to the one-period dividend and net rent in equity and real estate markets, respectively – one-period net earnings are the corresponding shipping cash flow. Specifically, shipping net earnings represent the one-period operating profit (that is, the revenue net of operating and maintenance expenses) realised by the investor – ship owner – from holding the asset – vessel – for one period.

Since, however, shipping investors know in advance the net earnings variable for the forthcoming period, we construct a “forward-looking” earnings yield which we believe is more consistent with reality; hence, more capable of exploiting the properties of our data. One could object that this is also the case in real estate markets – or, in general, in most real asset economies – since property

owners agree in advance with lessees upon the rent corresponding to the forthcoming period. Therefore, we believe that – conditional on data characteristics and availability – the proposed extension is more realistic and appropriate for the valuation of tangible assets within this framework.

Regarding the existing shipping literature, a first theoretical point towards this asset pricing direction was the suggestion by Greenwood and Hanson (2015) that the earnings yield – in the dry bulk Panamax sector – must strongly forecast low future earnings growth. The authors base this argument on the nature of competition in shipping but do not justify it empirically. Furthermore, a number of studies have incorporated a definition of the shipping earnings yield. Namely, Alizadeh and Nomikos (2007) identify a long-run cointegrating relationship between net earnings and vessel prices. Accordingly, they use the established relationship to develop investment technical trading strategies in the dry bulk sector. Their results suggest that shipping earnings yields comprise substantial information about future market conditions that can benefit agents when making their investment decisions. Their paper, though, incorporates a different, “lagged”, definition of the earnings yield compared to the one used in this chapter. In addition, the theoretical framework and the corresponding empirical methodology included in that paper do not account for the economic depreciation in the value of the vessel.

Papapostolou *et al* (2014) use the earnings yield as a sentiment proxy that captures market valuation in order to construct a shipping investors’ sentiment index. The paper by Papapostolou *et al* is closely related to this research for two reasons. First, the authors also define a “forward-looking” earnings yield.¹¹ Second, they argue that a high earnings yield is associated with positive sentiment which, in turn, serves as a contrarian indicator for future shipping conditions. Hence, similar to our findings, Papapostolou *et al* (2014) suggest that the earnings yield is negatively related to future cycle phases. The latter argument is also in line with Greenwood and Hanson (2015) who implicitly assume the same “forward-looking” definition of the earnings yield as we do. Nevertheless, none of those papers examine formally the relation between shipping earnings yields and future net earnings growth. Except for the purpose of filling this gap in the shipping literature, however, an extension of the Campbell-Shiller (1988b) asset pricing framework to shipping markets appears to be interesting for the following reasons.

As mentioned in Chapter 1, shipping industry is highly important to the world economic activity (Killian, 2009) since vessels transport roughly 90% of the world trade (Papapostolou *et al*, 2014). Furthermore, the price of a cargo carrying vessel amounts to tens of millions of dollars depending on

¹¹ Papapostolou *et al* (2014) define the earnings-price ratio as the ratio of annual earnings (revenue) to the corresponding 5-year old vessel price. Therefore, they do not incorporate net earnings in their definition since operating and maintenance expenses are not subtracted from the annual revenue.

the vessel's type and age and the prevailing market conditions. Accordingly, from an economic point of view, it is important to understand the pricing dynamics of this asset class. Moreover, in contrast to equity markets, shipping industry consists of tangible assets with limited economic lives. Hence, the value of a vessel is substantially affected by economic depreciation. While residential and commercial premises are also real assets, the existing literature does not incorporate any adjustment to account for this feature. From a technical perspective, therefore, the application to shipping is important because it enables us to extend the Campbell-Shiller variance decomposition and VAR frameworks (1988b and 1988a, respectively) to account for both forward-looking valuation ratios and economic depreciation in the value of the respective asset. To the best of our knowledge, this is the first time that these features are explicitly incorporated in this asset pricing framework. Accordingly, the proposed methodologies can be used for the valuation of assets in other real asset economies such as the real estate and airline industries.

Furthermore, shipping markets can be examined from a worldwide perspective since (i) sector-specific shipping services are perceived as homogeneous (Kalouptsidi, 2014), (ii) the total number of assets in this industry is relatively small, and (iii) the supply of and demand for vessels are derived globally. Therefore, our results are robust and apply to the entire dry bulk shipping industry. In contrast, this is neither the case for the real estate industry nor for equity markets. Since shipping is a capital-intensive industry with clear and directly observable supply and demand determinants and mechanism, it provides an ideal environment to build a bridge between this empirical asset pricing framework and the economics of the market. Subsequently, this reasoning can be extended to other real asset industries as commercial real estate.

Finally, the underlying economic principles and industrial organisation of this market are unique and lead to the well-documented shipping boom-bust cycles (Stopford, 2009) which, in turn, result in extremely volatile cash flows. In relative terms, shipping cash flows exhibit noticeably more volatile behaviour over time compared to vessel prices. This stylised fact coincides with evidence from a significant number of international equity markets (Rangvid *et al*, 2014). By contrast, in the post-WWII U.S. equity markets, while stock prices fluctuate significantly, the respective dividends appear to be relatively smooth¹² over time (Fama and French, 1988b). From a statistical point of view, therefore, in line with Chen *et al* (2012) and Rangvid *et al* (2014), we should a priori expect that in real asset economies – where the cash flows received by investors cannot be smoothed – and especially in industries in which cash flows are extremely volatile, the corresponding valuation ratios

¹² Following the definition by Chen *et al* (2012), dividend smoothing is the phenomenon where dividend payout is determined not only by current or permanent earnings but also by past dividend payout (Lintner 1956; Marsh and Merton, 1987).

will strongly predict future cash flow changes. Indeed, as we demonstrate in the following sections, this is precisely the case in the dry bulk shipping industry.

From an economic point of view, in order for valuation ratios to move due to expectations about future cash flows, the latter should be predictable by market agents using the current information set. Vice versa, if future cash flows are not predictable using current market information, then they can neither be predicted by the earnings yield. We argue that the major determinants of valuation ratios are the second-order effects (SOEs) that the current cash flows have on current prices through the future cash flow stream. If there are no profound SOEs, then there is no reason for future cash flows to be predictable by the current information filtration.

To the best of our knowledge, this is the first time in the shipping literature that the bulk of shipping earnings yields' variation is attributed formally and unambiguously to variation in future net earnings growth. From a vessel valuation point of view, our results imply that vessel prices vary – mainly – due to news related to expected net earnings, not due to news about expected returns, and not due to news about the terminal – scrap – price of the vessel. Furthermore, according to our empirical estimation, there is no consistent, significant evidence of time-varying expected returns or, equivalently, time-varying risk premia in the second-hand dry bulk shipping industry. The obtained results are both economically and statistically significant. In turn, this implies that, from a market participant's perspective (specifically, in terms of developing forecasting and trading strategies), the earnings yield can be incorporated as a significant predictor of future spot market conditions but not as an indicator related to future risk premia.

Moreover, by relating the obtained empirical results to economic principles, we provide a comparison with the main equity and real estate findings and explain theoretically the observed similarities and differences. Specifically, our results are diametrically opposed to the ones in the post-WWII U.S. equity markets and residential (housing) real estate markets but in line with the ones obtained from both the pre-WWII U.S. equity markets and the bulk of international equity markets as well as the majority of the commercial real estate industry – and the REIT index market. Therefore, this chapter provides strong evidence for further discussion regarding the economic principles that drive the forecasting properties of valuation ratios.

The remainder of this chapter is organised as follows. Section 2.II introduces the dataset employed and analyses the main variables of interest along with some preliminary results. Section 2.III illustrates the methodology and the main empirical findings of this chapter. Section 2.IV discusses the economic rationale behind the empirical results and provides a theoretical comparison with the stock and real estate markets in terms of the respective findings. Section 2.V concludes.

2.II. Data and Variables of Interest

The dataset employed consists of monthly and quarterly observations on newbuilding (N/B), second-hand (S/H), and scrap vessel prices and 6-month and 1-year time-charter rates for each of the four dry bulk sectors under consideration.¹³ In addition, we have obtained data for various supply and demand variables related to the dry bulk shipping industry.¹⁴ Our main shipping data source is Clarksons Shipping Intelligence Network.¹⁵ The operating and maintenance expenses are approximated through discussions with industry participants and the adopted values agree with estimated figures in the recent literature. In addition, data for the U.S. Consumer Price Index (CPI) are obtained from Thomson Reuters Datastream Professional.

In line with previous researches (Greenwood and Hanson, 2015), due to the construction lag between the ordering and delivery of a new vessel, we use the second-hand – 5-year old – vessel price instead of the newbuilding one. A major implication of this feature is that the owner of the newly ordered ship is not able to immediately exploit the prevailing net earnings. As a result, the Campbell-Shiller (1988b) present value relationship cannot be directly applied. In contrast, though, the owner of the second-hand vessel can lease the asset immediately; hence, the present value identity can be incorporated in a straightforward manner. Moreover, due to data availability, we particularly choose the price of a 5-year old vessel. Importantly, our results are not sensitive to this choice.

Table 2.1: Sample and sector characteristics for dry bulk shipping.

Sector	Sample period	T	Representative vessel (dwt)	Costs (\$)
Capesize	1/1992-12/2014	276	180,000	8,000
Panamax	1/1976-12/2011	432	76,000	7,000
Handymax	4/1986-6/2014	339	56,000	6,500
Handysize	1/1976-12/2014	468	32,000	5,500

Notes: The number of observations in the sample is denoted by T . Costs are expressed in – December 2014 –dollars per day and refer to the total operating and maintenance expenses of the vessel.

The main problem with the initial shipping dataset is that vessel prices and time-charter rates do not refer to vessels of the same cargo carrying capacity, at least for a significant time interval. Accordingly, we have constructed new time-charter rates series by adjusting the initial time-charter rates to the size of the vessel on which the vessel price time-series refer, using an appropriate scale

¹³ In particular, our second-hand price dataset consists of observations for 5, 10, 15, and 20-year old vessels.

¹⁴ Unfortunately, some of the supply and demand variables can only be provided on an annual basis.

¹⁵ Clarksons are considered as the most significant and widely used provider of shipping data to market participants (Greenwood and Hanson, 2015); this renders our dataset both accurate and easily retrievable.

factor.¹⁶ In line with previous researches (e.g., Greenwood and Hanson, 2015), we assume that vessels operate in consecutive one-year time-charter contracts. In this type of arrangement, only the operating and maintenance expenses are borne by the ship owner. As it is common in the literature, taxes and interest expenses are ignored from the analysis. After discussions with industry participants, we have approximated the summation of daily operating and maintenance costs for the representative 5-year vessel for each of the four dry bulk sectors (see Table 2.1). Following Stafford *et al* (2002), we assume that, for a given vessel, operating and maintenance costs increase with inflation. Since the analysis is conducted under real terms, we define the December 2014 nominal figures as our benchmark real values. In addition, we assume that vessels spend 10 days per annum in maintenance and repairs (Stopford, 2009). During this out-of-service period, ship owners do not receive the corresponding time-charter rates but bear the operating and maintenance expenses.¹⁷

An inherent characteristic of the shipping industry is that revenues for period $t \rightarrow t + 1$ are agreed and, in turn, determined at time t . Indeed, assuming that the charterer will not default on his payment obligations, owners know in advance and with certainty the amount that they will receive in one-period time.¹⁸ Implicitly, we also assume that the operating and maintenance expenses that refer to period $t \rightarrow t + 1$ do not deviate from the estimated figure at time t . Hence, one-period-ahead net earnings calculations are performed with the respective time-charter rates and expenses prevailing at time t – both quoted in dollars per day. Mathematically, the one-period-ahead net earnings variable in shipping is \mathcal{F}_t -measurable. This point is very important since in the equity and real estate markets literatures, the generated cash flows – i.e., dividends and rents, respectively – corresponding to period $t \rightarrow t + 1$ are assumed to be unknown at time t .

Furthermore, our net earnings estimation accounts also for the commission that the brokering house receives for bringing the ship owner and the charterer into an agreement. Once again, after

¹⁶ For illustrative purposes, consider the Panamax sector where 5-year old vessel prices correspond to a 76,000-dwt carrier while the time-charter rates to a 65,000 dwt one. Following Greenwood and Hanson (2015), we scale the time-charter rates by multiplying the initial time series by 76/65. While it was not feasible to retrieve the same sample size for all sectors, we have acquired the largest possible dataset for each one. Furthermore, all our samples include the most recent shipping crisis of 2008. Since the Panamax sample ends in December 2011, we have tested for robustness another subsample ending in December 2014 and the obtained results coincide.

¹⁷ Note that we have examined the robustness of our estimation for several values of both the operating and maintenance expenses and the duration of the out-of-service period; the obtained results are qualitatively the same and quantitatively very similar to the ones presented here.

¹⁸ In practice, ship owners and charterers agree upon the time-charter rate of the vessel – for the entire leasing period – before the corresponding leasing period begins. Accordingly, the agreed time-charter rates are typically received every 15 days – sometimes also in advance. As a result, the probability of default on the part of the charterer is reduced. Moreover, a wide brokering network and the fact that ship owners normally lease their vessels to solvent charterers assure transparency and low probability of default. In contrast, a time-charter lease with a less creditworthy charterer will incur higher rates to compensate the ship owner for the higher probability of default on the part of the charterer. Finally, additional contractual agreements included in the charter party ensure that the owner will receive the full time-charter rate agreed.

discussions with industry participants we have estimated this brokerage fee to be approximately equal to 2.5% of the daily time-charter rate.¹⁹ Accordingly, the one-period net earnings for the owner of a vessel, Π_{t+1} , are estimated through the following equation:

$$\Pi_{t+1} \equiv \Pi_{t \rightarrow t+1} = 355 \cdot 0.975 \cdot TC_{t \rightarrow t+1} - 365 \cdot OPEX_{t \rightarrow t+1}, \quad (2.1)$$

where $TC_{t \rightarrow t+1}$ and $OPEX_{t \rightarrow t+1}$ refer to the corresponding daily time-charter rates and the sum of daily operating and maintenance expenses, respectively. Due to the fact that taxes are excluded from the analysis, depreciation costs are not deducted from the gross revenue. Finally, our estimation procedure implicitly assumes that net earnings realised by a specific vessel are not a function of her age. This adjustment, however, does not have a qualitative impact on the results.

Since vessels are real assets with limited economic lives, we must account for economic depreciation. In particular, at each point in time, a 6-year old vessel is less valuable than an identical 5-year one for two reasons. Namely, the former has one less year of future economic life but also deteriorated performance compared to the latter. In line with previous researches, we assume that a newly built vessel has an economic life equal to 25 years. Since Clarksons only provide us with 5-, 10-, 15-, 20-year old, and scrap vessel prices, we need to adopt a depreciation scheme to approximate the price of $(5 + n)$ -year old vessels at each time t . We denote this price by $P_{5+n,t}$, for $0 \leq n \leq 20$, where n corresponds to an integer; thus, the price of a 5-year old vessel at time t is defined as $P_{5,t}$.

Accordingly, we first estimate the average price ratios of 10-year to 5-year, 15-year to 10-year, 20-year to 15-year, and scrap to 20-year old vessels. These ratios are approximately equal to 0.75 in all dry bulk categories. Consequently, the assumption of a straight-line depreciation scheme for each 5-year age window implies a 5% annual value reduction compared to the price of the youngest vessel in the interval at each corresponding time t . In order to illustrate the adopted depreciation mechanism, consider the prices of vessels between 5 and 10 years of age. At each t , these prices can be estimated using the formula:

$$P_{5+n,t} = (1 - 0.05n) \cdot P_{5,t}, \quad 1 \leq n \leq 5. \quad (2.2)$$

Accordingly, the one-year horizon raw return is estimated using the formula:

$$R_{n,t \rightarrow t+1} \equiv R_{n,t+1} = \frac{\Pi_{t+1} + P_{n+1,t+1}}{P_{n,t}}, \quad (2.3)$$

where $R_{n,t+1}$ is the – ex post – holding period raw return realised at time $t + 1$ from an investment made at time t for a n -year old vessel; thus, this variable quantifies the return realised by an investor

¹⁹ Roughly, this fee is the sum of the brokerage and address commissions.

who owns the vessel for one year. As analysed above, due to economic depreciation of the vessel, we set $P_t \equiv P_{n,t}$ and $P_{t+1} \equiv P_{n+1,t+1}$. Intuitively, this formula assumes that the investor at time t purchases the vessel at price $P_{n,t}$ and immediately leases her out for one year to earn the one-period profit, Π_{t+1} . In addition, at the end of period $t \rightarrow t + 1$ – that is, at $t + 1$ – the owner sells the vessel²⁰ at the prevailing market price, $P_{n+1,t+1}$. In the following, for expositional simplicity, we drop the age index from the return's notation.

In a pure empirical asset pricing context, we would like to examine the source of vessel price volatility; that is, investigate whether vessel prices vary due to changing forecasts about future net earnings, changing forecasts about future returns, and/or due to changing forecasts about the terminal –scrap – price of the vessel. In order to answer this question, we apply the familiar Shiller (1981) present value framework to the shipping case. Furthermore, in line with the related literature, we linearise this relationship by incorporating the log transformation of all variables of interest – that is, net earnings, prices, and one period returns.

Since – as is commonly the case in the empirical literature – the log transformation of net earnings and prices yields nonstationary variables – in particular, variables that are integrated of order one, $I(1)$ – we follow Campbell and Shiller (1988a and 1988b) and Cochrane (1992) by incorporating the established cointegrating relationship between log shipping net earnings and log vessel prices (Alizadeh and Nomikos, 2007). Specifically, we obtain a stationary variable by taking the first difference between the log net earnings variable and the corresponding log 5-year old vessel price – or, equivalently, the natural logarithm of the ratio of net earnings to the prevailing 5-year old vessel price. Accordingly, we define the ratio as the shipping earnings yield and we denote it by $\frac{\Pi_{t+1}}{P_{5,t}}$.

Importantly, the earnings yield can be interpreted as a valuation ratio that measures the profit from utilising the vessel for the period $t \rightarrow t + 1$ as a fraction of the prevailing price of the asset at t . From an investor's perspective, a high – low – earnings yield reflects the relative degree of undervaluation – overvaluation – in the price of the vessel (Papapostolou *et al*, 2014). This valuation ratio is the natural analogue in shipping of the dividend and rent yields in equity and real estate markets, respectively. This parallelism can be justified by the fact that, in all three cases, the numerator of the corresponding valuation ratio is the annual net income to the investor who holds the asset while the denominator is the current price of the asset.²¹

²⁰ When a vessel is sold in the second-hand market, the sale and purchase (S&P) broker usually receives commission equal to 1% of the resale price. However, in the context of this research, we ignore this transaction cost since it complicates the mathematical analysis in Section 2.III while it has a negligible effect on the empirical results.

²¹ Net earnings in shipping are the equivalent of dividend and net rent in equity and real estate markets, respectively. In other words, as in equity (real estate) markets the net income to an investor who holds the asset for one year comes in the form of the annual dividend (net rent), this net income in shipping consists of

Arguably, though, there exists a significant difference in the definition of our shipping valuation ratio compared to the ones incorporated in the existing asset pricing literature. Specifically, in the equity and real estate markets, the price figure used in the denominator of the ratio is net of the respective dividend – rent – value used in the numerator. Mathematically, the ratio $\frac{D_t}{P_t}$ corresponds to the last net income paid up to time t – that is, the one paid during period $t - 1 \rightarrow t$ – divided by the asset price at t . Campbell and Shiller (1988b) argue that dividends are lagged as to be ensured that they are \mathcal{F}_t -measurable. Hence, the buyer of the asset at time t is not entitled to dividend – rent – D_t but to the net income stream $\{D_{t+i}\}_{i \geq 1}$. Note that previous shipping researches (Alizadeh and Nomikos, 2007) have used a lagged interpretation of the earnings yield, defined as $\frac{\Pi_t}{P_{5,t}}$, which is analogous to the dividend – rent – yield used in the equity – real estate – markets.

Consequently, in the context of this research and in accordance with Papapostolou *et al* (2014), we suggest that the appropriate valuation ratio in shipping is “forward-looking”. This adjustment can be justified by the fact that, in shipping, Π_{t+1} is known in advance. Recall that the owner of the vessel at time t is also entitled to the deterministic, at time t , value of net earnings, Π_{t+1} – and in total to the one-period net earnings stream $\{\Pi_{t+i}\}_{1 \leq i \leq 20}$. Thus, in contrast to equity markets, the \mathcal{F}_t -measurable shipping cash flow Π_{t+1} does not only serve as a forecasting scheme for future cash flows; it is the first term of the forthcoming cash flow series. This is equivalent to saying that the asset is trading “cum dividend” in equity markets.


In support of this statement, Fama and French (1988b) argue that the most commonly incorporated definition of the dividend yield in equity markets, D_t/P_t , has the following drawback. While stock prices, P_t , are forward-looking, the incorporated dividend, D_t , is old relative to the dividend forecasts in P_t . Accordingly, positive news about future dividends results in a high price P_t relative to the last dividend D_t which, in turn, implies a low current dividend yield, D_t/P_t . Furthermore, this increase in P_t produces also a high return $r_{t-1 \rightarrow t}$ and, as a result, there is a negative correlation between the disturbance ε_{t-1} and the time t shock to D_t/P_t . Consequently, the slope coefficients in regressions of $r_{t \rightarrow t+1}$ on D_t/P_t tend to be upward-biased. On the other hand, the alternative measure – that is, D_t/P_{t-1} – does not incorporate the entire information set available at time t . Hence, it is expected to have lower forecasting ability (specifically, to be too conservative) compared to D_t/P_t . In general, however, the recent empirical asset pricing literature in equity markets uses the more “updated” definition of the dividend yield, D_t/P_t . Since in shipping both net earnings and prices are forward-looking, the proposed net earnings-to-price ratio, $\Pi_{t+1}/P_{5,t}$, is time-consistent.

the difference between the annual time-charter rates received and the annual operating and maintenance expenses paid by the owner (equation 2.1).


In conclusion, instead of focusing on vessel prices, the main variable of interest in our empirical estimation is the shipping earnings yield. Equivalently, using the Campbell-Shiller (1988b) variance decomposition framework, instead of detecting the primary source of vessel price volatility, we will examine the primary source of earnings yield's volatility (Cochrane, 2005).

Panel C: Handymax sector from 4/1986 to 6/2014.

Panel D: Handysize sector from 1/1976 to 12/2014.



Panel A: Capesize sector from 1/1992 to 12/2014.



Panel B: Panamax sector from 1/1976 to 12/2011.


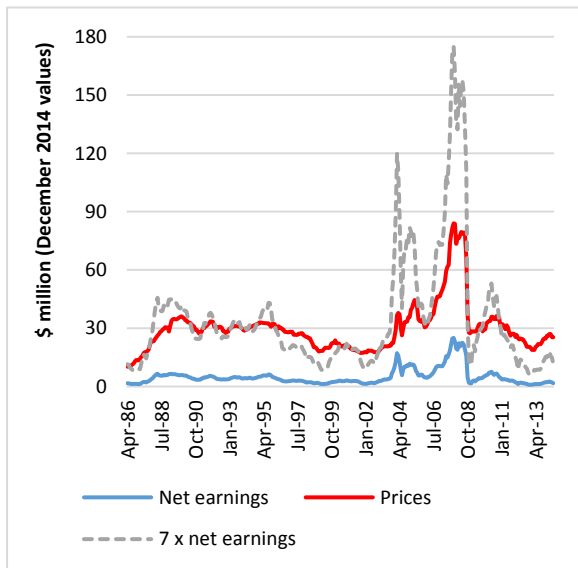



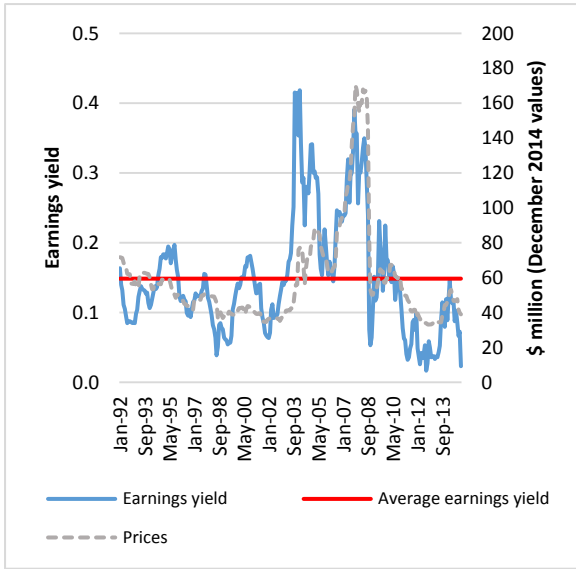


Figure 2.1: Net Earnings and Vessel Prices.

Panels A-D present real annual net earnings and real 5-year old vessel prices for the representative vessel of each dry bulk sector. Notice that, following Cochrane (2011), the dashed lines correspond to the product of the sector-specific average earnings yield times the corresponding prevailing net earnings.



Panel C: Handymax sector from 4/1986 to 6/2014.

Panel D: Handysize sector from 1/1976 to 12/2014.

Figure 2.2: Earnings Yields and Vessel Prices.

Panels A-D present the evolution of the earnings yield in each dry bulk sector. For illustrative purposes, we have also plotted the respective sample mean earnings yield and the evolution of the corresponding 5-year old vessel prices.

Table 2.2: Descriptive statistics for vessel prices, net earnings, and earnings yields.

Variable	T	Mean	SD	CV	Median	Max	Min	ρ_1	ρ_{12}	ρ_{24}
Panel A: Capesize Sector (from January 1992 to December 2014)										
Π (\$m)	276	10.46	11.66	1.12	6.35	60.91	0.57	0.97	0.41	0.15
P (\$m)	276	58.61	28.31	0.48	50.25	170.25	33.13	0.98	0.51	0.24
Π/P	276	0.15	0.08	0.57	0.13	0.42	0.02	0.95	0.47	0.18
Panel B: Panamax Sector (from January 1976 to December 2011)										
Π (\$m)	432	5.03	4.66	0.93	3.58	30.11	0.02	0.97	0.25	-0.06
P (\$m)	432	34.21	15.58	0.46	32.23	103.05	11.90	0.98	0.51	0.23
Π/P	432	0.13	0.06	0.47	0.13	0.35	0.00	0.94	0.24	-0.07
Panel C: Handymax Sector (from April 1986 to June 2014)										
Π (\$m)	339	4.97	4.39	0.88	3.81	24.98	0.86	0.97	0.39	0.13
P (\$m)	339	29.82	12.67	0.42	28.45	84.01	10.15	0.98	0.52	0.20
Π/P	339	0.15	0.07	0.45	0.14	0.47	0.04	0.97	0.40	0.09
Panel D: Handysize Sector (from January 1976 to December 2014)										
Π (\$m)	468	3.09	2.46	0.80	2.53	14.86	0.59	0.98	0.43	0.17
P (\$m)	468	22.16	8.46	0.38	21.32	58.64	5.61	0.98	0.60	0.31
Π/P	468	0.13	0.05	0.43	0.12	0.32	0.04	0.96	0.41	0.12

Notes: This table presents descriptive statistics related to real net earnings, real prices, and the corresponding net earnings yields. The included statistics are the number of observations, mean, standard deviation, coefficient of variation, maximum, minimum, 1, 12, and 24-month autocorrelation coefficients. Real net earnings, Π , refer to the one-year time-charter revenue minus the operating and maintenance expenses, all expressed in December 2014 million dollars. Real price, P , refers to the price of a 5-year old vessel, expressed in December 2014 million dollars, while Π/P denotes the earnings yield.

Table 2.2 summarises descriptive statistics related to annual net earnings, 5-year old vessel prices, and earnings yields for the four dry bulk sectors while Figures 2.1 and 2.2 illustrate the evolution of these variables. As Figure 2.1 depicts, both net earnings and vessel prices are characterised by highly volatile behaviour. Specifically, in line with Table 2.2, while both variables are highly persistent in the one-month horizon, the autocorrelation coefficients are rapidly reduced as the horizon increases.²² This finding verifies the boom-bust nature of the shipping industry. In a cross-sector comparison, we observe that the means, standard deviations, and coefficients of variation of annual net earnings increase with the size of the vessel. This result suggests that larger vessels generate larger but also more volatile cash flows (Alizadeh and Nomikos 2007). Interestingly, when we compare the stochastic behaviour of net earnings to the one of vessel prices, we observe that the former variable appears to be significantly more volatile –in relative terms – than the latter.

²² This finding is in line with our assumptions regarding the corresponding cash flow processes in the theoretical models of Chapters 3 and 4 of this thesis.

Namely, net earnings have more than two times higher coefficients of variation than prices in all sectors under consideration (Table 2.2).

Furthermore, vessel prices and net earnings exhibit a very high degree of co-movement (Table 2.3).

Therefore, in line with Greenwood and Hanson (2015), we can argue that second-hand vessel prices are highly responsive to the prevailing net earnings. The fact, though, that prices and the respective cash flows are highly correlated does not imply that these two variables change proportionately. If this were the case, net earnings yields would be constant across time. In contrast, as it becomes evident from Figure 2.2 and Table 2.2, they fluctuate significantly in all sectors exhibiting noticeable mean-reverting behaviour.²³ Interestingly, in a cross-industry comparison, shipping earnings yields are significantly less persistent than dividend yields and rent yields in the post-WWII U.S. equity (Cochrane, 2005)²⁴ and real estate markets (Ghysels *et al*, 2012), respectively. As demonstrated in Section 2.III, the empirical implication of this statistical feature is very important since it explains why the slope coefficients and R^2 s of shipping net earnings growth predictive regressions do not increase linearly with the forecasting horizon as in the case of the U.S. equity markets' returns regressions (Cochrane, 2005).

Table 2.3: Correlations between net earnings, vessel prices, and earnings yields.

Market	$Corr(P, \Pi)$	$Corr(\Delta p, \Delta \pi)$	$Corr(\Pi/P, \Pi)$	$Corr(\Pi/P, P)$
Capesize Sector	0.95	0.84	0.86	0.72
Panamax Sector	0.89	0.79	0.81	0.54
Handymax Sector	0.92	0.88	0.82	0.59
Handysize Sector	0.89	0.79	0.84	0.55

Notes: We denote asset prices by P , net earnings by Π , annual log price growth by Δp , annual log net earnings growth by $\Delta \pi$, and earnings yields by Π/P . Furthermore, by $Corr(X, Y)$ we indicate the corresponding correlation coefficient.

Moreover, these preliminary findings demonstrate another important feature of the dry bulk shipping industry. Namely, since the price-earnings ratio reflects the relative degree of overvaluation in the price of an asset, one could argue that vessels are overvalued during market troughs and vice versa. Graphically, following Cochrane (1011), this can also be observed by comparing the dashed lines in Figure 2.1 – which correspond to the product of the sector-specific average earnings yield times the corresponding prevailing net earnings – to the actual price of the vessel. This finding coincides with Papapostolou *et al* (2014) who mention that high price-earnings ratios are associated

²³ The Augmented Dickey-Fuller test suggests that all earnings yield variables satisfy the stationarity condition.

²⁴ Cochrane (2005) has estimated an AR(1) coefficient of 0.9 in annual U.S. equity markets data.

with low sentiment levels in shipping. As mentioned above, the Campbell-Shiller framework incorporates the log transformation of the variables of interest. Accordingly:

$$\pi_{t+1} - p_{5,t} = \ln(\Pi_{t+1}) - \ln(P_{5,t}) = \ln\left(\frac{\Pi_{t+1}}{P_{5,t}}\right), \quad (2.4)$$

where $\pi_{t+1} - p_{5,t} \sim I(0)$ is the shipping log net earnings yield. Furthermore, we estimate the n -period log net earnings growth rate using the formula:

$$\pi_{t+n} - \pi_t = \ln\left(\frac{\Pi_{t+n}}{\Pi_t}\right) = \sum_{i=1}^n \ln\left(\frac{\Pi_{t+i}}{\Pi_{t+i-1}}\right) = \sum_{i=1}^n \Delta\pi_{t+i}, \quad (2.5)$$

where $\Delta\pi_{t+1} = \pi_{t+1} - \pi_t \sim I(0)$ is the one-period log net earnings growth.

In addition, the one-year horizon log return is defined as $r_{t+1} = \ln(R_{t+1})$. Regarding the computation of the multi-year (cumulative) returns, recall that we assume that the vessel is employed in consecutive one-period time-charters. Thus, we estimate the multi-year raw log returns by summing the corresponding one-year ones. When adding, however, the corresponding one-period returns, we must also account for the economic depreciation in the price of the vessel:

$$r_{t+n} = \sum_{i=1}^n \ln\left(\frac{\Pi_{t+i} + P_{5+i,t+1}}{P_{5+i-1,t}}\right), \quad 1 \leq n \leq 20. \quad (2.6)$$

For statistical robustness, the Augmented Dickey-Fuller test was performed on all incorporated log earnings yields, annual net earnings growth rates, and annual log returns. As one would expect, for each of the three variables, the null hypothesis of non-stationarity is rejected.

Finally, the 1-year horizon vessel price growth refers to the growth in the price of a specific vessel – and not to the change in the price of a specific age class across time. Due to economic depreciation, the appropriate formula for the growth rate of the vessel-specific price is the following one:

$$\Delta p_{n+1,t+1} = p_{n+1,t+1} - p_{n,t} = \ln\left(\frac{P_{n+1,t+1}}{P_{n,t}}\right). \quad (2.7)$$


Essentially, equation 2.7 quantifies the annual change in the price of a given vessel. Evidently, this growth rate is closely related to the one-period return; hence, it is the one relevant to our empirical estimation. Note that, in the remaining of this chapter, by price growth rate we mean the change in the price of a specific vessel between the fifth and sixth years of her economic life.

Using equations 2.4-2.7, we compute log earnings yields, 1-, 2-, and 3-year horizon net earnings yields, 1-, 2-, and 3-year horizon net earnings growth rates, 1-, 2-, and 3-year horizon raw log returns, and 1-year horizon vessel-specific price growth rates for each of the four dry bulk sectors – for a representative vessel that at time t was 5 years old. Table 2.4 summarises descriptive statistics related to these variables in the 1-year horizon case while Figure 2.3 demonstrates the evolution of annual log returns, annual net




Panel B: Panamax sector from 1/1976 to 12/2011.

Panel A: Capesize sector from 1/1992 to 12/2014.



Panel C: Handymax sector from 4/1986 to 6/2014.



Panel D: Handysize sector from 1/1976 to 12/2014.

with, and vessel price growth.




Figure 2.3: Net Earnings Growth, Vessel Price Growth, and Returns.

Panels A-D present 1-year horizon annual real net earnings growth, 1-year horizon second-hand vessel real price growth (for a representative vessel that at time t was 5 years old), and 1-year horizon log returns for each dry bulk sector. All growth variables correspond to log differences.

Table 2.4. Descriptive statistics for earnings yields, returns, vessel price growth, and earnings growth.

Variable	T	Mean	Abs Mean	SD	Median	Max	Min	ρ_1	ρ_{12}
Panel A: Capesize Sector (from January 1992 to December 2014)									
$\pi - p$	264	-2.06	-	0.60	-2.01	-0.87	-4.09	0.94	0.50
r	264	0.09	0.20	0.25	0.06	0.70	-0.59	0.96	0.09
Δp	264	-0.07	0.24	0.33	-0.05	0.63	-1.26	0.96	-0.14
$\Delta \pi$	264	-0.02	0.69	0.87	0.05	1.88	-3.07	0.91	-0.21
Panel B: Panamax Sector (from January 1976 to December 2011)									
$\pi - p$	420	-2.12	-	0.53	-2.04	-1.06	-4.58	0.92	0.15
r	420	0.09	0.24	0.29	0.08	0.82	-0.63	0.96	0.01
Δp	420	-0.06	0.27	0.36	-0.04	0.74	-1.33	0.96	-0.08
$\Delta \pi$	420	-0.07	0.76	1.07	0.04	2.65	-5.04	0.93	-0.29
Panel C: Handymax Sector (from April 1986 to June 2014)									
$\pi - p$	327	-1.97	-	0.42	-1.95	-0.76	3.17	0.97	0.39
r	327	0.13	0.21	0.25	0.09	0.67	-0.54	0.97	0.17
Δp	327	-0.03	0.22	0.30	-0.04	0.56	-1.17	0.97	-0.03
$\Delta \pi$	327	0.01	0.54	0.69	0.01	1.74	-2.69	0.96	-0.20
Panel D: Handysize Sector (from January 1976 to December 2014)									
$\pi - p$	456	-2.13	-	0.43	-2.14	-1.14	-3.33	0.96	0.43
r	456	0.07	0.22	0.27	0.06	0.77	-0.54	0.96	0.19
Δp	456	-0.06	0.24	0.31	-0.05	0.72	-0.98	0.96	0.07
$\Delta \pi$	456	-0.03	0.51	0.65	-0.03	1.72	-2.84	0.96	-0.17

Notes: Panels A-D present descriptive statistics for log net earnings yields, $\pi - p$, 1-year horizon log returns, r , 1-year horizon real log S/H (5-year old) vessel price growth, Δp , and 1-year horizon real log net earnings growth, $\Delta \pi$, corresponding to the four dry bulk sectors. The included statistics are the number of observations, mean, mean of absolute values, standard deviation, maximum, minimum, 1, and 12-month autocorrelation coefficients.

Recall from the analysis above that shipping earnings yields fluctuate significantly over time. Mathematically, this finding is equivalent to the statement that during each 1-year period, $t \rightarrow t + 1$, the rates of change of net earnings and vessel prices are not equal. Indeed, as Panels A-D of Figure 2.3 demonstrate, this is precisely the case. Interestingly, returns and vessel price growth rates are relatively smooth compared to net earnings growth. Specifically, in all dry bulk markets and during the entire periods under examination, net earnings growth is significantly higher in absolute value than both returns and price growth.²⁵ Notably, when returns and price growth rates are positive, net earnings growth rates are also positive and vice versa, however, in a substantially greater magnitude. This finding is also supported by the results presented in Table 2.4. Namely, in all dry bulk sectors, the mean of the absolute values of net earnings growth rates is more than twice as high as the respective statistic for prices and returns. Equivalently, the ratio of returns volatility to net earnings volatility is lower than 0.5 in all sectors.

In conclusion, the preliminary analysis in this section suggests that net earnings are noticeably more volatile than vessel prices.²⁶ Interestingly, in the U.S. equity markets we observe the opposite phenomenon at an extreme level. In particular, Fama and French (1988b) find that, after 1940, stock returns are more than 2.4 times as volatile as dividend changes. Consequently, asset prices in the U.S. equity markets exhibit significantly higher volatility than the respective generated cash flows or, equivalently, dividends are relatively smooth compared to the corresponding stock prices. This stylised fact, known as “dividend smoothing”,²⁷ is well-documented in the corporate finance and equity markets empirical asset pricing literatures. As a result, Shiller (1981) and Cochrane (2005) argue that most of U.S. stock prices’ extraordinary volatility does not seem to be accompanied by any important news about fundamentals. In other words, stock price movements cannot be attributed to any objective new information about dividends.

However, as analysed in the Introduction of this chapter, the above U.S. equity markets’ stylised fact cannot be generalised to the international equity markets since, depending on the country, the ratio of returns volatility to dividend volatility can be significantly higher (e.g., Australia) or lower (e.g., Argentina) than one (Rangvid *et al*, 2014). More importantly, Rangvid *et al* (2014) demonstrate

²⁵ The results are qualitatively identical and quantitatively approximately the same when we use sector-specific price growth, that is, the change in the price of the 5-year old age class across time.

²⁶ Importantly, recall that this comparison is conducted under absolute terms.

²⁷ See Allen and Michaely (2003) and Tirole (2006).

that dividend predictability increases with dividend volatility or, equivalently, dividend predictability decreases with dividend smoothing. This significant result can explain the difference in the results across international equity markets. Noticeably, the fact that net earnings are the natural analogue in shipping of dividends in equity markets combined with the observation that net earnings are far from smooth, yield to the a priori expectation of significant cash flow predictability in shipping. Indeed, as we analyse in the following section, this is precisely the case.

2.III. Predictability of Net Earnings in Shipping

2.III.A. Predictive Regressions

In this section, we address the main question of interest of this chapter; that is, whether shipping earnings yield vary due to expectations about future returns, expectations about future net earnings growth or expectations about the terminal spread between the resale – scrap – price of the vessel and the respective net earnings prevailing in the market. Extending the Cochrane (2005) argument to shipping, a high price-net earnings ratio should forecast either high future net earnings growth or/and

low future returns or/and a high terminal ratio:

$$\frac{P_{5,t}}{\Pi_{t+1}} = E_t \left[R_{t+1}^{-1} + R_{t+1}^{-1} \cdot \sum_{i=1}^{19} \left(\prod_{j=1}^i R_{t+j+1}^{-1} \cdot \frac{\Pi_{t+j+1}}{\Pi_{t+j}} \right) + \left(\prod_{j=1}^{20} R_{t+21-j}^{-1} \cdot \frac{\Pi_{t+22-j}}{\Pi_{t+21-j}} \right) \cdot \frac{S_{t+20}}{\Pi_{t+21}} \right], \quad (2.8)$$

where $S_{t+20} \equiv P_{t+20}^{25}$ denotes the scrap price of the vessel 20 years ahead (i.e., at $t + 20$).

Equation 2.8 holds ex post as an identity and justifies the above present value statement in shipping (Appendix 2.A presents the analytical derivation of this equation). However, since the non-linearity of (2.8) renders simple time-series tools inappropriate for further analysis, we linearise it by incorporating the Campbell and Shiller (1988b) – and Cochrane (2005) – framework. Importantly, though, we extend the existing methodology by (i) accounting for the fact that our net earnings-price ratio is forward-looking and (ii) adjusting for economic depreciation in the value of the asset. An immediate consequence of the latter feature is that we do not have to impose the transversality or “no-bubbles” condition. Accordingly, – in Appendix 2.B – we derive the following equation:

$$\begin{aligned} \pi_{t+1} - p_{5,t} &\approx - \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) k_i \\ &+ E_t \left[- \sum_{i=1}^n \left(\prod_{j=1}^i \rho_j \right) \Delta \pi_{t+i+1} + \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) r_{t+i} + \left(\prod_{i=1}^n \rho_i \right) (\pi_{t+n+1} - p_{5+n,t+n}) \right], \end{aligned} \quad (2.9)$$

where $\rho_i = \frac{P_{5+i}/\Pi}{1+P_{5+i}/\Pi}$ for $i \in \{1, \dots, n\}$, while for $i = 0$ we set $\rho_0 = 1$.

In addition, $k_i = -(1 - \rho_i) \ln(1 - \rho_i) - \rho_i \ln(\rho_i)$, for $i \in \{1, \dots, n\}$. Notice that for $n = 20$ we obtain $p_{25,t+20} \equiv s_{t+20} = \ln(S_{t+20})$ which corresponds to the log scrap price of the vessel. From a technical perspective, due to economic depreciation of the vessel, the i^{th} subsequent Taylor expansion is taken around the corresponding age-varying approximation point, defined as:

$$p_{5+i} - \pi = \ln(P_{5+i}/\Pi), \text{ where } P_{5+i}/\Pi = 1 / \left[(1/T) \sum_{t=0}^{T-1} \Pi_{t+1} / P_{5+i,t+1} \right] \text{ and } i \in \{1, \dots, n\}.$$

Finally, recall from Section 2.II that all variables incorporated in (2.9) and, in turn, in the empirical estimation below are $I(0)$, that is, they satisfy the stationarity condition.

Similar to (2.8), equation 2.9 illustrates that a high “forward-looking” log net earnings yield is a consequence of either low expectations about future log net earnings growth or/and high expectations about future log returns or/and high expectations regarding the spread between the log resale – scrap – price of the vessel and the corresponding log net earnings prevailing in the market. Importantly, note that high expectations about future – log – returns can be interpreted as the existence of high risk premia in the shipping industry.

Accordingly, for the purpose of examining which of these three sources is the major driving force behind the observed shipping earnings yields, as a preliminary test, we run one- and multi-year horizon forecasting OLS regressions in the spirit of Fama and French (1988b) and Cochrane (2005 and 2011). In particular, we regress future real log returns, r_{t+n} , future real log net earnings growth rates, $\pi_{t+n+1} - \pi_{t+1}$, and future terminal spreads, $\pi_{t+n+1} - p_{5+n,t+n}$, on the current log net earnings-price ratio, $\pi_{t+1} - p_{5,t}$:

$$r_{t+n} = \alpha_{r,t+n} + \beta_{r,t+n} \cdot (\pi_{t+1} - p_{5,t}) + \varepsilon_{r,t+n}, \quad (2.10)$$

$$\pi_{t+n+1} - \pi_{t+1} = \alpha_{\Delta\pi,t+n} + \beta_{\Delta\pi,t+n} \cdot (\pi_{t+1} - p_{5,t}) + \varepsilon_{\Delta\pi,t+n}, \quad (2.11)$$

$$\pi_{t+n+1} - p_{5+n,t+n} = \alpha_{\pi-p,t+n} + \beta_{\pi-p,t+n} \cdot (\pi_{t+1} - p_{5,t}) + \varepsilon_{\pi-p,t+n}, \quad (2.12)$$

where $n \in \{1, \dots, 20\}$. Recall that current log earnings yields, real log net earnings' growth rates, and future real log returns are estimated through equations 2.4, 2.5 and 2.6, respectively. Since our dataset consists of monthly observations we must deal with the overlapping nature of returns and growth rates. Therefore, in line with the existing literature, we incorporate (i) Newey-West (1987) heteroskedasticity and autocorrelation consistent (HAC) standard errors and (ii) Hodrick (1992) standard errors to account for this feature.²⁸ Table 2.5 summarises the results from these predictive regressions for the 1-, 2-, and 3-year horizon cases.

The findings presented in Panels A-D of Table 2.5 clearly indicate that shipping earnings yields strongly and negatively forecast future net earnings growth. Specifically, we observe that all net earnings growth coefficients are noticeably large in absolute magnitude. Furthermore, the signs of the growth coefficients are negative in every sector and horizon; that is, their signs are consistent with the present value linearisation – as presented through equation 2.8. Importantly, the t-statistics indicate a high level of significance – at least at the 5% level – across all sectors and horizons. Additionally, the R^2 s of growth regressions in all sectors and horizons are substantial and well above 10% – in some cases they are even close to 30%. Hence, according to our empirical results, there is clear evidence of cash flow predictability in the dry bulk shipping industry. However, since this chapter examines a question of relative predictability, we should compare the results obtained from the growth regressions to the respective findings from the returns and terminal spread ones.

First, regarding the returns regressions, we observe that the corresponding slope coefficients and t-statistics – using both standard errors – are substantially smaller in absolute value compared to the growth ones. Moreover, the returns coefficients are mainly insignificant even at the 10% level. Notice that only in the Capesize and Handysize sectors – and solely in the 3-year horizon case – the returns coefficients are significant at the 5% level or higher. Finally, the R^2 s of all returns regressions are close to zero and at least four times smaller compared to the respective growth ones. Therefore, we can argue that there is no consistent evidence of substantial statistical relationship between shipping earnings yields and expected returns.

Table 2.5: Regressions of future earnings yield, returns, and earnings growth on current earnings yield.

n	T	Earnings yield			Return			Net earnings growth		
		β	t^{NW}	R^2	β	t^{NW}	R^2	β	t^{NW}	R^2
Panel A: Capesize Sector										
1	264	0.48***	3.57	0.23	0.07	1.25	0.03	-0.60***	-2.91	0.17

²⁸ For conciseness, Table 2.5 reports only the Newey-West (1987) t-statistics since the respective ones obtained using the Hodrick (1992) correction have very similar values.

2	252	0.24	1.03	0.05	0.08	0.57	0.01	-0.95***	-2.84	0.26
3	240	0.14	1.04	0.01	0.17***	2.62	0.03	-1.01***	-5.16	0.20
Panel B: Panamax Sector										
1	420	0.22	0.89	0.03	-0.02	-0.29	0.00	-0.93***	-3.27	0.21
2	408	0.08	0.67	0.00	-0.10	-0.98	0.02	-1.20***	-5.72	0.28
3	396	0.27*	1.70	0.04	-0.16	-1.17	0.03	-1.08***	-3.94	0.20
Panel C: Handymax Sector										
1	327	0.47**	2.50	0.20	0.09	1.32	0.03	-0.60**	-2.42	0.13
2	315	0.17	1.13	0.02	0.09	0.53	0.01	-1.01***	-3.94	0.20
3	303	0.25**	2.40	0.05	0.16	1.28	0.02	-0.87***	-4.28	0.14
Panel D: Handysize Sector										
1	456	0.47***	2.75	0.21	0.12*	1.85	0.03	-0.55**	-2.48	0.13
2	444	0.17	0.80	0.03	0.18	1.48	0.03	-0.84***	-3.66	0.17
3	432	0.18	1.26	0.03	0.22**	2.31	0.03	-0.81***	-4.59	0.13
Panel E: Dry Bulk Industry										
	<i>n</i>	β	t^W	R^2	β	t^W	R^2	β	t^W	R^2
	1	0.34***	5.57	0.10	0.06*	1.92	0.02	-0.70***	-7.64	0.17
	2	0.16***	4.64	0.04	0.04	0.62	0.01	-1.02***	-13.09	0.23
	3	0.21***	7.76	0.04	0.06	0.62	0.01	-0.96***	-16.21	0.17

Notes: Panels A-D report 1-, 2-, and 3-year horizon OLS regressions of future log earning yield, real log return, and real log net earnings growth on current log earnings yield for each dry bulk sector. To account for the overlapping nature of the variables, the t-statistics, t^{NW} , reported are estimated using the Newey-West (1987) HAC correction. The predictive coefficient, β , is accompanied by *, **, or *** when the absolute t^{NW} statistic indicates significance at the 10%, 5% or 1% levels, respectively. In addition, Panel E summarises the respective results from 1, 2, and 3-year horizon pooled-time series least squares regressions using the same set of variables. These regressions embody cross-section fixed effects while the incorporated sample is unbalanced. The corresponding t-statistics, t^W , are estimated using the "White period" method.

Second, regarding the terminal spread regressions, the obtained results can be explained by the rapid mean reversion that characterises the shipping earnings yields. Namely, as we analysed in Section 2.II, the autocorrelation coefficients of earnings yields fall from above 0.9 in the one-month horizon to less than 0.5 in the 1-year case and to less than 0.2 in the 2-year one (Table 2.2). This stochastic property is equivalent to the substantial reduction in the magnitude of the slope coefficients and t-statistics, the significance of the coefficients, and the R^2 s of the respective regressions in the 2-year horizon, compared to the 1-year case. Furthermore, the fact that in the 1-year horizon the Capesize, Handymax, and Handysize sectors' earning yields have moderate autocorrelation coefficients explains the substantial magnitude of the coefficients, t-statistics, and R^2 s, and the level of significance of the corresponding regressions.

Notice that the Panamax sector's earnings yield is much more mean-reverting (Tables 2.2 and 2.4); hence, the results of the 1-year horizon terminal spread regression are even less significant compared to the other three sectors. Consequently, only in the 1-year horizon case the results from the terminal spread regressions can be compared to the ones from the growth ones. However, as we

also demonstrate through the variance decomposition, as the horizon increases, the aggregate effect of the terminal spread's variation on the volatility of current net earnings yield becomes substantially small compared to the one of the net earnings growth variation.

A further implication of the rapid mean reversion of shipping earnings yields is the fact that we do not observe any clear, general pattern related to the magnitude of the slope coefficients and the R^2 s of the shipping growth and returns regressions across different horizons and sectors. As mentioned in Section 2.II, this result contrasts with the U.S. equity markets returns regressions and is an immediate consequence of the fact that the shipping forecasting variable (i.e., the earnings yield) is extremely less persistent than the corresponding U.S. equity markets one. Specifically, Cochrane (2005) illustrates mathematically that, within this empirical framework, when the

foreca
sting

Panel A: Capesize sector from 1/1992 to 12/2014. *Panel B: Panamax sector from 1/1976 to 12/2011.*

variable is highly persistent, the slope coefficients and the R^2 s of the forecasting regressions add up over longer horizons. In particular, they increase approximately linearly with the horizon in the beginning and then at a decreasing rate. Therefore, since in the U.S. equity markets the forecasting variable (i.e., the dividend yield) is not only strongly related with future returns but also highly persistent, the values of the slope coefficients and R^2 s in the returns regressions increase with the horizon.

In contrast, in shipping where the forecasting variable is related to future net earnings growth – and not to future returns – but at the same time is highly mean-reverting, neither the slope coefficients and R^2 s of growth regressions, nor the ones of returns regressions increase consistently with time. Finally, an interesting feature presented in Table 2.5 is that the slope coefficients and R^2 s of growth regressions increase substantially in the 2-year horizon compared to the 1-year case. From

Panel C: Handymax sector from 4/1986 to 6/2014. *Panel D: Handysize sector from 1/1976 to 12/2014.*

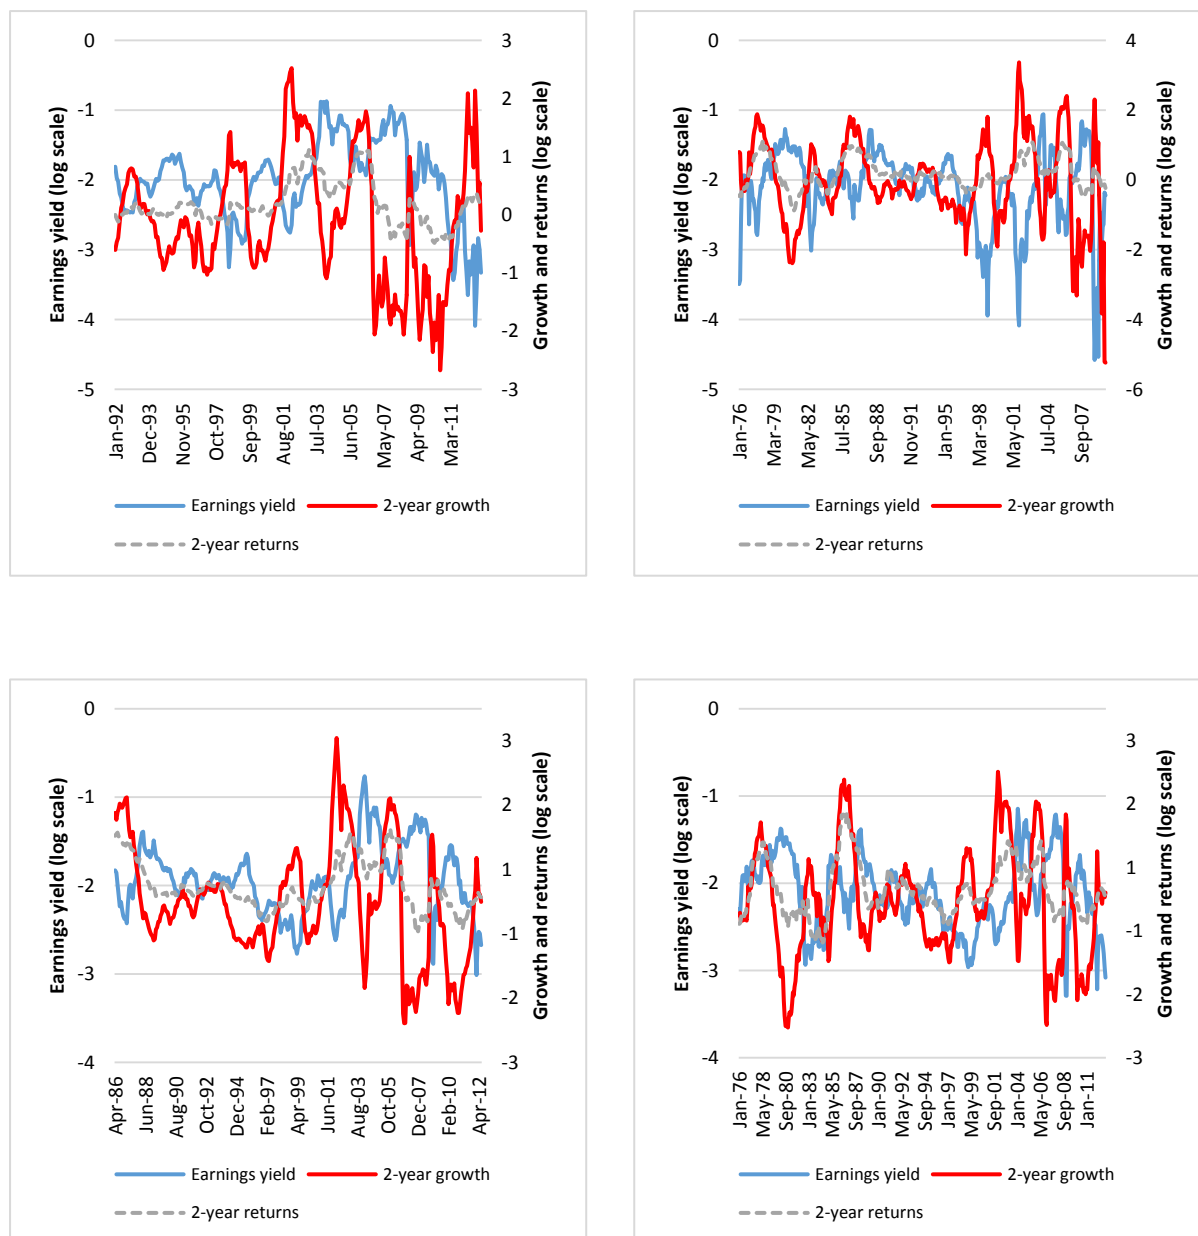


Figure 2.4: Current Earnings Yields, Future Net Earnings Growth, and Future Returns.

Panels A-D present current earnings yields and the corresponding 2-year horizon log net earnings growth and log returns for each dry bulk sector.

an economic perspective, this result may be related to the time-lag required for the delivery of a new vessel – which is approximately equal to 2 years on average.²⁹

Figure 2.4 illustrates the main empirical result discussed above. Specifically, Panels A-D depict the strong, negative relation between current shipping earnings yields and the corresponding 2-year

²⁹ The economic rationale behind the obtained empirical results is analysed further in Section 4.IV.

horizon net earnings growth across the four dry bulk sectors. In addition, they demonstrate the fact that earnings yields are essentially uncorrelated with the 2-year horizon returns.

Finally, in order to further test the robustness of our conclusions, we examine the dry bulk industry as a whole. Namely, we perform predictive pooled-time-series least squares regressions using the same explanatory and explained shipping variables as in the simple – sector-specific – time-series estimation. Specifically, we employ fixed effects in the cross-section to account for the differences across the four shipping sectors while we use an unbalanced sample based on the four dry bulk shipping sectors and all corresponding observations. Accordingly, we run the following set of regressions for the 1-, 2-, and 3-year horizons:

$$x_{i,t+n} = c + \alpha_{i,x} + \beta_x \cdot (\pi_{i,t+1} - p_{i,5,t}) + \varepsilon_{i,t+n}, \quad n \in \{1, 2, 3\}, \quad (2.13)$$

where $x_{i,t+n}$ alternately denotes $r_{i,t+n}$, $\pi_{i,t+n+1} - \pi_{i,t+1}$, and $\pi_{t+n+1} - p_{i,5+n,t+n}$. Moreover, $\alpha_{i,x}$ represents the cross-section fixed effects while by i we index the corresponding dry bulk sector. Note that we incorporate the “White period” method for standard errors which assumes that the errors within a cross-section suffer from heteroscedasticity and serial correlation. As one would expect, the results from these pooled-time-series predictive regressions – summarised in Panel E of Table 2.5 – indicate precisely the same patterns as the ones obtained from the simple time-series estimation.

In conclusion, according to this first set of predictive regressions, we can argue that shipping earnings yields appear to be strongly related to expected net earnings growth and not to expected returns and the expected terminal earnings yield.

2.III.B. Variance Decomposition

Furthermore, in order to quantify formally the relative magnitude of each of the three potential sources of variation, we decompose the variance of the shipping earnings yield using the following equation (see Appendix 2.C):

$$1 \approx -b_{\Delta\pi,n} + b_{r,n} + b_{\pi-p,n}, \quad (2.14)$$

where $b_{i,n}$ is the n -year horizon coefficient corresponding to the i^{th} element of the decomposition. Following Cochrane (1992, 2005, and 2011), these regression coefficients can be interpreted as fractions of net earnings-price ratio variation attributed to each of the three sources. In particular, $b_{\Delta\pi,n}$ corresponds to the fraction attributed to future net earnings growth, $b_{r,n}$ to the fraction attributed to future returns, and $b_{\pi-p,n}$ to the fraction attributed to the terminal spread. Notice that the elements of this decomposition do not have to be between 0 and 100%.

Accordingly, we examine formally to which of the three sources does the bulk of earnings yield volatility correspond by running the following set of exponentially weighted regressions for each of the four dry bulk shipping sectors:

$$\sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) r_{t+i} = \alpha_{r,n} + b_{r,n} \cdot (\pi_{t+1} - p_{5,t}) + \varepsilon_{r,t+n}, \quad (2.15)$$

$$\sum_{i=1}^n \left(\prod_{j=1}^i \rho_j \right) \Delta \pi_{t+i+1} = \alpha_{\Delta \pi,n} + b_{\Delta \pi,n} \cdot (\pi_{t+1} - p_{5,t}) + \varepsilon_{\Delta \pi,t+n}, \quad (2.16)$$

$$\left(\prod_{j=1}^n \rho_j \right) (\pi_{t+n+1} - p_{5+n,t+n}) = \alpha_{\pi-p,n} + b_{\pi-p,n} \cdot (\pi_{t+1} - p_{5,t}) + \varepsilon_{\pi-p,t+n}, \quad (2.17)$$

where $n \in \{1, \dots, 20\}$. Table 2.6 presents the results from the variance decomposition corresponding to the 5-year horizon case for each dry bulk sector under consideration.³⁰

Table 2.6: Variance decomposition of the earnings yield.

Sector	n	Returns	Net earnings Growth	Terminal spread
		$b_{r,5}$	$b_{\Delta \pi,5}$	$b_{\pi-p,5}$
Capesize	5	-0.04	-1.38	-0.23
Panamax	5	-0.22	-1.25	-0.08
Handymax	5	0.01	-1.28	-0.20
Handysize	5	-0.09	-1.30	-0.18

Notes: $b_{i,5}$ is the exponentially weighted 5-year horizon regression coefficient corresponding to the i^{th} element of the decomposition. See equations 2.14 – 2.17 of the main text.

As expected, our findings suggest that in the dry bulk shipping industry almost all variation in net earnings-price ratios corresponds to variation in expected net earnings growth. Therefore, we can argue that high vessel prices relative to current net earnings significantly forecast high future net earnings growth, not low future returns, and not a high “terminal earnings yield”. Note that, in our context, the “terminal earnings yield” is equivalent to the spread between the terminal price of the vessel and the last net earnings realised by the ship owner. Finally, for robustness, we have performed the same analysis presented in this section using numerous sub-periods of the sample (including and excluding the 2008 financial recession and the last shipping super-cycle) for each dry bulk sector and the obtained patterns coincide. As a result, we can argue that our empirical findings

³⁰ We have tested various horizons and the obtained results indicate precisely the same patterns.

characterise the entire second-hand dry bulk shipping industry for the period from January 1976 to December 2014.

2.III.C. An Extension of the Campbell and Shiller (1988a) VAR Framework to Shipping

Finally, in this subsection, we examine an argument closely related to the results presented above. Namely, the fact that – within the Campbell-Shiller (1988b) framework – there is no consistent, strong statistical evidence of time-varying one-period required returns implies that investors require a – relatively – constant 1-period return when valuing vessels. If, however, 1-period required returns are constant, one should expect that an unrestricted econometric vector autoregressive (VAR) model including the three remaining variables of equation 2.9 will provide an accurate description of the data; that is, it will be able to explain sufficiently well the observed price-net earnings (PE) ratios. Following Campbell and Shiller (1988a), this comparison can be achieved either in a formal statistical way or simply by collating the movements of the observed and generated PE ratios.

Lof (2015) follows the exact same procedure in order to compare the observed price-dividend (PD) ratios with the ones generated by the VAR model, using US equity markets data. However, while he finds a significantly high correlation between the observed and generated PD ratios (approximately equal to 0.799) the obtained volatility ratio is noticeably poor (roughly 0.135).³¹ In other words, while the generated ratio can explain sufficiently well the direction of the movements of the realised variable, its volatility is almost 8 times smaller compared to the actual one.³² In line with this finding, as analysed above, it has been well-established in the U.S. equity markets literature that PD ratios can strongly predict future returns. In contrast, dividend growth volatility appears to explain a substantially small proportion of PE ratio's volatility.

Consequently, one should a priori expect that imposing the assumption of constant returns – and, accordingly, excluding from the VAR specification the returns variable – would result in poor volatility matching or, equivalently, in a failure of the variance bound test. Of course, this does not imply a rejection of the present value model or a failure of the efficient markets hypothesis. It can only be interpreted as a failure of this specific log linear relationship and the two-variable VAR to describe efficiently the observed data. This, in turn, can lead to a rejection of the constant discount rate model in the case of the U.S. equity markets. In analogy, if we assumed constant net earnings growth in shipping while simultaneously let future returns to vary, the generated PE ratios from this

³¹ The volatility ratio is defined as the ratio between the standard deviations of the observed and model-generated variables.

³² At a first glance, one could interpret such failure of the variance ratio test as a rejection of the efficient market hypothesis (EMH).

alternative VAR specification would match the realised data significantly worse compared to the ones in the time-varying earnings growth and constant returns case.

Accordingly, following Campbell and Shiller (1988a) and Lof (2015), the series of model-implied log price-net earnings ratios (for a 5-year old vessel), δ'_t , can be generated through the following equation (see Appendix 2.D for the derivation):

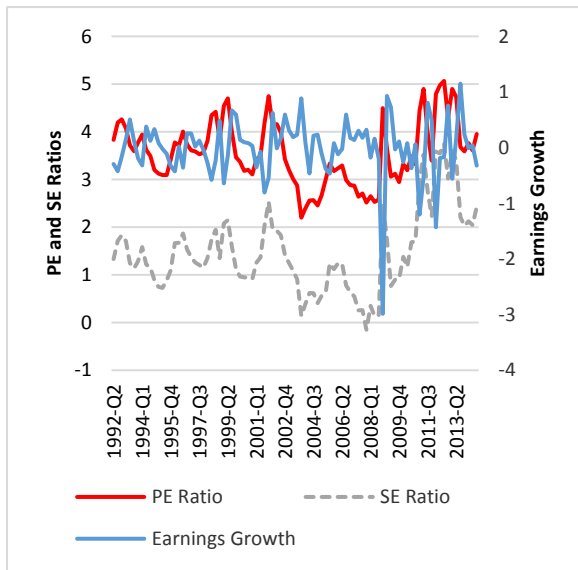
$$\delta'_t = \left[\sum_{i=1}^n \left(\prod_{j=1}^i \rho_j \right) \mathbf{e}2' \mathbf{A}^i + \left(\prod_{j=1}^n \rho_j \right) \mathbf{e}3' \mathbf{A}^n \right] \mathbf{z}_t, \quad (2.18)$$

where \mathbf{z}_t is a $3p \times 1$ matrix of state variables and p is the optimal number of lags corresponding to the incorporated VAR model. The state variables in this case are the actual log price-net earnings ratio, $\delta_t = p_t - \pi_{t+1}$, the one period log net earnings growth, $\Delta\pi_{t+1}$, and the log scrap-net earnings ratio, $\tau_t = s_t - \pi_{t+1}$, plus $(p - 1)$ lags of each state variable. Importantly, note that all variables in this equation are demeaned. Furthermore, \mathbf{A} is a $3p \times 3p$ matrix of constants and $\mathbf{e}2$, $\mathbf{e}3$ are selection vectors such that $\mathbf{e}2' \mathbf{z}_t = \Delta\pi_{t+1}$ and $\mathbf{e}3' \mathbf{z}_t = \tau_t$.


Table 2.7: Descriptive statistics for price-net earnings ratios, net earnings growth, and terminal spreads.

Variable	First Obs	Last Obs	T	Mean	SD	Median	Max	Min	ρ_1
Panel A: Capesize Sector									
δ	1992Q2	2014Q3	90	3.56	0.67	3.51	5.06	2.20	0.78
$\Delta\pi$	1992Q2	2014Q3	90	-0.01	0.53	0.05	1.15	-2.99	0.01
τ	1992Q3	2014Q3	90	1.47	0.86	1.27	3.75	-0.15	0.84
Panel B: Panamax Sector									
δ	1990Q1	2011Q4	88	3.54	0.64	3.44	6.49	2.42	0.65
$\Delta\pi$	1990Q1	2011Q4	88	-0.02	0.64	0.02	1.49	-4.75	-0.05
τ	1990Q1	2011Q4	88	1.40	0.73	1.29	4.45	-0.10	0.66
Panel C: Handymax Sector									
δ	1991Q4	2012Q2	83	3.15	0.38	3.22	3.88	2.25	0.85
$\Delta\pi$	1991Q4	2012Q2	83	-0.01	0.31	0.01	0.63	-1.99	0.20
τ	1991Q4	2012Q2	83	0.94	0.45	0.99	1.95	-0.54	0.73
Panel D: Handysize Sector									
δ	1992Q1	2014Q2	90	3.32	0.34	3.37	3.97	2.45	0.85
$\Delta\pi$	1992Q1	2014Q2	90	-0.01	0.27	0.04	0.57	-1.84	0.18
τ	1992Q1	2014Q2	90	1.08	0.47	1.09	2.21	-0.39	0.80

Notes: Panels A-D present descriptive statistics for log price-net earnings ratio, δ , 1-quarter horizon real log net earnings growth, $\Delta\pi$, and log scrap-net earnings ratio, τ , corresponding to the four dry bulk sectors. The included statistics are the number of observations, the dates of the first and last sample observations, the mean, standard deviation, median, maximum, minimum, and 1-quarter autocorrelation coefficients.



Panel A: Capesize sector, from 1992:Q2 to 2014:Q3.





Panel B: Panamax sector from 1990:Q1 to 2011:Q4.

Figure 2.5: Net Earnings Growth, S/H Price-Net Earnings Ratio, and Scrap Price-Net Earnings Ratio.

Panels A-D present the evolutions of the 1-quarter horizon real quarterly net earnings growth, current real second-hand (5-year old) vessel price-net earnings ratio, and real scrap price-net earnings ratio for each dry bulk sector. All variables correspond to log transformations.

Panel C: Handymax sector from 1991:Q4 to 2012:Q2.

Panel D: Handysize sector from 1992:Q1 to 2014:Q2.



From an economic perspective, if market agents value vessels requiring constant – expected – returns, the observed price-net earnings ratios should be the same as – or, in practice, very close to – the ones generated by (2.18). In order to examine this argument in dry bulk shipping, we estimate the time-series of the three variables of interest, that is, δ_t , $\Delta\pi_{t+1}$, and τ_t . Due to data limitations, mainly related to scrap prices time-series, we assume consecutive quarterly operating periods (i.e., frequency, f , equal to 4) as opposed to annual ones – therefore, the remaining economic life of a 5-year old vessel is equal to $n = 20 \cdot f = 80$ quarterly periods. Accordingly, after discussions with industry participants, we also adjust for the out-of-service period and the operating and maintenance expenses of the vessel to be consistent with reality.³³ Moreover, to obtain stationary variables – as required within this estimation framework, we slightly reduce the length of the sample in the cases of Capesize, Handymax, and Handysize sectors. Table 2.7 summarises descriptive statistics related to the three variables of interest and Figure 2.5 illustrates their evolution.

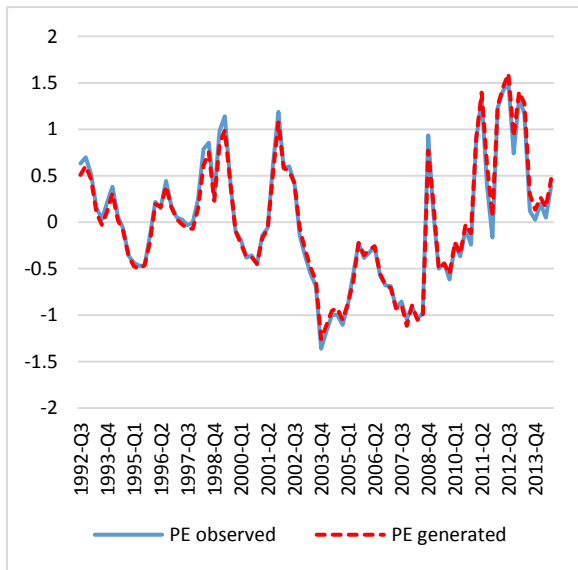
Table 2.8: Lag length information criterion.

1	2	3	4	5	6	7	8
Panel A: Capesize Sector							
-1.440	-1.448*	-1.414	-1.330	-1.249	-1.173	-1.042	-0.970
Panel B: Panamax Sector							
-1.549	-1.617*	-1.588	-1.546	-1.507	-1.370	-1.202	-1.044
Panel C: Handymax Sector							
-2.826	-3.118*	-3.044	-2.949	-2.858	-2.792	-2.668	-2.570
Panel D: Handysize Sector							
-3.357	-3.614	-3.618*	-3.580	-3.497	-3.386	-3.314	-3.208


Notes: Panels A-D present lag selection based on the Akaike Information Criterion (AIC) for the VAR(p) models corresponding to the four dry bulk sectors. The first row indicates the respective lag length while the optimal one is indicated by *.

³³ After discussions with industry participants, we allow for an out-of-service period between consecutive quarterly time-charter contracts in order to account for maintenance requirements, (delays in) chartering arrangements, travelling distance between lading and loading ports, and port congestion, among other factors. During this period, the ship owner does not receive the corresponding time-charter rate but pays for the operating and maintenance expenses. The estimation results presented here correspond to an out-of-service period of 13 days per quarter while the operating and maintenance expenses are assumed to be fixed for the entire quarterly leasing period – and might differ compared to the ones presented in Table 2.1. Importantly, however, we have examined the robustness of our estimation for several values of the operating and maintenance expenses and the duration of the out-of-service period. As expected, since we are mainly interested in the second moments of the variables of interest, the results are very similar to the ones presented here. Moreover, since there is no data for 3-month time-charter rates, we use the corresponding 6-month ones.


In order to select the optimal lag length, p , of the unrestricted VAR(p) model we incorporate the



Panel A: Capesize sector from 1992:Q3 to 2014:Q3.



Panel B: Panamax sector from 1990:Q2 to 2011:Q4.



Panel C: Handymax sector from 1992:Q1 to 2012:Q2.




Figure 2.6: Observed and Generated Price-Net Earnings Ratios.

Panels A-D present the observed price-net earnings ratios and the ones generated by the VAR model (that is, through equation 2.18 of the main text) for each dry bulk sector.

Akaike Information Criterion (AIC). Furthermore, we perform the VAR Stability Condition Test for each dry bulk sector and, according to the results, all models are stable – that is, there is no root lying outside the unit circle. Table 2.8 presents the optimal lag length for each sector. Accordingly, Figure 2.6 illustrates a comparison between the observed log price-net earnings ratios and the ones generated by equation 2.18 for each dry bulk sector. In addition, Table 2.9 presents the estimated correlation coefficients between the two variables and the ratios between their respective standard deviations; that is, the corresponding volatility ratios, denoted by $\frac{\sigma(\delta_t)}{\sigma(\delta'_t)}$. Evidently, our unrestricted VAR model with constant required returns matches sufficiently well the observed data in each dry bulk shipping sector.

Table 2.9: Comparison between the observed and generated PE ratios.

Sector	Correlation	Volatility Ratio
Capesize	0.99	1.01
<i>Panel D: Handysize sector from 1992:Q3 to 2014:Q2.</i>		
Panamax	0.96	0.84
Handymax	0.88	0.83
Handysize	0.95	0.97

Notes: This table illustrates a comparison between the observed and generated log price-net earnings ratios. The latter are estimated through equation 2.18 of the main text. Correlation refers to the correlation coefficient between the two variables while the volatility ratio corresponds to the fraction between the volatility of the observed price-net earnings ratio and the volatility of the generated one.

Unfortunately, as mentioned above, the empirical results obtained from the variance decomposition (subsections 2.III.A and 2.III.B) and the VAR (subsection 2.III.C) frameworks are not directly comparable since, due to data limitations, the incorporated operating periods of the vessel

and, in turn, the respective samples do not coincide. However, we can argue that both estimation procedures' results suggest that there does not appear to be any no consistent, significant evidence of time-varying required returns in the valuation of dry bulk vessels.

2.IV. Economic Interpretation and Discussion

As we demonstrated above, the bulk of shipping earnings yields' volatility is attributed to variation in expected net earnings growth. In other words, vessel valuation ratios vary due to expectations about future net earnings growth and not due to expectations about future returns. Specifically, high net earnings-price ratios strongly, negatively forecast future net earnings growth. Equivalently, during market peaks net earnings are high compared to vessel prices and, vice versa, during market troughs vessel prices are high compared to the prevailing net earnings.

Note that since in the Campbell-Shiller variance decomposition methodology the earnings yield is the sole state variable, it is assumed to be summarising the time t information filtration; namely, the historical and prevailing economic conditions (Fama and French, 1998a). As Campbell and Shiller (1988a) argue, while we cannot observe everything that shipping agents do, we observe the earnings yield which should summarise the market's relevant information. In particular, the net earnings yield variable is significantly more informative regarding both current and future market conditions – and, in turn, future net earnings growth – compared to the lagged net earnings growth variable. Accordingly, we verify this argument by running 1-, 2-, and 3-year horizon regressions of future net earnings growth on lagged net earnings growth. In turn, we compare the results from these regressions to the ones obtained from regressing future net earnings growth on current earnings yields (equation 2.16). Table 2.10 summarises these findings.

Table 2.10: Regressions of future net earnings growth on lagged net earnings growth and current earnings yields.

n	Lagged net earnings growth					Earnings yield				
	T	β	t^{NW}	t^H	R^2	T	β	t^{NW}	t^H	R^2
Panel A: Capesize Sector										
1	252	-0.22	-1.50	-1.52	0.04	264	-0.60***	-2.91	-2.94	0.17
2	240	-0.42**	-2.23	-2.37	0.10	252	-0.95***	-2.84	-2.89	0.26
3	228	-0.19**	-2.27	-2.29	0.02	240	-1.01***	-5.16	-5.06	0.20
Panel B: Panamax Sector										
1	408	-0.32***	-2.82	-2.76	0.08	420	-0.93***	-3.27	-3.17	0.21
2	396	-0.52***	-3.94	-3.90	0.16	408	-1.20***	-5.72	-5.67	0.28
3	384	-0.30	-1.56	-1.52	0.04	396	-1.08***	-3.94	-4.13	0.20
Panel C: Handymax Sector										
1	315	-0.19	-1.42	-1.44	0.04	327	-0.60**	-2.42	-2.43	0.13
2	303	-0.43**	-2.56	-2.59	0.13	315	-1.01***	-3.94	-4.07	0.20

3	291	-0.18***	-2.82	-2.92	0.02	303	-0.87***	-4.28	-3.92	0.14
Panel D: Handysize Sector										
1	444	-0.17	-1.05	-1.04	0.03	456	-0.55**	-2.48	-2.57	0.13
2	432	-0.40***	-2.87	-2.85	0.10	444	-0.84***	-3.66	-3.86	0.17
3	420	-0.26***	-2.72	-2.76	0.03	432	-0.81***	-4.59	-4.54	0.13

Notes: This table reports results from 1-, 2-, and 3-year horizon forecasting regressions of real log net earnings growth on lagged real net earnings growth. For expositional simplicity and direct comparison purposes we have included the results from the regressions of future net earnings growth on current earnings yield, presented in Table 2.5. To account for the overlapping nature of the variables, t-statistics are estimated using the Newey-West (1987) HAC and the Hodrick (1992) corrections and are denoted by t^{NW} , and t^H , respectively. The predictive coefficient, β , is accompanied by *, **, or *** when the absolute t^{NW} statistic indicates significance at the 10%, 5% or 1% levels, respectively.

Figure 2.7: Dry Supply and


Panel A provides a comparison between the total development (measured in million dwt) and the total dry bulk in billion tonnes). The data are in an

As analysed corresponding results illustrate

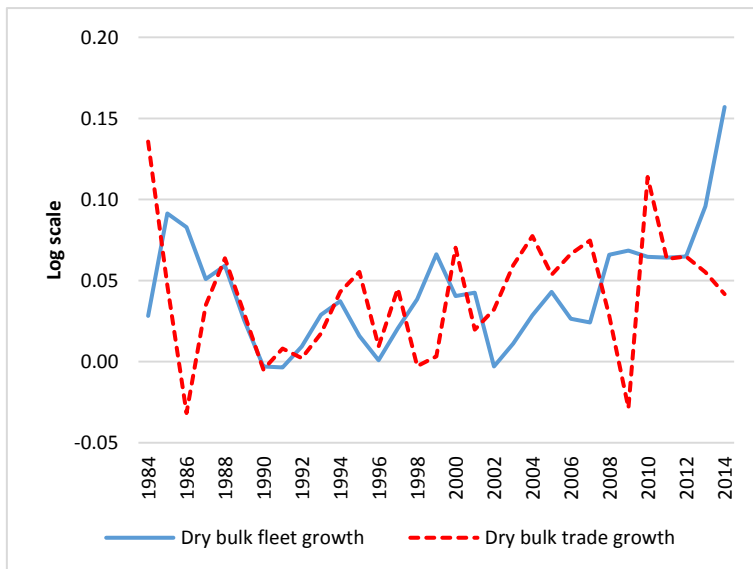
are

significantly more informative regarding future net earnings growth compared to the lagged net earnings growth variable. Note that we further reinforce this argument by performing bivariate forecasting regressions using both current earnings yields and lagged net earnings growth as explanatory variables and future net earnings growth as the explained one. In line with the univariate regressions' findings, in these bivariate regressions, only the coefficients of current earnings yields are economically and statistically significant.³⁴

The drawback, though, of the – reduced form – variance decomposition specification incorporated in this chapter is that it does not allow us to understand the economic principles



Panel A: Dry bulk fleet and trade development from 1983 to 2014.



Panel B: Dry bulk fleet and trade growth from 1983 to 2014.

Bulk Shipping Demand. comparison dry bulk fleet million evolution of the trade (measured Panel B compares the one-period the two variables. annual frequency. above, the regression that earnings

³⁴ The results from these bivariate regressions have not been included in this document, however, they can be provided upon request.

behind the obtained results (Ghysels *et al*, 2012). However, the shipping industry provides an ideal environment to explain the underlying mechanism and, accordingly, relate – at least theoretically – the Campbell-Shiller framework to economic rationale. Specifically, shipping is a capital-intensive industry with clear and well-documented supply and demand mechanism. Furthermore, due to data availability, we can directly observe the (dis)investment decisions of market participants and how these affect the industry equilibrium and, in turn, future shipping cash flows.

2.IV.A. The Shipping Supply and Demand Mechanism

To begin with, as analysed in Chapter 1, due to the time-to-build feature, shipping supply adjusts sluggishly to demand (Kalouptsi, 2014). As a result, while the aggregate supply and demand variables exhibit a high degree of co-movement (Panel A of Figure 2.7), their respective one-period growth rates are extremely less correlated (Panel B of Figure 2.7). Thus, since time-charter rates and, in turn, net earnings are the equilibrium outcome of the supply and demand mechanism, one should expect net earnings growth to be positively and negatively related to – growth in – demand and supply, respectively.³⁵ Indeed, by approximating shipping demand through the aggregate dry bulk seaborne trade,³⁶ we estimate a significant positive relationship between net earnings growth and shipping demand growth across all dry bulk sectors (ranging from 0.49 to 0.63). Equivalently, net earnings growth is expected to be negatively related to the spread – imbalance – between supply and demand growth rates. Accordingly, we quantify this spread corresponding to period $t \rightarrow t + 1$ through the following equation:³⁷

$$S_{t+1} = \ln\left(\frac{F_{t+1}}{F_t}\right) - \ln\left(\frac{D_{t+1}}{D_t}\right), \quad (2.19)$$

where D_t is the aggregate demand for shipping services during period $t \rightarrow t + 1$ while F_t is the aggregate fleet capacity (i.e., supply) at time t .³⁸ Consequently, we examine for each dry bulk sector the relation between 1-year horizon net earnings growth, $\Delta\pi_{t+1}$, and the spread, S_{t+1} . In line with our expectations, Figure 2.8 demonstrates the strong negative relation between net earnings growth

³⁵ In order to depict that statement, the reader can think of a simple inverse linear demand function.

³⁶ Unfortunately, due to data limitations, we could not quantify the sector-specific demand variables.

³⁷ Alternatively, we could define the spread variable by expressing current net earnings as a function of current supply and demand (e.g., through a linear inverse demand curve). However, the results obtained from both specifications indicate precisely the same patterns.

³⁸ In line with the existing literature (e.g., Kalouptsi, 2014; Greenwood and Hanson, 2015), we assume that fleet capacity remains constant during the period. However, this simplifying assumption does not have a major impact on the results.

and the spread across all dry bulk sectors (the estimated coefficients range between -0.79 and -0.87).³⁹

In summary, the dry bulk shipping supply and demand mechanism operates in the following way. Random shocks in demand drastically perturb the short-run equilibrium and, consequently, the prevailing net earnings (this can be defined as a first-order effect). In turn, this increase in current net earnings has an indirect dramatic effect on future net earnings through the current investment decisions of market participants – ship owners. More importantly, due to the time-to-build characteristic of the industry, this change in supply will not be realised immediately but during subsequent periods (this can be defined as a second-order effect). This fact, accompanied by the mean-reverting – around an upward trend – character of the exogenous demand (recall Chapter 1

and

Panel A: Capesize sector from 1992 to 2014.

Figure *Panel B: Panamax sector from 1989 to 2011.* 1.1)

result in extremely volatile shipping cash flows. Consequently, shipping net earnings are not exogenously but partially endogenously determined by the investment decisions of market participants (Stopford, 2009; Greenwood and Hanson, 2015).

Having analysed the relation between supply, demand, and net earnings, we will now examine from an economic point of view the interaction between net earnings, prices, and earnings yields.⁴⁰ For illustrational purposes, consider a discrete time, dynamic environment. At each time t , annual net earnings corresponding to period $t \rightarrow t + 1$ are determined through the previously analysed supply and demand mechanism. Assume further that, due to an unexpected positive demand shock, current net earnings, $\Pi_{t \rightarrow t+1}$, are significantly high. Therefore, the owner of a vessel at time t can immediately exploit the prosperous market conditions and, thus, realise a significant, deterministic inflow at time $t + 1$. In anticipation of this increased short-term net cash flow, current vessel prices,

Panel C: Handymax sector from 1989 to 2013.

P_t , *Panel D: Handysize sector from 1989 to 2014.*
jump

³⁹ As mentioned before, due to data limitations, we can estimate only the aggregate and not the sector-specific demand for dry bulk shipping services. Moreover, the aggregate demand data obtained are in an annual frequency. Hence, the supply variable is calculated using the aggregate fleet capacity of the dry bulk sector in an annual frequency. Therefore, we approximate the spread as the difference between the annual growth of the aggregate dry bulk fleet (in dwt) and the annual growth of the total dry bulk seaborne trade (in tonnes). Since shipping services are perceived as a homogeneous product (Stopford, 2009; Kalouptsi, 2014), we believe that this approximation does not have a qualitative impact on the results.

⁴⁰ This economic analysis is in line with Beenstock and Vergottis (1989), Stopford (2009), and the behavioural model of Greenwood and Hanson (2015).

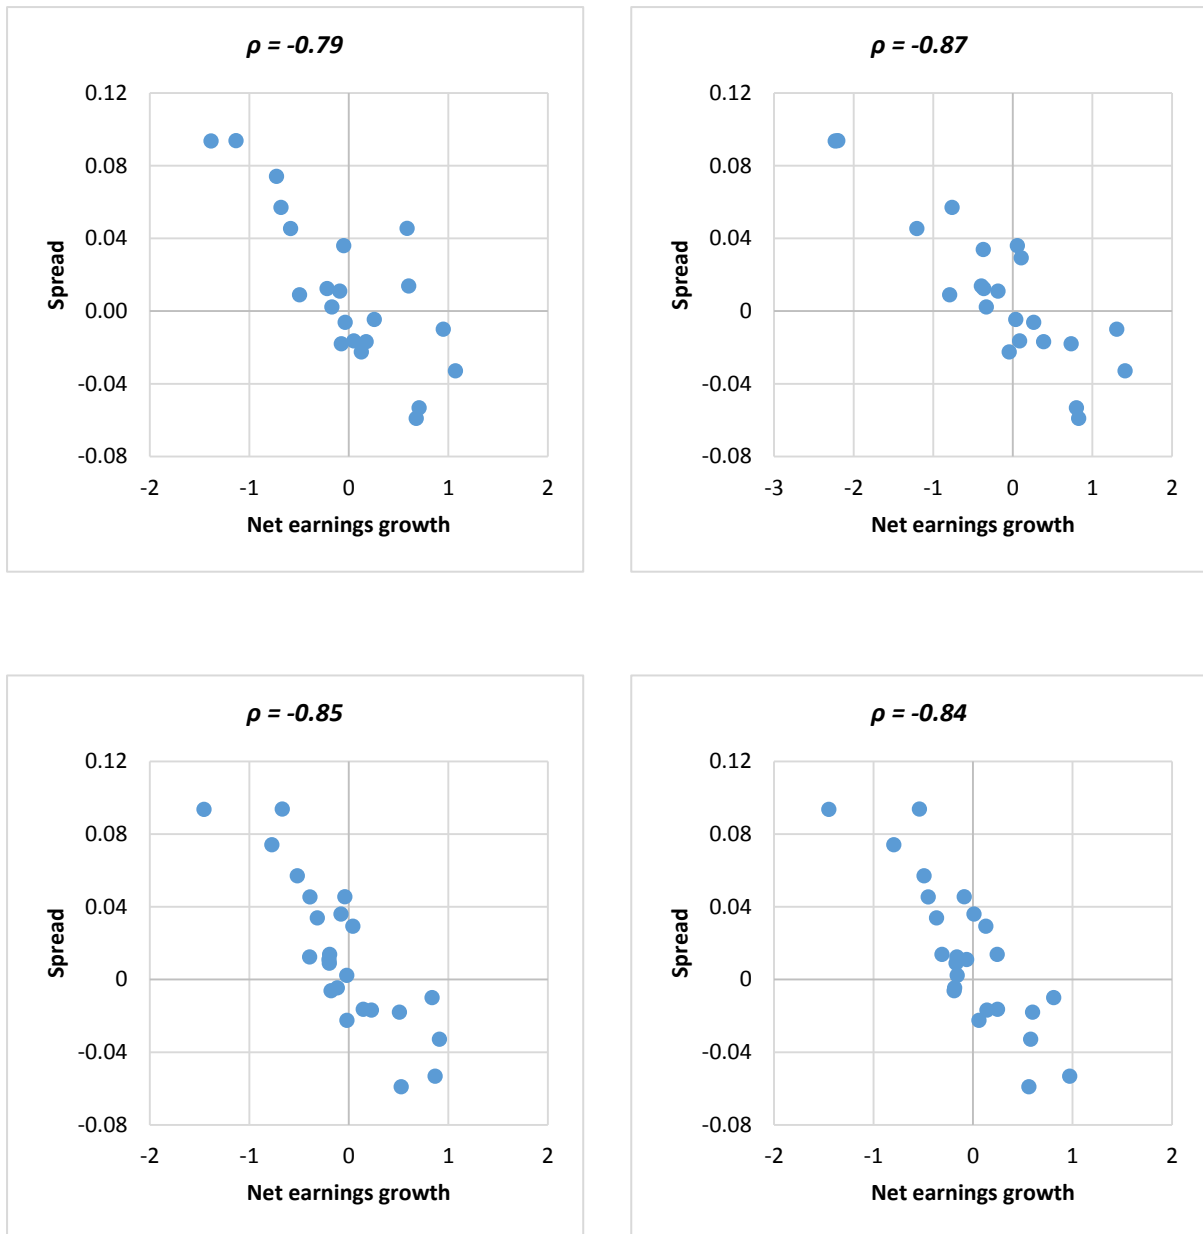




Figure 2.8: Correlation Between Net Earnings Growth, Supply Growth, and Demand Growth.


Panels A-D plot the annual net earnings growth for each dry bulk sector against the spread between the aggregate dry bulk fleet capacity growth and the aggregate dry bulk demand growth. Fleet capacity is measured in million dwt while demand in million tonnes. In addition, we report the correlation coefficient, ρ , between the spread and the respective net earnings growth.



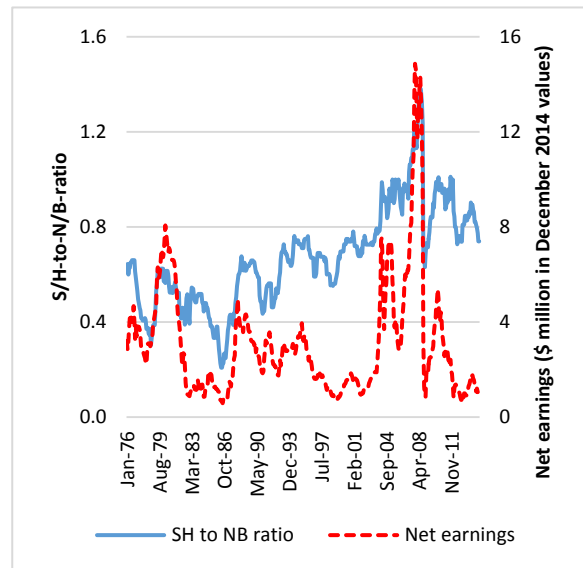
Panel A: Capesize sector from 1/1992 to 12/2014.



Panel B: Panamax sector from 1/1976 to 12/2011.



Panel C: Handymax sector from 4/1986 to 6/2014.



Panel D: Handysize sector from 1/1976 to 12/2014.

Figure 2.9: Net Earnings, Second-Hand Vessel Prices, and Newbuilding Vessel Prices.

Panels A-D depict the relation between the ratio of 5-year old-to concurrent newbuilding vessel prices and current net earnings for each dry bulk sector. Note that the second-hand and newbuilding prices capacity referring to the same dry bulk sector correspond to vessels of approximately the same cargo carrying capacity (dwt). Since we are interested in the cross-time relationship between price ratios and net earnings these discrepancies do not have an implication in the resulting patterns.

compared to their previous level, P_{t-1} .⁴¹ This substantial price increase is a positive first-order effect (FOE) of the current – increased – net earnings. Technically, this is an implication of the fact that net earnings for period $t \rightarrow t + 1$ are \mathcal{F}_t -measurable. This strong, positive relationship between current net earnings and vessel prices is depicted in Figure 2.1. In addition, it is justified by the extremely high correlation coefficients (both in levels and growth rates) between these two variables (Table 2.3).

Furthermore, in analogy to commodity markets literature, due to the time-to-build required for the delivery of a newbuilding order, this first-order effect can be interpreted as a “convenience yield” – for having the asset readily available for leasing. In turn, this is reflected in the ratio of the 5-year old to the concurrent newbuilding vessel prices.⁴² In particular, as we observe in Figure 2.9, this ratio is positively correlated with net earnings. Noticeably, during market upturns the ratio is significantly higher than one and vice versa; that is, during a market peak – trough – 5-year old vessels can be substantially more – less – expensive than newbuilding ones. This result becomes even more interesting if we consider that the latter have significantly longer economic lives compared to the former.

However, apart from this first-order effect, increased current net earnings result in increased current net investment. As Kalouptsidei (2014) argues, entry into dry bulk shipping markets is free subject to an entry cost and time-to-build delays. In order to demonstrate this argument, we define current monthly net investment similar to Papapostolou *et al* (2014):


$$NI_t = (order_{t+1} - order_t + del_t) - scrap_t, \quad (2.20)$$

where $order_t$ denotes the order book at the beginning of period $t \rightarrow t + 1$, del_t the delivery of newly built fleet capacity during period $t \rightarrow t + 1$, and $scrap_t$ the demolished fleet capacity during the same period. Equation 2.20 incorporates also cancelations in existing orders and scrapping activity which are both regarded as negative investment decisions. Note that all variables in (2.20) are measured in dwt. Accordingly, we scale the net investment variable by the respective fleet size at the beginning of period $t \rightarrow t + 1$. Moreover, notice that the whole analysis below is conducted using data in monthly frequency. Therefore, the time index, t , corresponds to months and not years as in the previous analysis. Figure 2.10 demonstrates the fact that current net earnings and current scaled net investment are significantly positively correlated in every dry bulk sector.⁴³


⁴¹ For expositional simplicity, we have dropped the age subscript from the price notation.

⁴² Papapostolou *et al* (2014) define it as “delivery premium”.


⁴³ Specifically, the correlation coefficient in the Capesize sector is equal to 0.77, in the Panamax one is 0.52, in the Handymax one is 0.60, and in the Handysize one is 0.71.




Panel A: Capesize sector from 1/1996 to 12/2014.



Panel B: Panamax sector from 1/1996 to 12/2011.



Panel C: Handymax sector from 1/1996 to 6/2014.



Panel D: Handysize sector from 1/1996 to 12/2014.

n

turn, increased net investment results in increased future fleet capacity which, ceteris paribus, leads to decreased future net earnings. Notice that this decrease can be highly exacerbated due to the

Figure 2.10: Net Earnings and Net Investment.

Panels A-D illustrate the evolutions of the net earnings and scaled net investment variables for each dry bulk sector. Net investment is defined in equation 2.20 of the main text. Accordingly, it is scaled by the capacity of

the fleet at the beginning of the corresponding period. We measure net investment in dwt as opposed to the number of vessels. Net earnings are measured in December 2014 million dollars. The data are in a monthly frequency.

mean-reverting – around a time trend – character of demand. In general, this is the case when future demand is lower compared to the corresponding demand expectations formed at time t . We justify this argument formally by performing 1-, 2-, and 3-year horizon predictive OLS regressions of future log net earnings growth on current scaled net investment, for each of the four dry bulk sectors:

$$\pi_{t+12n} - \pi_t = \alpha_{NI,n} + \beta_{NI,n} \cdot NI_{t,n} + \varepsilon_{NI,t+12n}, \quad n \in \{1, 2, 3\}, \quad (2.21)$$

where $\pi_{t+12n} - \pi_{t+1}$ is the n -year log net earnings growth and $12n$ is the forecasting horizon measured in months.

However, the problem with the net investment variable as defined in (2.20) is that the order book data become available from January of 1996. As a result, the net investment variable time series are relatively small. Therefore, in order to examine the relation between net investment and net earnings over a longer period of time, we incorporate an additional investment variable, defined as “realised net investment”. This variable is solely based on data related to deliveries and scrapping activity. Specifically, similar to Greenwood and Hanson (2015), we assume that current newbuilding contracting is realised within the next 13 to 24 months, that is, during period $t + 13 \rightarrow t + 24$. Furthermore, we assume that current demolitions take place over the period $t \rightarrow t + 12$. Accordingly, we define realised net investment, RI_t , as:

$$RI_t = del_{t+13 \rightarrow t+24} - scrap_{t \rightarrow t+12}. \quad (2.22)$$

Once again, we scale the investment variable by the corresponding fleet size at the beginning of period $t \rightarrow t + 1$. Accordingly, we perform a second set of regressions, similar to the ones in equation 2.21, using this time the realised investment as the explanatory variable:

$$\pi_{t+12n} - \pi_t = \alpha_{RI,n} + \beta_{RI,n} \cdot RI_{t,n} + \varepsilon_{RI,t+12n}, \quad n \in \{1, 2, 3\}, \quad (2.23)$$

As one would expect, the results in Table 2.11 suggest that current net investment negatively predicts future net earnings growth. In particular, the slope coefficients are negative across all sectors and horizons. Importantly, we observe that the absolute magnitude and significance of the slope coefficients increase noticeably in the 2-year horizon. As we have already analysed in Section

2.III, the latter result can be explained by the time lag required for the delivery of a newbuilding order which is on average equal to approximately 2 years.

Accordingly, one may question the negative sign of the 1-year regression slope coefficients. The explanation to this objection can be divided into two parts. First, future net earnings growth depends also on demand which, as analysed, is exogenous and reverts rapidly around an upward sloping long-term drift. Hence, a negative shock in demand may lead to a decrease in future net earnings even if the supply had not increased in the meantime. Second, recall that the net contracting variable includes also cancellations and scrapping activity. In practice, when freight rates are high cancellations and scrapping are at significantly low levels. Thus, the combined result of these two facts can be the decrease of the 1-period ahead net earnings. Of course, as the horizon increases, the newbuilding investment decisions made at t will be realised and the net earnings decrease will be amplified.

Table 2.11: Regressions of future net earnings growth on net investment.

n	Net Investment					Realised Net Investment				
	T	β	t^{NW}	t^H	R^2	T	β	t^{NW}	t^H	R^2
Panel A: Capesize Sector										
1	216	-13.97*	-1.66	-1.68	0.07	264	-3.78*	-3.18	-3.27	0.09
2	204	-26.71***	-10.78	-11.23	0.16	252	-7.73***	-3.72	-3.69	0.24
3	192	-32.01***	-7.45	-7.09	0.20	240	-10.29***	-3.50	-3.40	0.38
Panel B: Panamax Sector										
1	180	-47.00***	-3.40	-3.35	0.12	420	-5.73***	-3.29	-3.34	0.08
2	168	-53.91***	-5.15	-5.60	0.11	408	-11.62***	-3.64	-3.70	0.22
3	156	-52.74***	-7.03	-7.19	0.13	396	-10.86***	-5.72	-5.76	0.18
Panel C: Handymax Sector										
1	210	-14.41	-1.31	-1.34	0.05	327	-2.75***	-2.69	-2.58	0.06
2	198	-33.87***	-5.40	-5.46	0.18	315	-6.60***	-5.11	-5.15	0.21
3	186	-32.04***	-5.26	-5.56	0.16	303	-7.85***	-9.30	-9.30	0.31
Panel D: Handysize Sector										
1	216	-30.33	-1.60	-1.60	0.10	456	-2.94*	-1.92	-2.00	0.03
2	204	-46.71***	-4.43	-4.54	0.16	444	-10.30***	-4.01	-4.09	0.18
3	192	-50.69***	-4.79	-5.51	0.18	432	-13.53***	-4.56	-4.55	0.28

Notes: This table reports results from 1-, 2-, and 3-year horizon regressions of real log net earnings growth on net and realised net investment. The data for the net and realised net investment starts from January 1996 and January 1976, respectively. To account for the overlapping nature of the variables, t-statistics are estimated using the Newey-West (1987) HAC, and the Hodrick (1992) corrections and are denoted by t^{NW} , and t^H , respectively. The predictive coefficient, β , is accompanied by *, **, or *** when the absolute t^{NW} statistic indicates significance at the 10%, 5% or 1% levels, respectively.

Consequently, market participants at time t anticipate – on average and up to a certain degree – this mechanism and, in turn, value second-hand vessels as if they expect future net earnings to be decreased compared to the prevailing ones.⁴⁴ Hence, current net earnings – through current investment – have a negative second-order effect (SOE) to current second-hand prices. Therefore, in a market upturn the growth rate of net earnings is significantly higher compared to the one of prices (Figure 2.2). Vice versa, during market downturns, current net earnings decrease substantially more than vessel prices because investors anticipate – the mean reversion of net earnings and, thus – that future net earnings will be higher. Specifically, low net earnings result in low (even negative) current net investment which, in conjunction with an expected increase in future demand, results in expectations of higher future net earnings. In this case, current net earnings have a negative first-order effect on current prices but a positive second-order one.

This explanation is in accordance with Greenwood and Hanson (2015) who argue that investors recognise – on average and up to a certain degree – the mean-reverting character of net earnings. This, in turn, results in a much more conservative – less naïve – valuation of vessels compared to the extreme case in which investors would assume that current earnings will also prevail in the future.

A first implication of this mechanism is the fact that net earnings are substantially more volatile than prices (as illustrated in Tables 2.2 and 2.4). A second one is that earnings yields are strongly positively related with net earnings and vessel prices (Table 2.3). Since in financial markets valuation ratios are used as indicators of fundamental value of the generated cash flow relative to corresponding price of the asset (Campbell and Shiller, 1988b), we can argue that, in shipping, during market peaks (troughs) vessels are undervalued (overvalued) compared to their respective generated cash flows (Figures 2.1 and 2.2). The third and most important implication of this mechanism is that high shipping earning yields strongly reflect market expectations about deteriorating future market conditions (i.e., negative net earnings growth).

Note that the above economic argument and its implications are further analysed in Chapter 3 where we develop a structural microeconomic model that directly relates net earnings, second-hand vessel prices, and trading activity in the sale and purchase market for vessels.

2.IV.B. Comparison to Other Markets

From a statistical perspective, we have answered a question of relative predictability. Specifically, according to the Campbell-Shiller variance decomposition framework (1988b), variability of valuation ratios must be due to either predictability of future returns or/and predictability of future cash flow growth or/and predictability of the terminal cash flow-price ratio. As Chen *et al* (2012)

⁴⁴ Note that this point is further analysed in Chapter 3 of this thesis.

argue, this relative predictability reveals which component is more important in driving price movements since, essentially, there is a trade-off between cash flow growth and return predictability. Furthermore, Chen *et al* (2012) and Rangvid *et al* (2014) illustrate that cash flow predictability is positively related to cash flow volatility. Therefore, one should expect that in dry bulk shipping where the generated cash flows are extremely volatile – in particular, extremely mean-reverting in longer horizons – due to boom-bust cycles, the variability of valuation ratios will be attributed mainly to cash flow predictability. Indeed, as we have demonstrated, this is precisely the case.

Accordingly, this argument can be applied to other real asset economies as well; in particular, this is also the case for the bulk of the U.S. residential and commercial real estate markets (Ghysels *et al*, 2012). Noticeably, though, the results are even more profound in the commercial part of the industry (Plazzi, Torous, and Valkanov, 2010) which, in terms of fundamentals, is arguably closer to shipping markets.

Specifically, Hamilton and Schwab (1985) examine 49 urban housing markets and find a strong negative relation between the rent yield and future rent growth. This result is in line with Gallin (2008) who incorporates a longer forecasting horizon equal to 4 years. In addition, Gallin illustrates that while there exists a positive relation between the rent yield and future returns, it is statistically insignificant. Ghysels *et al* (2012) estimate predictive regressions of future returns on the current rent yield using data for residential and commercial properties. In both cases, their results suggest that the returns coefficients, while being positive, are statistically insignificant. In addition, they incorporate in their estimation the Real Estate Investment Trust (REIT) index series and distinguish further between companies that mainly hold industrial buildings and offices, retail properties, and apartment buildings. Accordingly, they run future returns and cash flow growth regressions on the corresponding valuation ratio for various horizons. Once again, the results from the future cash flow regressions are higher in absolute magnitude and much more significant than the ones from the returns regressions. In line with Plazzi, Torous, and Valkanov (2010), this predictability is even more substantial in the case of industrial and office properties.

In contrast to these findings, Campbell *et al* (2009), applying the dynamic Gordon growth model to the U.S. housing markets, demonstrate that risk premia account for a substantial proportion of rent yields' volatility. When they split their sample into two subperiods, however, their findings suggest that, during the period 1997-2007, variation in expected rent growth was the principle source of variation in the rent-price ratio.

In addition, while our results coincide with recent findings from the majority of international (Rangvid, *et al*, 2014) and the pre-WWII U.S. (Chen 2009) equity markets they are diametrically

opposed to the ones from the post-WWII U.S. equity markets (Cochrane, 2011). Namely, in the latter case, the dividend yield is strongly positively related with future returns while future dividend growth appears to be unpredictable. A further finding, though, related to the U.S. equity markets is that future dividend growth is highly predictable when examining firm-level data (Vuolteenaho, 2002). The reason is that when using the – aggregated – U.S. market-level data this idiosyncratic predictability is “washed out” and, as a result, the aggregate market-wide dividend growth is unpredictable (Rangvid *et al*, 2014).

Thus, the question of interest is what drives the observed similarities and differences in the obtained results across these different industries – but also, in the equity markets case, the differences

both across time for a given country (as in the U.S. equity markets) and in a cross-country setting. To this end, recall that cash flow predictability appears to be positively related with cash flow volatility. Moreover, in order for asset prices at time t to move due to expectations about future cash flows there should be news about the latter; that is, cash flows must be economically predictable by market agents at time t . Vice versa, if future cash flows are not predictable using the time t information filtration then they cannot be predicted by valuation ratios. Namely, as mentioned above, according to Fama and French (1988a), the forecasting variable is implicitly assumed to be summarising the time t historical and prevailing economic conditions.

Therefore, from an economic perspective, as we illustrated for the case of dry bulk shipping, the major determinant of cash flow (net earnings) growth predictability by the valuation ratio (net earnings yield) is the significance of second-order effects of current cash flows (net earnings) on current prices. If there are no profound SOEs then there is no reason for future cash flows to be predictable using the current information filtration. Naturally, this explanation can be directly applied to the real estate industry. In support of this argument, Abraham and Hendershott’s (1996) findings suggest that rent growth predictability in residential markets is related to local supply elasticity measures (e.g., the availability of desirable land). Furthermore, Wheaton and Torto (1988) illustrate a strong relationship between future rent growth and current excess vacancy.

Finally, in order to relate the shipping and real estate results to the stock markets ones, it is fruitful to incorporate the well-known corporate finance notion of “dividend smoothing” and the corresponding equity markets’ empirical findings.⁴⁵ Namely, Chen *et al* (2009) show that there exists statistically significant evidence of dividend smoothing on the aggregate U.S. level in the post-WWII

⁴⁵ Notice that in equity markets there is information asymmetry between the issuer and the holder of the asset. In contrast, in real asset industries, such as shipping, there is, essentially, no information asymmetry between the asset holder (ship owner) and the cash flow payer (charterer). As a result, future cash flows can be more accurately predicted – there is no signalling hypothesis.

period. Furthermore, dividends are unambiguously more smoothed during this period compared to the pre-WWII era during which, the evidence of dividend smoothing is statistically insignificant. Accordingly, using simulation analysis, the authors demonstrate that dividend smoothing causes two substantial effects.

First it expunges the predictability of future dividend growth. This result should be a priori expected on a theoretical basis since dividend smoothing disentangles dividends from fluctuations in dividend-price ratios. In line with this finding, Rangvid *et al* (2014) show that in equity markets that experience less future dividend growth predictability, dividends are indeed more smoothed. The general conclusion demonstrated by Rangvid *et al* is that predictability is stronger in countries where dividends

are less smooth, the typical firm is small, and volatility is higher; that is, in relatively small and less developed markets.⁴⁶ Second, dividend smoothing results in substantially high persistence of the dividend yield⁴⁷ which, in turn, has a great implication on the empirical results. Specifically, as analysed in Sections 2.II and 2.III, high dividend yield persistence, accompanied by strong return predictability, causes the slope coefficients and R^2 s of future returns regressions to increase with the forecasting horizon (Fama and French, 1988b).

In line with these arguments, Chen *et al* (2012) extend the equity markets' literature by exploring and decomposing two alternative valuation ratios which are substantially less affected by dividend smoothing; namely, the earnings and net payout yields. The results from this decomposition are remarkable since they suggest that in the U.S. equity markets, both in the pre- and post-WWII era, the bulk of earnings (net payout) yield's volatility is attributed to earnings (net payout) growth and not to returns. This finding implies that news about future cash flows have a much more significant role than news about future returns in the determination of stock prices.

As we analysed in this chapter, however, cash flows in real economies – particularly, in shipping and commercial real estate – characterised by severe boom-bust investment cycles are far from smooth. Thus, since the degree of cash flow smoothing is negatively related to cash flow volatility (Chen *et al*, 2012; Rangvid *et al*, 2014) we should expect to evidence similar patterns in the shipping and real estate industries – but also in the pre-WWII U.S. stock markets and the bulk of international equity markets. In other words, when cash flows are significantly smoothed, asset prices do not move due to news about future cash flows because, essentially, there is no news. Accordingly, there

⁴⁶ The link between the size and the level of maturity of a firm (or market) and the degree of dividend smoothing was first established by Leary and Michaely (2011) using firm-level U.S. data.

⁴⁷ Interestingly, the estimated AR(1) coefficients for the pre-war and post-war periods are equal to 0.557 and 0.956, respectively.

is no significant second-order effect implied by the current cash flow on current asset prices; hence, future cash flow growth predictability is substantially low.

As a result, since we examine a question of relative predictability, when future cash flow growth predictability is low, predictability of future returns will be increased. This is precisely the case in the post-WWII U.S. equity markets. In contrast, the less smoothed cash flows are, the greater the SOEs of current cash flows on current asset prices become because investors can anticipate the effect of the most recently paid cash flow on future ones. In particular, when dividends are not smoothed, they are highly dependent upon the corresponding earnings. In turn, future earnings volatility and predictability imply future dividend volatility and predictability, respectively. Accordingly, market participants are more capable of predicting future dividends by forecasting future earnings. In line with the corporate finance literature, investors may perceive a large dividend increase as a lack of investment and growth opportunities by the firm (Mozes and Rapaccioli, 1998). Consequently, investors may believe that this increase is associated with negative future net earnings growth and, in turn, with negative future dividend growth – due to lack of dividend smoothing. As a result, in markets where firms smooth their dividends less, dividend yields strongly and negatively predict future dividend growth. In this case, a high current dividend has a negative second-order effect on current prices. Importantly, this argument is perfectly aligned with Chen *et al* (2012) who decompose the variance of earnings yield – in addition to the dividend-price ratio – in the U.S. stock markets. Namely, they find a significant negative relationship between the current earnings yield and future net earnings growth.

2.V. Conclusion

This chapter analyses the relation between second-hand vessel prices, net earnings, and holding period returns in the Capesize, Panamax, Handymax, and Handysize sectors of the dry bulk shipping industry. Namely, we examine whether earnings yields move due to changing expectations about future net earnings growth or/and due to changing expectations about future returns or/and due to changing expectations about the terminal spread between the resale – or scrap – price of the vessel and the corresponding prevailing net earnings in the market – in our context, this is referred to as “terminal earnings yield”.

Specifically, through the Campbell-Shiller (1988b) variance decomposition framework, we provide strong statistical evidence that almost the entire volatility of earnings yields can be attributed to variation in expected net earnings growth; almost none to expected returns variation and almost none to varying expectations about the terminal earnings yield. Therefore, we demonstrate formally that vessel valuation ratios mainly move due to news about net earnings

growth and not due to time-varying expected returns. Equivalently, dry bulk vessel prices – mainly – move due to news about future net earnings and not due to news about future returns.

In particular, shipping net earnings-price ratios are negatively and significantly related to future net earnings growth. Furthermore, there is no consistent, strong statistical evidence supporting the existence of time-varying risk premia in the valuation of dry bulk shipping vessels. This latter argument is further reinforced using the Campbell-Shiller (1988a) VAR framework. Specifically, we illustrate that actual price-net earnings ratios can be replicated sufficiently well through a VAR model with constant required returns. To the best of our knowledge, these stylised facts had never been documented before in the shipping literature as authors have mainly focused on the predictability of future returns.

From a technical perspective, we contribute to the empirical asset pricing literature by extending the familiar Campbell-Shiller variance decomposition (1988b) and VAR (1988a) frameworks to account for both a “forward-looking” valuation ratio and economic depreciation in the value of the asset. Note that this extension can also be incorporated in other real economies where the respective assets have limited economic lives, such as the commercial real estate and airplane industries.

In addition – by examining a real, capital intensive industry with distinct supply and demand determinants – we provide an economic interpretation for the obtained empirical results. To this end, we examine and incorporate the well-known shipping supply and demand freight rate mechanism. Accordingly, from an economic point of view, we argue that in order for valuation ratios to significantly predict future cash flows, current cash flows must have a profound second-order effect on the current price of the asset through the future cash flow stream. Therefore, we provide a bridge between the existing empirical asset pricing theory and the basic microeconomic principles that characterise a real asset industry like shipping. Furthermore, we extend this argument to explain and justify the observed similarities and differences in the respective results across different markets. In particular, our shipping results agree with recent findings from the pre-WWII U.S. equity markets, the bulk of international equity markets, and the majority of the U.S. real estate industry. They are diametrically opposed, however, to the corresponding findings in the post-WWII U.S. equity markets literature.

Finally, from a statistical perspective – and in line with recently obtained evidence – we argue that the significant predictability of earnings growth by the earnings yield is driven by the extreme volatility of shipping net earnings.

2.V.A. Connection to Chapter 3

As illustrated above, Chapter 2 analyses the determination mechanism of the net earnings yield—that is, the main valuation ratio of the second-hand shipping market – and, in turn, the interrelation between vessel prices, net earnings, and holding period returns. Accordingly, Chapter 3 integrates and concludes the examination of this market by incorporating in the analysis the trading activity related to second-hand vessels. For this purpose, we develop and estimate empirically a heterogeneous expectations asset pricing model with microeconomic foundations that can account for some distinct characteristics of the market.

Namely, among other features, our partial equilibrium model reproduces the actual behaviour of vessel prices and the positive correlation between net earnings and second-hand vessel transactions. Moreover, our model implicitly captures the fact that second-hand markets for vessels are rather illiquid while it also accounts for the main findings of Chapter 2; namely, for the fact that net earnings yields are highly positively correlated with the prevailing market conditions and, in turn, strongly negatively forecast future net earnings growth, but also for the finding that the bulk of the yield's volatility is attributed to expected cash flow variation and not to time-varying expected returns.

Our discrete-time economy consists of two agent types, conservatives and extrapolators, who form heterogeneous expectations about future net earnings and at the same time under (over) estimate the future demand responses of their competitors. Interestingly, formal estimation of the model suggests that in order to simultaneously match the empirical regularities, the average investor expectations in the second-hand market for ships must be “near-rational”. In particular, the investor population must consist of a very large proportion of agents with totally – or very close to – rational beliefs while the remaining fraction must hold highly extrapolative beliefs; thus, there must exist significant heterogeneity of beliefs in the market.

From an economic perspective, this finding is in accordance with the nature of the shipping industry; namely, the large fraction of conservative investors corresponds to the large number of established shipping companies that operate in the industry. In some instances, ship owning families have been present in the market for more than a century (Stopford, 2009) and, consequently, have strong prior experience and expertise about the key supply and demand drivers of the industry, analysed subsection 2.IV. In turn, their superior knowledge translates into more rational forecasts about future market conditions compared to relatively new investors.

Extrapolators, on the other hand, reflect new entrants such as diversified investors (e.g., private equity firms) with little or no previous experience of the market. It is well-documented that during prosperous periods, new entrants, impressed by the high prevailing earnings and short-term returns, are eager to buy vessels which, subsequently, are more than keen to sell as conditions begin to

deteriorate. In contrast, there are many cases where traditional, established owners have realised significant returns by selling vessels at the peak of the market and buying at the trough – a strategy known as “playing the cycles” (Stopford, 2009). Finally, note that while the empirical analysis in Chapter 3 focuses on the Handysize sector, our results have been tested to the remaining dry bulk sectors and are both qualitatively and quantitatively robust; thus, our conclusions are representative of the entire dry bulk industry.

2.V.B. Connection to Chapter 4

Having analysed in Chapters 2 and 3 the physical shipping market for second-hand vessels – that is, real assets – Chapter 4 examines the derivative market for Forward Freight Agreements (FFAs) – that is, financial instruments – related to the dry bulk sector of the shipping industry. Among other stylised facts, we illustrate formally that the bulk of volatility of the FFA basis – that is, the main valuation ratio in the FFA market – can be attributed to expectations about future physical market conditions rather than expectations about future risk premia. Interestingly, this result is perfectly aligned with the main finding of Chapter 2 that the bulk of earnings yields’ volatility can be attributed to variation in future market conditions rather than expected returns. More importantly, to provide an interpretation for our finding, we incorporate and, in turn, extend the economic and statistical arguments developed in Chapter 2 in conjunction with the related arguments in the seminal commodity markets literature. Namely, as analysed in Section 2.IV, cash flow predictability by valuation ratios is positively related to cash flow volatility and, thus, inversely related to the degree of cash flow smoothing. Accordingly, in Chapter 4, we attribute the role of dividend smoothing in equity markets to inventories and the cost of storage in commodity – and, in turn, shipping – markets.

Appendix 2

A.2.A. Exact Present Value Relation

We begin by deriving equation 2.8 of the main text. The first step towards this direction is to use the identity:

$$1 = R_{t+1}^{-1} \cdot R_{t+1} = R_{t+1}^{-1} \cdot \left[\frac{\Pi_{t+1} + P_{6,t+1}}{P_{5,t}} \right]. \quad (2. A1)$$

Subsequently, we substitute in (2. A1) the definition of the one-period return, illustrated in equation (2.3) of the main text:

$$1 = R_{t+1}^{-1} \cdot \left[\frac{\Pi_{t+1} + P_{6,t+1}}{P_{5,t}} \right]. \quad (2. A2)$$

Accordingly, multiplying both sides of (2. A2) by $P_{5,t}/\Pi_{t+1}$, and doing some algebra yields:

$$\frac{P_{5,t}}{\Pi_{t+1}} = R_{t+1}^{-1} \cdot \left[1 + \frac{P_{6,t+1}}{\Pi_{t+1}} \right]. \quad (2. A3)$$

Equation 2. A3 can be generalised to:

$$\frac{P_{5+n,t+n}}{\Pi_{t+n+1}} = R_{t+n+1}^{-1} \cdot \left[1 + \frac{P_{5+n+1,t+n+1}}{\Pi_{t+n+1}} \right], \quad (2. A4)$$

where $5 + n \leq 24$ denotes the age of the vessel when acquired by the owner. The problem with (2. A3) is that we cannot iterate it forward like Campbell and Shiller (1988a) and Cochrane (2005). Fortunately, though, since vessels have limited economic lives we can apply backward iteration in order to obtain P_{t+1}^6/Π_{t+1} . Namely, assuming an economic life of 25 years, at the end of which the vessels is scrapped⁴⁸ – and adjusting for economic depreciation of the asset – the terminal – scrap – price of the vessel 20 years ahead (i.e., at $t + 20$) is denoted by $S_{t+20} \equiv P_{25,t+20}$. Using (2. A4) with $n = 19$ and multiplying both sides of the equation by Π_{t+20}/Π_{t+19} , yields:

$$\frac{P_{24,t+19}}{\Pi_{t+19}} = R_{t+20}^{-1} \cdot \left[1 + \frac{S_{t+20}}{\Pi_{t+20}} \right] \cdot \frac{\Pi_{t+20}}{\Pi_{t+19}}.$$

Accordingly, iterating backwards, we observe that the ratio $P_{5+n,t+n}/\Pi_{t+n}$ can be obtained at any $n \in \{1, \dots, 19\}$ using the formula:

$$\frac{P_{5+n,t+n}}{\Pi_{t+n}} = \sum_{i=1}^{20-n} \left(\prod_{j=1}^i R_{t+n+j}^{-1} \cdot \frac{\Pi_{t+n+j}}{\Pi_{t+n+j-1}} \right) + \left(\prod_{j=1}^{20-n} R_{t+21-j}^{-1} \cdot \frac{\Pi_{t+21-j}}{\Pi_{t+20-j}} \right) \cdot \frac{S_{t+20}}{\Pi_{t+20}}. \quad (2. A5)$$

For $n = 1$, equation 2. A5 implies:

$$\frac{P_{6,t+1}}{\Pi_{t+1}} = \sum_{i=1}^{19} \left(\prod_{j=1}^i R_{t+j+1}^{-1} \cdot \frac{\Pi_{t+j+1}}{\Pi_{t+j}} \right) + \left(\prod_{j=1}^{19} R_{t+21-j}^{-1} \cdot \frac{\Pi_{t+21-j}}{\Pi_{t+20-j}} \right) \cdot \frac{S_{t+20}}{\Pi_{t+20}}. \quad (2. A6)$$

In turn, by substituting (2. A6) into (2. A3) we obtain:

$$\frac{P_{5,t}}{\Pi_{t+1}} = R_{t+1}^{-1} \cdot \left[1 + \sum_{i=1}^{19} \left(\prod_{j=1}^i R_{t+j+1}^{-1} \cdot \frac{\Pi_{t+j+1}}{\Pi_{t+j}} \right) + \left(\prod_{j=1}^{19} R_{t+21-j}^{-1} \cdot \frac{\Pi_{t+21-j}}{\Pi_{t+20-j}} \right) \cdot \frac{S_{t+20}}{\Pi_{t+20}} \right].$$

Equivalently, multiplying the last term within the square brackets by $\frac{\Pi_{t+21}}{\Pi_{t+21}}$ and performing some algebraic manipulation results in:

⁴⁸ In contrast to equity markets, the asset's economic life is limited in shipping; hence, we do not need to impose the transversality condition. The transversality or "no-bubbles" condition in equity markets is defined as:

$$\lim_{i \rightarrow \infty} E_t \left[\left(\prod_{j=1}^i R_{t+j}^{-1} \frac{D_{t+j}}{D_{t+j-1}} \right) \frac{P_{t+i}}{D_{t+i}} \right] = 0.$$

$$\frac{P_{5,t}}{\Pi_{t+1}} = R_{t+1}^{-1} + R_{t+1}^{-1} \cdot \sum_{i=1}^{19} \left(\prod_{j=1}^i R_{t+j+1}^{-1} \cdot \frac{\Pi_{t+j+1}}{\Pi_{t+j}} \right) + \left(\prod_{j=1}^{20} R_{t+21-j}^{-1} \cdot \frac{\Pi_{t+22-j}}{\Pi_{t+21-j}} \right) \cdot \frac{S_{t+20}}{\Pi_{t+21}}.$$

Finally, taking conditional expectations and exploiting the fact that Π_{t+1} is \mathcal{F}_t -measurable yields equation 2.8 of the main text:

$$\frac{P_{5,t}}{\Pi_{t+1}} = E_t \left[R_{t+1}^{-1} + R_{t+1}^{-1} \cdot \sum_{i=1}^{19} \left(\prod_{j=1}^i R_{t+j+1}^{-1} \cdot \frac{\Pi_{t+j+1}}{\Pi_{t+j}} \right) + \left(\prod_{j=1}^{20} R_{t+21-j}^{-1} \cdot \frac{\Pi_{t+22-j}}{\Pi_{t+21-j}} \right) \cdot \frac{S_{t+20}}{\Pi_{t+21}} \right]. \quad (2.A7)$$

A.2.B. Linearisation of the Present Value Relation

Equation 2.9 of the main text corresponds to the linearisation of equation 2.A7. For this purpose, we follow Campbell and Shiller (1988a), Cochrane (2005), and Alizadeh and Nomikos (2007). In addition, we extend the existing framework by (i) accounting for the fact that our net earnings-price ratio is forward-looking and (ii) adjusting for economic depreciation in the value of the asset – which, in turn, results in not imposing the transversality or “no-bubbles” condition.

Specifically, starting from equation 2.A3:

$$P_{5,t} = R_{t+1}^{-1} \cdot \left[1 + \frac{P_{6,t+1}}{\Pi_{t+1}} \right] \cdot \Pi_{t+1}$$

and taking logs on both sides the equation, we obtain:

$$p_{5,t} = -r_{t+1} + \pi_{t+1} + \ln(1 + e^{p_{6,t+1} - \pi_{t+1}}),$$

where $p_{n,t} = \ln(P_{n,t})$, $\pi_t = \ln(\Pi_t)$ and $r_t = \ln(R_t)$.

Applying a first-order Taylor expansion of the last term around a point $p_6 - \pi = \ln(P_6/\Pi)$ yields:

$$\begin{aligned} p_{5,t} &\approx -r_{t+1} + \pi_{t+1} + \ln\left(1 + \frac{P_6}{\Pi}\right) + \frac{P_6/\Pi}{1 + P_6/\Pi} \cdot [p_{6,t+1} - \pi_{t+1} - (p_6 - \pi)] \\ &\Rightarrow p_{5,t} \approx -r_{t+1} + \pi_{t+1} + \rho_1(p_{6,t+1} - \pi_{t+1}) + k_1, \end{aligned} \quad (2.B1)$$

where

$$\rho_1 = \frac{P_6/\Pi}{1 + P_6/\Pi} \text{ and } k_1 = -(1 - \rho_1) \ln(1 - \rho_1) - \rho_1 \ln(\rho_1).$$

As illustrated by Campbell and Shiller (1988a), the higher-order terms of the Taylor expansion that are neglected from (2.B1) create an approximation error and as a result (2.B1) does not hold

exactly.⁴⁹ Furthermore, in the equity markets' asset pricing literature, the point of expansion is usually assumed to be the natural logarithm of the sample mean price-dividend ratio. However, as Cochrane (2011) argues, this does not need to be the case. For instance, Lof (2015) approximates ρ by the sample mean of the ratio $\frac{P_t}{P_t+D_t}$. Alternatively, we can use as an approximation point the natural logarithm of the inverse of the sample mean dividend-price ratio – similar to (Cochrane, 2005) – or the natural logarithm of the fraction of the geometric mean of prices to the geometric mean of the corresponding cash flows (Alizadeh and Nomikos, 2007). Accordingly, for the first Taylor expansion, we set $P_6/\Pi = P_6/\Pi = 1/[(1/T)\sum_{t=0}^{T-1}\Pi_{t+1}/P_{6,t+1}]$. Notice that the choice of the expansion point and, consequently, of ρ , does not have a major implication on the results. Specifically, for robustness, we have incorporated a variety of expansion points and the empirical results remain approximately the same while the obtained conclusions are identical.⁵⁰

Accordingly, iterating equation 2. B1 forward yields:

$$p_{5,t} \approx \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) (1 - \rho_i) \pi_{t+i} - \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) r_{t+i} + \left(\prod_{i=1}^n \rho_i \right) p_{5+n,t+n} + \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) k_i, \quad 1 \leq n \leq 20, \quad (2. B2)$$

where $p_{5,t}$ is the current log price of a 5-year old vessel while $p_{5+n,t+n} = \ln(P_{5+n,t+n})$ is the log price of a $(5 + n)$ -year old vessel after n years. In addition, for $n = 20$ we obtain $p_{25,t+20} \equiv s_{t+20} = \ln(S_{t+20})$ which corresponds to the log scrap price of the vessel.

In the context of this research, however, we also have to account for economic depreciation in the value of the asset. Consequently, the i^{th} subsequent Taylor expansion is taken around the corresponding age-varying approximation point, defined as:⁵¹

$$p_{5+i} - \pi = \ln(P_{5+i}/\Pi), \text{ where } P_{5+i}/\Pi = 1/[(1/T)\sum_{t=0}^{T-1}\Pi_{t+1}/P_{5+i,t+1}] \text{ and } i \in \{1, \dots, n\}.$$

Subsequently,

$$\rho_i = \frac{P_{5+i}/\Pi}{1+P_{5+i}/\Pi} \quad \text{and} \quad k_i = -(1 - \rho_i) \ln(1 - \rho_i) - \rho_i \ln(\rho_i), \quad i \in \{1, \dots, 20\}.$$

⁴⁹ Campbell and Shiller (1988a) show that this error is in practice small and almost constant. Moreover, they argue that a constant approximation error does not have any implication on the empirical results when no restrictions on the means of the data are tested.

⁵⁰ Campbell and Shiller (1988b) demonstrate that letting ρ vary within a plausible range does not have a significant impact on the results and the conclusions.

⁵¹ Specifically, we construct new net earnings-price ratios variables, the numerators of which are equal to Π_{t+1} , while the denominators are equal to the prices of the corresponding $(5 + i)$ -year old vessels one-period ahead, $P_{5+i,t+1}$. In turn, we find the sample (arithmetic) means of these ratios.

Finally, notice that we have set $\rho_0 = 1$ since $i = 0$ corresponds to no Taylor expansion.

In contrast, in equity markets' asset pricing literature, where assets are assumed to be infinitely lived, the approximation points are constant and not age-varying. Consequently, ρ and k are also constant and (2.B2) is simplified to the well-known Campbell and Shiller (1988b) linear present-value formula. Thus, in this chapter we provide a generalisation of the existing framework that can cover in a mathematical rigorous manner the class of real assets with limited economic lives (e.g., vessels, houses, and airplanes).

The intuition behind formula 2.B2 is straightforward: high vessel prices are related to either high future net earnings or/and low future returns or/and high future vessel prices. However, since log prices and log net earnings are – usually – nonstationary variables, it is not appropriate to apply variance-bounds tests to equation 2.B2. Following Alizadeh and Nomikos (2007), a natural solution to this problem is to capitalise the cointegrating relationship between log vessel prices and log net earnings. Specifically, this can be achieved by subtracting the corresponding net earnings from both sides of equation 2.B2. However, since our definition of the net earnings-price ratio is forward-looking, we deviate from the existing asset pricing literature by subtracting π_{t+1} instead of π_t . Accordingly, we obtain:

$$p_{5,t} - \pi_{t+1} \approx \sum_{i=1}^{n-1} \left(\prod_{j=1}^i \rho_j \right) \Delta \pi_{t+i+1} - \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) r_{t+i} + \left(\prod_{j=1}^n \rho_j \right) (p_{5+n,t+n} - \pi_{t+n}) \\ + \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) k_i.$$

Equivalently,

$$\pi_{t+1} - p_{5,t} \approx - \sum_{i=1}^{n-1} \left(\prod_{j=1}^i \rho_j \right) \Delta \pi_{t+i+1} + \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) r_{t+i} \\ + \left(\prod_{j=1}^n \rho_j \right) (\pi_{t+n} - p_{5+n,t+n}) - \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) k_i, \quad (2.B3)$$

where $\pi_{t+1} - p_{5,t}$ is the forward-looking log net earnings-price ratio and $\Delta \pi_{t+1} = \pi_{t+1} - \pi_t$ is the 1-year horizon (log) net earnings growth. Notice that $\Delta \pi_{t+i}$ and r_{t+i} do not enter (2.B3) symmetrically since the log-net earnings growth series has one less term compared to the log-returns one. However, as we have analysed, both $P_{5+n,t+n}$ and Π_{t+n+1} are \mathcal{F}_{t+n} -measurable.

Therefore, we modify (2. B3) by adding and subtracting $(\prod_{j=1}^n \rho_j) \pi_{t+n+1}$ to and from the right-hand side of the equation:

$$\begin{aligned} \pi_{t+1} - p_{5,t} \approx & - \sum_{i=1}^n \left(\prod_{j=1}^i \rho_j \right) \Delta \pi_{t+i+1} + \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) r_{t+i} \\ & + \left(\prod_{j=1}^n \rho_j \right) (\pi_{t+n+1} - p_{5+n,t+n}) - \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) k_i. \end{aligned} \quad (2. B4)$$

Finally, since equation 2. B4 holds ex post, we can take conditional expectations at time t :

$$\begin{aligned} \pi_{t+1} - p_{5,t} \approx & - \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) k_i \\ & + E_t \left[- \sum_{i=1}^n \left(\prod_{j=1}^i \rho_j \right) \Delta \pi_{t+i+1} + \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) r_{t+i} + \left(\prod_{j=1}^n \rho_j \right) (\pi_{t+n+1} - p_{5+n,t+n}) \right]. \end{aligned} \quad (2. B5)$$

This corresponds to equation 2.9 of the main text.

Of course, equation 2. B5 can be easily extended to capture cases where each chartering period (i.e., each operating period of the vessel) corresponds to less than one year. In this case, the number of remaining operating periods, n , is estimated through $n = 20 \cdot f$, where f is the number of equal, consecutive time-charter contracts within a year (by definition, for annual contracts $f = 1$). Accordingly, in this general case, ρ_i is given by:

$$\rho_i = \frac{P_{5+i/f}/\Pi}{1 + P_{5+i/f}/\Pi}.$$

A.2.C. Variance Decomposition

In order to decompose the variance of the shipping net earnings-price ratio (equation 2.14 of the main text), we start by multiplying both sides of (2. B4) by $[(\pi_{t+1} - p_{5,t}) - E(\pi_{t+1} - p_{5,t})]$. Accordingly, taking expectations at time t on both sides, we obtain:

$$\begin{aligned}
var(\pi_{t+1} - p_{5,t}) &\approx -cov \left[\pi_{t+1} - p_{5,t}, \sum_{i=1}^n \left(\prod_{j=1}^i \rho_j \right) \Delta \pi_{t+i+1} \right] \\
&+ cov \left[\pi_{t+1} - p_{5,t}, \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) r_{t+i} \right] \\
&+ cov \left[\pi_{t+1} - p_{5,t}, \left(\prod_{i=1}^n \rho_i \right) (\pi_{t+n+1} - p_{5+n,t+n}) \right].
\end{aligned} \tag{2. C1}$$

The three terms in the right-hand side of (2. C1) are numerators of exponentially weighted long-run regression coefficients. Finally, dividing both sides of (2. C1) by $var(\pi_{t+1} - p_{5,t})$ yields:

$$1 \approx -b_n^{\Delta \pi} + b_n^r + b_n^{\pi-p}, \tag{2. C2}$$

where b_n^i is the n-year horizon coefficient corresponding to the i^{th} element of the decomposition.

A.2.D. Extension of the Campbell and Shiller (1988a) VAR Framework to Shipping

We begin from the log linear relation between the one-period holding return, the one-period net earnings, and the current and future prices for a 5-year old vessel (see Appendix 2.B):

$$p_{5,t} \approx -r_{t+1} + \pi_{t+1} + \rho_1(p_{5+1/f,t+1} - \pi_{t+1}) + k_1. \tag{2. D1}$$

For expositional simplicity, the age subscript will be dropped in the analysis below. Therefore, p_t corresponds to the price of a 5-year old vessel at time t , while p_{t+1} to the price of the same vessel after one period – at which point the asset will be 6 years old. In addition, as stated in the main text (Subsection 2.III.C), we impose the assumption that expected returns from holding the vessel for one period are constant; hence, $r_{t+1} = E_t[r_{t+1}] = r$. Incorporating these modifications in equation 2. D1, we obtain:

$$p_t \approx -r + \pi_{t+1} + \rho_1(p_{t+1} - \pi_{t+1}) + k_1. \tag{2. D2}$$

Iterating (2. D2) forward yields:

$$\begin{aligned}
p_t &\approx \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) (1 - \rho_i) \pi_{t+i} - \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) r + \left(\prod_{i=1}^n \rho_i \right) p_{t+n} \\
&+ \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) k_i,
\end{aligned}$$

where $n = 20 \cdot f$. Equivalently,

$$p_t \approx \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) (1 - \rho_i) \pi_{t+i} + \left(\prod_{i=1}^n \rho_i \right) p_{t+n} + \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) (k_i - r). \quad (2. D3)$$

Once again, due to the limited economic life of the vessel, we did not have to impose the transversality or “no-bubbles” condition when iterating forward the difference equation 2. D2. Next, in order to create the forward-looking log price-net earnings ratio, we subtract π_{t+1} from both sides of (2. D3):

$$\delta_t \approx \sum_{i=1}^{n-1} \left(\prod_{j=1}^i \rho_j \right) \Delta \pi_{t+i+1} + \left(\prod_{j=1}^n \rho_j \right) (p_{t+n} - \pi_{t+n}) + \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) (k_i - r), \quad (2. D4)$$

where, for expositional simplicity, we denote the price-net earnings ratio for the 5-year old vessel by $\delta_t = p_t - \pi_{t+1}$. Finally, adding and subtracting $(\prod_{j=1}^n \rho_j) \pi_{t+n+1}$ to and from the right-hand side of equation 2. D4 results in:

$$\delta_t \approx \sum_{i=1}^n \left(\prod_{j=1}^i \rho_j \right) \Delta \pi_{t+i+1} + \left(\prod_{j=1}^n \rho_j \right) \tau_{t+n} + \sum_{i=1}^n \left(\prod_{j=1}^i \rho_{j-1} \right) (k_i - r), \quad (2. D5)$$

where $\tau_{t+n} = s_{t+n} - \pi_{t+n+1}$ is the terminal – scrap – spread between the log of the scrap price of the vessel and the log of the prevailing net earnings at time t . In contrast to (2. B5), equation 2. D5 suggests that price-net earnings ratios’ movements are attributed to either future net earnings growth volatility or/and volatility of the terminal spread.

Redefining all variables as deviations from their means enables us to drop the constant term and, thus, simplify further equation 2. D5:

$$\delta_t \approx \sum_{i=1}^n \left(\prod_{j=1}^i \rho_j \right) \Delta \pi_{t+i+1} + \left(\prod_{j=1}^n \rho_j \right) \tau_{t+n}. \quad (2. D6)$$

Accordingly, parameters k_i and r are omitted from the analysis below.

Importantly, in this exercise, the age-varying approximation points, ρ_i , are estimated through:⁵²

$$\rho_i = \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{1 + \Pi_t / (1 - \frac{i}{100}) P_{5,t}}, \quad i \in \{1, \dots, 80\}.$$

Therefore, following the procedure described in Campbell and Shiller (1988a) we can test the model in equation 2.D6 using a log-linear Vector Autoregressive Model with p lags. In particular, we compare the observed log price-net earnings ratio, δ_t , with the forecast of the net earnings growth and scrap spread generated by the VAR(p) model, δ'_t .

To begin with, consider the case where at the beginning of period $t \rightarrow t + 1$ all market agents observe a vector of state variables denoted by \mathbf{y}_t which is assumed to summarise the current state of the economy. This vector includes the log price-net earnings ratio, δ_t , the log net earnings growth, $\Delta\pi_{t+1}$, and the terminal spread, τ_t . Equivalently, $\mathbf{y}_t = [\delta_t, \Delta\pi_{t+1}, \tau_t]'$. At this point, recall that the net earnings variable corresponding to period $t \rightarrow t + 1$ is \mathcal{F}_t -measurable. In addition, assume that all market participants at time t have access to precisely the same information set; that is, the history of state vectors, $\{\mathbf{y}_t, \mathbf{y}_{t-1}, \mathbf{y}_{t-2}, \dots\}$, denoted by the information filtration \mathcal{F}_t . Specifically, the state vector is assumed to follow a linear stochastic process with constant coefficients which are known to all market agents. This feature is very important since it implies that all market agents are symmetrically informed. Mathematically, the stochastic linear process that characterises the evolution of \mathbf{y}_t is expressed as a VAR(p):

$$\mathbf{y}_t = A_1 \mathbf{y}_{t-1} + A_2 \mathbf{y}_{t-2} + \dots + A_p \mathbf{y}_{t-p} + \boldsymbol{\varepsilon}_t, \quad (2.D7)$$

where A_i with $i = 1, 2, \dots, p$ are 3×3 matrices of coefficients known to market participants. Therefore, we can denote by A_{ijk} the slope coefficient of the j^{th} variable in the state vector \mathbf{y}_t on the k^{th} variable with a time lag equal to i . Accordingly, A_{ijk} is the (j, k) element of matrix A_i . Furthermore, $\boldsymbol{\varepsilon}_t$ is a 3×1 matrix consisting of error terms (white noises). In order to illustrate this notation, consider the equations for the three state variables at time t :

$$\begin{aligned} \delta_t &= \sum_{i=1}^p A_{i11} \delta_{t-i} + \sum_{i=1}^p A_{i12} \Delta\pi_{t+1-i} + \sum_{i=1}^p A_{i13} \tau_{t-i} + \varepsilon_{1,t} \\ \Delta\pi_{t+1} &= \sum_{i=1}^p A_{i21} \delta_{t-i} + \sum_{i=1}^p A_{i22} \Delta\pi_{t+1-i} + \sum_{i=1}^p A_{i23} \tau_{t-i} + \varepsilon_{2,t} \end{aligned}$$

⁵² Recall that the choice of the approximation point has negligible effect on the results.

$$\tau_t = \sum_{i=1}^p A_{i31} \delta_{t-i} + \sum_{i=1}^p A_{i32} \Delta\pi_{t+1-i} + \sum_{i=1}^p A_{i33} \tau_{t-i} + \varepsilon_{3,t}.$$

Following Sargent (1979), we can write this VAR(p) model in companion form (as a first-order autoregressive model) to take advantage of the convenient conditional expectations formula. Namely, we define a new vector, \mathbf{z}_t , which consists of $3p$ elements instead of 3; that is, apart from the 3 initial variables, $\delta_t, \Delta\pi_{t+1}$, and τ_t , it also includes $(p - 1)$ lags of each state variable. Similar to Campbell and Shiller (1988a), we can demonstrate this conversion by considering the VAR(2) model. In this case, $\mathbf{z}_t = [\delta_t, \delta_{t-1}, \Delta\pi_{t+1}, \Delta\pi_t, \tau_t, \tau_{t-1}]'$ and $\boldsymbol{\varepsilon}_t = [\varepsilon_{1,t}, 0, \varepsilon_{2,t}, 0, \varepsilon_{3,t}, 0]'$. Furthermore, the evolution of \mathbf{z}_t is characterised by a first-order VAR written in the following form:

$$\begin{bmatrix} \delta_t \\ \delta_{t-1} \\ \Delta\pi_{t+1} \\ \Delta\pi_t \\ \tau_t \\ \tau_{t-1} \end{bmatrix} = \begin{bmatrix} A_{111} & A_{211} & A_{112} & A_{212} & A_{113} & A_{213} \\ 1 & 0 & 0 & 0 & 0 & 0 \\ A_{121} & A_{221} & A_{122} & A_{222} & A_{123} & A_{223} \\ 0 & 0 & 1 & 0 & 0 & 0 \\ A_{131} & A_{231} & A_{132} & A_{232} & A_{133} & A_{233} \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} \delta_{t-1} \\ \delta_{t-2} \\ \Delta\pi_t \\ \Delta\pi_{t-1} \\ \tau_{t-1} \\ \tau_{t-2} \end{bmatrix} + \begin{bmatrix} \varepsilon_{1,t} \\ 0 \\ \varepsilon_{2,t} \\ 0 \\ \varepsilon_{3,t} \\ 0 \end{bmatrix}.$$

In general, a VAR(p) in companion form can be expressed using the following equation:

$$\mathbf{z}_t = \mathbf{A}\mathbf{z}_{t-1} + \boldsymbol{\varepsilon}_t, \tag{2.D8}$$

where \mathbf{z}_t and $\boldsymbol{\varepsilon}_t$ are $3p \times 1$ matrices and \mathbf{A} is a $3p \times 3p$ matrix of constants. Noticeably, the rows describing the initial state variables are stochastic, while the remaining ones deterministic. As mentioned above, the VAR(p) written in the form of equation 2.D8 has the following, very convenient, property:

$$E[\mathbf{z}_{t+1} | \mathcal{F}_t] = E_t[\mathbf{z}_{t+1}] = \mathbf{A}\mathbf{z}_t \Rightarrow E_t[\mathbf{z}_{t+n}] = \mathbf{A}^n \mathbf{z}_t, \tag{2.D9}$$

which implies that once matrix \mathbf{A} is estimated it can be incorporated to forecast n periods ahead, simply by multiplying \mathbf{z}_t by the n^{th} power of \mathbf{A} . Finally, following the notation in Campbell and Shiller (1988a), we define the selection vectors $\mathbf{e1}$, $\mathbf{e2}$, and $\mathbf{e3}$ such that $\mathbf{e1}'\mathbf{z}_t = \delta_t$, $\mathbf{e2}'\mathbf{z}_t = \Delta\pi_{t+1}$, and $\mathbf{e3}'\mathbf{z}_t = \tau_t$, respectively. As an example, in the VAR(2) case illustrated above, these vectors correspond to $\mathbf{e1}' = [1,0,0,0,0,0]'$, $\mathbf{e2}' = [0,0,1,0,0,0]'$, and $\mathbf{e3}' = [0,0,0,0,1,0]'$.

Importantly, the VAR(p) model described above is tightly restricted by the log-linear present-value model in equation 2.D6. Specifically, taking conditional expectations at time t – that is, expectations conditional on \mathcal{F}_t – on both sides of equation 2.D6 and exploiting the fact that δ_t is \mathcal{F}_t -measurable yields:

$$\delta_t \approx E_t \left[\sum_{i=1}^n \left(\prod_{j=1}^i \rho_j \right) \Delta \pi_{t+i+1} + \left(\prod_{j=1}^n \rho_j \right) \tau_{t+n} \right] \equiv \delta'_t, \quad (2. D10)$$

where δ'_t is the unrestricted VAR forecast of $\sum_{i=1}^n (\prod_{j=1}^i \rho_j) \Delta \pi_{t+i+1} + (\prod_{j=1}^n \rho_j) \tau_{t+n}$. Multiplying the selection vectors by (2. D9) and iterating forward results in:

$$E_t[\delta_{t+1}] = \mathbf{e1}' \mathbf{A} \mathbf{z}_t \Rightarrow E_t[\delta_{t+i}] = \mathbf{e1}' \mathbf{A}^i \mathbf{z}_t$$

$$E_t[\Delta \pi_{t+2}] = \mathbf{e2}' \mathbf{A} \mathbf{z}_t \Rightarrow E_t[\Delta \pi_{t+i+1}] = \mathbf{e2}' \mathbf{A}^i \mathbf{z}_t$$

$$E_t[\tau_{t+1}] = \mathbf{e3}' \mathbf{A} \mathbf{z}_t \Rightarrow E_t[\tau_{t+i}] = \mathbf{e3}' \mathbf{A}^i \mathbf{z}_t.$$

Accordingly, equation 2. D10 can be written as:

$$\begin{aligned} \delta_t = \mathbf{e1}' \mathbf{z}_t &\approx \sum_{i=1}^n \left(\prod_{j=1}^i \rho_j \right) \mathbf{e2}' \mathbf{A}^i \mathbf{z}_t + \left(\prod_{j=1}^n \rho_j \right) \mathbf{e3}' \mathbf{A}^n \mathbf{z}_t \\ &= \left[\sum_{i=1}^n \left(\prod_{j=1}^i \rho_j \right) \mathbf{e2}' \mathbf{A}^i + \left(\prod_{j=1}^n \rho_j \right) \mathbf{e3}' \mathbf{A}^n \right] \mathbf{z}_t \equiv \delta'_t \end{aligned} \quad (2. D11)$$

Therefore, in order for the left- and right-hand sides of equation 2. D11 to be equal, the following condition has to be satisfied:

$$\mathbf{e1}' = \sum_{i=1}^n \left(\prod_{j=1}^i \rho_j \right) \mathbf{e2}' \mathbf{A}^i + \left(\prod_{j=1}^n \rho_j \right) \mathbf{e3}' \mathbf{A}^n, \quad (2. D12)$$

Equation 2. D12 imposes a set of $3p$ nonlinear restrictions on the coefficients of the VAR model. In conclusion, the series of model-implied log price-net earnings ratios, δ'_t , can be generated through the following equation – which corresponds to equation 2.18 of the main text:

$$\delta'_t = \left[\sum_{i=1}^n \left(\prod_{j=1}^i \rho_j \right) \mathbf{e2}' \mathbf{A}^i + \left(\prod_{j=1}^n \rho_j \right) \mathbf{e3}' \mathbf{A}^n \right] \mathbf{z}_t. \quad (2. D13)$$

Chapter 3: Heterogeneous Expectations and the Second-Hand Market for Dry Bulk Ships

Abstract. This chapter investigates the joint behaviour of vessel prices, net earnings, and second-hand activity in the dry bulk shipping industry. We develop and estimate empirically a behavioural asset pricing model with microeconomic foundations that can account for some distinct characteristics of the market. Namely, among other features, our partial equilibrium model reproduces the actual volatility and behaviour of vessel prices, the average trading activity in the market, and the positive correlation between net earnings and second-hand transactions. To explain these findings, we depart from the rational expectations benchmark of the model, incorporating extrapolative beliefs – mainly – on a part of the investor population. In contrast to the majority of financial markets’ behavioural models, however, in our environment agents extrapolate fundamentals, not past returns. Accordingly, we introduce two types of agents who hold heterogeneous beliefs regarding the cash flow process. Formal estimation of the model indicates that a heterogeneous beliefs environment where a small fraction of market agents highly extrapolates fundamentals compared to the rest of the population – while both agent types simultaneously under (over) estimate their competitors’ future demand responses – can explain the positive relation between net earnings, prices, and second-hand vessel transactions. To the best of our knowledge, the second-hand market for vessels has never been examined from the perspective of a structural, behavioural economic model in the shipping literature before.

Keywords: Asset Pricing, Vessel Valuation, Biased Beliefs, Cash Flow Extrapolation, Heterogeneous Agents, Trading Activity

3.1. Introduction

As it is well-established in the asset pricing literature, most rational expectations models fail to explain numerous empirical regularities related to asset prices. Among others, two prominent examples are the “excess volatility puzzle” (Leroy and Porter, 1981) and the positive correlation between trading volume and asset prices (Barberis *et al*, 2015b). To explain these findings, one of the tools that researchers have developed are heterogeneous beliefs economic models that incorporate behavioural biases, termed as heuristics (Barberis *et al*, 2015a).

In this chapter, we extend the application of this type of models to real assets and, specifically, vessels. As analysed in Chapter 1, shipping is a very important sector of the world economy since 90% of the world trade is transported by sea and it is justifiably considered as a leading indicator of world economic activity (Killian, 2009). Hence, it is important to understand the pricing and trading dynamics of this asset class. To the best of our knowledge, this is the first time that a structural heterogeneous beliefs asset pricing model is applied to a real asset economy. Thus, we provide a

framework that can be incorporated and, in turn, empirically evaluated in other markets with similar characteristics, such as the airplane and the commercial real estate industries.

Accordingly, we develop a heterogeneous beliefs model that can provide a plausible economic interpretation for numerous empirical findings related to the sale and purchase market for vessels of the dry bulk sector of the shipping industry. While the empirical analysis focuses on the Handysize sector, our results have been tested to the remaining dry bulk ones and are both qualitatively and quantitatively robust; thus, our conclusions are representative of the entire dry bulk industry. Therefore, the main motivation for this Chapter is to construct an economic model able to simultaneously explain (i.e., *ex post*) several empirical regularities observed in the shipping industry – that is, the aim is not to develop a forecasting framework (i.e., *ex ante*).

Namely, the proposed partial equilibrium framework explains the observed behaviour of second-hand vessel prices: in particular, the actual price volatility and the high correlation between prices and prevailing net earnings. In addition, our model reproduces and justifies the stylised fact that trading activity is positively related to both market conditions and absolute changes in net earnings between two consecutive periods. In our sample, the two correlation coefficients are equal to 0.53 and 0.65, respectively, which implies that investors trade more aggressively during prosperous market conditions but also when net earnings have significantly changed compared to the previous period. Moreover, our model implicitly captures the fact that second-hand markets for vessels are rather illiquid: during the period 1995-2014, the average annual sale and purchase turnover was approximately 5.8% of the corresponding fleet size. Finally, the proposed framework also accounts for the stylised features presented in Chapter 2 of this thesis; namely, for the finding that net earnings yields are highly positively correlated with the prevailing market conditions and, in turn, strongly negatively forecast future net earnings growth but also for the fact that the bulk of the earnings yield's volatility is attributed to expected cash flow variation and not to time-varying expected returns.

Our discrete time environment consists of two agent types: “conservatives” and “extrapolators”. Annual shipping net earnings are the sole state variable – observed at each period by the entire investor population – and, when valuing the asset at each period, agents maximise recursively a constant absolute risk aversion (CARA) utility function defined over next period's wealth. In accordance with the nature of the industry, both agents face short-sale constraints. Importantly, both types of agent value vessels based on fundamentals – that is, shipping net earnings – however, they form heterogeneous expectations regarding their evolution and at the same time under (over) estimate the future demand responses of their competitors. Specifically, while conservatives have totally rational or “near-rational” beliefs about the cash flow process, extrapolators hold highly

extrapolative expectations. From a psychological perspective, the extrapolation of fundamentals can be the result of several heuristic-driven biases, the most frequent being the “representativeness heuristic” according to which, individuals believe that small samples are representative of the entire population (Tversky and Kahneman, 1974).

In addition, each agent’s investment strategy is independent of the other’s. Namely, both types of agent assume that in all future periods the other type will maintain his per-capita fraction of the risky asset supply (Barberis *et al*, 2015b). From a psychological point of view, this misbelief can be driven by a bias known as “competition neglect” (Camerer and Lovo, 1999; Kahneman, 2011) which leads agents to form forecasts about competitors’ reactions incorporating a simplified economic framework instead of a more elaborate model of the market (Glaeser, 2013; Greenwood and Hanson, 2015). A first-order effect of the proposed framework is that, in the presence of extrapolative expectations, vessel prices become more sensitive to the prevailing cash flow. As a result, the extrapolative-model generated price deviates from the asset’s fundamental value whenever the corresponding cash flow variable deviates from its steady state. This fact implies an immediate over – or under – valuation of the vessel which, in turn, generates “excess” price volatility.

While there can be alternative – “rational” – explanations for the observed patterns in either trading activity (such as limits to arbitrage) or vessel price behaviour (such as time-varying risk preferences), the proposed model has the advantage of simultaneously explaining in a sufficient manner numerous stylised facts. For instance, while a homogeneous-agent setting with extrapolative expectations could capture the observed price behaviour, it would not be sufficient to justify the second-hand market transactions. Therefore, trading activity in our framework is the consequence of heterogeneous beliefs and, in turn, valuations of the asset by market participants.

Furthermore, in line with Cochrane (2011), most of the potential alternative “rational” explanations incorporate “exotic preferences” rendering them almost indistinguishable from behavioural ones. Equivalently, their predictions stem from auxiliary assumptions and not from the rationality assumption *per se* (Arrow, 1986). The fact, however, that almost any biased beliefs model can be re-expressed as a rational expectations’ one with time-varying preferences/discount factors (Cochrane, 2011) does not validate the latter approach or invalidate the former one. Specifically, as Lof (2015) argues, biased beliefs models are very appealing when modelling boom-bust cycles as the ones documented in the shipping industry (Greenwood and Hanson, 2015). More importantly, as we illustrate in the following, the economic interpretation of the model and the respective results are plausible and in line with the nature of the shipping industry.

To the best of our knowledge, this is the first time in the shipping literature that a structural economic model incorporates the coexistence of heterogeneous beliefs agents to explain the joint behaviour of observed vessel prices, net earnings, and second-hand vessel transactions. Regarding the existing shipping literature, Beenstock (1985), Beenstock and Vergottis (1989), and Kalouptside (2014) construct and estimate rational expectations general equilibrium models in a homogeneous agents' setting which, however, does not allow for the explanation of the second-hand market activity. Greenwood and Hanson (2015) develop a microeconomic model in which agents extrapolate current demand conditions while simultaneously neglect their competitors' supply responses. The behavioural mechanism proposed here is similar to that of Greenwood and Hanson (2015), however, to be able to capture vessel trading activity, we focus on the market for second-hand vessels instead of the new-building and demolition ones. Furthermore, the introduction of two types of agents allows us to simultaneously capture the observed behaviour of prices and the relation between net earnings and second-hand activity in the market. Finally, recall that while Chapter 2 of this thesis analyses the behaviour of vessel valuation ratios, it does not explicitly model the underlying mechanism of the behaviour of asset prices per se or the relation between prices, earnings, and second-hand activity.

This chapter looks at the main features of heterogeneous agents' models but also introduces important modifications which are required to capture stylised features of the shipping markets. Recent articles – mainly in equity but also in commodity markets (Ellen and Zwinkels, 2010) – have attempted to explain empirical asset pricing findings using heterogeneous beliefs models in which a fraction of the population forms biased expectations about future returns. Barberis *et al* (2015a) develop an extrapolative capital asset pricing model (X-CAPM) that explains the volatility of the aggregate stock market. Furthermore, Barberis *et al* (2015b) incorporate a heterogeneous beliefs extrapolative model of returns in order to analyse the formation of asset bubbles in equity markets. Some key features of their model are very closely related to the one presented in this chapter.

However, in contrast to these papers and the bulk of the behavioural equity markets literature, in our model there is cash flow and not return extrapolation. The motivation for this is based on actual market practice and the economics of the industry. Namely, shipping industry participants characterise market conditions based on the prevailing – and forecasts of future – net earnings and not on realised returns. Thus, it is much more plausible for investors to form biased expectations regarding fundamentals rather than returns. In contrast, in equity markets, recent evidence from surveys (Greenwood and Shleifer, 2014) suggests that many investors extrapolate stock market returns. In addition to this argument, to be able to capture simultaneously and in a sufficient manner some key stylised features of the shipping industry – among which, the fact that the bulk of earnings

yield volatility is attributed to net earnings and not returns – there must be cash flow and not return extrapolation in the market. As Barberis *et al* (2015a) explain, the opposite is true in equity markets since models that incorporate cash flow extrapolation (Choi and Mertens, 2013; Hirshleifer and Yu, 2013; Alti and Tetlock, 2014) struggle to match key empirical findings like the survey evidence.

Apart from this significant difference, we further depart from the frameworks incorporated in Barberis *et al* (2015a; 2015b) by examining an asset with finite life that is significantly affected by economic depreciation due to wear and tear. This fact provides different challenges in the economic modelling of the market compared to the case of an infinitely lived financial asset. Finally, our model is flexible enough to allow market agents to hold distorted beliefs at different degrees. This feature enables us to simultaneously capture, more sufficiently, a number of stylised features of the market among which, asset undervaluation during market troughs, the positive correlation between market conditions and trading activity, and the relatively low liquidity of the shipping markets. In addition, it renders our framework simple enough so that it can be easily extended and applied to other real asset markets characterised by alternative forms of – biased – investor behaviour.

Our simulation results suggest that even a small fraction of extrapolators – that is, less than 10% of the population – can reproduce the observed findings. This result is of interest since the model of Barberis *et al* (2015a) suggests that, to match the “excess volatility” in the U.S. equity markets, extrapolators must constitute 50% of the population. While the results in the two models are not directly comparable, we can draw two interesting conclusions. First, from a mathematical perspective, the cash flow extrapolative expectations mechanism incorporated in our model is very direct as even modest one-period cash flow shocks are immediately translated into significant vessel price fluctuations. In contrast, Barberis *et al* assume a much slower extrapolative expectations process regarding the price – return – variable. Second, from an economic perspective, due to the fundamental differences in the structures of the shipping and equity markets, it is much more plausible for extrapolators to be a substantially larger fraction of the latter market compared to the former one.

The remainder of this chapter is organised as follows. Section 3.II introduces the environment of our economy and the solution of the theoretical model. Section 3.III presents the dataset employed along with the empirical estimation of the model. It also provides an economic interpretation of the results. Section 3.IV examines several alternative hypotheses regarding the investor population composition. Section 3.V concludes.

3.II. Environment and Model Solution

Consider a discrete-time environment where the passage of time is denoted by t . The economy consists of two asset classes: the first one is risk-free while the second one is risky. The risk-free asset can be thought of as an infinitely lived financial instrument in perfectly elastic supply, earning an exogenously determined constant rate of return equal to R_f . The risky asset class consists of otherwise identical assets (i.e., vessels) which are further categorised by their age. All age classes have fixed per capita supply over time equal to Q .⁵³ In what follows, we restrict our attention to the modelling of the market for 5-year old vessels. However, the same principles apply for the valuation of the other age classes. Following market practice, we assume that a newly built vessel has an economic life of 25 years after which is scrapped and exits the economy. Accordingly, setting the time-step of the model, Δt , equal to one year implies that a 5-year old asset has $T = 20$ periods of remaining economic activity.

As analysed in Chapter 2 of this thesis, an inherent characteristic of the shipping industry is that next period's net earnings are known in advance. Accordingly, assuming no default on the part of the charterer, the ship owner at time t knows precisely his net earnings for the period $t \rightarrow t + 1$, defined as Π_t .⁵⁴ This is equivalent to saying that the asset is trading "cum dividend" in equity markets. Therefore, the owner at time t is entitled to an exogenously determined stream of annual net earnings, $\{\Pi_n\}_t^{t+T}$. In the context of our model, net earnings are the sole state variable. In line with the data (see also Chapter 2) and the existing literature (Greenwood and Hanson, 2015), annual net earnings are assumed to be following a mean-reverting process in discrete time:

$$\Pi_{t+1} = (1 - \rho_0)\bar{\Pi} + \rho_0\Pi_t + \varepsilon_{t+1}, \quad (3.1)$$

where $\bar{\Pi}$ is the long-term mean, $\rho_0 \in [0,1)$, and $\varepsilon_{t+1} \sim N(0, \sigma_\varepsilon^2)$, *i. i. d.* over time. Importantly, in contrast to $\bar{\Pi}$, parameters ρ_0 and σ_ε^2 are not public information.

The economy consists of two investor types, i : "conservatives" and "extrapolators", denoted by c and e , respectively. We normalise the investor population related to each asset age-class to a unit measure and further assume that the fractions of conservatives, μ^c , and extrapolators, μ^e , are fixed, both across all age classes and through each specific asset's life. In what follows, we set $\mu^c = \mu$;

⁵³ This is justified by the fact that we are interested in the modelling of a real asset with economic depreciation. Hence, the supply of the age-specific asset cannot increase over time. Furthermore, scrapping very rarely occurs before the 20th year of a vessel's life; thus, we assume that supply cannot be reduced either. Since accidents, losses, and conversions constitute an insignificant proportion of the fleet, they are not considered here. Finally, while supply may differ across age classes, this feature does not affect the predictions of our model.

⁵⁴ In practice, ship owners and charterers agree upon the time-charter rate of the vessel before the corresponding leasing period begins. Accordingly, the agreed rates are typically received every 15 days – sometimes also in advance. Ships are chartered via an extensive network of competitive brokers using established contractual agreements in the charter-party contract which provide some guarantee that the owner will receive the full rate agreed. Thus, one can plausibly assume that next period's net earnings are \mathcal{F}_t -measurable (see also Chapter 2).

hence, $\mu^e = 1 - \mu$. The difference between the two types lies in the alternative ways in which they form expectations about future cash flows. Specifically, compared to extrapolators, conservatives' perception is significantly closer – in principle, it might even be identical – to (3.1). We assume that in agent i 's mind, net earnings related to the valuation of the 5-year old vessel evolve according to

$$\Pi_{t+1} = (1 - \rho_i)\bar{\Pi} + \rho_i\Pi_t + \varepsilon_{t+1}^i, \quad (3.2)$$

in which $\rho_0 \leq \rho_c < \rho_e < 1$ and $\varepsilon_{t+1}^i \sim N(0, \vartheta_5^i \sigma_\varepsilon^2)$, *i. i. d.* over time, where $0 < \vartheta_5^e < \vartheta_5^c < 1$. The strictly positive parameter ϑ_5^i adjusts the – true – variance of the cash flow shock according to agent's

i perspective while the subscript denotes the current age-class of the vessel being valued.

The conservative agent parameters – μ , ρ_c , and ϑ_5^c – characterise completely the information structure of our model. When $\mu = 1$, $\rho_c = \rho_0$, and $\vartheta_5^c = 1$, all agents have perfect information about the economy. This case is defined as the benchmark “rational” economy of our model and we term this agent type as fundamentalist, f ; hence, $\rho_f = \rho_0$ and $\vartheta_5^f = 1$. When $\mu = 1$, $\rho_c \neq \rho_0$, and $\vartheta_5^c \neq 1$ or, $\mu = 0$, all agents have imperfect information about the economy. However, in all cases above, there is no information asymmetry among agents and, in turn, no trading activity in the market. Finally, when $\mu \in (0,1)$, information is both imperfect and asymmetric (Wang, 1993) and, as a result, there is trading activity in the economy.

The timeline of the model is as follows. At each point t , Π_t is realised and observed by all market participants. Furthermore, the 25-year old age class is scrapped and replaced by newly built vessels. Accordingly, both agent types determine their time t demands for each age class asset with the aim of maximizing a constant absolute risk-aversion (CARA) utility function, defined over next period's wealth. For the 5-year old vessel, this corresponds to

$$\max_{N_{5,t}^i} E_t^i \left[-e^{-\alpha^i w_{t+1}^i} \right], \quad (3.3)$$

where α^i and $N_{5,t}^i$ are investor i 's coefficient of absolute risk-aversion and time t per-capita demand for the 5-year old vessel, respectively. Agent i 's next period's wealth, w_{t+1}^i , is given by

$$w_{t+1}^i = (w_t^i - N_{5,t}^i P_{5,t})(1 + R_f) + N_{5,t}^i (\Pi_t + P_{6,t+1}), \quad (3.4)$$

in which $P_{5,t}$ and $P_{6,t+1}$ are the prices of the 5- and 6-year old vessel at t and $t + 1$, respectively.⁵⁵

⁵⁵ In principle, each agent could invest a fraction of his wealth in every age-class of the risky asset. However, to obtain closed-form solutions for the demand functions, we assume that – at each t – a new unit mass of investors solely interested in 5-year old vessels enters the industry. In turn, at $t + 1$ this investor population

In what follows, we normalise the rate of return of the risk-free asset to zero (Wang, 1993). Therefore, investor i 's objective becomes

$$\max_{N_{5,t}^i} E_t^i \left[-e^{-a^i(w_t^i + N_{5,t}^i(\Pi_t + P_{6,t+1} - P_{5,t}))} \right]. \quad (3.5)$$

Accordingly, the time t price of the 5-year old vessel is endogenously determined through the market-clearing condition

$$\mu N_{5,t}^c + (1 - \mu) N_{5,t}^e = Q. \quad (3.6)$$

Following the same principles, the time t per-capita demand of agent i for the 6-year old vessel, $N_{6,t}^i$, and the corresponding 6-year old vessel price, $P_{6,t}$, are determined (Appendix 3.A). Finally, trading activity corresponding to time t – that is, to period $t - 1 \rightarrow t$ – takes place in the market. In shipping, this activity refers to the sale and purchase market for second-hand vessels. Since this is a discrete-time model, we impose the assumption that trading occurs instantaneously at each point t (Barberis *et al*, 2015b). Note that because vessels are real assets with limited economic life, their values are affected by economic depreciation due to wear and tear.⁵⁶ Thus, a 5-year old vessel acquired at time $t - 1$ will be a 6-year old one – when sold – at time t . Accordingly, we define as trading activity the agent-specific change in demand for the – same – risky asset between points $t - 1$ and t , multiplied by the respective population fraction:

$$V_{t-1 \rightarrow t} \equiv V_t = \mu^i |N_{6,t}^i - N_{5,t-1}^i|, \quad (3.7)$$

Figure 3.1 summarises the timeline of the model.

Consistent with the nature of the industry, we impose short-sale constraints (Barberis *et al*, 2015b). Appendix 3.A shows that the time t per-capita demand of agent i for the 5-year old vessel is

$$N_{5,t}^i = \max \left\{ \frac{\frac{1 - \rho_i^{21}}{1 - \rho_i} (\Pi_t - \bar{\Pi}) + 21\bar{\Pi} - X_5^i \sigma_\varepsilon^2 Q - P_{5,t}}{Y_5^i \sigma_\varepsilon^2}, 0 \right\}, \quad (3.8a)$$

with


will be solely interested in the 6-year old class while a new unit mass related to the 5-year old class will enter the market.

⁵⁶ Hence, at any t , a 6-year old vessel is less valuable than an identical 5-year one.

$$\begin{cases} X_5^i = \left[\frac{20}{(1 - \rho_i)^2} - \frac{(1 - \rho_i^{20})(1 + 2\rho_i - \rho_i^{20})}{(1 + \rho_i)(1 - \rho_i)^3} \right] \alpha^i \vartheta_5^i \\ Y_5^i = \left(\frac{1 - \rho_i^{20}}{1 - \rho_i} \right)^2 \alpha^i \vartheta_5^i \end{cases}, \quad (3.8b)$$

where both X_5^i and Y_5^i are strictly positive constants. Equation 3.8a along with the market-clearing condition 3.6 determine the equilibrium 5-year old vessel price at each t . From an economic perspective, the fraction in (3.8a) reflects the expected one-period net income for investor i scaled by the product of investor's risk aversion times the risk he is bearing according to his perception. Note that, to derive the agent-specific demand functions, we have assumed that agent i makes the simplifying assumption that his counterpart, $-i$, will hold his fraction of the risky asset constant at $\mu^{-i}Q$, irrespective of the corresponding future net earnings variable. Equivalently, agent i 's $\mu^{-i}Q$ is constant, which is the "naïve" effect of his competitor's valuation of the asset (Barberis et al, 2015b).

Figure 3.1: Timeline of the Model.



⁵⁷ As analysed in Appendix 3.A, in principle, investors could understand that their beliefs about either the cash flow process and/or their competitors' strategy are inaccurate (Barberis et al, 2015a). We do not incorporate an explicit learning process, however, since this would gradually eliminate both the "excess price volatility" feature and – the observed patterns related to – second-hand activity in the market. Accordingly, we adopt a rather indirect learning mechanism. Specifically, we assume that agents become more "suspicious" as the specific asset's age grows and they indirectly respond by increasing the perceived risk associated with their investment.

Since extrapolators have “more incorrect” beliefs about the net earnings process, it might be the case that in the long-run their wealth will be significantly reduced, if not depleted.⁵⁸ Notice though that the use of exponential utility implies that the demand function is independent of the respective wealth level. This, in turn, allows us to abstract from the “survival on prices” effect (Barberis *et al*, 2015a) and focus solely on the pricing and trading implications of the heterogeneous agents’ economy. In reality, even if extrapolators are not able to invest due to limited wealth, it is plausible to assume that they will be immediately replaced by a new fraction of investors with exactly the same characteristics. In shipping, this cohort could correspond to diversified investors with substantial cash availability – such as private equity firms – but little or no prior experience of the industry.

Proposition: Equilibrium price for 5-year old vessels. In the environment presented above, a market-clearing – or equilibrium – price for the 5-year old vessel, $P_{5,t}^*$, always exists. The equilibrium price of the vessel depends on the prevailing market conditions. We denote the net earnings thresholds at which extrapolators and conservatives related to the 5-year old vessel class exit the market by Π_5^e and Π_5^c , respectively.

First, when

$$\Pi_5^e = \bar{\Pi} + \frac{(X_5^e - X_5^c - \frac{Y_5^c}{\mu})\sigma_\varepsilon^2 Q}{\frac{1 - \rho_e^{21}}{1 - \rho_e} - \frac{1 - \rho_c^{21}}{1 - \rho_c}} < \Pi_t < \bar{\Pi} + \frac{(X_5^e - X_5^c + \frac{Y_5^e}{1 - \mu})\sigma_\varepsilon^2 Q}{\frac{1 - \rho_e^{21}}{1 - \rho_e} - \frac{1 - \rho_c^{21}}{1 - \rho_c}} = \Pi_5^c, \quad (3.9a)$$

both agents are present in the market, and the market-clearing price, $P_{5,t}^{*c+e}$, is equal to

$$P_{5,t}^{*c+e} = 21\bar{\Pi} + \frac{\mu Y_5^e \frac{1 - \rho_c^{21}}{1 - \rho_c} + (1 - \mu) Y_5^c \frac{1 - \rho_e^{21}}{1 - \rho_e}}{\mu Y_5^e + (1 - \mu) Y_5^c} (\Pi_t - \bar{\Pi}) - \frac{\mu Y_5^e X_5^c + (1 - \mu) Y_5^c X_5^e + Y_5^c Y_5^e}{\mu Y_5^e + (1 - \mu) Y_5^c} \sigma_\varepsilon^2 Q. \quad (3.9b)$$

Second, in the case where $\Pi_t \leq \Pi_5^e$, extrapolators exit the market and the market-clearing price, $P_{5,t}^{*c}$, is given by

⁵⁸ Appendix 3.D illustrates that – for the basic parameterisation of the model – while extrapolators’ one-period changes in wealth are significantly more volatile than conservatives’ ones, both types of agent realise approximately the same mean change. Furthermore, for reasons that become apparent in Section 3.III, extrapolators have a positively skewed distribution of wealth changes in contrast to conservatives who have a normal one. Therefore, there is no formal indication that extrapolators “suffer” (on average) by limitations of wealth more than conservatives do. This auxiliary result is similar to the one in Barberis *et al* (2015a).

$$P_{5,t}^{*c} = 21\bar{\Pi} + \frac{1 - \rho_c^{21}}{1 - \rho_c} (\Pi_t - \bar{\Pi}) - \left[X_5^c + \frac{Y_5^c}{\mu} \right] \sigma_\varepsilon^2 Q. \quad (3.10)$$

Third, in the scenario where $\Pi_5^c \leq \Pi_t$, conservatives exit the market and the equilibrium price, $P_{5,t}^{*e}$, is given by

$$P_{5,t}^{*e} = 21\bar{\Pi} + \frac{1 - \rho_e^{21}}{1 - \rho_e} (\Pi_t - \bar{\Pi}) - \left[X_5^e + \frac{Y_5^e}{1 - \mu} \right] \sigma_\varepsilon^2 Q. \quad (3.11)$$

■

As the first term of equations 3.9b, 3.10, and 3.11 indicates, the price of the vessel depends on the long-term mean of the cash flow variable multiplied by the total number of payments to be received until the end of the asset's economic life. The second term corresponds to the effect of the product of the perceived persistence of the net earnings variable times its current deviation from the long-term mean. Essentially, this term measures the main bulk of over (under) valuation in the price of the risky asset.⁵⁹ Finally, the last term corresponds to the aggregate discounting by which future cash flows are reduced in order for investors to be compensated for the risk they bear (Wang, 1993).⁶⁰

Benchmark rational economy. It is also useful to examine the benchmark rational economy, denoted by f , in which the market consists solely of agents who know the actual stochastic process that governs the evolution of net earnings. The equilibrium price of the 5-year old vessel in this case is

$$P_{5,t}^f = \frac{1 - \rho_0^{21}}{1 - \rho_0} (\Pi_t - \bar{\Pi}) + 21\bar{\Pi} - [X_5^f + Y_5^f] \sigma_\varepsilon^2 Q. \quad (3.12)$$

As equations 3.8a and 3.8b indicate, fundamentalists' perception of the risk they are bearing is given by the product $\left(\frac{1 - \rho_0^{20}}{1 - \rho_0} \right)^2 \sigma_\varepsilon^2$. In this benchmark case, this perception is correct. In the presence of extrapolators, though, it is just an approximation since future asset prices will also depend on extrapolators' future demand responses and not just on the riskiness of cash flows.

Moreover, the unconditional volatility of the fundamental price is given by

⁵⁹ Note that due to the assumed form of extrapolation and the structure of our economy, a substantial over (under) valuation of the asset can occur – and, accordingly, disappear or even revert – within one period; that is, a single cash flow shock suffices. In contrast, in the model of Barberis *et al* (2015b), for an overvaluation to occur (referred as a “bubble”), we need to have a series of positive cash flow shocks and, in turn, a 3-stage displacement process (in line with Kindleberger, 1978). Furthermore, in contrast to Barberis *et al* (2015b), our model can also account for severe undervaluation of the risky asset.

⁶⁰ Extending the proof of the Proposition, it is straightforward to show that a vessel age-specific market-clearing price always exists (Appendices 3.A and 3.B).

$$\sigma(P_{5,t}^f) = \frac{1 - \rho_0^{21}}{1 - \rho_0} \sigma(\Pi_t). \quad (3.13)$$

Finally, taking unconditional expectations on both sides of equation 3.12 and setting the unconditional mean of the net earnings variable equal to its long-term mean, $\bar{\Pi}$, yields

$$E[P_{5,t}^f] = 21\bar{\Pi} - [X_5^f + Y_5^f] \sigma_\varepsilon^2 Q. \quad (3.14)$$

Corollary 1: Steady state equilibrium. We define the “steady state” of our economy as the one in which the net earnings variable is equal to its long-term mean, $\bar{\Pi}$. As equation 3.1 indicates, the economy reaches this state after a sequence of zero cash flow shocks. In the steady state, the price of the risky asset is equal to its respective fundamental value. Furthermore, both types of agent are present in the market and each type holds the risky asset in analogy to his fraction of the total population. Accordingly, the “steady state” equilibrium price of the 5-year old vessel, \bar{P}_5^* , is given by

$$\bar{P}_5^* = 21\bar{\Pi} - [X_5^i + Y_5^i] \sigma_\varepsilon^2 Q, \quad (3.15a)$$

under the restriction

$$X_5^c + Y_5^c = X_5^e + Y_5^e = X_5^f + Y_5^f = \frac{\mu Y_5^e X_5^c + (1 - \mu) Y_5^c X_5^e + Y_5^c Y_5^e}{\mu Y_5^e + (1 - \mu) Y_5^c}. \quad (3.15b)$$

In a similar manner, the “steady state” equilibrium price of the 6-year old vessel is

$$\bar{P}_6^* = 20\bar{\Pi} - [X_6^i + Y_6^i] \sigma_\varepsilon^2 Q, \quad (3.16a)$$

under the restriction

$$X_6^c + Y_6^c = X_6^e + Y_6^e = X_6^f + Y_6^f. \quad (3.16b)$$

Therefore, if in two consecutive periods the net earnings variable is equal to its long-term mean, the change in the price of the asset is

$$\bar{P}_6^* - \bar{P}_5^* = -\bar{\Pi} - [X_6^i + Y_6^i - (X_5^i + Y_5^i)] \sigma_\varepsilon^2 Q. \quad (3.17)$$

The right-hand side of (3.17) is negative and corresponds to the one-year economic depreciation in the value of the vessel. Finally, in this scenario, there is no activity in the second-hand market, since the change in share demand of each agent is equal to zero. ■

Corollary 2: Deviation from the fundamental value. Whenever the value of the net earnings variable deviates from its long-term mean, the model-generated price of the 5-year old vessel deviates from its fundamental value. In the following, we denote by $D_{5,t}$ the degree of deviation;

namely, a positive (negative) value of $D_{5,t}$ corresponds to over (under) valuation of the asset relative to its fundamental analogue, $P_{5,t}^f$. Note that, in the following, we define as prosperous (adverse) market conditions the case where the net earnings variable is above (below) its steady state value, $\bar{\Pi}$.

First, in the case where both agents are present in the market, the deviation, D_t^{c+e} , is given by

$$D_{5,t}^{c+e} = \frac{\mu \left(\frac{1 - \rho_c^{21}}{1 - \rho_c} - \frac{1 - \rho_0^{21}}{1 - \rho_0} \right) Y_5^e + (1 - \mu) \left(\frac{1 - \rho_e^{21}}{1 - \rho_e} - \frac{1 - \rho_0^{21}}{1 - \rho_0} \right) Y_5^c}{\mu Y_5^e + (1 - \mu) Y_5^c} (\Pi_t - \bar{\Pi}). \quad (3.18)$$

Since the fraction is always positive, the sign of price deviation solely depends on the sign of net earnings deviation. Thus, during prosperous market conditions the asset is overpriced and vice versa.

Second, when only conservatives exist in the market the deviation, $D_{5,t}^c$, is estimated through

$$D_{5,t}^c = \left(\frac{1 - \rho_c^{21}}{1 - \rho_c} - \frac{1 - \rho_0^{21}}{1 - \rho_0} \right) (\Pi_t - \bar{\Pi}) - \left[X_5^c + \frac{Y_5^c}{\mu} - X_5^f - Y_5^f \right] \sigma_\varepsilon^2 Q, \quad (3.19)$$

which is always negative. Thus, during adverse market conditions the vessel is undervalued.

Third, when only extrapolators are present, the discrepancy, $D_{5,t}^e$, is

$$D_{5,t}^e = \left(\frac{1 - \rho_e^{21}}{1 - \rho_e} - \frac{1 - \rho_0^{21}}{1 - \rho_0} \right) (\Pi_t - \bar{\Pi}) - \frac{\mu Y_5^e}{1 - \mu} \sigma_\varepsilon^2 Q, \quad (3.20)$$

which is always positive⁶¹ and for significantly prosperous market conditions the degree of overvaluation becomes severe. ■

Corollary 3: Sensitivity of exit points to the fraction of conservatives. As the expressions for Π_5^c and Π_5^e suggest (i.e., equation 3.9a), the agent-specific exit points differ due to the quantities $-\frac{Y_5^c}{\mu}$ and $\frac{Y_5^e}{1-\mu}$. The implication of this fact is that whenever $Y_5^c/\mu \neq Y_5^e/(1-\mu)$ there is no symmetry around $\bar{\Pi}$ between the two points. As a result, the positive and negative shock cases are not mirror images of each other. Taking the first partial derivative of the extrapolators' 5-year exit point with respect to the fraction of conservatives yields

$$\frac{\partial \Pi_5^e}{\partial \mu} = \frac{1}{\mu^2} \cdot \frac{Y_5^c \sigma_\varepsilon^2 Q}{\frac{1 - \rho_e^{21}}{1 - \rho_e} - \frac{1 - \rho_c^{21}}{1 - \rho_c}} \quad (3.21)$$

⁶¹ It is straightforward to verify this by plugging in (3.20) the expression for Π_5^c from 3.9a.

which is strictly positive. Hence, the higher the fraction of conservatives, the more prone extrapolators are to exit from the market during adverse conditions. Similarly, the first partial derivative of conservatives' exit point with respect to their relative fraction is equal to

$$\frac{\partial \Pi_5^c}{\partial \mu} = \frac{1}{(1-\mu)^2} \cdot \frac{Y_5^e \sigma_\varepsilon^2 Q}{\frac{1-\rho_e^{21}}{1-\rho_e} - \frac{1-\rho_c^{21}}{1-\rho_c}} \quad (3.22)$$

which is strictly positive. Thus, the higher the fraction of conservatives, the less prone they are to exit from the market during prosperous conditions. The same principles apply for the 6-year old vessel valuation. Hence, the asymmetry increases as μ deviates from the midpoint 0.5. ■

Trading volume and net earnings. Appendix 3.B shows that trading activity is quantified through

$$V_t = \mu^i \left| \max \left\{ \frac{\frac{1-\rho_i^{20}}{1-\rho_i} (\Pi_t - \bar{\Pi}) + 20\bar{\Pi} - X_6^i \sigma_\varepsilon^2 Q - P_{6,t}}{Y_6^i \sigma_\varepsilon^2}, 0 \right\} - \max \left\{ \frac{\frac{1-\rho_i^{21}}{1-\rho_i} (\Pi_{t-1} - \bar{\Pi}) + 21\bar{\Pi} - X_5^i \sigma_\varepsilon^2 Q - P_{5,t-1}}{Y_5^i \sigma_\varepsilon^2}, 0 \right\} \right| \quad (3.23)$$

Due to the short-sale constraints, the agent-specific demand functions are not strictly monotonic with respect to the net earnings variable in the entire Π_t domain. As a result, trading activity depends on the realisation of the net earnings variable during the two corresponding consecutive dates, $t-1$ and t . In Appendix 3.B, we examine all possible scenarios. Note that in the absence of constraints, absolute net earnings changes would be almost perfectly correlated with trading activity. Due to the existence of short-sale constraints, however, the correlation between the two variables is much lower.⁶²

Moreover, Corollary 3 demonstrates that both exit points increase (decrease) with the fraction of conservatives (extrapolators) and the perceived persistence on behalf of extrapolators (conservatives). Hence, the higher the values of the exit points, the more the two types of agent coexist during prosperous market conditions and the less they interact during adverse ones. Thus, a high value of μ , along with a significant spread between ρ_c and ρ_e , will result in both positive correlation between current net earnings and trading activity and less than perfect correlation

⁶² In order to illustrate this point, let's define trading activity as in the equity markets literature where there is no depreciation in the value of the asset due to wear and tear; namely, we set $N_{6,t}^i = N_{5,t}^i$. Equivalently, we substitute A_5^c for A_6^c in (3.B23). Thus, in the absence of short-sale constraints, trading activity, V_t , would always be equal to $\mu |A_5^c| |\Pi_t - \Pi_{t-1}|$ and, in turn, $\text{corr}(|\Pi_t - \Pi_{t-1}|, V_t) = 1$.

between absolute net earnings changes and trading activity. These theoretical predictions are analysed in the next section.

3.III. Empirical Estimation of the Model in the Dry Bulk Shipping Industry

In this section, the dataset employed and the construction of the variables of interest are discussed. Accordingly, we evaluate empirically the predictions of our model by performing a large number of simulations. We also provide a deeper intuition of the results by implementing impulse response and sensitivity analyses. Finally, we discuss our findings from an economic and practical perspective.

3.III.A. Data on Net Earnings, Prices, and Trading Activity

The dataset employed consists of annual observations on second-hand vessel prices, 1-year time-charter rates, fleet capacity, and second-hand vessel transactions related to the Handysize dry bulk sector. Our main source of shipping data is Clarksons Shipping Intelligence Network. In addition, data for the U.S. Consumer Price Index (CPI) are obtained from Thomson Reuters Datastream Professional. Note that while the empirical estimation focuses on the Handysize sector, our results have been tested to the remaining dry bulk sectors and are both qualitatively and quantitatively robust; thus, our conclusions are representative of the entire dry bulk industry.

In line with Chapter 2 and the existing literature (Greenwood and Hanson, 2015), we assume that vessels operate in consecutive one-year time-charter contracts. In this type of arrangement, only the operating and maintenance costs are borne by the ship owner. Since these costs increase with inflation (Stafford *et al*, 2002), we use the December 2014 nominal figures as the benchmark real values – after discussions with industry participants, we arrived at a figure of \$5,500 (see also Chapter 2). In addition, we assume that vessels spend 10 days per annum off-hire for maintenance and repairs (Stopford, 2009). During this period, ship owners do not receive the corresponding time-charter rates but bear the operating and maintenance costs. We also consider the commission that the brokering house receives for bringing the ship owner and the charterer into an agreement; this is equal to 2.5% of the daily time-charter rate. Finally, as it is common in the literature, interest and tax expenses are ignored from the analysis. Thus, similar to Chapter 2, the annual net earnings variable is given by

$$\Pi_t \equiv \Pi_{t \rightarrow t+1} = 355 \cdot 0.975 \cdot TC_{t \rightarrow t+1} - 365 \cdot OPEX_{t \rightarrow t+1}, \quad (3.24)$$

where $TC_{t \rightarrow t+1}$ and $OPEX_{t \rightarrow t+1}$ refer to the corresponding daily time-charter rates and total daily operating and maintenance costs, respectively. Moreover, the one-year horizon log return is given by

$$r_{t+1} = \ln \left(\frac{\Pi_t + P_{6,t+1}}{P_{5,t}} \right), \quad (3.25)$$

where $P_{5,t}$ and $P_{6,t+1}$ refer to the current and next period's price of the 5 and 6-year old vessel, respectively. As analysed in Chapter 2, since generic 6-year old vessel prices are not readily available, we set $P_{6,t} = 0.95P_{5,t}$ to estimate the actual one-period returns.⁶³

In order to construct the annual trading activity variable, V_t , we scale the total number of second-hand transactions taking place within the period of interest by the fleet size in the beginning of the respective period.⁶⁴ Table 3.1 summarises descriptive statistics related to annual net earnings, 5-year old vessel prices, and annual trading activity, from 1989 to 2014. Panels A and B of Figure 3.2 illustrate the relation between trading activity and net earnings and trading activity and absolute one-year changes in net earnings, respectively. Evidently, trading activity is significantly positively correlated with both variables. Namely, the correlation coefficients are equal to 0.53 and 0.65, respectively. Note that, as analysed in Section I, these two key empirical findings are the main motivation for this model.

Table 3.1: Descriptive statistics for vessel prices, net earnings, and trading activity.

Variable	L	Mean	SD	Median	Max	Min	ρ_0
Π (\$m)	26	3.10	2.39	2.42	9.96	0.91	0.58
P (\$m)	26	22.86	7.65	22.32	50.23	13.43	0.49
V	20	0.058	0.020	0.054	0.099	0.031	0.11

Notes: This table presents the number of observations (L), mean, standard deviation (SD), median, maximum, minimum, and 1-year autocorrelation coefficient, (ρ_0), for net earnings, Π , 5-year old vessel prices, P , and trading activity, V . Shipping data are provided by the Clarksons Shipping Intelligence Network. The sample is annual, covering the period from 1989 to 2014, apart from trading activity which becomes available in 1995. Net earnings and prices are expressed in December 2014 million dollars through the U.S. Consumer Price Index (CPI), obtained from Thomson Reuters Datastream Professional. Since the 5-year old vessel price time series refers to a 32,000-dead weight tonnage (dwt) carrier while the time-charter rate series to a 30,000 one, we multiply the initial rate series by 32/30. Trading activity is scaled by the respective size of the fleet.

⁶³ We have estimated the average ratio of 10- to 5-year old vessel prices to be approximately equal to 0.75. Accordingly, adopting a straight-line depreciation scheme implies $P_{6,t} = 0.95P_{5,t}$.

⁶⁴ Since we only have data regarding the total number of transactions realised during each corresponding period, we assume that this scaled figure is representative of each vessel-age interval.


3.III.B. Simulation Methodology and Results

In this subsection, we evaluate empirically the predictions of our model for several combinations of the three main parameters of interest, $\{\mu, \rho_c, \rho_e\}$, using numerical simulations. Accordingly, we compare the model-generated moments to the actual ones. Using equation 3.1, we generate 10,000 sample paths for our economy where each path corresponds to 100 periods.⁶⁵ Finally, we estimate the average of each statistic under consideration across all valid paths (Barberis *et al*, 2015a).


To conduct the simulations, we calibrate two sets of model parameters. The first set contains the asset-level parameters, $\{\bar{\Pi}, \rho_0, \sigma_\varepsilon^2, Q, T, R^f\}$, and remains the same irrespective of the population composition and characteristics. We set $\bar{\Pi}$ equal to the long-term mean of the net earnings variable. The coefficient of persistence, ρ_0 , is approximated through the actual 1-year autocorrelation coefficient of the variable. We set the standard deviation of the error term, σ_ε^2 , equal to 1 to reduce the number of discarded paths but at the same time ensure a sufficient degree of net earnings volatility.⁶⁶ We set the remaining economic life of the 5-year old vessel, T , equal to 20. Finally, we

⁶⁵ In the simulation, we discard the paths that lead to negative values either for net earnings or vessel prices. We impose this restriction to be able to perform the predictive regressions which use log quantities as variables. Even if we do not discard these paths, the remaining results remain essentially the same.

⁶⁶ This value per se has no direct qualitative impact on the estimation.



Panel A: Relation between Trading Activity and Net Earnings.



Panel B: Relation between Trading Activity and Absolute Changes in Net Earnings. Figure 3.2: Net Earnings and Trading Activity.

Panel A depicts the relation between annual trading activity and annual net earnings. Panel B depicts the relation between annual trading activity and absolute changes in annual net earnings. The sample runs from 1995 to 2014. Annual trading activity is expressed as a percentage of the fleet in the beginning of the corresponding period. Prices and net earnings are expressed in December 2014 million dollars.

normalise the fixed per capita supply, Q , to one and the risk-free rate of return, R^f , to zero.

The second set includes the agent-specific parameters $\mu, \rho_i, \vartheta_5^i, \vartheta_6^i, \alpha^i$, and w_0^i for $i \in \{f, c, e\}$. Regarding the parameter μ , we choose values within the interval $[0,1]$. While fundamentalists' characteristics are fixed by definition, the ones related to both conservatives and extrapolators are recalibrated each time depending on the scenario choice. Recall that, by assuming an exponential utility, our results are independent from the initial level of wealth; hence, we do not have to assign a value to w_0^i . Since fundamentalists form expectations about future net earnings based on the true stochastic process, ρ_f, ϑ_5^f , and ϑ_6^f are assigned the values of 0.58, 1, and 1, respectively. The last parameter is the coefficient of absolute risk aversion, α^f . Appendix 3.B shows that α^f and the steady

Table 3.2: Parameter values.

Parameter	Assigned Value
\overline{P}_5^*	22.86
$\overline{\Pi}$	3.1
ρ_0	0.58
σ_ε^2	1
Q	1
T	20
R_f	0
μ	{0.1,0.5,0.95}
α^f	0.42
ρ_f	0.58
ϑ_5^f	1
ϑ_6^f	1
α^c	0.35
ρ_c	{0.58,0.65,0.75}
α^e	0.15
ρ_e	{0.9,0.99}

Notes: The table summarises the assigned values regarding the long-term means of the 5-year old vessel prices, \overline{P}_5^* , and the net earnings variable, $\overline{\Pi}$; the actual autocorrelation of net earnings, ρ_0 ; the variance of the net earnings shock, σ_ε^2 ; the vessel supply, Q ; the remaining economic life of the 5-year old vessel, T ; the risk-free rate, R_f ; the fraction of conservatives in the investor population, μ ; the coefficient of absolute risk aversion of fundamentalists, α^f ; the perceived persistence of fundamentalists, ρ_0 ; the 5- and 6-year variance adjustment coefficients of fundamentalists, that is, ϑ_5^f and ϑ_6^f , respectively; the coefficient of absolute risk aversion of conservatives, α^c ; the perceived persistence of conservatives, ρ_c ; the coefficient of absolute risk aversion of extrapolators, α^e ; and the perceived persistence of extrapolators, ρ_e . Note that we list parameters ϑ_5^i and ϑ_6^i only for the fundamentalist since in the cases of conservatives and extrapolators these depend solely on the choice of ρ_i .

state equilibrium prices of the 5- and 6-year old vessels, \overline{P}_5^* and \overline{P}_6^* , respectively, are nested. In line with (3.11), we set \overline{P}_5^* equal to the sample arithmetic mean of the 5-year old vessel prices (Table 3.1). In turn, this yields $\alpha^f = 0.42$ and $\overline{P}_6^* = 22.14$.

The conservative and extrapolator agent-specific parameters are estimated in a similar manner. Since these parameters are nested, for any chosen value of the key parameter of interest ρ_i , the values of the products $\alpha^i \vartheta_5^i$ and $\alpha^i \vartheta_6^i$ are endogenously determined (Appendix 3.B). Hence, it suffices to arbitrarily fix either the parameter ϑ_5^i or ϑ_6^i or α^i . Notably, this choice does not have any qualitative or quantitative implication on the results since only the value of the product matters. We choose to set conservatives' and extrapolators' coefficients of absolute risk aversion equal to 0.35 and 0.15, respectively. Finally, depending on the choice of agent i 's perceived persistence, ρ_i , the equilibrium conditions assign the corresponding values to ϑ_5^i and ϑ_6^i . Table 3.2 summarises the model parameters. Accordingly, we estimate the moments of interest for each scenario under consideration. Table 3.3 presents our model's predictions for several combinations of the three agent-specific parameters of interest $\{\mu, \rho_c, \rho_e\}$. In addition, the right-most column illustrates the actual values of the moments of interest.

To begin with, an apparent feature of the simulation results is that the average prices and average earnings yields are very close to their actual values, irrespective of the selected parameterisation. This was expected since the steady state equilibrium price – recall the equilibrium restrictions in Corollary 1 – has been set equal to the sample mean of the actual vessel prices. In addition, recall that net earnings are exogenously determined and, thus, independent of the chosen parameterisation. Furthermore, equations 3.9b, 3.10b, and 3.11b imply a high positive correlation between net earnings and vessel prices (Alizadeh and Nomikos, 2007). As a result, the autocorrelations of net earnings and 5-year old prices are closely related, irrespective of the scenario. Taken together, these facts explain why the latter statistic has the same value across all scenarios and is also very close to the actual one.

The price volatility statistic is defined as the ratio of the standard deviation of 5-year old vessel prices in the extrapolative heterogeneous-agent economy to the standard deviation of the fundamental value of the 5-year old asset, for a given net earnings shock sequence. When this ratio is higher than one, the heterogeneous-agent model prices are more volatile than the ones in the benchmark rational economy (Barberis *et al*, 2015); hence, this ratio⁶⁷ captures the “excess volatility”

⁶⁷ We assign a benchmark value to this statistic by considering the volatility of vessel prices in a counterfactual fully rational economy, given by equation 3.13. Substituting in this formula the actual volatility of net earnings from Table 3.1, we estimate the fundamental value for our actual data; that is, $\sigma(P_{5,t}^f) = 5.71$. However, this

value is arguably lower compared to the actual volatility of vessel prices in the data, $\sigma(P_{5,t}) = 7.65$. Specifically, the price volatility ratio is approximately equal to 1.34.

– that is, the finding that actual vessel prices are more volatile than those obtained by optimally forecasted net earnings (Greenwood and Hanson, 2015). As expected, our simulation results suggest that the vessel price volatility statistic is positively related to the perceived autocorrelation coefficient of both types of agent and negatively related to the relative fraction of conservatives. Thus, the higher the average degree of net earnings extrapolation in the market, the higher the volatility of vessel prices.

Evidently, a parameterisation close to the ones in columns D and E of Table 3.3 is able to generate a price volatility statistic that approaches the respective empirical value. From an economic perspective, the market must consist from a very large fraction of investors holding totally rational or “near-rational” beliefs (i.e., conservatives) and a very small fraction of participants with extremely extrapolative expectations (i.e., extrapolators). Therefore, the “average investor” must hold “near-rational” beliefs regarding the evolution of the net earnings variable or, equivalently, we can argue that the degree of cash flow extrapolation in the market must be relatively low in equilibrium (and not as is the case in columns A, B, and F of Table 3.3). Importantly, note that this result is in line with the analysis in Chapter 2 of this thesis where we argued that the average shipping investor appears to anticipate – up to a certain degree – the mean-reverting character of net earnings (Greenwood and Hanson, 2015). As a result, vessels are not highly over (under) valued in equilibrium and, in turn, vessel price volatility – while being higher – is not extremely higher than the respective one in the benchmark fully rational economy.

Furthermore, the results related to the correlation between net earnings and net earnings yields and the net earnings yields regressions confirm a well-analysed argument in the recent shipping literature. Namely, while vessel prices and net earnings are highly correlated they do not change proportionately over time (Greenwood and Hanson, 2015). Consequently, net earnings yields fluctuate significantly over time. Specifically, recall that earnings yields are highly correlated with the prevailing net earnings and strongly and negatively forecast future net earnings growth. What is more, the bulk of net earnings yield volatility is attributed to expected net earnings growth variation and not to time-varying expected returns. Our model’s explanation for these facts follows from the analysis in the previous paragraph and Chapter 2 of this thesis.

Specifically, assume that – due to an unexpected positive demand shock – current net earnings, Π_t , are significantly high. Accordingly, the owner of a vessel at time t can immediately exploit the prosperous market conditions and realise a significant, deterministic operating profit during the forthcoming period. Due to this increase in current net earnings, current vessel prices jump compared to their previous level. The mean-reverting character of the net earnings variable, though, implies that future net earnings will – are expected to – be decreased compared to their current

level. Since, however, the average investor has “near-rational” expectations, he values second-hand vessels anticipating – up to a significant degree – the mean reversion of net earnings. As a result, vessels do not become highly overvalued in equilibrium and, in turn, the growth rate of vessel prices is significantly lower compared to the corresponding one of net earnings. Vice versa, after an unexpected negative demand shock. Consequently, net earnings yields are high (low) during prosperous (adverse) market conditions.

The previous analysis explains also why only a small proportion of net earnings yield volatility is attributed to future returns. Once again, assume that at time t there is a positive shock in net earnings. In turn, this will result in a high 5-year old vessel price but, more importantly, in a high earnings yield, $\Pi_t/P_{5,t}$. Accordingly, due to mean reversion, net earnings at time $t + 1$ will be reduced and, as a result, the six-year old vessel price will also be decreased compared to the 5-year old price one period before; thus, the ratio $P_{6,t+1}/P_{5,t}$ will be low. Equation 3.25, however, indicates directly that these two facts have an offsetting effect on the one-period return; hence, the volatility of the return variable is significantly lower compared to the one of net earnings growth.⁶⁸ Consequently, only a small fraction of earnings yield volatility is attributed to returns.

In contrast, if the average degree of extrapolation in the market were much higher, changes in vessel prices would be – in the same direction and, thus, of the same sign but – of a larger magnitude than the corresponding ones in net earnings and, as a result, net earnings yields and net earnings would be negatively correlated. In turn, due to the mean reversion of net earnings, the earnings yield would be strongly positively related with future net earnings growth. What is more, a substantial fraction of the earnings yield volatility would be attributed also to future returns.⁶⁹ This scenario – illustrated in Column A of Table 3.3 – is in sharp contrast with reality. As more conservative participants enter the market, however, average investor expectations become closer to rational and the model-implied predictive regressions results approach the respective empirical values. Finally, we observe that the R-squared of the net earnings growth regressions – and the slope coefficients’ p-values – are significantly high in all cases. Since net earnings are – exogenously – generated through equation 3.1, they exhibit the same highly volatile behaviour irrespective of the parameterisation. Hence, a significant portion of net earnings yields variation is always attributed to variation in future net earnings growth.

Since trading is the result of heterogeneous beliefs in the market, one should expect that average

⁶⁸ For thorough analysis of this point see Chapter 2 of this thesis.

⁶⁹ As the presence of extrapolative beliefs in the market increases, the volatility of the earnings yield becomes significantly reduced. This is because changes in vessel prices weaken the effect of net earnings changes on the earnings yield. To illustrate this argument, in scenarios A and C of Table 3.3, the earnings yield volatilities – scaled by the earnings yield volatility in the benchmark rational economy – were 0.32 and 0.97, respectively.

trading activity would increase with the degree of heterogeneity and decrease with the difference in the population fractions. Our numerical results suggest that this is precisely the case; when both types have a strong presence in the market and a noticeable belief disagreement, trading activity is high (column B of Table 3.3) and vice versa. As we illustrate in the following, the market exit points of the agents and, in turn, the correlation between net earnings and trading activity are extremely sensitive to the choice of parameter μ . Keeping the values of ρ_c and ρ_e constant, we see that for μ equal to 0.1, 0.5, and 0.95, the respective correlation coefficients are negative, approximately zero, and positive, respectively (columns A-C of Table 3.3). Similarly, for fixed $\mu = 0.95$, columns C-F of Table 3.3 suggest that the correlation coefficient is positively related to ρ_e and negatively to ρ_c . Finally, due to the short-sale constraints and, in turn, the asymmetry in investors' market exit points, the correlation between absolute net earnings changes and trading activity – while being very high across all parameterisations – is not perfectly positive.

In conclusion, parameterisations $\{0.95, 0.58, 0.99\}$ and $\{0.95, 0.65, 0.90\}$ – that is, columns D and E of Table 3.3, respectively – appear to be able to capture sufficiently almost all stylised facts under consideration.⁷⁰ Therefore, our empirical estimation suggests that conservatives must have totally rational or “near-rational” beliefs, extrapolators must hold highly extrapolative beliefs (thus, there must exist significant heterogeneity of beliefs among the two types of investor), and the fraction of conservative investors must be very high. In turn, these prerequisites imply that the average investor expectations must be “near-rational”.

3.III.C. Economic Interpretation

From an economic perspective, this finding is in accordance with the nature of the shipping industry; namely, the large fraction of conservative investors corresponds to the large number of established – either publicly-owned or privately-held – shipping companies that operate in the industry. In some instances, ship owning families have been present in the market for more than a century (Stopford, 2009) and, consequently, have strong prior experience and expertise about the key supply and demand drivers of the industry – these were analysed in Chapters 1 and 2 of this thesis. In turn, their superior knowledge translates into more rational forecasts about future market conditions compared to relatively new investors. Extrapolators, on the other hand, reflect new

⁷⁰ Table 3.3 presents the results for 6 scenarios, however, by conducting numerous simulations using alternative parameterisations, we observe that the main statistics under consideration are strictly monotonous functions of the respective population parameter in the intervals between the examined cases, for a given net earnings sequence. For example, keeping the values of ρ_c and ρ_e equal to 0.58 and 0.09, respectively, vessel price volatility is a strictly decreasing function of μ in the interval $[0.1, 0.5]$ (columns A and B). Furthermore, while one can obtain values closer to the actual moments through finer adjustment of the set of parameters, the results and the realised patterns will be qualitatively very similar.

entrants such as diversified investors (e.g., private equity firms) with little or no previous experience of the market. It is well-documented that during prosperous periods, new entrants, impressed by the high prevailing earnings and short-term returns, are eager to buy vessels which, subsequently, are more than keen to sell as conditions begin to deteriorate. In contrast, there are many cases where traditional, established owners have realised significant returns by selling vessels at the peak of the market and buying at the trough – a strategy known as “playing the cycles” (Stopford, 2009). As analysed in the following subsections, our model accounts for this fact through the two market exit points; namely, extrapolators exit during adverse market conditions while conservatives during extremely prosperous ones. Finally, as mentioned above, our simulation results suggest that, in order to simultaneously match the empirical regularities, the average investor expectations must be “near-rational”. From an industrial and microeconomic point of view, this conjecture is plausible since the distinct supply and demand determinants and freight rate mechanism of the shipping markets – as analysed in Chapters 1 and 2 of this thesis – in conjunction with the established presence of experienced, traditional investors, render future cash flows predictable up to a highly significant degree (recall the main findings of Chapter 2 and the relevant discussion).

In a cross-industry comparison, the finding that even a small fraction of extrapolators can reproduce the observed findings – in particular, the volatility of vessel prices – is of interest since the model of Barberis *et al* (2015a) suggests that to produce the “excess volatility” in the U.S. equity markets, extrapolators must constitute 50% of the population. While the results in the two models are not directly comparable, since Barberis *et al* (2015a) examine the “excess volatility” of the price-dividend ratio and not the asset price volatility per se, as is the case here, we can draw two interesting conclusions. First, from a mathematical perspective, the cash flow extrapolative expectations mechanism incorporated in our model is very direct as even a one-period positive (negative) cash flow shock is immediately translated into an over (under) valuation of the vessel. In contrast, Barberis *et al* (2015a) assume a much slower price – return – extrapolative expectations process; namely, in their model, a substantial over (under) valuation of the asset requires consecutive periods of positive (negative) shocks.

Second, while U.S. stock prices are – in relative terms – more volatile than their respective dividends, vessel prices fluctuate relatively less compared to the corresponding net earnings. What is more, vessel prices exhibit – in relative terms – significantly less “excess volatility” compared to U.S. stock prices (recall that this point is extensively analysed in Chapter 2 of this thesis). Therefore, it should be directly expected that a smaller fraction of extrapolators – with, nevertheless, highly extrapolative beliefs – would be able to reproduce the observed vessel price volatility. Vice versa, from an economic perspective, due to both the distinct supply and demand mechanism of the

shipping market as opposed to the U.S. equity one (once again, recall the discussion in Chapter 2) and the fundamental differences in the structures of the two markets (e.g., liquidity concerns, proportion of traditional investors, entry conditions, etc.), it is much more plausible that extrapolators are – can be – a substantially larger fraction of the U.S. equity market population compared to the shipping industry. Finally, note that since the purpose of this model is to simultaneously explain (i.e., *ex post*) several empirical regularities observed in the shipping industry – and not to be applied as a forecasting framework (i.e., *ex ante*) – it does not offer any new tangible trading strategy implication; importantly, however, it does strongly explain and verify the established “playing the cycles” one.

3.III.D. Sensitivity Analysis


To provide a deeper intuition of the mechanism that creates the positive correlation between net earnings and trading activity, this subsection examines the sensitivity of agents’ exit points to the key model parameters. In each case, we allow the relevant parameter of interest to vary while keeping the remaining ones fixed. The corresponding fixed parameters are based on the parameterisation $\{0.95, 0.65, 0.9\}$, that is, column E of Table 3.3.

As Corollary 3 suggests, both agents’ exit points are strictly increasing functions of conservatives’ fraction. Panel A of Figure 3.3 depicts this relation for $\rho_c = 0.65$, $\rho_e = 0.9$, and $\mu \in [0.05, 0.95]$. Evidently, as μ deviates from the midpoint 0.5, the asymmetry between the two exit points increases. Specifically, when μ reaches the value of 0.95, extrapolators exit the market even for slightly adverse net earnings values while conservatives remain active in the market even for significantly high ones. The opposite phenomenon is observed when the fraction of conservatives in the market is low. Panels B and C plot the sensitivity of both agents’ exit points to the perceived persistence of extrapolators and conservatives, respectively. Namely, Panel B suggests that conservatives’ exit point is a strictly decreasing function of the perceived persistence of extrapolators while extrapolators’ exit point is a strictly increasing one. Finally, as Panel C illustrates, the inverse is true for conservatives’ exit point. The implications of these features are illustrated in the following.


Specifically, a large fraction of conservatives combined with a high ρ_e and a low ρ_c result in an exit point for extrapolators that is very close to the steady state equilibrium, $\bar{\Pi}$.⁷¹ On the other hand, these population characteristics yield a conservatives’ exit threshold that is significantly above the steady state. Essentially, this implies that conservatives are always present in the market while extrapolators’ optimal investment policy is to remain inactive even during slightly adverse market

⁷¹ Specifically, the parameterisation $\{0.95, 0.65, 0.9\}$ yields $\Pi_5^e = 2.44$ and $\Pi_5^c = 15.75$, while $\bar{\Pi} = 3.1$.


conditions. Figure 3.4 illustrates the immediate effect of conservatives' fraction and extrapolators' persistence on market trading activity. Namely, we examine the dependence of the – relative – magnitude of trading



Panel A: Sensitivity of exit points to the fraction of conservatives.




Panel B: Sensitivity of exit points to extrapolators' persistence.




Panel C: Sensitivity of exit points to conservatives' persistence.

Figure 3.3: Sensitivity of Market Exit Points to Parameter Values.

Figure 3.3 depicts the relation between agents' exit points and the key parameters of the model. Panel A illustrates the sensitivity to the fraction of conservatives for $\rho_c = 0.65$ and $\rho_e = 0.9$. Panel B shows the sensitivity to extrapolators' perceived persistence for $\mu = 0.95$ and $\rho_c = 0.65$. Panel C demonstrates the sensitivity to conservatives' perceived persistence for $\mu = 0.95$ and $\rho_e = 0.9$. The horizontal solid black line in each panel shows the steady state value of the net earnings variable.



Panel A: Sensitivity of trading activity to the fraction of conservatives.



Panel B: Sensitivity of trading activity ratio to extrapolators' persistence.

Figure 3.4: Sensitivity of Trading Activity to Key Model Parameters.

Panel A presents the relation between the fraction of conservatives and trading activity following positive and negative two standard-deviation shocks for $\rho_c = 0.65$ and $\rho_e = 0.9$. Panel B presents the relation between extrapolators' persistence and trading activity following positive and negative two standard-deviation shocks for $\mu = 0.95$ and $\rho_c = 0.65$. The arrow indicates the limiting value of extrapolators' perceived persistence, ρ_e^* .

activity during prosperous and adverse market conditions on these parameters.⁷² Accordingly, we perturb the steady state equilibrium with a positive and a negative two standard-deviation shock, respectively. As before, we allow each time the corresponding parameter of interest to vary while holding the other ones fixed.

Panel A of Figure 3.4 illustrates the relationship between trading activity and the fraction of conservatives in the market. Notice that for $\mu = 0$ and $\mu = 1$ there is no heterogeneity among agents; hence, there is no trading activity. Furthermore, for $\mu = 0.5$ trading activity is approximately the same after the two shocks. For values of μ higher than 0.5, trading activity is much higher following a positive shock. It is straightforward to interpret this pattern. In line with Panel A of Figure 3.3, for large values of μ , extrapolators exit relatively quickly following a negative net earnings shock. As a result, extrapolators' demand and, in turn, their holdings of the risky asset become zero. In addition, from Corollary 1, in the steady state equilibrium both agents hold the risky asset according to their population fractions. Therefore, as Appendix 3.B illustrates, for large values of μ , trading activity after the negative shock equals $(1 - \mu)Q$. Since μ is large, in this case, the resulting activity is relatively small. In contrast, after a positive shock, both agents are present in the market and for this set of parameter values trading activity is significantly higher than $(1 - \mu)Q$. It is this mechanism that creates the asymmetry between the two cases. Note that if we had perturbed the steady state equilibrium with shocks of smaller absolute value than the one inducing the exit of extrapolators, trading activity in the positive and negative cases would have been essentially the same.

Panel B illustrates the relation between trading activity and extrapolators' perceived persistence. As the latter variable deviates from ρ_c , the heterogeneity of beliefs and, in turn, trading activity in the market increases. Up to a limiting value of extrapolators' persistence, denoted by ρ_e^* , trading activity in the positive and negative shock cases is approximately the same. In the interval $(\rho_e^*, 1)$, however, trading activity after the positive shock is higher compared to the one after the negative shock. In line with Panel B of Figure 3.3, this follows from the fact that extrapolators' exit point is a strictly increasing function of ρ_e . Accordingly, trading activity after the negative shock is equal to $(1 - \mu)Q$ which for our chosen setting is 0.05. As ρ_e increases further, extrapolators' exit point increases as well, however, the trading activity after the negative shock is bounded since it cannot be higher than 0.05.

3.III.E. Impulse Response Functions

Having conducted the sensitivity analysis, we now examine the effect on the economy of a one-time shock in the net earnings variable. In what follows, we present model-implied impulse response

⁷² The corresponding results for the degree of conservatives' persistence are the inverse of the ones for extrapolators and can be directly inferred from those (Appendix 3.C).

functions for the parameterisation $\{\mu, \rho_c, \rho_e\} = \{0.95, 0.65, 0.9\}$, that is, column E of Table 3.3. Figures 3.5 and 3.6 illustrate the behaviour of net earnings, 5-year old vessel prices, vessel demand, and trading activity after a two standard-deviation positive and negative shock, respectively. In panel B of each figure, apart from the model-generated 5-year old price – i.e., the equilibrium price of the risky asset – we also present the respective agent-specific valuations. The latter refer to the “fair” value of the asset from each agent’s perspective.⁷³ For comparative purposes, we also plot the fundamental value of the asset. As Corollary 1 suggests, in the steady state equilibrium all four valuations coincide.

In the following, the negative shock case is analysed. The positive shock case can be directly inferred from this one. At $t = 0$, net earnings are equal to their long-term mean, $\bar{\Pi}$, and the model is in its steady state (Panel A of Figure 3.6). Hence, the agent-specific valuations of the vessel coincide (Panel B); thus, all agents have the same per capita demand for the asset (Panel C). Furthermore, assuming that in the previous period the model was also in its steady state, there is no trading activity in the market (Panel D). At $t = 1$, we perturb the steady state equilibrium by generating a negative 2 standard-deviation (i.e., \$2 million) shock. The immediate first-order effect is the decrease of current net earnings by this amount. Due to the mean-reverting property of net earnings, this shock is completely attenuated within roughly 10 years. However, extrapolators expect net earnings to revert to their steady state value after more than 20 periods while conservatives in about 12 (Panel A).

As a result, agent-specific valuations of the risky asset are lower compared to the fundamental one. Nevertheless, extrapolators consider the asset to be overvalued compared to the prevailing market-clearing price while conservatives consider the asset to be undervalued – with respect to their subjective “fair” valuation (Panel B).⁷⁴ Essentially, agents compare their valuation of the asset to its equilibrium price and not to the fundamental price of the asset – which by not being fundamentalists, they totally ignore.⁷⁵ Consequently, extrapolators’ (conservatives’) demand for the 5-year old vessel is lower (higher) compared to the steady state of the economy. The same applies for the valuation of the corresponding 6-year old vessel at $t = 1$.⁷⁶ Therefore, extrapolators’ (conservatives’) demand for the 6-year old vessel is lower (higher) compared to their demand for the


⁷³ Specifically, this corresponds to the expression $\frac{1-\rho_i^{21}}{1-\rho_i}(\Pi_t - \bar{\Pi}) + 21\bar{\Pi} - X_5^i \sigma_\varepsilon^2 Q$ in equation 3.8a.

⁷⁴ Unfortunately, the scale of the graphs does not allow us to distinguish between the two values.


⁷⁵ To be precise, agents compare their expected one-period income from the asset to its equilibrium price. This comparison is quantified for the 5- and 6-year old vessels through the numerator of the fraction inside the maximum function in equations 3.8a.

⁷⁶ The scale of the graphs does not allow us to distinguish between the agent-specific demands for the 5- and 6-year old vessels at the same t , that is, $N_{5,t}^i$ and $N_{6,t}^i$, respectively. However, it allows us to illustrate clearly the difference between the two variables at two consecutive points in time, $N_{5,t}^i$ and $N_{6,t+1}^i$ – which is the relevant one for trading activity.


respective 5-year old vessel one period before (Panel C). In particular, extrapolators' demand for both the 5- and 6-year old vessels




Panel A: Net Earnings.



Panel B: 5-Year Old Vessel Prices.





Panel C: Demand.





Panel D: Trading Activity.

Figure 3.5: Model-Implied Impulse Response Functions Following a Positive Shock.

Figure 3.5 displays model-implied impulse response functions following a positive two standard-deviation (\$2 million) shock to net earnings, for the parameterisation $\{\mu, \rho_c, \rho_e\} = \{0.95, 0.65, 0.9\}$. Panel A illustrates the actual evolution of net earnings and the evolution perceived by each extrapolator type. Panel B shows the model-generated 5-year old vessel prices and the fundamental and agent-specific valuations. Panel C demonstrates the agent-specific share demands for the 5- and 6-year old vessels. Finally, Panel D plots the



Panel A: Net Earnings Panel B: 5-Year Old Vessel Prices.



Panel C: Demand. Panel D: Trading Activity.

trading activity in the market. The horizontal solid black line in each panel shows the steady state value of the corresponding variable.

Figure 3.6: Model-Implied Impulse Response Functions Following a Negative Shock.

Figure 3.6 displays model-implied impulse response functions following a negative two standard-deviation (\$2 million) shock to net earnings, for the parameterisation $\{\mu, \rho_C, \rho_e\} = \{0.95, 0.65, 0.9\}$. Panel A illustrates the

actual evolution of net earnings and the evolution perceived by each extrapolator type. Panel B shows the model-generated 5-year old vessel prices and the fundamental and agent-specific valuations. Panel C demonstrates the agent-specific share demands for the 5- and 6-year old vessels. Finally, Panel D plots the trading activity in the market. The horizontal solid black line in each panel shows the steady state value of the corresponding variable.

is equal to zero because the two exit points are higher than the corresponding net earnings variable at $t = 1$.⁷⁷ Thus, extrapolators exit from the market. Due to this rapid change in demand, there is significant trading activity in the second-hand market for vessels (Panel D). Specifically, extrapolators reduce their relative fractions of the risky asset while conservatives increase it.

However, since the short-sale constraints bind, trading activity is much lower compared to the one in the positive shock case (Panel D of Figure 3.5). In year 2, net earnings revert towards their long-term mean, although, they are still below both exit thresholds, Π_5^e and Π_6^e . Therefore, extrapolators remain out of the market and there is no trading activity during this date. In year 3, net earnings are slightly higher than Π_6^e but still below Π_5^e .⁷⁸ Accordingly, there is rather small trading activity. In year 4, though, activity becomes noticeably higher since the demand for the 6-year old vessel is substantially higher than the demand for the 5-year old vessel in year 3 – which was equal to zero (Panel C of Figure 3.6). Finally, from this point onwards, both agents are present in the market and trading activity strictly decreases with time until it becomes zero – when net earnings converge to their long-term mean. In contrast, in the positive shock case, both agents are always present in the market (Panel C of Figure 3.5) and, as a result, there is positive – and, thus, significantly higher compared to the negative shock case – trading activity during the corresponding period (Panel D of Figure 3.5).

3.III.F.Expectations of Returns and Realised Returns

This subsection examines the agent-specific expectations of future returns and the corresponding realised returns. In line with Section 3.II, agent i 's one-period expected return from operating the vessel for the interval between her fifth and sixth years of economic life is given by⁷⁹

$$R_t^i \equiv R_{t \rightarrow t+1}^i = \frac{\frac{1 - \rho_i^{21}}{1 - \rho_i} (\Pi_t - \bar{\Pi}) + 21\bar{\Pi} - X_5^i \sigma_\varepsilon^2 Q - P_{5,t}}{P_{5,t}} = \frac{E_t^i [P_{6,t+1}] + \Pi_t - P_{5,t}}{P_{5,t}}. \quad (3.26)$$

⁷⁷ Namely, for the parameter values incorporated in this section, $\Pi_{t=1} \cong 1.1$, $\Pi_5^e \cong 2.44$, and $\Pi_6^e \cong 2.42$.

⁷⁸ Specifically, $\Pi_{t=2} \cong 1.94$ and $\Pi_{t=3} \cong 2.43$ while $\Pi_5^e \cong 2.44$ and $\Pi_6^e \cong 2.42$.

⁷⁹ Recall that the numerator of the fraction in (3.8a) reflects the expected one-period net income for investor i .

Since there is one market-clearing price at each t , agent i 's expected return depends on his specific beliefs and the current realisation of the net earnings variable. Specifically, the numerator in 3.26 is – ceteris paribus – an increasing function of ρ_i . Thus, during prosperous market conditions extrapolators have higher expected returns compared to conservatives and, in turn, are more eager to invest compared to the latter, and vice versa.

In order to compare which investor type's expectations are on average closer to the realised return

one period hence, we define the agent-specific prediction error, Z_t^i , as the absolute deviation between agent i 's expected return and the realised (actual) return:

$$Z_t^i = |R_t^i - R_t^a|, \quad (3.27)$$

where the realised return, R_t^a , is estimated through

$$R_t^a \equiv R_{t \rightarrow t+1}^a = \frac{P_{6,t+1} + \Pi_t - P_{5,t}}{P_{5,t}}. \quad (3.28)$$

Plugging (3.26) and (3.28) in equation 3.27 yields

$$Z_t^i = \frac{\left| \frac{1 - \rho_i^{21}}{1 - \rho_i} (\Pi_t - \bar{\Pi}) + 21\bar{\Pi} - X_5^i \sigma_\varepsilon^2 Q - \Pi_t - P_{6,t+1} \right|}{P_{5,t}}. \quad (3.29)$$

In the heterogeneous-agent economy, the prediction error, Z_t^i , depends on the stochasticity of the net earnings variable and the determination mechanism of the equilibrium market price; recall that each type of agent neglects the strategy of the other type and, in turn, both agents' investment strategies are based on the misbelief that the price of the vessel will revert to its fair – according to their beliefs – value within one period. Supposing that conservatives explicitly incorporated in their valuation the strategy of extrapolators, then they would have always formed more accurate returns expectations and, as a result, they would have been able to exploit their “more correct” beliefs. Due to competition neglect, however, the equilibrium price and the realised returns depend on a complex weighted average of both agents' beliefs where the weights correspond to the population fractions in the economy.

We further clarify this argument by examining the heterogeneous-agent economy for three different values of μ . The estimation procedure and the remaining parameter values are as in the previous subsections. The statistics under consideration are the mean and standard deviation of each agent i 's expected return; the mean and standard deviation of the realised returns; and the mean and standard deviation of the agent-specific prediction error. In addition, we estimate the

expected returns, R_t^f , the realised returns, and the prediction error, Z_t^f , in the counterfactual rational economy. In this case, the expected return formula is simplified to

$$R_t^f \equiv R_{t \rightarrow t+1}^f = \frac{Y_5^f \sigma_\varepsilon^2 Q}{\frac{1 - \rho_f^{21}}{1 - \rho_f} (\Pi_t - \bar{\Pi}) + 21\bar{\Pi} - [X_5^f + Y_5^f] \sigma_\varepsilon^2 Q}. \quad (3.30)$$

By construction, if no shock occurs between two consecutive periods, the rationally expected return is equal to the realised one.

Furthermore, in the steady state equilibrium of the rational economy, the expected return, R_t^{*f} , is

$$R_t^{*f} = \frac{Y_5^f \sigma_\varepsilon^2 Q}{21\bar{\Pi} - [X^f + Y^f] \sigma_\varepsilon^2 Q}.$$

For our parameter values, this is approximately equal to 0.1044. Thus, we should expect that, after 10,000 simulations, both the realised and rationally expected returns will converge to this value. Table 3.4 summarises the statistics of interest for the three parameterisations of the heterogeneous-agent economy (Panels A-C) and the benchmark rational economy (Panel D).

Table 3.4: Expected Returns, Realised Returns, and Prediction Error.

Variables	Mean	Standard Deviation
Panel A: Expected Returns, Realised Returns and Prediction Error for $\{\mu, \rho_c, \rho_e\} = \{0.95, 0.65, 0.9\}$		
Conservatives' Expected Return	0.1093	0.0241
Extrapolators' Expected Return	0.1375	0.2906
Realised Return	0.0922	0.1350
Conservatives' Prediction Error	0.1057	0.0796
Extrapolators' Prediction Error	0.2838	0.2133
Panel B: Expected Returns, Realised Returns and Prediction Error for $\{\mu, \rho_c, \rho_e\} = \{0.5, 0.65, 0.9\}$		
Conservatives' Expected Return	0.1083	0.1472
Extrapolators' Expected Return	0.1020	0.1950
Realised Return	0.1079	0.2546
Conservatives' Prediction Error	0.1914	0.1511
Extrapolators' Prediction Error	0.2806	0.2279
Panel C: Expected Returns, Realised Returns and Prediction Error for $\{\mu, \rho_c, \rho_e\} = \{0.1, 0.65, 0.9\}$		
Conservatives' Expected Return	0.3927	0.8369
Extrapolators' Expected Return	0.1773	0.0712
Realised Return	0.2237	0.7411
Conservatives' Prediction Error	0.3747	0.4769

Extrapolators' Prediction Error	0.3927	0.5599
Panel D: Expected Returns, Realised Returns and Discrepancy in the Benchmark Rational Economy		
Expected Return	0.1054	0.0131
Realised Return	0.1083	0.1052
Prediction Error	0.0828	0.0627

Notes: This table summarises the mean and standard deviation of the quantities of interest presented in the left column, for three different populations compositions, across 10,000 simulations. Panel A presents the case where conservatives constitute a very large fraction of the population. Panel B illustrates the case where each agent type constitutes half of the population. Panel C summarises the case where extrapolators constitute a very large fraction of the population. Finally, Panel D presents the corresponding results for the benchmark rational economy.

Evidently, when the market is dominated by conservatives (Panel A), their average prediction error is extremely smaller than that of extrapolators.⁸⁰ In line with this, conservatives' average expected return is also closer to the average realised return while the standard deviations of both expected return and prediction error are among the lowest across the cases considered. In accordance with the previous analysis, conservatives' discrepancy can be mainly attributed to the stochasticity of the error term and their slight extrapolative expectations and to a lesser extent to competition neglect – since extrapolators constitute a very small fraction of the population. The inverse is true for extrapolators.

In contrast, when extrapolators constitute the largest fraction of the population, agent-specific prediction errors are quite high (Panel C). In the case of conservatives, this error is mainly attributed to competition neglect – since they constitute a very small fraction of the population – and secondarily to the stochasticity of the error term and their extrapolative expectations. In the case of extrapolators, the inverse is true. Accordingly, the high standard deviation of realised returns is mainly attributed to agents' extrapolative expectations and competition neglect and secondarily to the stochasticity of net earnings. In contrast to the previous case, the model-generated average realised return in the market substantially deviates from the empirical value of the average one-period return. The case where each agent-type constitutes half of the populations lies somewhere in the middle.

Finally, Panel D of Table 3.4 shows that while the average expectations of rational investors converge to the average realised returns, there still exists an average prediction error between the two values. This relatively small discrepancy is solely attributed to the volatility of the cash flow shock. In line with this argument, the standard deviations of both the expected return and the prediction error are substantially low. Regarding the former, we observe that expected returns have

⁸⁰ Figure 3.C2 in Appendix 3.C presents the probability density function related to each agent's prediction error corresponding to the economy in Panel A of Table 3.4.

almost zero standard deviation; thus, in the benchmark rational economy, investors have essentially constant required returns.

3.IV. Robustness

We now proceed to test the robustness of our model's predictions by examining five alternative hypotheses regarding the characteristics of the investor population. Namely, we allow our economy to consist of (i) contrarians and fundamentalists, (ii) contrarians and extrapolators, (iii) fundamentalists, (iv) extrapolators, and (v) contrarians. Accordingly, we compare the findings to both the empirical values and the results from our basic setting.

We introduce contrarian investors, denoted by x , in a straightforward manner. Specifically, we assume that they hold irrational beliefs regarding the net earnings process in the opposite way to that of extrapolators; that is, they overestimate the mean reversion of net earnings. Accordingly, their perceived persistence of net earnings (in equation 3.2), ρ_x , lies in the interval $[0, \rho_0)$. Apart from this feature, contrarians behave exactly as the other agent types. In particular, they also neglect the future demand responses of the other types and they upgrade the perceived riskiness of their investments as they grow older. Therefore, the Proposition and Corollaries 1-3 can be directly extended to capture this alternative specification.

Table 3.5 summarises the results obtained from these alternative hypotheses for a variety of investor population characteristics, $\{\mu, \rho_x, \rho_i\}$.⁸¹ For reasons of brevity, we present only the statistics related to the main quantities of interest. The estimation procedure and the basic parameter values are as in Section 3.III. Evidently, the results suggest that these alternative hypotheses are not able to simultaneously match sufficiently the empirical values. To begin with, in the heterogeneous agents scenaria (Panels B and C), we observe that the main effect of contrarians' presence in the market is the attenuation of vessel price volatility. This should be a priori expected since vessel price volatility is an increasing function of the perceived persistence. Therefore, in terms of price volatility, contrarians have the opposite effect in the market compared to extrapolators; that is, they generate less volatility than the benchmark rational economy does.

Extending the analysis of Section 3.II, in an economy consisting of contrarians and fundamentalists (Panel B), the latter will exit from the market during adverse market conditions. Accordingly, it is straightforward to interpret the remaining results in Table 3.5. Namely, a very small fraction of extrapolators combined with a sufficient degree of heterogeneity of beliefs results in low average trading activity and positive correlation between trading activity and net earnings. However, due to the large presence of contrarians, the volatility of vessel prices is significantly reduced by up

⁸¹ We have set contrarians' coefficient of absolute risk aversion, a_x , equal to 0.55.

to 50% compared to the fundamental economy. On the other hand, the specifications that generate “excess volatility” cannot simultaneously approach the significant positive correlation between trading activity and net earnings. Finally, in the homogeneous-agent economy (Panel D) there is no trading activity in the market since beliefs’ heterogeneity is what motivates trading in our model.

In conclusion, we have illustrated that none of those alternative hypotheses, regarding the investor composition, can reproduce the stylised facts under consideration. Of course, there exists a variety of alternative model extensions that can be considered. As an example, it is straightforward to model the coexistence of contrarians, fundamentalists, and extrapolators in the economy. However, the results obtained from these extensions lie somewhere between the ones illustrated in Tables 3.3 and 3.5; thus, they are not able to either improve the fit of the model regarding the main quantities of interest or to alter the economic interpretation of the results.

It is of utmost importance to note that this chapter provides a plausible explanation for several stylised facts observed in the second-hand market for vessels. As analysed in the Introduction, while there can be alternative – “rational” – explanations for the observed patterns in either trading activity (e.g., limits to arbitrage) or vessel price behaviour (e.g., time-varying risk preferences), the proposed model has the advantage of simultaneously explaining in a sufficient manner numerous empirical regularities.

Furthermore, in line with Cochrane (2011), most of the potential alternative “rational” explanations incorporate “exotic preferences” rendering them almost indistinguishable from behavioural ones. Equivalently, their predictions stem from auxiliary assumptions and not from the rationality assumption per se (Arrow, 1986). The fact, however, that almost any biased beliefs model can be re-expressed as a rational expectations’ one with time-varying preferences/discount factors (Cochrane, 2011) does not validate the latter approach or invalidate the former one. Specifically, as Lof (2015) argues, biased beliefs models are very appealing when modelling boom-bust cycles as the ones documented in shipping (Greenwood and Hanson, 2015). More importantly, as illustrated above, the economic interpretation of the model and the respective results are plausible and in line with the nature of the shipping industry.

3.V. Conclusion

This chapter examines the market for second-hand vessels related to the dry bulk sector of the shipping industry. Specifically, our partial equilibrium framework investigates the joint behaviour of vessel prices, net earnings, and second-hand trading activity. For this purpose, we develop and, accordingly, estimate empirically a behavioural asset pricing model with microeconomic foundations that can account for some distinct characteristics of the market.

Namely, among other features, our partial equilibrium model reproduces the actual behaviour of vessel prices, the average trading activity in the market, and the positive correlation between net earnings and second-hand vessel transactions. In addition, the proposed framework also accounts for the stylised features presented in Chapter 2 of this thesis; namely, for the finding that net earnings yields are highly positively correlated with the prevailing market conditions and, in turn, strongly negatively forecast future net earnings growth but also for the fact that the bulk of the net earnings yield’s volatility is attributed to expected cash flow variation and not to time-varying expected returns. While the empirical analysis focuses on the Handysize sector, our results have been tested and are representative of the entire dry bulk industry.

Our discrete-time economy consists of two agent types, conservatives and extrapolators, who form heterogeneous expectations about future net earnings and at the same time under (over)

estimate the future demand responses of their competitors. Formal estimation of the model indicates that a heterogeneous beliefs environment where extrapolators have highly extrapolative expectations while conservatives hold totally rational or “near-rational” beliefs and constitute a very high fraction of the investor population can explain the positive relation between net earnings, prices, and second-hand vessel transactions.

From an economic perspective, this finding is in accordance with the nature of the shipping industry; namely, the large fraction of conservative investors corresponds to the large number of established shipping companies that operate in the industry. In some instances, ship owning families have been present in the market for more than a century and, consequently, have strong prior experience and expertise regarding the industry which, in turn, translates into more accurate forecasts about future market conditions compared to relatively new investors. Extrapolators, on the other hand, reflect new entrants such as diversified investors (e.g., private equity firms) with little or no previous experience of the market.

In conclusion, the proposed partial equilibrium framework provides a first step towards the explicit modelling of the joint behaviour of net earnings, vessel prices, and trading activity which, to the best of our knowledge, had never been examined from the perspective of a structural, behavioural economic model in the literature before. Furthermore, to the best of our knowledge, this is the first time that a heterogeneous beliefs asset pricing model with microeconomic foundations is applied to a real asset economy. Therefore, we provide a framework that can be incorporated and, accordingly, empirically evaluated in other markets with similar characteristics, such as the airplane and the commercial real estate industries.

3.V.A. Connection to Chapter 4

Having concluded the analysis of the physical shipping market for second-hand vessels, in Chapter 4, we turn to the derivative market for Forward Freight Agreements (FFAs) related to the dry bulk sector of the shipping industry. Specifically, we begin by illustrating that the bulk of volatility in the main FFA valuation ratio can be attributed to expectations about future physical market conditions rather than expectations about future risk premia. This stylised fact and, more importantly, its economic justification are perfectly aligned with the main findings of Chapter 2.

Despite this result, though, there appears to be a statistically significant bias in FFA rates in the form of both a strong momentum effect and substantial predictability of risk premia by lagged price-based signals and economic variables that reflect recent changes in the physical market conditions. Accordingly, we develop a dynamic asset pricing model that can account for these findings. Namely,

our framework incorporates both the familiar “hedging pressure” feature – the rational dimension – and a heterogeneous beliefs explanation – the irrational dimension.

The distinct feature of the model is that, apart from having different objective functions, agents – that is, ship owners, charterers, and speculators – might also differ in the way they form expectations about future market conditions. Specifically, speculators are assumed to have distorted beliefs for two reasons: due to asymmetric and imperfect information but mainly due to a behavioural bias known as “the law of small numbers” or “gambler’s fallacy. In contrast, ship owners and charterers are totally rational investors.

The assumption of asymmetric and imperfect information can be justified by the fact that ship owners and charterers – who participate also in the physical market and, thus, have “inside” information regarding the actual future market conditions – are expected to be able to form more accurate forecasts about future spot rates than speculators – who participate only in the FFA market. Regarding the behavioural bias, speculators are assumed to believe that a realised shock in current spot prices will be followed by one of the opposite sign in the next period and, as a result, they adopt a contrarian investment strategy.

In practice, traders frequently follow contrarian strategies – which can be – influenced by behavioural biases such as the “gambler’s fallacy”. In particular, Kaniel *et al* (2008) provide evidence that numerous traders indeed select contrarian strategies while laboratory experiments, conducted by Bloomfield *et al* (2009), suggest that mainly uninformed investors usually adopt contrarian behaviour. What is more, Grinblatt and Keloharju (2000) show that, in Finnish markets, inexperienced investors frequently act as contrarians while more sophisticated ones tend to follow momentum strategies (Lof, 2015). Those findings are particularly related to our model since speculators correspond to financial investors who, as non-participants in the physical market, are assumed to be less sophisticated and informed regarding future spot market conditions compared to traditional physical shipping market agents.

At this point, recall that our empirical analysis in Chapter 3 suggests that the average investor expectations regarding future market conditions must be – slightly extrapolative but – “near-rational”. In turn, note that the “average investor” of Chapter 3 corresponds to the “ship owner” agent type in Chapter 4. Furthermore, charterers can be plausibly assumed to form rational expectations since they participate in the physical market as well. Therefore, the average physical investor expectations in Chapter 4 can be plausibly assumed to be “near-rational” as well – for simplicity and without loss of generality, we assume that physical players are totally rational.⁸²

⁸² Note that it is straightforward to account for slightly extrapolative beliefs on behalf of ship owners in our framework. Even if we do so, however, the qualitative predictions and conclusions of our model are not affected.

Since, there are no surveys regarding shipping industry participants' beliefs and investment strategies as in the equity markets literature (Greenwood and Shleifer, 2014), we further justify our behavioural explanation by contradiction, that is, using both theoretical predictions and numerical simulations of the proposed framework. Note that a similar justification is followed in the model of Lof (2015) who motivates the presence of contrarian investors empirically by illustrating that the observed regularities can be more sufficiently approached when incorporating contrarian expectations on behalf of a population fraction.

Accordingly, we begin by illustrating that a "fully-rational" model in its simplest form is not able to explain the documented empirical regularities. Then, we add the – time-varying – "hedging pressure" dimension and, in turn, examine the generated results. Having shown that neither this model can simultaneously generate the stylised facts, we incorporate the heterogeneous beliefs dimension to test whether we can qualitatively reproduce our findings. Specifically, the simulation results suggest that, to simultaneously match all observed regularities sufficiently well, one must depart from the rational benchmark of the model. While the predictions are not particularly sensitive to the degree of information asymmetry this is not true for the behavioural bias feature; namely, a fraction of investors must suffer from the "gambler's fallacy" and, in turn, follow a contrarian investment strategy.

Appendix 3

A.3.A. Derivation of the Demand Functions for the Age-Specific Vessels

5-Year Old Vessel

We begin by estimating the time t demand function for the 5-year old vessel for each agent type, $N_{5,t}^i$. In the following, $N_{5+n,t+n}^i$ and $P_{5+n,t+n}$ refer to the time $t+n$ agent i 's demand for and price of the $(5+n)$ -year old vessel, respectively. Since vessels are real assets with limited economic lives, we can estimate this demand recursively. Specifically, assuming that a newly built vessel has an economic life of 25 years, at the terminal date – that is, at $T = t + 20$ – the price of the 25-year old asset must be equal to the cash flow realised on that date which, in turn, its scrap price. However, since this scrap price is correlated with the net earnings variable, we impose the simplifying assumption that it is equal to the net earnings variable corresponding to period T ; that is, $P_{25,t+20} = \Pi_{t+20}$.⁸³ From equation 3.5 of the main text, agent i 's objective at time $T - 1 = t + 19$ is

$$\max_{N_{24,t+19}^i} E_{t+19}^i \left[-e^{-\alpha^i (w_{t+19}^i + N_{24,t+19}^i (\Pi_{t+19} + P_{25,t+20} - P_{24,t+19}))} \right]. \quad (3.A1)$$

Using the fact that $P_{25,T} = \Pi_T$ and, accordingly, incorporating (3.2) of the main text, results in

$$\max_{N_{24,t+19}^i} -e^{-\alpha^i (w_{t+19}^i + N_{24,t+19}^i ((1+\rho_i)\Pi_{t+19} + (1-\rho_i)\bar{\Pi} - P_{24,t+19}))} + \frac{(\alpha^i N_{24,t+19}^i)^2}{2} \vartheta_5^i \sigma_\varepsilon^2. \quad (3.A2)$$

Hence, agent i 's first-order condition implies that

$$N_{24,t+19}^i = \frac{(1 + \rho_i)\Pi_{t+19} + (1 - \rho_i)\bar{\Pi} - P_{24,t+19}}{\alpha^i \vartheta_5^i \sigma_\varepsilon^2}. \quad (3.A3)$$

The market-clearing condition at $T - 1$, $\mu N_{24,t+19}^C + (1 - \mu) N_{24,t+19}^e = Q$, along with (3.A3) yield

$$\Rightarrow P_{24,t+19} = (1 + \rho_i)\Pi_{t+19} + (1 - \rho_i)\bar{\Pi} - \frac{\alpha^i \vartheta_5^i \sigma_\varepsilon^2}{\mu^i} [Q - (1 - \mu^i) N_{24,t+19}^e]. \quad (3.A4)$$

In a similar manner, at time $T - 2 = t + 18$, trader i 's objective is

$$\max_{N_{23,t+18}^i} \left\{ -e^{-\alpha^i (w_{t+18}^i + N_{23,t+18}^i (\Pi_{t+18} - P_{23,t+18}))} E_{t+18}^i \left[e^{-\alpha^i N_{23,t+18}^i P_{23,t+18}} \right] \right\}. \quad (3.A5)$$

⁸³ It is straightforward to assume a scrap value given by an AR(1) process where the long-term mean is equal to the average scrap value in our sample and the random (white noise) term is highly correlated with the error term in (3.1) and (3.2). Alternatively, we could also assume a zero-terminal value of the asset.

Incorporating equation 3. A4, the expectation in (3. A5) can be expressed as

$$e^{-\alpha^i N_{23,t+18}^i \left[(1-\rho_i)\bar{\Pi} - \frac{\alpha^i \vartheta_5^i \sigma_\varepsilon^2}{\mu^i} Q \right]} E_{t+18}^i \left\{ e^{-\alpha^i N_{23,t+18}^i \left[(1+\rho_i)\Pi_{t+19} + \frac{\alpha^i \vartheta_5^i \sigma_\varepsilon^2}{\mu^i} (1-\mu^i) N_{24,t+19}^- \right]} \right\}.$$

At this point, we assume that each agent is characterised by an additional form of bounded rationality in the following sense. Agent i , instead of explicitly considering the strategy of agent $-i$, that is, trying to forecast the evolution of $-i$'s demand, makes the simplifying assumption that, in all future periods, $-i$ will just hold his per-capita fraction of the risky asset supply constant at $\mu^{-i}Q$ (Barberis *et al*, 2015b). Thus, using (3.2) of the main text, the objective function 3. A5 is simplified to

$$\max_{N_{23,t+18}^i} \left\{ -e^{-\alpha^i [w_{t+18}^i + N_{23,t+18}^i ((1+\rho_i + \rho_i^2)\Pi_{t+18} + (2+\rho_i)(1-\rho_i)\bar{\Pi} - \alpha^i \vartheta_5^i \sigma_\varepsilon^2 Q - P_{23,t+18})] + \frac{[\alpha^i(1+\rho_i)]^2}{2} \vartheta_5^i \sigma_\varepsilon^2} \right\}.$$

Therefore, agent i 's first-order condition implies

$$N_{23,t+18}^i = \frac{(1 + \rho_i + \rho_i^2)\Pi_{t+18} + (2 + \rho_i)(1 - \rho_i)\bar{\Pi} - \alpha^i \vartheta_5^i \sigma_\varepsilon^2 Q - P_{23,t+18}}{\alpha^i \vartheta_5^i (1 + \rho_i)^2 \sigma_\varepsilon^2}.$$

Similar to the previous two maximisation problems, agent i 's first-order condition at time $T - 3 = t + 17$, yields

$$N_{22,t+17}^i = \frac{(1 + \rho_i + \rho_i^2 + \rho_i^3)\Pi_{t+17} + (3 + 2\rho_i + \rho_i^2)(1 - \rho_i)\bar{\Pi} - \alpha^i \vartheta_5^i \sigma_\varepsilon^2 [1 + (1 + \rho_i)^2]Q - P_{22,t+17}}{\alpha^i (1 + \rho_i + \rho_i^2)^2 \vartheta_5^i \sigma_\varepsilon^2}.$$

Extending the above pattern up to 20 periods before the end of the vessels' economic life – that is, at time $T - 20 = t -$ and applying basic properties of geometric series, we obtain:

$$N_{5,t}^i = \frac{\frac{1 - \rho_i^{21}}{1 - \rho_i} (\Pi_t - \bar{\Pi}) + 21\bar{\Pi} - X_5^i \sigma_\varepsilon^2 Q - P_{5,t}}{Y_5^i \sigma_\varepsilon^2}, \quad (3. A6)$$

where

$$\begin{cases} X_5^i = \left[\frac{20}{(1 - \rho_i)^2} - \frac{(1 - \rho_i^{20})(1 + 2\rho_i - \rho_i^{20})}{(1 + \rho_i)(1 - \rho_i)^3} \right] \alpha^i \vartheta_5^i \\ Y_5^i = \left(\frac{1 - \rho_i^{20}}{1 - \rho_i} \right)^2 \alpha^i \vartheta_5^i \end{cases}. \quad (3. A7)$$

Finally, in order to be consistent with the nature of the industry, we impose short-sale constraints for each investor type. Following Barberis *et al* (2015b), equation 3.A6 becomes

$$N_{5,t}^i = \max \left\{ \frac{\frac{1 - \rho_i^{21}}{1 - \rho_i} (\Pi_t - \bar{\Pi}) + 21\bar{\Pi} - X_5^i \sigma_\varepsilon^2 Q - P_{5,t}}{Y_5^i \sigma_\varepsilon^2}, 0 \right\}. \quad (3.A8)$$

This corresponds to equation 3.8a of the main text.

6-Year Old Vessel

Following the same procedure, it is straightforward to derive the demand functions for the 6-year old vessel, $N_{6,t}^i$. At this point, it is of utmost importance to note that, in principle, investors could understand that their beliefs about either the cash flow process and/or their competitors' strategy are inaccurate; that is, to learn from their misperception and, accordingly, try to correct it (Barberis *et al*, 2015a). In the context of this framework, however, we do not incorporate an explicit learning process for the following reason. If investors could directly correct-update their beliefs, the main observed regularities would not be reproduced by this environment as the valuation of the asset would be approaching the fundamental one in the benchmark rational case. As a result, there would be neither "excess price volatility" – nor heterogeneity of beliefs and, in turn – nor the observed patterns related to second-hand activity in the market.

Accordingly, we adopt a rather indirect learning mechanism. Specifically, we assume that agents become more "suspicious" – or, equivalently, more risk averse – as the specific asset's age grows. This "suspicion" stems from the fact that they realise that the evolution of net earnings (and prices) does not evolve precisely in the way they expected in the previous period. As a result, agents indirectly respond by increasing the perceived risk associated with their investment. In order not to overcomplicate things, we model the update in agents' beliefs in a straightforward manner. Namely, we assume that agent i at t increases the value of the perceived cash flow shock variance corresponding to the valuation of the 6-year old vessel, $\vartheta_6^i \sigma_\varepsilon^2$, compared to the one incorporated for the valuation of the 5-year old one at $t - 1$, $\vartheta_5^i \sigma_\varepsilon^2$; therefore, $\vartheta_5^i < \vartheta_6^i$.⁸⁴ Thus, for a given t , investors related to different vessel-age classes have different beliefs about the variance of the error term. Of course, in the special case where conservatives are fundamentalists, this specific agent knows the precise stochastic process; hence, no variance update occurs between periods $t - 1$ and t .⁸⁵

⁸⁴ Apart from the economic justification, this result is also an indirect implication of the model solution.

⁸⁵ Alternatively, we could have assumed that agent i becomes more risk averse, which would imply an increase of the CARA coefficient from period t to $t + 1$. Both methods yield exactly the same results. We impose this condition in order for the steady state equilibrium of our economy to be well-defined from a

Thus, according to agent i , net earnings related to the valuation of the 6-year old vessel evolve as

$$\Pi_{t+1} = (1 - \rho_i)\bar{\Pi} + \rho_i\Pi_t + \varepsilon_{t+1}^i,$$

in which $\rho_0 \leq \rho_c < \rho_e < 1$ and $\varepsilon_{t+1}^i \sim N(0, \vartheta_6^i \sigma_\varepsilon^2)$, *i. i. d.* over time, where $0 < \vartheta_6^e < \vartheta_6^c$. Despite their increased “suspicion”, however, agents remain irrational since they still do not form unbiased forecasts of either the cash flow process or their competitors’ demand responses. Following precisely the same procedure as for the 5-year old asset, agent i ’s time t demand for the 6-year old one is

$$N_{6,t}^i = \max \left\{ \frac{\left(\frac{1 - \rho_i^{20}}{1 - \rho_i} (\Pi_t - \bar{\Pi}) + 20\bar{\Pi} - X_6^i \sigma_\varepsilon^2 Q - P_{6,t} \right)}{Y_6^i \sigma_\varepsilon^2}, 0 \right\}, \quad (3.A9)$$

where $P_{6,t}$ refers to the time t price of the 6-year old vessel and

$$\begin{cases} X_6^i = \left[\frac{19}{(1 - \rho_i)^2} - \frac{(1 - \rho_i^{19})(1 + 2\rho_i - \rho_i^{19})}{(1 + \rho_i)(1 - \rho_i)^3} \right] \alpha^i \vartheta_6^i \\ Y_6^i = \left(\frac{1 - \rho_i^{19}}{1 - \rho_i} \right)^2 \alpha^i \vartheta_6^i \end{cases}. \quad (3.A10)$$

Note that, for our parameter values, the fact that agents adjust upwardly the perceived riskiness of the cash flow shock implies that $Y_6^i \sigma_\varepsilon^2 > Y_5^i \sigma_\varepsilon^2$. Thus, the expected one-period net income for the 6-year old investment is scaled by a higher quantity compared to the respective 5-year old one.

A.3.B. Proposition and Corollaries

Proof of Proposition

In order to prove the Proposition, it is convenient to define the aggregate demand at time t as $N_{5,t} = \mu N_{5,t}^c + (1 - \mu) N_{5,t}^e$, where the agent-specific demands are given by equation 3.8a. To begin with, we can directly observe that the lower the price of the vessel, the higher the value of aggregate demand. On the other hand, demand can be equal to zero for a sufficiently high value of the vessel price variable. Formally, aggregate demand is a continuous function of the vessel price, $P_{5,t}$.⁸⁶ Moreover, it is a strictly decreasing function of $P_{5,t}$ (as a sum of strictly decreasing functions) with a minimum value of zero. Accordingly, since the market supply of vessels cannot be negative,

mathematical perspective. Even if we do not impose this assumption, however, the steady state equilibrium restrictions will hold approximately and our results will be essentially the same.

⁸⁶ As a sum of continuous functions. Notice that $\max(f(x), 0)$ is continuous for all continuous f and in our case, $f(P_{5,t})$ – which is given by plugging (3.8a) in equation 3.6 – is a continuous function of $P_{5,t}$.

there always exists a vessel price at which the aggregate demand for the risky asset at time t is equal to the aggregate supply of the vessel, Q . Due to monotonicity of the aggregate demand function, this price is unique. We call this value “market-clearing price” or “equilibrium price” of the 5-year old vessel at each t and we denote it by $P_{5,t}^*$.

Accordingly, we determine this equilibrium price by proceeding in a similar fashion to Barberis *et al* (2015b). In particular, we begin by defining the price at which investor i 's short-sale constraint binds at time t

$$\widetilde{P}_{5,t}^i = \frac{1 - \rho_i^{21}}{1 - \rho_i} (\Pi_t - \bar{\Pi}) + 21\bar{\Pi} - X_5^i \sigma_\varepsilon^2 Q. \quad (3. B1)$$

Since $\frac{1 - \rho_i^{21}}{1 - \rho_i}$ is an increasing function of the perceived net earnings' persistence, ρ_i , there exists a net earnings threshold, denoted by $\widehat{\Pi}_5$ and given by

$$\widehat{\Pi}_5 = \bar{\Pi} + \frac{(X_5^e - X_5^c) \sigma_\varepsilon^2 Q}{\frac{1 - \rho_e^{21}}{1 - \rho_e} - \frac{1 - \rho_c^{21}}{1 - \rho_c}} \quad (3. B2)$$

such that

$$\Pi_t \leq \widehat{\Pi}_5 \Leftrightarrow \widetilde{P}_{5,t}^e \leq \widetilde{P}_{5,t}^c. \quad (3. B3)$$

Namely, when shipping net earnings are –equal to or – below this threshold, the cut-off price of extrapolators is –equal to or – lower compared to the one of conservatives and vice versa.

In order to simplify the illustration, we denote the highest and lowest cut-off prices at time t by $\widetilde{P}_{5,t}^1$ and $\widetilde{P}_{5,t}^0$, respectively, so that $\widetilde{P}_{5,t}^1 \geq \widetilde{P}_{5,t}^0$. Furthermore, we define the aggregate demand when the price is equal to $\widetilde{P}_{5,t}^i$ as $N_{\widetilde{P}_{5,t}^i}$. The fact that demand is strictly decreasing in vessel price implies

$$\widetilde{P}_{5,t}^1 \geq \widetilde{P}_{5,t}^0 \Leftrightarrow 0 = N_{\widetilde{P}_{5,t}^1} \leq N_{\widetilde{P}_{5,t}^0}.$$

Accordingly, we distinguish between two scenarios. First, assume that $N_{\widetilde{P}_{5,t}^0} < Q$, that is, the aggregate demand at the lowest cut-off price at time t is lower than the market supply of vessels. Due to market-clearing, however, total demand will adjust to be equal to total supply at each point in time. Therefore, aggregate demand at time t , $N_{5,t}$, will increase and, accordingly, will become higher than $N_{\widetilde{P}_{5,t}^0}$. In order, though, for demand to increase, price must decrease beyond $\widetilde{P}_{5,t}^0$ which is the lowest cut-off price at this point – since aggregate demand is a strictly decreasing function of the price. In turn, this price decrease implies that the demand of the trader with the lowest cut-off price becomes positive as well. Hence, in this scenario, all traders in the market have strictly positive

demand. Thus, substituting equation 3.8a in the market-clearing condition 3.6 and rearranging for $P_{5,t}$, we obtain the equilibrium price of the vessel:

$$P_{5,t}^{*c+e} = 21\bar{\Pi} + \frac{\mu Y_5^e \frac{1-\rho_c^{21}}{1-\rho_c} + (1-\mu)Y_5^c \frac{1-\rho_e^{21}}{1-\rho_e}}{\mu Y_5^e + (1-\mu)Y_5^c} (\Pi_t - \bar{\Pi}) - \frac{\mu Y_5^e X_5^c + (1-\mu)Y_5^c X_5^e + Y_5^c Y_5^e}{\mu Y_5^e + (1-\mu)Y_5^c} \sigma_\varepsilon^2 Q. \quad (3. B4)$$

This corresponds to equation 3.9b of the main text.

Second, assume that $N_{P_{5,t}^{\bar{1}}} \leq Q \leq N_{P_{5,t}^{\bar{0}}}$.⁸⁷ Due to the fact that aggregate demand is a strictly decreasing function of the price, it follows that the equilibrium price belongs in the interval defined by the lowest and the highest cut-off prices; that is, $\widetilde{P}_{5,t}^0 \leq P_{5,t}^* \leq \widetilde{P}_{5,t}^1$. Accordingly, in equilibrium, only the agents with the highest cut-off price will have strictly positive demand for the vessel. Intuitively, once again, due to the market-clearing condition, the aggregate demand for the risky asset must be equal to the aggregate supply. As a result, the price must be lower than $\widetilde{P}_{5,t}^1$ and higher than $\widetilde{P}_{5,t}^0$. When, however, price is lower than the highest cut-off price, the corresponding agents' demand becomes positive; thus, they are in the market. At the same time, though, the price while being lower than the highest cut-off price, remains higher than the lowest one. Therefore, the corresponding agent type has zero demand and, in turn, stays out of the market. In conclusion, in this second scenario, only one type of agent is active in the market. Which type is this and, thus, the determination of the equilibrium price, depends on the prevailing market conditions.

Specifically, when net earnings are below the threshold $\widehat{\Pi}_5$, then

$$\widetilde{P}_{5,t}^0 = \widetilde{P}_{5,t}^e \leq \widetilde{P}_{5,t}^c = \widetilde{P}_{5,t}^1 \Leftrightarrow N_{P_{5,t}^e} \leq N_{P_{5,t}^c}.$$

Namely, when market conditions are sufficiently adverse, the demand of extrapolators becomes zero and only conservatives have strictly positive demand. Therefore, for $N_{5,t}^e = 0$, from equation 3.8a along with the market-clearing condition 3.6, we obtain the equilibrium price of the vessel in the scenario where only conservatives hold the risky asset

$$P_{5,t}^{*c} = 21\bar{\Pi} + \frac{1-\rho_c^{21}}{1-\rho_c} (\Pi_t - \bar{\Pi}) - \left[X_5^c + \frac{Y_5^c}{\mu} \right] \sigma_\varepsilon^2 Q. \quad (3. B5)$$

⁸⁷ Obviously, the aggregate supply of the risky asset cannot be negative; thus, we cannot observe the scenario where $Q < N_{P_{5,t}^{\bar{1}}}$ – since, by definition, $N_{P_{5,t}^{\bar{1}}} = 0$.

This equation corresponds to (3.10) of the main text. Furthermore, it is straightforward to find the critical point at which extrapolators exit the market. Namely, at this point, the short-sale constraint of extrapolators is binding; hence, the equilibrium price of the market is given also by equation 3.B1. Since, the equilibrium price at each t is unique, by equating 3.B1 to 3.B5, we can obtain the value of the net earnings variable at which extrapolators exit from the market, Π_5^e . Accordingly,

$$\Pi_5^e = \bar{\Pi} + \frac{(X_5^e - X_5^c - \frac{Y_5^c}{\mu})\sigma_\varepsilon^2 Q}{\frac{1 - \rho_e^{21}}{1 - \rho_e} - \frac{1 - \rho_c^{21}}{1 - \rho_c}}. \quad (3. B6)$$

As expected, since Y_5^c is positive, Π_5^e is lower than the threshold $\widehat{\Pi}_5$. This suggests that – depending on the sign of the fraction in condition 3. B6 – even during adverse market conditions extrapolators can be present in the market. Namely, the higher the fraction of extrapolators in the market, the more tolerant they are to unfavourable net earnings conditions. From an economic point of view, this result is straightforward; the stronger the fraction of extrapolators, the more difficult it becomes to be entirely driven out of the market, that is, to trade their aggregate holdings with the other part of the investor population.

In a similar manner, when net earnings are above the threshold $\widehat{\Pi}_5$, then

$$\widetilde{P}_{5,t}^0 = \widetilde{P}_{5,t}^c \leq \widetilde{P}_{5,t}^e = \widetilde{P}_{5,t}^1 \Leftrightarrow N_{\widetilde{P}_{5,t}^c} \leq N_{\widetilde{P}_{5,t}^e}.$$

Specifically, when market conditions are significantly prosperous, the demand of conservatives becomes zero and only extrapolators have strictly positive demand. Therefore, for $N_{5,t}^c = 0$, from equation 3.8a along with the market-clearing condition 3.6, we obtain the equilibrium price of the vessel in the scenario where only extrapolators hold the risky asset:

$$P_{5,t}^{*e} = 21\bar{\Pi} + \frac{1 - \rho_e^{21}}{1 - \rho_e} (\Pi_t - \bar{\Pi}) - \left[X_5^e + \frac{Y_5^e}{1 - \mu} \right] \sigma_\varepsilon^2 Q. \quad (3. B7)$$

This equation corresponds to (3.11) of the main text.

Following the same line of reasoning (i.e., equating 3.B1 with 3.B7), the value of the net earnings variable at which conservatives exit from the market, Π_t^c , is

$$\Pi_5^c = \bar{\Pi} + \frac{(X_5^e - X_5^c + \frac{Y_5^e}{1 - \mu})\sigma_\varepsilon^2 Q}{\frac{1 - \rho_e^{21}}{1 - \rho_e} - \frac{1 - \rho_c^{21}}{1 - \rho_c}}. \quad (3. B8)$$

Since Y_5^e is positive, Π_5^c is higher than $\widehat{\Pi}_5$.

In conclusion, the necessary and sufficient condition for agents to coexist in the market is

$$\Pi_5^e = \bar{\Pi} + \frac{(X_5^e - X_5^c - \frac{Y_5^c}{\mu})\sigma_\varepsilon^2 Q}{\frac{1 - \rho_e^{21}}{1 - \rho_e} - \frac{1 - \rho_c^{21}}{1 - \rho_c}} < \Pi_t < \bar{\Pi} + \frac{(X_5^e - X_5^c + \frac{Y_5^e}{1 - \mu})\sigma_\varepsilon^2 Q}{\frac{1 - \rho_e^{21}}{1 - \rho_e} - \frac{1 - \rho_c^{21}}{1 - \rho_c}} = \Pi_5^c. \quad (3. B9)$$

Condition 3. B9 corresponds to (3.9a) of the main text. Furthermore, for our parameter values, 3. B9 implies that when $\Pi_t = \bar{\Pi}$ both agents are present in the market. ■

Equilibrium Price for the 6-Year Old Vessel

Extending the arguments illustrated above, it is straightforward to prove that a vessel age-specific market-clearing price always exists. Below, we state the equilibrium price conditions for the 6-year old vessel.

First, in the case where both agents are present in the market, that is, when

$$\Pi_6^e = \bar{\Pi} + \frac{(X_6^e - X_6^c - \frac{Y_6^c}{\mu})\sigma_\varepsilon^2 Q}{\frac{1 - \rho_e^{20}}{1 - \rho_e} - \frac{1 - \rho_c^{20}}{1 - \rho_c}} < \Pi_t < \bar{\Pi} + \frac{(X_6^e - X_6^c + \frac{Y_6^e}{1 - \mu})\sigma_\varepsilon^2 Q}{\frac{1 - \rho_e^{20}}{1 - \rho_e} - \frac{1 - \rho_c^{20}}{1 - \rho_c}} = \Pi_6^c, \quad (3. B10)$$

the price is given by

$$P_{6,t}^{*c+e} = 20\bar{\Pi} + \frac{\mu Y_6^e \frac{1 - \rho_c^{20}}{1 - \rho_c} + (1 - \mu) Y_6^c \frac{1 - \rho_e^{20}}{1 - \rho_e}}{\mu Y_6^e + (1 - \mu) Y_6^c} (\Pi_t - \bar{\Pi}) - \frac{\mu Y_6^e X_6^c + (1 - \mu) Y_6^c X_6^e + Y_6^c Y_6^e}{\mu Y_6^e + (1 - \mu) Y_6^c} \sigma_\varepsilon^2 Q. \quad (3. B11)$$

Second, when only conservatives hold the vessel, that is, when

$$\Pi_t \leq \bar{\Pi} + \frac{(X_6^e - X_6^c - \frac{Y_6^c}{\mu})\sigma_\varepsilon^2 Q}{\frac{1 - \rho_e^{20}}{1 - \rho_e} - \frac{1 - \rho_c^{20}}{1 - \rho_c}} = \Pi_6^e, \quad (3. B12)$$

the price is given by

$$P_{6,t}^{*c} = 20\bar{\Pi} + \frac{1 - \rho_c^{20}}{1 - \rho_c} (\Pi_t - \bar{\Pi}) - \left[X_6^c + \frac{Y_6^c}{\mu} \right] \sigma_\varepsilon^2 Q. \quad (3. B13)$$

Third, in the scenario where only extrapolators hold the risky asset; namely, when

$$\Pi_6^c = \bar{\Pi} + \frac{(X_6^e - X_6^c + \frac{Y_6^e}{1-\mu})\sigma_\varepsilon^2 Q}{\frac{1-\rho_e^{20}}{1-\rho_e} - \frac{1-\rho_c^{20}}{1-\rho_c}} \leq \Pi_t, \quad (3. B14)$$

the price equals

$$P_{6,t}^{*e} = 20\bar{\Pi} + \frac{1-\rho_e^{20}}{1-\rho_e}(\Pi_t - \bar{\Pi}) - \left[X_6^e + \frac{Y_6^e}{1-\mu} \right] \sigma_\varepsilon^2 Q. \quad (3. B15)$$

■

Proof of Corollary 1

From the Proposition and the definition of the “steady state” equilibrium, it is straightforward to derive equations 3.15a and 3.15b of the main text. Specifically, equation 3.8 combined with the fact that in the steady state $\Pi_t = \bar{\Pi}$ result in

$$\bar{N}_5^i = \max \left\{ \frac{21\bar{\Pi} - X_5^i \sigma_\varepsilon^2 Q - \bar{P}_5^*}{Y_5^i \sigma_\varepsilon^2}, 0 \right\}. \quad (3. B16)$$

Since, however, in the steady state, both agents coexist in the market, type i 's time t demand becomes

$$\bar{N}_5^i = \frac{21\bar{\Pi} - X_5^i \sigma_\varepsilon^2 Q - \bar{P}_5^*}{Y_5^i \sigma_\varepsilon^2}. \quad (3. B17)$$

Substituting (3. B17) in the market-clearing condition 3.6, we obtain

$$\bar{P}_5^* = 21\bar{\Pi} - \frac{\mu Y_5^e X_5^c + (1-\mu) Y_5^c X_5^e + Y_5^c Y_5^e}{\mu Y_5^e + (1-\mu) Y_5^c} \sigma_\varepsilon^2 Q. \quad (3. B18)$$

Moreover, both types hold the risky asset in analogy to their fraction of the total population if and only if

$$\bar{N}_5^i = \frac{21\bar{\Pi} - X_5^i \sigma_\varepsilon^2 Q - \bar{P}_5^*}{Y_5^i \sigma_\varepsilon^2} = Q \Leftrightarrow \bar{P}_5^* = 21\bar{\Pi} - (X_5^i + Y_5^i) \sigma_\varepsilon^2 Q.$$

However, since the steady state equilibrium price is unique

$$X_5^c + Y_5^c = X_5^e + Y_5^e = X_5^f + Y_5^f = \frac{\mu Y_5^e X_5^c + (1 - \mu) Y_5^c X_5^e + Y_5^c Y_5^e}{\mu Y_5^e + (1 - \mu) Y_5^c}, \quad (3. B19)$$

which corresponds to condition 3.15b of the main text.

Vice versa, restrictions 3.B19 ensure that in the steady state both agents are present in the market. Namely, the parenthesis and, in turn, the second terms on the right-hand side of conditions 3.B6 and 3.B8 are negative and positive, respectively; hence, $\Pi_5^e < \bar{\Pi} < \Pi_5^c$. Following the same procedure, we obtain the steady state equilibrium conditions for the 6-year old case. ■

Agent- and Age-Specific Parameters

The steady state equilibrium conditions 3.15a and 3.16a imply that our model's parameters are nested. This interrelationship can be illustrated through the following system of equations

$$\alpha^i = \frac{21\bar{\Pi} - \bar{P}_5^*}{\left[\frac{20 + (1 - \rho_i^{20})^2}{(1 - \rho_i)^2} - \frac{(1 - \rho_i^{20})(1 + 2\rho_i - \rho_i^{20})}{(1 + \rho_i)(1 - \rho_i)^3} \right]} \vartheta_5^i \sigma_\varepsilon^2 Q, \quad (3. B20)$$

and

$$\alpha^i = \frac{20\bar{\Pi} - \bar{P}_6^*}{\left[\frac{19 + (1 - \rho_i^{19})^2}{(1 - \rho_i)^2} - \frac{(1 - \rho_i^{19})(1 + 2\rho_i - \rho_i^{19})}{(1 + \rho_i)(1 - \rho_i)^3} \right]} \vartheta_6^i \sigma_\varepsilon^2 Q. \quad (3. B21)$$

The implications of this fact are analysed in the empirical estimation of the model.

Trading Volume and Net Earnings

The general expression for the trading activity variable is

$$V_t = \mu^i \left| \max \left\{ \frac{\left(\frac{1 - \rho_i^{20}}{1 - \rho_i} (\Pi_t - \bar{\Pi}) + 20\bar{\Pi} - X_6^i \sigma_\varepsilon^2 Q - P_{6,t} \right)}{Y_6^i \sigma_\varepsilon^2}, 0 \right\} \right. \\ \left. - \max \left\{ \frac{\left(\frac{1 - \rho_i^{21}}{1 - \rho_i} (\Pi_{t-1} - \bar{\Pi}) + 21\bar{\Pi} - X_5^i \sigma_\varepsilon^2 Q - P_{5,t-1} \right)}{Y_5^i \sigma_\varepsilon^2}, 0 \right\} \right|. \quad (3. B22)$$

Due to the short-sale constraints, however, the agent-specific demand functions are not strictly monotonic with respect to the net earnings variable in the entire Π_t domain; namely, strict monotonicity disappears whenever the constraints are binding. As a result, the precise equation

quantifying the trading activity variable – and, therefore, its value – depends on the realisation of the net earnings variable during the two corresponding consecutive dates, $t - 1$ and t . In the following, we examine all possible scenarios.

In the first scenario, both agents are present in the market for two consecutive periods. Equivalently, conservative agents' demands for 5- and 6-year old vessels are positive. Incorporating the equilibrium prices from (3. B4) at $t - 1$ and (3. B11) at t in equation 3. B22 results in

$$V_t = \mu^i |A_6^i \Pi_t - A_5^i \Pi_{t-1} + (A_6^i - A_5^i) \bar{\Pi}|, \quad (3. B23)$$

where

$$A_6^i \Pi_t - A_5^i \Pi_{t-1} + (A_6^i - A_5^i) \bar{\Pi} = N_{6,t}^i - N_{5,t-1}^i \quad (3. B24)$$

is agent i 's change in demand for the asset between periods $t - 1$ and t . The agent-specific constants are given by

$$\begin{cases} A_5^c = \frac{(1 - \mu) \left(\frac{1 - \rho_c^{21}}{1 - \rho_c} - \frac{1 - \rho_e^{21}}{1 - \rho_e} \right)}{[\mu Y_5^e + (1 - \mu) Y_5^c] \sigma_\varepsilon^2} < 0 \\ A_6^c = \frac{(1 - \mu) \left(\frac{1 - \rho_c^{20}}{1 - \rho_c} - \frac{1 - \rho_e^{20}}{1 - \rho_e} \right)}{[\mu Y_6^e + (1 - \mu) Y_6^c] \sigma_\varepsilon^2} < 0 \end{cases}, \quad (3. B25)$$

and

$$\begin{cases} A_5^e = \frac{\mu \left(\frac{1 - \rho_e^{21}}{1 - \rho_e} - \frac{1 - \rho_c^{21}}{1 - \rho_c} \right)}{[\mu Y_5^e + (1 - \mu) Y_5^c] \sigma_\varepsilon^2} > 0 \\ A_6^e = \frac{\mu \left(\frac{1 - \rho_e^{20}}{1 - \rho_e} - \frac{1 - \rho_c^{20}}{1 - \rho_c} \right)}{[\mu Y_6^e + (1 - \mu) Y_6^c] \sigma_\varepsilon^2} > 0 \end{cases}, \quad (3. B26)$$

Since trading volume in the market is the same irrespective of the agent type's perspective – from which we analyse it – in the following we examine this variable from the conservative agent's point of view. Accordingly, equation 3. B23 becomes

$$V_t = \mu |A_6^c \Pi_t - A_5^c \Pi_{t-1} + (A_6^c - A_5^c) \bar{\Pi}|. \quad (3. B27)$$

The second scenario is when both agents are present at time $t - 1$ but conservatives exit at t , that is, $N_{6,t}^c$ equals zero. Incorporating the equilibrium prices from (3. B4) at $t - 1$ and (3. B15) at t in equation 3. B22 yields

$$V_t = \mu |A_5^c(\Pi_{t-1} - \bar{\Pi}) + Q|. \quad (3. B28)$$

In the third scenario, conservatives are not present in the market at time $t - 1$ but both agent types are active at t . Proceeding in a similar fashion to before, equation 3. B22 becomes

$$V_t = \mu |A_6^c(\Pi_t - \bar{\Pi}) + Q|. \quad (3. B29)$$

The fourth scenario refers to the case where both agents are present in the market at time $t - 1$ but extrapolators exit at t . In this case, trading activity is given by

$$V_t = \mu \left| A_5^c(\Pi_{t-1} - \bar{\Pi}) - \frac{(1 - \mu)}{\mu} Q \right|. \quad (3. B30)$$

The fifth scenario is when only conservatives are present in the market at time $t - 1$ but both types at t . Therefore, equation 3. B22 becomes

$$V_t = \mu \left| A_6^c(\Pi_t - \bar{\Pi}) - \frac{(1 - \mu)}{\mu} Q \right|. \quad (3. B31)$$

In the sixth (seventh) scenario, only agents of type i ($-i$) are present in the market at time $t - 1$ and only of type $-i$ (i) at t . Namely, (3. B22) simplifies to

$$V_t = Q. \quad (3. B32)$$

Furthermore, if in two consecutive dates agents i are out of the market, there is no trading activity. Finally, if $\mu^i = 0$ or, equivalently, $\rho_c = \rho_e$, the market-clearing condition along with equations 3.8a and 3. A9 suggest that there are no second-hand transactions in the economy. ■

A.3.C. Omitted Figures

The following figures correspond to the results presented in the sensitivity analysis (Subsection 3.III.C) and the expectations of returns and realised returns analysis (Subsection 3.III.E) conducted in the main body of this chapter.



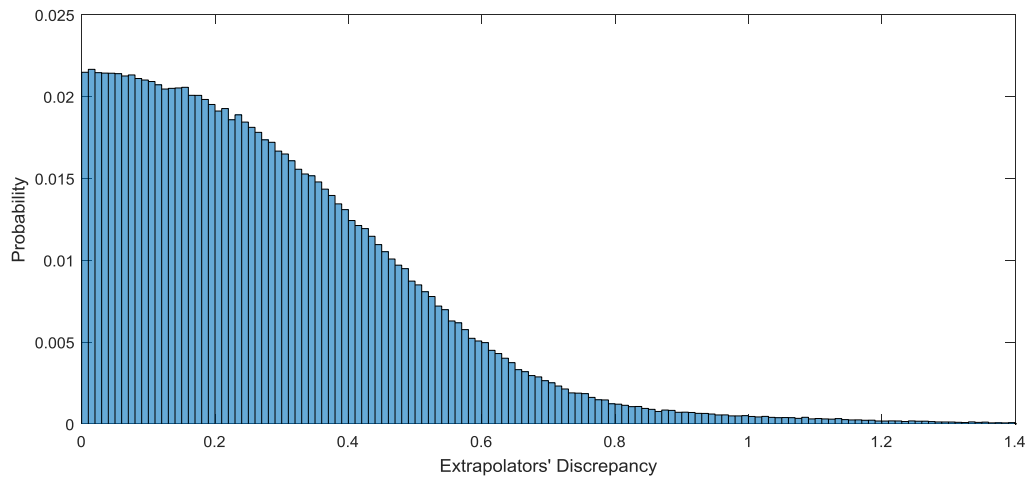


Figure 3.C1: Sensitivity of Trading Activity to Conservatives' Persistence.

Figure 3.C1 presents the relation between conservatives' persistence and trading activity following (positive and negative) two standard-deviation shocks for $\mu = 0.95$ and $\rho_e = 0.9$. The arrow indicates the limiting value of conservatives' perceived persistence, ρ_C^* .



Panel A: Probability Density Function of Conservatives' Discrepancy.



Panel B: Probability Density Function of Extrapolators' Discrepancy.

Figure 3.C2: Probability Density Functions of Agent-Specific Prediction Discrepancy.

Figure 3.C2 presents probability density functions of agent-specific discrepancies between the expected and the realised returns. The incorporated parameterisation is $\mu = 0.95$, $\rho_e = 0.65$, and $\rho_e = 0.9$. Panel A illustrates the case for conservatives and Panel B for extrapolators.


A.3.D. Changes in Investor Wealth

The results in Subsection 3.III.E of the main text suggest that extrapolators have both more volatile expected returns and less accurate expectations compared to conservatives. Therefore, one should expect that the former will have significantly more skewed distribution of one-period wealth changes than the latter. We examine this prediction by estimating agent i 's one-period change in wealth, Δw_{t+1}^i , through

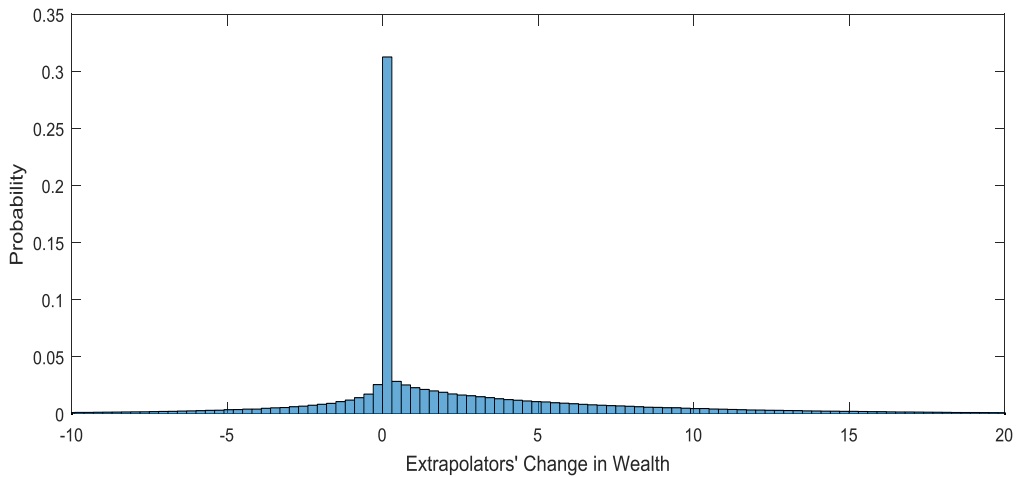
$$\Delta w_{t+1}^i = N_{5,t}^i (\Pi_t + P_{6,t+1} - P_{5,t}). \quad (3.D1)$$

Namely, the one-period change in wealth of agent i equals his time t holdings of the risky asset multiplied by the realised net income at t . Figure 3.D1 illustrates the probability density functions of both agents' one-period changes in wealth after 10,000 simulated paths. The most striking feature of these simulations is that extrapolators realise zero change in one-period wealth with a probability approximately equal to 27.5%. This is an immediate consequence of the fact that extrapolators exit from the market rapidly even during slightly adverse market conditions. When they are present, however, the fact that they form less accurate and more volatile expectations results in very volatile one-period wealth changes. Accordingly, the probability distribution of their one-period wealth change is significantly skewed (Panel B of Figure 3.D1). Namely, the mean, standard deviation, skewness, and excess kurtosis of the distribution are equal to 2.29, 6.28, 1.19, and 6.9, respectively.

In contrast, conservatives' change in wealth closely resembles a normal distribution with mean, standard deviation, skewness, and excess kurtosis equal to 2 and 2.94, 0.15, and -0.04, respectively. As analysed in subsection 3.III.E, conservatives' realised returns heavily depend on the stochasticity of the error term. In turn, the normally distributed error term, combined with fact that they are always present in the market, result in this probability density function (Panel A of Figure 3.D1). In conclusion, while extrapolators' one-period changes in wealth are significantly more volatile than conservatives' ones (this fact is in line with our assumption that extrapolators are in general more risk tolerant compared to conservatives), both types of agent realise approximately the same mean change. Therefore, there is no formal indication that extrapolators "suffer" (on average) by limitations of wealth more than conservatives do. This auxiliary result is similar to the one in Barberis *et al* (2015a).



Panel A: Probability Density Function of Conservatives' Change in Wealth.



Panel B: Probability Density Function of Extrapolators' Change in Wealth.

Figure 3.D1: Probability Density Functions of Agent-Specific One-Period Change in Wealth.

Figure 3.D1 presents probability density functions of agent-specific one-period changes in wealth. The incorporated parameterisation is $\mu = 0.95$, $\rho_e = 0.65$, and $\rho_e = 0.9$. Panel A illustrates the case for conservatives and Panel B for extrapolators.

Chapter 4: The Formation of FFA Rates in Dry Bulk Shipping: Spot Rates, Risk Premia, and Heterogeneous Expectations

Abstract. This chapter examines the formation of FFA rates in the dry bulk shipping industry. We illustrate that the bulk of volatility in the FFA basis can be attributed to expectations about future physical market conditions rather than expectations about future risk premia, as is commonly suggested in the commodity markets literature. Despite this finding, though, there appears to be a bias in FFA rates in the form of both a strong momentum effect and significant predictability of risk premia by lagged price-based signals and economic variables that reflect recent changes in the physical market conditions. An additional interesting finding is the evidence of “contango” in the FFA market. The evidence of bias in FFA rates is also supported by the results of three econometric tests which suggest rejection of the unbiased expectations hypothesis. We further contribute to the literature by developing an asset pricing framework that can explain both the existence of momentum and the documented sort of predictability of future risk premia. Importantly, our dynamic framework can simultaneously account for both the familiar “hedging pressure” feature – the rational dimension – and a heterogeneous beliefs explanation – the irrational dimension. The proposed model incorporates three types of traders: ship owners, charterers, and speculators. The distinct feature of our framework is that, apart from having – as is standard in the literature – different objective functions, agents might also differ in the way they form expectations about future market conditions. Specifically, we develop an asymmetric information environment where speculators suffer from a behavioural bias known as “the law of small numbers” – or, equivalently, “reversion to the mean” or “gambler’s fallacy”. Accordingly, we illustrate formally that, to simultaneously match the observed empirical regularities, one must depart from the rational expectations benchmark of the model. To the best of our knowledge, the FFA market had never been examined from the perspective of a structural, behavioural economic model before. In addition, we contribute to the generic commodity finance literature by incorporating explicitly the behavioural dimension in the formation of derivative contracts rates.

Keywords: Asset Pricing, FFA Rates, Speculation, Biased Beliefs, Law of Small Numbers, Heterogeneous Agents, Unbiased Expectations Hypothesis, Asymmetric Information

4.1. Introduction

This chapter investigates the formation of Forward Freight Agreement (FFA) rates in the dry bulk shipping industry. Specifically, our empirical analysis concentrates upon the Capesize BCI 4TC and Panamax BPI 4TC monthly contracts. Our contribution to the literature is threefold. First, by applying

a variance decomposition framework for the first time in the FFA market, we illustrate the significant forecasting power of FFA contracts regarding future market conditions. Second, we document, for the first time in the literature, several noticeable empirical regularities related to FFA rates and risk premia. Third, we develop a theoretical, heterogeneous agents', behavioural asset pricing model that can account for these stylised facts.

We begin by analysing empirically the formation of the most frequently incorporated FFA valuation ratio, that is, the FFA basis – defined as the log ratio of the FFA rate to the respective prevailing spot price. As it is well-analysed in the empirical asset pricing literature, but also in Chapter 2 of this thesis, predictability of future cash flow growth and/or predictability of future returns constitute the rational benchmark for the interpretation of variation in assets' valuation ratios (Bansal and Yaron, 2007). In the case of futures and forward contracts, however, this question of relative predictability becomes highly important for market participants due to the risk transfer and price discovery roles of these derivative markets.

In their seminal papers, Fama (1984a and 1984b) and Fama and French (1987) illustrate that the variance of the basis of any futures – forward – contract can be decomposed into the sum of the covariance between the basis and the expected change in the spot price plus the covariance between the basis and an expected premium over the spot price at maturity. This premium can be interpreted as the bias in the futures price as a forecast of the future spot price or, equivalently, as the excess return for an investor who goes short on the futures contract. Therefore, through this variance decomposition, it is straightforward to examine which of the two sources is the major determinant of the observed futures bases – prices.

Accordingly, applying this framework for the first time to shipping derivative markets, we illustrate formally that the bulk of volatility in the FFA basis can be attributed to expectations about future physical market conditions, rather than expectations about future risk premia, as is commonly suggested in the commodity markets literature (Fama and French, 1987). More importantly, we provide both an economic interpretation of this result and a comparison with the ones obtained from commodity futures and forward markets. Namely, our results validate and extend the economic arguments presented in the seminal commodity market papers (Hazuka, 1984; French, 1986; Fama and French, 1987) that examine the forecasting power of derivative contracts. In addition, our empirical results and the economic explanation provided are in line with the analysis in Chapter 2 regarding the formation of vessel valuation ratios.

Briefly, our line of reasoning is as follows. Those seminal commodity market articles illustrate that predictability of spot rates is an increasing function of the commodity cost of storage. Equivalently, since inventories tend to smooth predictable adjustments in spot prices, the “more

storable” the commodity the lower the predictability of future spot rates. In shipping, however, the commodity – defined as the mode of seaborne transport – is a service, that is, a non-storable one. Accordingly, the fact that the shipping industry is subject to significant supply and demand shocks which cannot be smoothed through adjustments of the short-term supply, as is the case with storable commodities – the reader can parallelise this to a lack of inventory and, thus, lack of spot price smoothing – results in predictable variation of spot rates and, in turn, substantial forecasting ability of FFA rates.

While, however, the bulk of FFA basis’ volatility is attributed to future spot growth, we cannot exclude the existence of – time-varying – risk premia. Accordingly, for the first time in the shipping literature, we provide evidence of three stylised features that might be of interest to both academic researchers and market participants. First, in contrast to most futures and forward commodity markets, there is no sign of “backwardation” in any type of contract or maturity in the dry bulk FFA market. What is more, we find strong statistical evidence of “contango” in the one-month contracts; that is, the realised risk premia – defined as the log ratio of the FFA rate to the respective settlement price of the contract – appear to be on average positive. Second, we demonstrate the existence of a momentum effect in the FFA market; namely, lagged risk premia positively forecast future risk premia in a strong statistical manner. Third, in line with the previous feature, we provide evidence that there exists – both economically and statistically – significant predictability of future risk premia in this derivative market. The documented predictability is more robust for the Panamax contracts but also for shorter maturities. In particular, FFA risk premia can be forecasted by both price-based signals and economic indicators related to commodity trade and shipping demand. Regarding the former, apart from the momentum effect there appears to be strong predictability using two additional types of price-based indices; namely, lagged spot market indicators and the FFA basis. Regarding the latter, we illustrate that lagged realisations of – changes in – economic variables such as commodity prices (e.g., iron ore) and trade indicators (such as the quantities of imported and exported dry bulk commodities) strongly negatively forecast future risk premia. Note that, as analysed in the following sections, this finding is the key motivation for the development of our heterogeneous beliefs model.

In addition, we provide evidence that future risk premia can also be – negatively – forecasted by past trading activity in the sale and purchase market for second-hand vessels. Interestingly, note that trading activity has been used as an indicator of market liquidity (recall Chapter 3 of this thesis), but also as an investor sentiment index in the shipping literature (Papapostolou *et al*, 2014). Finally, we test whether future market conditions and risk premia can be predicted by market activity variables that incorporate the FFA trading volume and open interest figures related to the

corresponding contracts. While there appears to exist some sort of predictability, mainly in the Capesize sector, the results cannot yet be generalised given the small size of the employed dataset.

From an economic perspective, the existence of statistically significant predictability of future risk premia contradicts the unbiased expectations hypothesis and, in turn, the efficiency of the FFA market. We further examine the validity of the hypothesis by performing three frequently incorporated econometric tests, including Johansen's (1988) cointegration approach. Despite the sensitivity of these tests to the model specification, the results suggest that there exists a bias in the formation of the 1-month FFA rates in both contracts. Regarding the 2-month contracts, our findings point towards the existence of bias, especially in the Panamax BPI 4TC case. Consequently, our empirical estimation results are robust and consistent. Therefore, we demonstrate formally, for the first in the shipping literature, the existence of bias in the dry bulk FFA market. Note that, from an industry participant's perspective (such as "shipping commodity hedge funds"), the existence of a bias in FFA rates suggests that it would be interesting to further examine – in the context of future research – potential profitable trading strategies that incorporate the documented stylised facts

Accordingly, in order to justify economically and, in turn, reproduce our main empirical findings – in particular, the existence of future risk premia predictability, that is, the bias in FFA rates – we develop a theoretical model of FFA price determination. While the proposed framework draws its main features from the last generation of structural economic models in the commodity futures literature (Gorton *et al*, 2012; Acharya *et al*, 2013), we modify and extend the basic setting in two, quantitatively simple but conceptionally important, manners.

First, since shipping services are a non-storable commodity, the "cost-of-carry model" cannot be applied; hence, in contrast to most commodity futures models, our framework departs from the "theory of storage" explanation of "time-varying" risk premia – as is the case in Gorton *et al* (2012) and Ekeland *et al* (2016). An immediate consequence of this fact is the expansion of the common in the existing literature 2-period economic environment to an infinite horizon model. This fact significantly simplifies the empirical evaluation of the generated framework. Accordingly, we are able to test and validate the theoretical predictions of our model through a large number of numerical simulations.

Second, as it is well-established in the asset pricing literature, the majority of rational expectations models fail to explain numerous empirical regularities related to asset prices. Among others, prominent examples are the "excess volatility puzzle" (Leroy and Porter, 1981; Shiller, 1981), the "equity premium puzzle" (Mehra and Prescott, 1985), the positive correlation between trading volume and asset prices (Barberis *et al*, 2015b), and the strong positive relation between the aggregate dividend yield and future returns in the post-WWII U.S. equity markets (Campbell and

Shiller, 1988a; Fama and French, 1988b; Cochrane, 2011). In order to explain these findings, one of the tools that researchers have developed are heterogeneous beliefs economic models that incorporate behavioural biases, termed as heuristics (Barberis *et al*, 2015a). In line with this growing body of research, recall that Chapter 3 of this thesis proposes a heterogeneous agents', behavioural asset-pricing model that can reproduce several key empirical regularities related to the – physical – shipping market for second-hand dry bulk vessels.

Accordingly, this chapter applies the heterogeneous beliefs framework to a derivative shipping market, for the first time in the literature. Specifically, we aim to explain the stylised facts observed in the FFA market by extending the “mean-variance optimisation” rational expectations models to incorporate the existence of distorted beliefs on behalf of a fraction of the investor population. Our discrete-time economy consists of three types of agent; ship owners, charterers, and speculators. The distinct feature of the proposed framework is that, apart from having – as is standard in the commodity markets literature – different objectives, speculators also differ in the way they form expectations about future market conditions for two reasons. Namely, due to asymmetric and imperfect information but mainly due to a behavioural bias known as “representativeness”. In contrast, ship owners and charterers are assumed to be totally rational investors.

The assumption of asymmetric and imperfect information can be justified by the fact that ship owners and charterers, also defined as “physical hedgers”, are expected to be more experienced and better informed – since, by participating also in the physical market, they have “inside” information regarding the actual future market conditions – than speculators. As a result, they are assumed to form more accurate forecasts of future spot market conditions than the latter.

Regarding the behavioural bias assumption, Tversky and Kahneman (1974) state that “representativeness” is a heuristic-driven bias according to which individuals believe that small samples are representative of the entire population. In the context of our model, it is assumed that speculators suffer from a variation of “the law of small numbers” bias which is also known as “regression – reversion – to the mean” and “gambler’s fallacy”. In line with Shefrin (2000), “the law of small numbers” arises “because people misinterpret the law of averages, technically known as ‘the law of large numbers’. They think the law of large numbers applies to small as well as to large samples” or, equivalently, “they exaggerate how likely it is that a small sample resembles the parent population from which is drawn” (Tversky and Kahneman, 1971; Terrell, 1994; Rabin, 2002). As a result, individuals that suffer from this misperception inappropriately predict – rapid – reversal of a trend or shock.

We introduce this irrationality in a rather straightforward manner. Namely, speculators in our model believe that spot price shocks tend to cancel out each other rapidly; thus, they expect that a

realised shock in current spot prices will be followed by one of the opposite sign in the next period. Equivalently, they believe that the spot price variable tends to revert rapidly to its previous level, that is, the one before the last realised shock. As Rabin (2002) argues, an individual suffering from the “gambler’s fallacy” believes that draws of one signal – a spot price shock in our case – increase the odds of next drawing other signals – that is, a spot price shock of the opposite sign. A natural consequence of this bias is a contrarian investment behaviour on behalf of speculators.

There is a large body in the financial markets literature modelling explicitly the existence of – both rational and irrational – contrarian investors to either explain puzzling empirical results or examine potential trading strategies (Lakonishok *et al*, 1994; Jegadeesh and Titman, 1995; Park and Sabourian, 2011). Regarding the commodity markets literature, Ellen and Zwinkels (2010) adopt a behavioural finance approach with heterogeneous speculators to explain the oil price dynamics. Despite some similarities regarding the investor composition, commodity demand in their framework is not derived explicitly through a structural economic model as in our case.

In practice, traders frequently form expectations about future market conditions and, in turn, devise investment strategies following simple technical analysis rules that are based on contrarian beliefs – which can be – influenced by behavioural biases such as the “gambler’s fallacy”. In particular, Kaniel *et al* (2008) provide evidence that numerous traders indeed select contrarian strategies while laboratory experiments, conducted by Bloomfield *et al* (2009), suggest that mainly uninformed investors usually adopt contrarian behaviour. What is more, Grinblatt and Keloharju (2000) show that, in Finnish markets, inexperienced investors frequently act as contrarians while more sophisticated ones tend to follow momentum strategies (Lof, 2015). Those findings are particularly related to our model since speculators correspond to financial investors who, as non-participants in the physical market, are assumed to be less sophisticated and informed regarding future shipping market conditions compared to traditional physical market agents.

Importantly, recall that Chapter 3 of this thesis concluded that to simultaneously match the empirical regularities related to the physical market for second-hand vessels, the average investor expectations regarding future market conditions must be – slightly extrapolative but – “near-rational”. In turn, note that the “average investor” of Chapter 3 corresponds to the “ship owner” agent type in Chapter 4. Furthermore, charterers can be plausibly assumed to form rational expectations since they participate in the physical market as well. Therefore, also the average physical investor expectations in Chapter 4 are expected to be “near-rational” as well. In the following, for simplicity and without loss of generality, we assume that they are totally rational.⁸⁸

⁸⁸ Note that it is straightforward to account for slightly extrapolative beliefs on behalf of ship owners in our framework. Even if we do so, however, the qualitative predictions and conclusions of our model are not affected.

Since, however, there are no surveys regarding shipping industry participants' beliefs and investment strategies as in the equity markets literature (Greenwood and Shleifer, 2014), we further justify our behavioural explanation by contradiction, that is, using both theoretical predictions and numerical simulations of the proposed framework. Note that a similar justification is followed in the model of Lof (2015) who motivates the presence of contrarian investors empirically by illustrating that the observed regularities can be more sufficiently approached when incorporating contrarian expectations on behalf of a population fraction.

Specifically, we begin by illustrating that a “fully-rational” model in its simplest form is not able to explain the documented empirical regularities. Then, we add the – time-varying – “hedging pressure” dimension and, in turn, examine the generated results. Since the “theory of storage” does not apply in this market, however, we cannot determine endogenously the hedging pressure variable. Thus, we impose a reasonable assumption to estimate it exogenously based on the corresponding spot market conditions. Having shown that neither this model can simultaneously generate the stylised facts, we incorporate the heterogeneous beliefs dimension to test whether we can qualitatively reproduce our findings. Accordingly, the simulation results suggest that, to simultaneously match all observed regularities sufficiently well, one must depart from the rational benchmark of the model since the hedging pressure dimension alone cannot capture the negative predictability of risk premia by lagged market conditions. While the predictions are not particularly sensitive to the degree of information asymmetry this is not true for the behavioural bias feature; namely, a fraction of investors must suffer from the “gambler’s fallacy” and, in turn, follow a contrarian investment strategy.

Consequently, our model lies on the intersection of empirical asset pricing and behavioural finance. To the best of our knowledge, the FFA market had never been examined from the perspective of a structural, heterogeneous beliefs, economic model before. In addition, we contribute to the generic commodity finance literature by incorporating explicitly the behavioural dimension in the formation of derivative contracts rates. Thus, the proposed framework could be further empirically evaluated in commodity markets for which exists microstructure data regarding the composition of traders.

The remainder of this chapter is organised as follows. Section 4.II describes briefly the FFA market and discusses the employed dataset. Section 4.III performs the empirical analysis and examines the main findings from an economic perspective. Section 4.IV presents the environment of our economy and develops the theoretical model. Accordingly, it provides the results from the numerical simulation of the model. Section 4.V concludes.

4.II. Data and Variables of Interest

In the context of this chapter, we examine the Capesize BCI 4TC and Panamax BPI 4TC dry bulk FFA contracts since – as illustrated in Chapter 1 of this thesis – they constitute by far the most liquid instruments. These contracts correspond to the equally weighted average of the four trip-charters of the Baltic Capesize Index and the Baltic Panamax Index, respectively. Furthermore, due to significantly higher data availability, we focus on the 1- and 2-month maturity contracts. In Appendix 4, however, we present some key empirical findings related to the 3- and 4-month maturities as well. On average, the volume related to these maturities is almost 50% of the total one in the respective

Table 4.1: Descriptive statistics for the variables of interest.

x	T	n	Start	End	\bar{x}	MD	SD	CV	Max	Min	ρ_1	ρ_2	ρ_{12}
Panel A: Variables in Levels (in '000 \$) for the Capesize Sector (BCI 4TC)													
Spot	-	116	1.07	8.16	37.2	15.8	48.4	1.30	222.8	0.9	0.93	0.85	0.45
Set	-	116	2.07	9.16	36.6	14.0	47.1	1.29	201.1	0.7	0.94	0.85	0.45
FFA1	1	116	1.07	8.16	37.0	16.7	46.9	1.27	201.3	1.6	0.95	0.87	0.47
FFA2	2	115	1.07	7.16	36.4	16.9	45.1	1.24	170.1	2.3	0.95	0.89	0.45
Panel B: Variables in Levels (in '000 \$) for the Panamax Sector (BPI 4TC)													
Spot	-	116	1.07	8.16	19.9	11.8	20.7	1.04	93.4	2.3	0.96	0.89	0.39
Set	-	116	2.07	9.16	20.1	11.4	20.8	1.03	86.1	2.6	0.96	0.90	0.40
FFA1	1	116	1.07	8.16	20.3	11.8	20.8	1.02	88.3	2.7	0.97	0.90	0.40
FFA2	2	115	1.07	7.16	20.7	11.5	20.7	1.00	87.3	3.3	0.96	0.91	0.41
Panel C: Variables in Log Differences for the Capesize Sector (BCI 4TC)													
b	1	116	1.07	8.16	0.09	0.06	0.24	-	0.75	-1.14	0.25	0.03	-0.04
Δs	1	116	2.07	9.16	0.02	0.01	0.37	-	1.63	-0.98	-0.08	-0.09	-0.08
r	1	116	2.07	9.16	0.07	0.03	0.32	-	0.80	-0.88	0.21	0.10	-0.02
b	2	115	1.07	7.16	0.13	0.07	0.39	-	1.15	-1.08	0.46	0.16	0.18
Δs	2	115	3.07	9.16	0.00	0.04	0.72	-	1.78	-2.33	0.28	-0.17	-0.02
r	2	115	3.07	9.16	0.13	0.12	0.60	-	2.39	-1.24	0.48	0.10	-0.08
Panel D: Variables in Log Differences for the Panamax Sector (BPI 4TC)													
b	1	116	1.07	8.16	0.05	0.02	0.13	-	0.52	-0.24	0.37	0.02	0.08
Δs	1	116	2.07	9.16	0.02	0.03	0.18	-	0.58	-0.48	-0.06	-0.10	0.01
r	1	116	2.07	9.16	0.03	0.00	0.16	-	0.72	-0.35	0.32	0.22	-0.11
b	2	115	1.07	7.16	0.09	0.02	0.23	-	0.84	-0.36	0.56	0.17	0.07
Δs	2	115	3.07	9.16	0.01	0.03	0.36	-	0.87	-1.44	0.37	-0.06	-0.05
r	2	115	3.07	9.16	0.08	0.06	0.35	-	1.73	-0.63	0.60	0.34	-0.21

Notes: Panels A-B present descriptive statistics for the levels of the spot, settlement, and FFA rates corresponding to the 1- and 2-month BCI 4TC and BPI 4TC FFA contracts. These variables (x) are expressed in thousand U.S. dollars. Panels C-D present descriptive statistics for the basis, b , the spot growth, Δs , and the risk premium, r , corresponding to the 1- and 2-month BCI 4TC and BPI 4TC FFA contracts. These variables (x) are expressed in log differences. The maturity of the contract and the number of observations are denoted by T and n , respectively. The first and last months of the variable in our

sample analysis are indicated by columns 4 and 5 (labelled “Start” and “End”), respectively (e.g., 1.07 refers to January 2007). The included statistics are the mean (\bar{x}), median (MD), standard deviation (SD), coefficient of variation (CV), maximum (max), minimum (min), 1-month (ρ_1), 2-month (ρ_2), and 12-month (ρ_{12}) autocorrelation coefficients.

contract. Since our empirical analysis focuses on the most liquid dry bulk FFA contracts, we believe that our findings are as robust as possible given the availability of data but also of high interest from an industry participant’s perspective.

Accordingly, our dataset consists of monthly observations on spot prices, settlement rates, and FFA rates for the BCI 4TC and BPI 4TC contracts with 1- and 2-month maturities,⁸⁹ obtained from the Baltic Exchange.⁹⁰ Incorporating the industry convention,⁹¹ settlement rates are estimated as the arithmetic average of the respective spot rates – of the corresponding trip-charter (TC) routes – over all trading days of the contract month. Furthermore, in line with Kavussanos and Nomikos (1999), we sample FFA rates and spot prices at the last trading day of each month.⁹² Table 4.1 summarises descriptive statistics related to these variables while Figure 4.1 illustrates the evolution of spot prices, settlement rates, and 1-month FFA rates. Note that the spot and settlement rates are the prices observed at issuance and maturity of the corresponding FFA 1-month contract, respectively. Moreover, Table 4.2 summarises the correlation coefficients among the variables of interest.

Table 4.2: Correlation matrix.


Variable	Levels				Log Differences
	Spot	Settlement	FFA1	FFA2	$Corr(\Delta s, r)$
Panel A: Capesize Sector (BCI 4TC)					
Spot	1.00				
Settlement	0.98	1.00			
FFA1	0.99	0.98	1.00		-0.77
FFA2	0.98	0.98	0.99	1.00	-0.84
Panel B: Panamax Sector (BPI 4TC)					
Spot	1.00				
Settlement	0.99	1.00			
FFA1	1.00	0.99	1.00		-0.72
FFA2	0.99	0.98	1.00	1.00	-0.80

⁸⁹ Since the 2-month FFA rates in the initial Baltic Exchange dataset start in September 2009, we extract the missing observations for the 2-month rates from the corresponding quarterly and 1-month contracts.


⁹⁰ The FFA rates are based on the Baltic Exchange Forward Assessments (BFA) which represent the mid-price of bids and offers for the dry bulk market, submitted and published every trading day at 17:30, London time.

⁹¹ This happens to ensure that settlement rates are neither subject to market manipulation on any given date nor vulnerable to extreme fluctuations due to the highly volatile nature of the industry (Alizadeh and Nomikos, 2010b).

⁹² For robustness, we have also performed the estimation procedure using the first trading day of each month. The corresponding results are both qualitatively and quantitatively very similar.



Panel A: BCI 4TC 1-month contract.



Notes: Panels A and B of this table correspond to the BCI 4TC and BPI 4TC FFA contracts, respectively. Columns 2-5 present the correlation coefficients for spot, settlement, and 1- and 2-month FFA rates. All these variables are in levels. The last column presents the corresponding correlation coefficients for the log spot growth, Δs , and the log risk premium, r , for the 1- and 2-month BCI 4TC and BPI 4TC FFA contracts. The latter two variables are expressed in log differences.

Figure 4.1: Spot, Settlement, and FFA Rates.

Panels A-B plot the evolutions of spot, settlement, and FFA rates from January 2007 to August 2016 for the 1-month BCI 4TC and BPI 4TC contracts, respectively. The spot and settlement rates are the prices observed at the issuance and maturity of the corresponding FFA 1-month contract, respectively. Accordingly, the first and last observations for the FFA and spot rates are in January 2007 and August 2016, respectively. For settlement rates, the first and last observations are in February 2007 and September 2016, respectively.

Figure 4.1 and the results presented in Table 4.1 suggest that all variables exhibit significantly volatile behaviour over time. Furthermore, while in the short-run variables appear to be highly persistent – as indicated by the 1- and 2-month autocorrelation coefficients – this persistence substantially decays as the horizon increases – as indicated by the 12-month autocorrelation coefficients. This result verifies the well-documented “boom-bust nature” of the shipping industry (Stopford, 2009). Importantly, note that this finding is in line with Chapters 2 and 3 of this thesis that examine the behaviour of shipping cash flows in the 1-year horizon.⁹³

Noticeably, spot, settlement, and FFA rates exhibit very similar descriptive statistics. In addition, as it becomes evident from Table 4.2 and Figure 4.1, these variables are extremely correlated. Taken together, these features suggest that these variables closely track each other. The fact, however, that – in contrast to settlement rates – both spot and FFA rates are \mathcal{F}_t -measurable yields two conjectures. First, it indicates that FFA rates are significantly affected by current physical market conditions. Second, it suggests that FFA rates have strong predictive power over the future settlement rates.

In line with the commodity markets literature (Fama and French, 1987), the difference between the FFA rate and the current spot price can be expressed as

$$F(t, T) - S(t) = S(t + T) - S(t) + F(t, T) - S(t + T), \quad (4.1a)$$

Panel B: BPI 4TC 1-month contract.

where $F(t, T)$ corresponds to the FFA rate at time t for a contract expiring in T

periods, $S(t)$ to the current spot price, and $S(t + T)$ to the settlement rate at maturity of the contract. As it is common in the empirical literature, we work with the log transformation of the variables of interest; hence, equation 4.1a can be re-expressed as

⁹³ Namely, while monthly spot rates appear to be non-stationary (Tables 4.2 and 4.4), annual cash flows exhibit a mean-reverting behaviour in the one-year horizon (Table 2.2). As a result, in the theoretical model of Chapter 3, where the time-period corresponds to one year, we assume a mean-reverting process for shipping cash flows. In contrast, in the theoretical model of this chapter, where the time-period corresponds to one month, we impose the random walk assumption for the spot rate process.

$$f(t, T) - s(t) = s(t + T) - s(t) + f(t, T) - s(t + T), \quad (4.1b)$$

where the lowercase letters correspond to the natural logarithm of the respective variable.

Since (4.1b) is an ex post identity, it also holds when incorporating the expectations operator conditional on any information set (Cochrane, 2011). Therefore, following Fama and French (1987), we can decompose the difference between the log FFA rate and the current log spot price into the sum of the expected change in the log spot – settlement – price and an expected premium over the log settlement price at maturity

$$f(t, T) - s(t) = E_t[s(t + T) - s(t)] + E_t[f(t, T) - s(t + T)], \quad (4.1c)$$


Note that this decomposition is equivalent to the log-linearisation approach followed in Chapter 2 of this thesis which is very common in the empirical asset pricing literature.⁹⁴ The quantity on the left-hand side of (4.1c) is termed the “(log) basis” of the FFA contract and constitutes a frequently incorporated valuation ratio. Furthermore, we define the term inside the first expectation on the left-hand side of (4.1c) as the “(log) spot growth” even if, strictly speaking, $s(t + T)$ corresponds to the settlement rate at maturity of the contract. Finally, the term inside the second expectation is defined as the “(log) risk premium”. This variable can be interpreted as the bias in the FFA rate as a forecast of the future settlement price or, equivalently, as the excess return for an investor who goes short on the FFA contract. From an economic point of view, equation (4.1c) illustrates that a high basis is a consequence of either high expectations about future log spot growth or/and high expectations of future log risk premia. Table 4.1 presents summary statistics related to these three variables for all contracts under consideration while Figure 4.2 illustrates their evolution in the 1-month horizon case.

To begin with, Table 4.1 suggests that FFA bases have on average been positive for both contracts. In addition, the mean basis strictly increases with the maturity of the contract. This implies that FFA rates have on average exceeded contemporaneous spot prices – recall that we have defined the basis as the log of the FFA rate minus the log of the respective spot price; this situation is often described by practitioners as “negative roll yield” (we further analyse this below). However, the standard deviations and 1-month autocorrelation coefficients in Table 4.1, along with Figure 4.2, suggest that FFA bases exhibit highly volatile behaviour over time; namely, all bases are characterised by significant mean-reversion. Therefore, to assess the statistical validity of this finding, we also estimate the respective t-statistics for the mean bases. Accordingly, as Table 4.3 indicates, there is strong statistical evidence of positive mean bases in all cases.


⁹⁴ For more on this topic, the reader can refer to Chapter 2 of this thesis and the broad literature by Campbell and Shiller, Cochrane, and Fama and French.

Furthermore, Table 4.1 provides evidence of positive mean risk premia in both contracts and across all horizons. According to our definition of the risk premium variable, this suggests that, on average, the FFA rate is higher than the corresponding realised settlement rate at maturity of the contract. Since also the realised risk premia are highly volatile, once again we estimate the corresponding t-statistics to assess the statistical significance of this result. Note that, due to the overlapping nature of observations, the estimated t-statistics are based on Newey-West corrected standard errors (Szymanowska *et al*, 2014).

Noticeably, the magnitudes of these premia are very high, implying annualised mean returns of above 30% in all cases (Table 4.3). More importantly, in the 1-month horizon these premia are also statistically higher than zero. Therefore, on average, 1-month FFA rates appear to provide an upwardly



Panel A: BCI 4TC 1-month contract.



Panel B: BPI 4TC 1-month contract.

Figure 4.2: FFA Bases, Spot Growth Rates, and Risk Premia.

Panels A-B plot the evolutions of the basis, spot growth, and risk premium variables for the 1-month BCI 4TC and BPI 4TC contracts, respectively. The sample runs from January 2007 to August 2016. All variables correspond to log differences.

Table 4.3: Significance of FFA bases and risk premia.

T	n	Mean Basis	t of Basis	An. Mean Premium	An. SD Premium	t^{NW} of Premium
Panel A: Capesize Sector (BCI 4TC)						
1	116	9.02%	4.07	79.63%	111.60%	2.22
2	115	12.90%	3.58	76.58%	146.59%	1.43
Panel B: Panamax Sector (BPI 4TC)						
1	116	4.83%	4.07	31.98%	56.85%	1.75
2	115	8.88%	4.20	49.95%	86.01%	1.25

Notes: This table presents descriptive statistics related to FFA bases and risk premia for the Capesize BCI 4TC and Panamax BPI 4TC 1- and 2-month contracts. The maturity of the contract and the number of observations are denoted by T and n , respectively. The included statistics are the mean and t-statistic of the basis and the annualised mean, standard deviation, and t-statistic, t^{NW} , of the risk premium. To deal with the overlapping nature of risk premia, the corresponding t-statistics are estimated using the Newey-West (1987) HAC correction. When the t-statistic indicates significance at least at the 10% level, the respective mean statistic appears in bold.

biased measure of future settlement rates (that is, there is statistical evidence of “contango”). In the 2-month horizon, however, while both coefficients are positive, they are not statistically significant. From an investor’s perspective, these results suggest that a FFA trader who does not participate in the physical market – that is, a “non-hedger” or “speculator” as is commonly termed in the commodity markets literature – should rather take the short position in the FFA market.

In conclusion, in the dry bulk FFA market there is neither economic nor statistical evidence of “normal backwardation”. This finding is of great importance since it contrasts with many commodity markets where futures prices are set at a premium to expected future spot prices. According to the prevailing conjecture in the literature, this happens in order for the long side of the futures agreement – which in most of the respective cases is assumed to be taken by “non-hedgers” – to be compensated for providing price insurance to the commodity producers (Gorton *et al*, 2012). Thus, following this argument, our results suggest that it is the short side of the FFA agreement the one being– on average – rewarded for providing price insurance. A natural explanation for this finding would be that there exists on average higher demand for the long position on the FFA contract (mainly from long hedgers) compared to the short one in this market. However, since we don’t have data on the investor composition, we cannot further examine this argument. Furthermore, those

stylised facts combined appear to verify the common view of practitioners that a positive basis – or, equivalently, a negative “roll yield”⁹⁵ – is a requirement for the existence of a positive risk premium to a short position in futures – forward – markets (Gorton *et al*, 2012). In the next section, however, we illustrate formally that only in the Panamax sector – and not in the Capesize one – there is documented a strong positive relationship between the FFA basis and the corresponding risk premium.

A third noticeable stylised fact presented in Table 4.3 is that the 1- and 2-month risk premia appear to be moderately positively autocorrelated in both contracts – as indicated by the columns labelled ρ_1 and ρ_2 for the 1- and 2-month contracts, respectively. This autocorrelation, however, is attenuated as the horizon and the maturity of the contract increase. Importantly, as analysed in the following section, this feature indicates the existence of a momentum effect. Moreover, Figure 4.2 suggests that risk premia and spot growth rates are substantially negatively correlated. This observation is further validated by the respective correlation coefficients (Table 4.2). A straightforward explanation for this feature is that an unexpected positive (negative) shock in spot rates will result in a negative (positive) realised risk premium. Finally, we observe that spot growth rates exhibit high volatility; namely, in all cases, the spot growth standard deviation is higher than both the respective basis and risk premium ones. As discussed in the following, this feature is closely related to – can explain – the obtained variance decomposition results. From an economic perspective, the high volatility and, in turn, the uncertainty regarding spot market conditions justifies the existence of the FFA market as a hedging instrument for physical market participants but also attracts the trading interest of investors outside the shipping markets, such as hedge funds and investment banks (Alizadeh and Nomikos, 2009).

4.III. Predictable Variation in the FFA Market

As it is well-analysed in the literature, futures and forward markets should, ideally, serve two significant social roles (French, 1986). First, they should act as a hedging instrument for participants in the physical market – the risk-transfer role – and, second, they should provide accurate forecasts of expected spot prices – the price-discovery role. While the first role is unambiguous, much controversy has been concentrated around the second one since the documented empirical results depend on both the market and the period under consideration but also on the incorporated econometric framework. Thus, we begin this section by examining the forecasting power of FFA rates regarding future physical market conditions.

⁹⁵ Practitioners define the roll yield as the ratio of the spot price over the contemporaneous futures contract rate. Therefore, it is equivalent to the inverse of the basis definition adopted in this chapter.

In addition, we test whether there exists significant predictability of returns – that is, risk premia – in the dry bulk FFA market. While this question has been analysed thoroughly in the financial and commodity markets empirical asset pricing literatures, we are the first – to the best of our knowledge – to examine it explicitly in the FFA market. Importantly, apart from the FFA market participants' perspective (e.g., for developing potential trading strategies), this research question appears to be of high interest also from an economist's point of view. Namely, according to Fama (1991), future asset returns should not be predicted by the current information filtration. Equivalently, if FFA markets are efficient, FFA risk premia should not be predicted by \mathcal{F}_t -measurable variables. Accordingly, by incorporating a large set of both economic and financial predictors, we address this question.

4.III.A. Predictability of Future Market Conditions and Risk Premia from the Basis

We begin by examining and quantifying the forecasting ability of the FFA market for ships and, in particular, of the BCI 4TC and BPI 4TC contracts. To this end, this subsection applies the variance decomposition framework (Fama, 1984a and 1984b; Fama and French, 1987). Note that, to the best of our knowledge, we are the first to apply this framework to the FFA market. While Kavussanos and Nomikos (1999) have already examined this question using cointegration techniques, our empirical analysis aims to fill certain gaps in the literature.

First, the incorporated sample corresponds to the most recent available data – including the extreme shipping cycle of the period 2008 to 2010 – regarding the futures and forward shipping markets. Second, and most important, the incorporated estimation procedure allows us not only to quantify the predictive power of the FFA contracts but also to provide an economic interpretation for the results. Namely, we explain our findings by performing a comparison with the results obtained from other commodity futures and forward markets – after applying the same framework.⁹⁶ Specifically, Hazuka (1984), French (1986), and Fama and French (1987) show that, among other factors, the forecasting ability of futures contracts is directly related to the importance of “seasonals” in supply and demand as well as the storage cost of the commodity.

In the following, we illustrate how their arguments can be extended to shipping where the corresponding commodity is a service. In addition, this decomposition allows us to quantify precisely the FFA basis variation that can be attributed to expectations about future market conditions and risk premia. Finally, in contrast to the cointegration approach where the estimates are highly

⁹⁶ Apart from commodity markets, the variance decomposition framework has also been applied to forward exchange rates and forward interest rates markets. Due to the significantly different economic principles that characterise these markets, however, we cannot directly compare the respective findings to the ones related to commodity markets.

sensitive to the specification of the Vector Error Correction Model (Kavussanos and Nomikos, 1999), the results from the variance decomposition framework are robust since the model cannot be misspecified.

4.III.A.i. A Simple Variance Decomposition Framework

As illustrated in Section 4.II, all FFA bases exhibit highly volatile behaviour. Following Cochrane (2011), it is straightforward to decompose the variance of the basis into two parts. Namely, multiplying both sides of (4.1b) by $f(t, T) - s(t) - E_t[f(t, T) - s(t)]$ and taking expectations yields

$$\begin{aligned} \text{var}[f(t, T) - s(t)] &= \text{cov}[f(t, T) - s(t), s(t + T) - s(t)] \\ &\quad + \text{cov}[f(t, T) - s(t), f(t, T) - s(t + T)]. \end{aligned} \quad (4.2a)$$

Therefore, the variance of the basis is exactly equal to the covariance between the basis and the future spot growth plus the covariance between the basis and the future risk premium. Dividing both sides of (4.2a) by the basis variance yields

$$\begin{aligned} \frac{\text{cov}[f(t, T) - s(t), s(t + T) - s(t)]}{\text{var}[f(t, T) - s(t)]} + \frac{\text{cov}[f(t, T) - s(t), f(t, T) - s(t + T)]}{\text{var}[f(t, T) - s(t)]} &= 1 \\ \Rightarrow \beta_{\Delta s, T} + \beta_{r, T} &= 1, \end{aligned} \quad (4.2b)$$

where $\beta_{i, T}$ is the T -period contract coefficient corresponding to the i^{th} element of the decomposition. Incorporating (4.1b) in (4.2b), these two coefficients are further analysed into

$$\beta_{\Delta s, T} = \frac{\text{var}[s(t + T) - s(t)] + \text{cov}[s(t + T) - s(t), f(t, T) - s(t + T)]}{\text{var}[f(t, T) - s(t)]} \quad (4.2c)$$

and

$$\beta_{r, T} = \frac{\text{var}[f(t, T) - s(t + T)] + \text{cov}[s(t + T) - s(t), f(t, T) - s(t + T)]}{\text{var}[f(t, T) - s(t)]}. \quad (4.2d)$$

The expressions above suggest that, apart from the individual variances of the spot growth and the risk premium, the variance of the basis and, in turn, the two slope coefficients also consist of a cross-term; that is, the covariance between those two components. Consequently, the relation between spot growth and risk premium has a major implication for this variance decomposition. Notice, though, that the contribution of this covariation to the variance of the basis is equally split

between the spot growth and risk premium coefficients (Fama, 1984a). Hence, $\beta_{\Delta s,T}$ and $\beta_{r,T}$ still contain the proportion of basis variance attributed to the variance of spot growth and risk premium, respectively. Thus, we can directly examine which of these two sources is the major determinant of the observed basis volatility by running forecasting OLS regressions in the spirit of Fama (1984a and 1984b), Fama and French (1987), and Cochrane (2011). Namely, we regress future log spot growth and future log risk premia on the current log basis:

$$s(t+T) - s(t) = \alpha_{\Delta s,T} + \beta_{\Delta s,T} \cdot [f(t,T) - s(t)] + \varepsilon_{\Delta s,t+T}, \quad (4.3a)$$

$$f(t,T) - s(t+T) = \alpha_{r,T} + \beta_{r,T} \cdot [f(t,T) - s(t)] + \varepsilon_{r,t+T}. \quad (4.3b)$$

In line with Fama and French (1987), statistical evidence that $\beta_{\Delta s,T}$ is positive means that the basis at t has forecasting power regarding the future change in the spot price which, in turn, implies that the FFA contract is a reliable predictor of the future spot rate. Statistical evidence that $\beta_{r,T}$ is different than zero implies that the basis at t has forecasting power regarding the future premium realised at T . Notice that equations 4.1c and 4.2b impose the restrictions $\alpha_{\Delta s,T} + \alpha_{r,T} = 0$, $\varepsilon_{\Delta s,t+T} + \varepsilon_{r,t+T} = 0$, and, most importantly, $\beta_{\Delta s,T} + \beta_{r,T} = 1$. The last equation implies that regressions 4.3a and 4.3b will always allocate all basis variation to either the expected spot growth or the expected risk premium or some combination of the two; thus, in analogy to the variance decomposition in Chapter 2, (4.3a) and (4.3b) examine a question of relative predictability through the magnitudes of the two slope coefficients, $\beta_{\Delta s,T}$ and $\beta_{r,T}$.

Table 4.4: Phillips-Perron unit root test.

Variable	Levels		Log Differences		
	Settlement	FFA	Basis	Growth	Risk Premium
Panel A: Capesize Sector (BCI 4TC)					
1-month Contract	-2.205 [0.206]	-1.811 [0.374]	-8.197 [0.000]	-14.996 [0.000]	-8.387 [0.000]
2-month Contract	-2.205 [0.206]	-1.593 [0.483]	-6.041 [0.000]	-8.994 [0.000]	-5.698 [0.000]
Panel B: Panamax Sector (BPI 4TC)					
1-month Contract	-1.845 [0.357]	-1.748 [0.405]	-7.189 [0.000]	-11.318 [0.000]	-7.780 [0.000]
2-month Contract	-1.845 [0.357]	-1.430 [0.565]	-5.471 [0.000]	-6.902 [0.000]	-5.361 [0.000]

Notes: Panels A-B report results from the Phillips-Perron (1988) unit root test for the Capesize BCI 4TC and Panamax BPI 4TC contracts, respectively, for the 1- and 2-month horizons. Namely, we test the null hypothesis that the series are non-stationary. The series of interest are the settlement and FFA rates, the basis, the spot growth, and the risk premium (all

expressed in logs). For each series, we present the adjusted t-statistics and the exact significance level in square brackets [·].

Before performing those regressions however, we examine formally whether the variables of interest satisfy the necessary stationarity condition. Table 4.4 presents the results from the Phillips-Perron (1988) unit root test that examines the null hypothesis of non-stationarity. Evidently, the null hypothesis is rejected for the three variables, for both contracts and maturities; therefore, all incorporated variables are $I(0)$.⁹⁷ Accordingly, we perform the predictive regressions (4.3) for the Capesize BCI 4TC and Panamax BPI 4TC contracts, for both maturities. In line with the existing literature, for the 2-month maturity contracts we incorporate Newey-West (1987) heteroskedasticity and autocorrelation consistent (HAC) standard errors to deal with the overlapping nature of risk premia and growth rates.

As it becomes evident from Table 4.5, all spot growth coefficients are noticeably large in magnitude while the signs are positive in every sector and horizon; hence, they are consistent with equation 4.1c. What is more, the respective t-statistics indicate significance at the 1% level in every case. Finally, the R^2 s of growth regressions are at least 14%. Therefore, the forecasting power of the log basis regarding future spot growth appears to be strong. Since this framework examines a question of relative predictability, we also compare the results obtained from the growth regressions to the respective findings from the risk premia ones. As Table 4.5 suggests, in both sectors, the slope coefficients and the respective t-statistics from the risk premia regressions are significantly smaller in magnitude compared to the ones from the growth regressions. In addition, the R^2 s of premia regressions are below 10% in all cases. Arguably, therefore, the bulk of variation in the FFA basis can be attributed to variation in expected spot growth and not to time-varying expected risk premia.

Table 4.5: Regressions of future risk premia and spot growth on current FFA bases.

Variable	T	n	α	t^{NW}	β	t^{NW}	R^2
Panel A: Capesize Sector (BCI 4TC)							
Δs	1	116	-0.05	-1.52	0.80***	6.39	0.26
r	1	116	0.05	1.52	0.20	1.61	0.02
Δs	2	115	-0.13	-1.40	1.04***	7.83	0.31
r	2	115	0.13	1.40	-0.04	-0.28	0.00
Panel B: Panamax Sector (BPI 4TC)							
Δs	1	116	-0.01	-0.57	0.63***	5.48	0.21
r	1	116	0.01	0.57	0.37***	3.20	0.08
Δs	2	115	-0.05	-1.04	0.59***	2.84	0.14

⁹⁷ This finding is also verified by the results from the Augmented Dickey-Fuller (ADF) test.

r	2	115	0.05	1.97	0.41*	1.97	0.07
-----	---	-----	------	------	-------	------	------

Notes: Panels A-B report results from 1- and 2-month horizon OLS regressions of future spot growth, Δs , and risk premia, r , on the current basis for the Capesize BCI 4TC and Panamax BPI 4TC contracts, respectively. To deal with the overlapping nature of the variables, t-statistics are estimated using the Newey-West (1987) HAC correction. The maturity of the contract and the number of observations are denoted by T and n , respectively. The intercept, α , and the slope coefficient, β , are accompanied by *, **, or *** when the absolute t^{NW} statistic indicates significance at the 10%, 5% or 1% level, respectively.

In a cross-sector comparison, we observe that the magnitudes of the Capesize risk premia slope coefficients are significantly smaller than the spot growth ones while there does not appear to be statistically significant predictability of risk premia from the basis. In contrast, in the Panamax sector, all slope coefficients are positive and, thus, consistent with (4.1c) while their magnitudes are around 0.4 for both maturities. What is more, we observe that for both maturities there appears to be strong statistical evidence of future risk premia predictability from the FFA basis. Noticeably, the 1-month horizon slope coefficient is statistically significant at the 1% level. Hence, this stylised fact indicates that, in short horizons, a high BPI 4TC FFA basis positively predicts the corresponding future risk premium. This finding is worth emphasising also for the following reason.

Namely, while the variance decomposition framework examines a question of relative predictability it does not impose any restrictions on either the spot rate process or the rationality of expectations. Specifically, Fama and French (1987) argue that any potential irrational forecasts of future spot prices are allocated by the regression 4.3b; equivalently, an irrational forecast of the future spot price in the futures/forward price will appear as a time-varying risk premium, that is, as a non-zero value of $\beta_{r,T}$. Therefore, Fama and French do not exclude the existence of irrational forecasts as a potential explanation for time-varying risk premia in futures/forwards markets; nevertheless, they do not examine it.

4.III.A.ii. Interpretation of the Results and Comparison to Other Markets

Our variance decomposition results clearly suggest that there exists strong predictability of future spot price changes from the FFA basis. In turn, this implies that FFA rates exhibit substantial forecasting ability regarding future spot rates. This finding is important since there is a long-standing debate in asset pricing regarding the forecasting ability of futures and forward markets. As analysed above, futures rates in many markets do not appear to possess statistically significant forecasting power while, in some cases, they do not even provide better forecasts compared to the current spot price. Having demonstrated that the former is not the case in the FFA market, we now also show

that the FFA rates are significantly better predictors of future market conditions compared to the current spot rates.

Accordingly, Table 4.6 compares the results from regressions of future spot growth on the first lag of the 1-month spot growth to the ones obtained from regressing future spot growth on the current FFA basis, that is, predictive regression 4.3a. Evidently, irrespective of the maturity and the sector under consideration, the magnitudes of the slope coefficients, the t-statistics, and the R^2 s of the bases regressions are significantly higher than the ones from the respective lagged spot growth regressions.⁹⁸ What is more (as expected from Tables 4.1 and 4.4), lagged spot growth does not have any significant forecasting power regarding future market conditions. Therefore, we can argue that FFA rates contain substantially more information compared to the concurrent spot and settlement rates.

From an economic point of view, the most interesting questions are, first, why do we obtain these results in shipping and, second, how can they be related to the ones from other commodity markets. Since these two questions are interrelated, however, we examine them in conjunction. To begin with, it is fruitful to restate French (1986) who argues that “if the current spot price equals the expectation of the future spot price, the futures price cannot provide a better forecast of the future spot price. Equivalently, the futures market cannot predict changes in the spot price unless the spot price is expected to change”. In simple words, for futures markets to be able to forecast future spot rates there must be something to be predicted. While this statement appears to be trivial, it is very subtle and important for the interpretation of the results related to forecasting questions and frameworks of this type (recall also the analysis in Chapter 2 of this thesis).

Table 4.6: Regressions of future spot growth on lagged spot growth and current FFA basis.

<i>T</i>	Lagged Spot Growth				FFA Basis			
	<i>n</i>	β	t^{NW}	R^2	<i>n</i>	β	t^{NW}	R^2
Panel A: Capesize Sector (BCI 4TC)								
1	115	-0.08	-0.83	0.01	116	0.80***	6.39	0.26
2	114	-0.24	-1.31	0.01	115	1.04***	7.83	0.31
Panel B: Panamax Sector (BPI 4TC)								
1	115	-0.06	-0.62	0.00	116	0.63***	5.48	0.21
2	114	-0.10	-0.46	0.00	115	0.59***	2.84	0.14

Notes: Panels A-B report 1- and 2-month horizon OLS forecasting regressions of future spot growth, Δs , on one period lagged 1-month spot growth and the current basis for the Capesize BCI 4TC and Panamax BPI 4TC contracts, respectively.

⁹⁸ This argument is also verified by the results of bivariate forecasting regressions using both current basis and lagged spot growth as the explanatory variables and future spot growth as the explained one. These results can be provided by the author upon request.

Spot growth is defined as the log of the ratio of the settlement rate to the spot price at the end of the previous month. To deal with the overlapping nature of returns and growth rates, t-statistics are estimated using the Newey-West (1987) HAC correction. The maturity of the contract and the number of observations are denoted by T and n , respectively. The slope coefficient, β , is accompanied by *, **, or *** when the absolute t^{NW} statistic indicates significance at the 10%, 5% or 1% level, respectively.

Therefore, one should expect that in markets where – the realised – spot rates exhibit significantly volatile behaviour there will be strong predictability of future spot rates from the futures contracts. Specifically, if investors know – up to a certain degree – the underlying economics of the market, they will be able to predict – up to a certain degree – the future spot rate; in the presence of futures markets, however, the expectation of spot rates is reflected – at least partially – on futures rates. *Ceteris paribus*, in such cases, spot rate volatility results in futures rate volatility. Consequently, in line with French (1986) and Fama and French (1987), futures and forward prices cannot provide reliable forecasts of future spot rates unless, on one hand, realised spot price changes exhibit substantially volatile behaviour and, on the other hand, the variance of expected spot price changes, as quantified by the basis, is comparable to – that is, of the same order of magnitude as – the one of the realised spot price changes.

Importantly, while the variance of the spot rate depends on the economics of the physical market under consideration, the variance of the basis depends also on the risk preferences, objectives, and beliefs – or, equivalently, expectations – of the participants in the derivative market. These three factors can significantly affect the formation of the futures rate and, in turn, the futures basis and its variance. Hence, in certain cases, even if the physical market is characterised by highly volatile cash flows, the futures basis can be a biased predictor of the expected spot rate. We relegate this discussion, however, to Section 4.IV of this chapter. In conclusion, from a statistical perspective, the necessary condition for the existence of spot growth predictability from the futures basis is that both variables are highly volatile.

More important, however, are the underlying economic principles related to this statistical observation. Namely, French (1986) shows that the forecasting power of the basis is an increasing function of seasonals in supply of and demand for the commodity and of the commodity cost of storage. Furthermore, Hazuka (1984) and Fama and French (1987) verify the direct relationship between the “theory of storage” and predictability of future spot rates for a variety of commodities. Specifically, the theory suggests that for commodities that are sensitive to supply and demand (seasonal) shocks, the degree of predictable variation in future spot prices should be an increasing function of the cost of storage – or, equivalently, a decreasing function of the inventory level.

The reason is that inventories tend to smooth predictable adjustments in spot prices in response to these shocks and, thus, tend to reduce the volatility of both realised and expected spot rates.

Since high storage costs relative to the commodity value deter storage, they also reduce the degree of spot price smoothing and, in turn, increase the amount of predictable spot price variation. As a result, for commodities characterised by high storage costs relative to value – that is, for commodities that are non-storable due to either perishability (e.g., animal products as broilers, eggs, hogs, live cattle, and pork bellies) or/and volume (e.g., cotton, oats, soybeans, and soymeal) – the respective futures prices exhibit significant forecasting power. In contrast, for commodities with low storage costs relative to value – that is, for storable commodities such as precious metals (e.g., gold, silver, and platinum) – futures prices are not informative regarding future market conditions.

Following this analysis, the results obtained from the FFA market should be a priori expected. First, from a statistical perspective, we observe that the necessary conditions stated by French (1986) are certainly met in the dry bulk FFA market. Namely, as Table 4.1 indicates, realised spot growth rates are highly volatile and, furthermore, FFA bases' volatility is comparable to the one of spot growth; note that the ratio of the basis standard deviation to the respective one of spot growth ranges from 0.54 to 0.72.

Second, from an economic perspective, it is well-documented that the shipping industry is highly sensitive to supply and demand shocks. Specifically, the notorious boom-bust shipping cycles – generated by the inelastic character of the exogenous demand for shipping services combined with the inelastic (highly elastic) supply of vessels in the short-run (long-run) – result in very volatile cash flows. Due to the nature of the industry, however, and the characteristics of physical hedgers,⁹⁹ future spot rates can on average be predicted – up to a certain degree – based on the time t public information filtration and/or investors' private information.¹⁰⁰ Accordingly, FFA rates are expected to reflect – up to a certain level – the economic predictions of market participants.

In contrast, if future spot rates could not be predicted using \mathcal{F}_t -measurable economic variables, the FFA basis at time t would have no forecasting power about future market conditions. In line with the analysis in Chapter 2 of this thesis, in the applied variance decomposition methodology, the FFA basis is the sole state variable. Thus, it is assumed to be summarising the time t information filtration; that is, the historical and prevailing market conditions (Fama and French, 1988a). As Campbell and Shiller (1988a) argue, while we cannot observe everything that shipping agents do,

⁹⁹ In line with Chapter 3, there is a large number of established shipping companies that operate in the industry. In some instances, these firms have been present in the market for more than a century (Stopford, 2009); consequently, they have strong prior experience and expertise about the key supply and demand drivers of the shipping industry. As a result, they can perform “near-rational” and, thus, relatively accurate forecasts about future market conditions. An analogous argument holds for large trading houses – that is, charterers.

¹⁰⁰ As analysed in the Introduction of this chapter, since ship owners and charterers participate also in the physical market they are expected to be better informed – that is, to have “inside” information regarding the actual future market conditions – than potential investors who trade only in the FFA market.

fortunately, we observe the FFA basis which summarises the market's relevant information. Of course, as mentioned above, the degree of FFA rates' prediction accuracy depends also on other factors among which the risk preferences, objectives, and rationality of beliefs of the FFA market participants.

Moreover, the relation between the "theory of storage" and the FFA results is straightforward. We analysed above that, in commodity markets, predictability of spot rates appears to be an increasing function of the commodity cost of storage. Equivalently, the "more storable" the commodity is the lower the predictability of future spot rates is expected to be. In shipping, however, the commodity is a service, thus, a non-storable one.¹⁰¹ Therefore, the fact that the industry is subject to significant supply and demand shocks which cannot be attenuated through adjustments of the short-term supply – the reader can parallelise this to a lack of inventory and, thus, lack of spot price smoothing – results in predictable variation of spot rates and, in turn, in substantial forecasting ability of FFA rates.

As illustrated in Chapter 2, the arguments presented above apply, not only to futures and forward contracts but also, to a variety of non-derivative – both financial (e.g., stocks) and real (e.g., real estate and vessels) – assets. Namely, Chen *et al* (2012) and Rangvid *et al* (2014) show that, in equity markets, cash flow predictability by valuation ratios (where, in this case, the dividend yield is the most frequently incorporated measure) is positively related to cash flow volatility; hence, inversely related to the degree of dividend smoothing. Therefore, we can relate the role of dividend smoothing in equity markets to inventories and the cost of storage in the commodity ones. More importantly, our variance decomposition results in Chapter 2 clearly indicate that vessel valuation ratios (namely, the earnings yield) have strong predictive power over future market conditions. Therefore, our physical – real asset – and FFA markets results are in line and, in turn, can be generalised to the entire dry bulk shipping industry. In addition, the analyses conducted in Chapters 2 and 4 reinforce and extend the economic justification in the related commodity markets literature.

Finally, recall that we examine a question of relative predictability: since $\beta_{\Delta s,T} + \beta_{r,T} = 1$ basis variation must be due to either predictability of future risk premia or/and predictability of future spot growth. Accordingly, the fact that spot growth changes are more volatile than the respective risk premia ones predisposes us for the allocation of the basis variability – that is, through equations 4.2c and 4.2d. In conclusion, in line with Chapter 2 and the analysis above, we argue that FFA bases move *mainly* due to expectations about future spot growth because the latter can be predicted – up

¹⁰¹ We have formally tested whether the theory of storage holds in the shipping industry by regressing the FFA basis on the nominal interest rate. Following Fama and French, the storage equation hypothesis is that the slope coefficient of the regression should be equal to one for any continuously storable commodity. The obtained coefficients in our case, however, are negative and statistically insignificant; thus, the theory is rejected. The corresponding results can be provided by the author upon request.

to a certain degree – by market agents at time t through the shipping supply and demand mechanism. For the *residual* proportion of basis variability in the FFA market, however, that is, the one attributed to time-varying risk premia, there can be two plausible economic justifications; a “rational” and an “irrational” one.

Regarding the former, there exist two – usually interconnected (Gorton *et al*, 2012; Ekeland *et al*, 2016) – “rational” theories for the existence of risk premia predictability in the commodity markets literature; namely, the “theory of storage” and the “theory of normal backwardation”. These theories justify the predictability of risk premia through the existence of – inventories which, in turn, result in – time-varying hedging pressure (usually on the part of commodity producers). Regarding the latter explanation, as analysed previously and in line with Fama and French (1987), the variability of the risk premia component can be attributed to irrational forecasts of future market conditions. In Section 4.IV, we illustrate formally why the latter explanation appears to be more plausible in the FFA market. Accordingly, incorporating the “irrational” explanation, the fact that the bulk of volatility is attributed to spot growth changes – and a smaller proportion to time-varying risk premia – implies that, while distorted expectations can justify the observed bias, the “average degree of expectations’ irrationality” is not extreme.

4.III.B. Predictability of Risk Premia from Lagged Risk Premia and Spot Market Indicators

As we illustrated in Subsection 4.III.A, predictability of future risk premia from the basis appears to depend on both the sector of the industry and the maturity of the contract. Specifically, in the Panamax BPI 4TC contracts there is statistically significant predictability of future risk premia – which appears to be stronger in the 1-month horizon. The fact, however, that in the Capesize BCI 4TC contracts there is no strong statistical evidence of time-varying risk premia in the formation of the basis does not imply either that expected premia are zero or that future risk premia cannot be predicted in general.

Regarding the first argument, regressions 4.3a and 4.3b are designed to detect variation in expected risk premia; hence, failure to identify time-varying expected premia does not imply that expected premia are zero (Fama and French, 1987). Indeed, as analysed in Section 4.II, there appears to be statistical evidence of positive mean risk premia in both contracts (Table 4.3). Regarding the second argument, the remainder of this subsection examines the predictive power of an additional, frequently incorporated price-based signal – namely, the lagged risk premium – and two spot market indicators – that is, the lagged spot growth and the lagged Baltic Dry Index (BDI). Specifically, the former predictor aims to examine whether there exists a momentum effect in risk premia while the latter two whether – recent changes in – physical market conditions forecast future risk premia.

To begin with, Table 4.7 summarises the results from regressions of 1- and 2-month risk premia on past realisations of the variable. The first three rows of each panel present the results from bivariate regressions where the lagged one-month risk premium is the predictor; that is, in the first row the regressor is the first lag of the risk premium variable related to the 1-month contract, in the second row is the second lag, and so on and so forth up to the third row. In the fourth row, the regressor is the corresponding previously *realised risk premium for each contract*; that is, for the 2-month contract expiring in $t + 2$ months, the predictor is the realised risk premium related to the 2-month contract that expired at t .¹⁰²

The figures in Table 4.7 indicate that there exists statistically significant predictability of future risk premia from lagged realisations of the variable in both contracts. Specifically, both the 1- and 2-month risk premia can be strongly positively forecasted by the first lag of the 1-month risk premium. Figure 4.3 depicts this positive relation for the 1-month maturity. Notably, in the Panamax sector, the slope coefficients are significant at the 1% level and, also, the second lag of the 1-month risk premium strongly positively predicts the 1- and 2-month risk premia – at the 5% and 1% levels of significance, respectively. Moreover, the 2-month BPI 4TC risk premium can be positively predicted – with statistical significance at the 1% level – also from the first lag of the 2-month premium. Notice that in both contracts – when incorporating lags of the 1-month risk premium as regressors – the magnitudes of the statistically significant coefficients increase with the horizon of the contract. More importantly, when we use higher lags as regressors, the values of the slope coefficients strictly decrease and become less significant.

Table 4.7: Regressions of future risk premia on lagged risk premia.

Variable	$f(t, 1) - s(t + 1)$				$f(t, 2) - s(t + 2)$			
	n	β	t^{NW}	R^2	T	β	t^{NW}	R^2
Panel A: Capesize Sector (BCI 4TC)								
$f(t - 1, 1) - s(t)$	115	0.21**	2.29	0.04	114	0.32*	1.74	0.03
$f(t - 2, 1) - s(t - 1)$	114	0.10	1.01	0.01	113	0.19	1.20	0.01
$f(t - 3, 1) - s(t - 2)$	113	0.00	0.04	0.00	112	-0.25	-1.64	0.02
$f(t - T, T) - s(t)$	115	0.21**	2.29	0.04	113	0.10	1.10	0.01
Panel B: Panamax Sector (BPI 4TC)								
$f(t - 1, 1) - s(t)$	115	0.32***	3.62	0.10	114	0.72***	2.66	0.11
$f(t - 2, 1) - s(t - 1)$	114	0.22**	2.36	0.05	113	0.60***	1.92	0.08
$f(t - 3, 1) - s(t - 2)$	113	0.15	1.62	0.02	112	0.21	1.18	0.01
$f(t - T, T) - s(t)$	115	0.32***	3.62	0.10	113	0.34***	3.69	0.12

Notes: Panels A-B report 1- and 2-month horizon OLS forecasting regressions of future risk premia, $f(t, T) - s(t + T)$, on lagged risk premia, for the Capesize BCI 4TC and Panamax BPI 4TC contracts, respectively. Namely, in the first three rows of

¹⁰² Note that, for the one-month contract, the first and fourth rows of the respective panel coincide.

each panel the predictor is the lagged one-period risk premium, $f(t-l, 1) - s(t-l+1)$; that is, the lagged risk premium related to the one-month contract where the number of lags, l , varies from 1 to 3. In the fourth row, the predictor is the corresponding previous risk premium for each contract, $f(t-T, T) - s(t)$; e.g., for the 2-month contract expiring in $t+2$ months, the predictor is the realised risk premium related to the two-month contract that expired at t . Note that, for the 1-month contract, the first and fourth rows of the respective panel coincide. The maturity of the contract and the number of observations are denoted by T and n , respectively. To deal with the overlapping nature of the variables, t-statistics are estimated using the Newey-West (1987) HAC correction. The slope coefficient, β , is accompanied by *, **, or *** when the absolute t^{NW} statistic indicates significance at the 10%, 5% or 1% level, respectively.

Thus, the results from these predictive regressions indicate that there is significant positive predictability of future risk premia from – at least – the first lag of the 1-month risk premium. This positive predictability is substantially more robust in the Panamax sector. From an economic point of view, this finding indicates the existence of a momentum effect in risk premia. Namely, a high realised risk premium appears to forecast high future premia or, equivalently, a high realised excess return from a position on a FFA contract positively predicts future short-term excess returns from taking the same position on the analogous contract. More importantly, the fact that this sort of predictability is





Figure 4.3: Risk Premia, Lagged Risk Premia, and Lagged Spot Growth.

Panels A-B plot the evolutions of risk premia, lagged risk premia, and lagged spot growth, for the 1-month BCI 4TC and BPI 4TC contracts, respectively. All variables correspond to log differences. The sample runs from February 2007 to August 2016. Lagged risk premia and lagged spot growth correspond to the first lags of the 1-month risk premium and 1-month spot growth, respectively. Spot growth is defined using the corresponding daily spot rate at maturity.

Panel A: BCI 4TC 1-month contract.

Panel B: BPI 4TC 1-month contract.

attenuated as the lag of the regressor increases, reinforces the argument for the existence of momentum.

Table 4.8: Regressions of future risk premia on past physical market indicators.

Variable	$f(t, 1) - s(t + 1)$				$f(t, 2) - s(t + 2)$			
	n	β	t^{NW}	R^2	T	β	t^{NW}	R^2
Panel A: Capesize Sector (BCI 4TC)								
$s(t) - s(t - 1)$	115	-0.18***	-3.88	0.12	114	-0.11	-1.14	0.01
$s(t - 1) - s(t - 2)$	114	0.02	0.36	0.00	113	-0.09	-1.04	0.01
$s(t - 2) - s(t - 3)$	113	-0.03	-0.68	0.00	112	0.02	0.25	0.00
$s(t) - s(t - T)$	115	-0.18***	-3.88	0.12	113	-0.10	-1.60	0.02
$BDI(t) - BDI(t - 1)$	116	-0.23*	-1.95	0.03	115	-0.32	-1.34	0.02
$BDI(t - 1) - BDI(t - 2)$	116	-0.09	-0.74	0.00	115	-0.27	-0.81	0.01
Panel B: Panamax Sector (BPI 4TC)								
$s(t) - s(t - 1)$	115	-0.15***	-3.06	0.08	114	-0.28	-1.40	0.06
$s(t - 1) - s(t - 2)$	114	-0.07	-1.41	0.02	113	-0.29**	-2.28	0.07
$s(t - 2) - s(t - 3)$	113	-0.13***	-2.74	0.06	112	-0.25***	-3.29	0.05
$s(t) - s(t - T)$	115	-0.15***	-3.06	0.08	113	-0.26**	-2.33	0.11
$BDI(t) - BDI(t - 1)$	116	-0.16***	-2.69	0.06	115	-0.38*	-1.95	0.07
$BDI(t - 1) - BDI(t - 2)$	116	-0.15**	-2.56	0.05	115	-0.35***	-2.84	0.06

Notes: Panels A-B report 1- and 2-month horizon OLS forecasting regressions of future risk premia, $f(t, T) - s(t + T)$, on past physical market conditions for the Capesize BCI 4TC and Panamax BPI 4TC contracts, respectively. Namely, in the first three rows of each panel the predictor is the lagged one-period spot growth $s(t - l) - s(t - l - 1)$ where the number of lags, l , varies from 1 to 3. In the fourth row, the predictor is the corresponding previous spot growth for each contract, $s(t) - s(t - T)$; e.g., for the 2-month contract expiring in $t + 2$ months, the predictor is the realised spot growth related to the two-month contract that expired at t , that is, the one corresponding to period $t - 2$ to t . Note that spot growth is estimated using the respective daily spot rate, $s(t)$, as the final spot price instead of the current settlement rate. In rows five and six of each panel, the predictors are the first and second lags of the first difference of the BDI variable, respectively. The maturity of the contract and the number of observations are denoted by T and n , respectively. To deal with the overlapping nature of the variables, t-statistics are estimated using the Newey-West (1987) HAC correction. The slope coefficient, β , is accompanied by *, **, or *** when the absolute t^{NW} statistic indicates significance at the 10%, 5% or 1% level, respectively.

We now examine whether future risk premia can be forecasted by – changes in – realised spot market conditions. To begin with, we perform OLS predictive bivariate regressions of 1- and 2-month risk premia on lagged spot growth. At this point, recall that FFA contracts are settled based on a monthly average; as a result, both spot growth and risk premia are estimated using the corresponding settlement rate as the final spot price, $s(t + T)$. When performing, however, this set of regressions, we incorporate an alternative spot growth variable. Namely, we estimate spot growth using the respective daily spot rate at maturity of the contract as the final spot price. The

motivation for this adjustment is that the daily spot price may be more informative regarding current market conditions

– that is, the ones prevailing during the initiation of the FFA contract – compared to the monthly average settlement rate.¹⁰³

In analogy to the risk premia regressions, in the first three rows of each panel of Table 4.8 the predictor is the lagged 1-period spot growth with the number of lags varying from one to three while in the fourth row the predictor is the respective previous spot growth for each contract; that is, for the 2-month contract expiring in $t + 2$ months, the predictor is the realised 2-month spot growth corresponding to period $t - 2$ to t .¹⁰⁴ For robustness, in addition to the spot growth regressions, we examine the predictability of risk premia by lags of the – first difference of the – Baltic Dry Index (BDI).

Following Alizadeh and Nomikos (2009), “the BDI is a composite index... widely used by practitioners as a general market indicator reflecting the movements in the dry-bulk market. It is in other words the ‘barometer’ of dry-bulk shipping”. Using this variable, therefore, we want to examine whether realised market conditions in the aggregate dry-bulk market can predict sector-specific future risk premia. Accordingly, in the fifth and sixth rows of each panel the predictors are the first and second lags of the first difference of the BDI variable, respectively. Since the sector-specific conditions are highly correlated in the dry bulk industry (see also Chapter 2 of this thesis), one should expect that the results from these regressions would closely assembly the ones using the sector-specific spot growth.

The results presented in Table 4.8 indicate that – in many cases – there exists statistically significant predictability of future risk premia from realised physical market conditions. Specifically, in both contracts the first lag of each spot market indicator variable negatively predicts 1-month future risk premia. This predictability is statistically significant at the 1% level when using lagged spot growth as a regressor. In the Panamax contract, this is also true for the BDI index. Figure 4.3 depicts the negative relationship between the 1-month risk premium and the first lag of the 1-month spot growth.

Therefore, a recent improvement in realised spot market conditions strongly predicts a decrease in future risk premia. While there is strong statistical evidence of risk premia predictability – from both regressors – also in the 2-month Panamax contract this is not the case for the 2-month Capesize one. In conclusion, similar to the basis and lagged risk premia regressions, also for this set

¹⁰³ We have also performed the same set of regressions incorporating as a predictor the spot growth variable defined in equation 4.1b and the results are qualitatively very similar, albeit, less significant compared to the ones presented here.

¹⁰⁴ Note that, for the one-month contract, the first and fourth rows of the respective panel coincide.

of forecasting variables, predictability is more profound for both the Panamax sector and the 1-month horizon. Note that the finding that spot market indicators have significant predictive power regarding future risk premia becomes more interesting if we recall that realised spot market conditions cannot predict future spot growth (Table 4.6). In turn, this result implies that these variables may affect in an “irrational” manner the formation of current FFA rates.

Finally, from a “non-hedger” industry participant’s perspective, our findings in subsections 4.III.A and 4.III.B may have useful implications for devising profitable investment strategies. Namely, current FFA bases, lagged risk premia, and lagged changes in physical market conditions can be incorporated as signals/indicators for taking the short or long position in the FFA market. For example, the lagged risk premia regression results suggest that taking the short position on the FFA contract after a positive risk premium is realised – and vice versa – might be a profitable investment strategy. More importantly, from an economist’s perspective, it is interesting to examine the potential drivers of this sort of predictability and momentum in the FFA market. To this end, in Section 4.IV we develop a theoretical model that can justify and reproduce those empirical results.

4.III.C. Predictability of Risk Premia from Economic Variables

As illustrated in the previous subsection, there is strong evidence of risk premia predictability by realised physical market conditions. Specifically, recent realised changes in spot market conditions negatively predict future risk premia. Since spot rates are determined in equilibrium through the freight rate mechanism (this topic is extensively analysed in Chapter 2 of this thesis), we further examine the predictability of FFA risk premia by economic variables related to the supply of and demand for shipping services. In particular, since we are interested in short-run predictability – that is, risk premia corresponding to the 1- and 2-month contracts – we focus on predictors that reflect current and recent short-term changes in supply and demand conditions.

We begin by incorporating shipping supply variables related to the capacity and availability of the fleet. Specifically, as indicators of the former and the latter, we use the 1-month log change in fleet capacity and the monthly congestion in main dry bulk ports as a proportion of the corresponding fleet capacity, respectively. For robustness, in addition to the sector-specific variables we also examine the ones related to the aggregate dry bulk fleet. Accordingly, we perform OLS bivariate regressions of 1- and 2-month risk premia on the first two lags of those variables. The obtained results, however, do not suggest the existence of any sort of predictability.¹⁰⁵ A potential explanation regarding the fleet capacity variable is that it is extremely slow-moving and highly inelastic in the

¹⁰⁵ The results from these “supply regressions” are omitted since they are statistically insignificant. However, they can be provided upon request.

short-run and, thus, not representative of the monthly market movements. Unfortunately, we do not have access to satellite data – related to, *exempli gratia*, the position of vessels and the utilisation of the fleet – that would be much more informative regarding short-term market conditions.

Table 4.9: Regressions of Capesize future risk premia on economic variables.

Variable	$f(t, 1) - s(t + 1)$				$f(t, 2) - s(t + 2)$			
	n	β	t^{NW}	R^2	T	β	t^{NW}	R^2
Steel Index (-1)	116	-0.21	-0.32	0.00	115	-0.66	-0.45	0.00
Steel Production (-1)	116	0.10	0.16	0.00	115	-0.92	-1.06	0.01
Steel Production (-2)	116	-0.37	-0.58	0.00	115	0.02	0.02	0.00
Iron Ore Spot (-1)	116	-0.80***	-3.13	0.08	115	-1.47*	-1.94	0.08
Iron Ore Spot (-2)	116	-0.29	-1.08	0.01	115	-0.37	-0.76	0.00
BFI (-1)	116	0.13	0.21	0.00	115	-0.88	-1.03	0.01
BFI (-2)	116	-0.34	-0.55	0.00	115	-0.19	-0.22	0.00
DRI (-1)	116	-0.63	-1.01	0.01	115	0.05	0.06	0.00
Iron Ore Exports (-2)	116	-0.30	-1.01	0.01	115	-0.07	-0.14	0.00
C. Coal Imports (-1)	116	-0.23	-0.98	0.01	115	-0.52*	-1.88	0.01
S. Coal Imports (-1)	116	-0.22	-0.77	0.01	115	-0.63**	-2.29	0.01
S. Coal Exports (-1)	115	-0.06	-0.18	0.00	115	-0.66*	-1.73	0.01
Chinese Imports (-1)	116	-0.17	-0.94	0.01	115	-0.46**	-2.07	0.02
Dry Bulk Exports (-2)	116	-0.47	-1.19	0.01	115	-0.05	-0.07	0.00
Spread (-2)	116	0.47	1.20	0.01	115	0.07	0.11	0.00
Brent Spot (-1)	116	-0.43	-1.37	0.02	115	-1.14	-1.24	0.03
Brent Spot (-2)	116	-0.19	-0.59	0.00	115	-0.72	-0.84	0.01
Propane Spot (-1)	116	-0.28	-1.13	0.01	115	-0.78	-1.22	0.03
Gasoline Spot (-1)	116	-0.38	-1.42	0.02	115	-0.37	-0.65	0.00
Gasoline Spot (-2)	116	0.17	0.61	0.00	115	0.15	0.21	0.00

Notes: This table reports 1- and 2-month horizon OLS forecasting regressions of Capesize risk premia, $f(t, T) - s(t + T)$, on numerous demand- and trade-related variables. BFI and DRI refer to the Blast Furnace Iron and Directly Reduced Iron indices, respectively. C. Coal and S. Coal refer to coking coal and steaming coal, respectively. The variable spread denotes the spread between the one-month log growth of dry bulk fleet capacity and the one-month log growth of dry bulk exports. To establish the stationarity condition, we incorporate the first differences of these predictors. Note that, by (-1) and (-2) we denote the first and second corresponding lags of the predictor, respectively. The number of observations is denoted by n . To deal with the overlapping nature of the variables, t-statistics are estimated using the Newey-West (1987) HAC correction. The slope coefficient, β , is accompanied by *, **, or *** when the absolute t^{NW} statistic indicates significance at the 10%, 5% or 1% level, respectively. Accordingly, a predictor appears in bold whenever it is statistically significant at any of the three conventional levels.

We now turn to the demand variables: these consist of trade and demand indicators related to the dry bulk industry. Namely, we incorporate the world steel production, the trade-weighted steel production index, the Blast Furnace Iron (BFI) and Directly Reduced Iron (DRI) indices, the iron ore

spot price, iron ore exports, coking coal imports, steaming coal imports and exports, Chinese imports, and global total dry bulk exports. In addition, we employ as a regressor the spread between the 1-month growth rates of dry bulk fleet supply and commodity demand (as quantified by total dry bulk exports). This spread variable, defined in Chapter 2 of this thesis, aims to capture imbalances between shipping supply and demand. Finally, we also use as predictors commodity prices related to the tanker (i.e., gasoline and crude oil prices) and LPG trades (i.e., propane prices) to account for trade and demand conditions in the entire shipping industry.¹⁰⁶

Table 4.10: Regressions of Panamax future risk premia on economic variables.

Variable	$f(t, 1) - s(t + 1)$				$f(t, 2) - s(t + 2)$			
	n	β	t^{NW}	R^2	T	β	t^{NW}	R^2
Steel Index (-1)	116	-0.70**	-2.15	0.04	115	-1.51*	-1.69*	0.04
Steel Production (-1)	116	-0.34	-1.05	0.01	115	-1.74**	-2.29	0.06
Steel Production (-2)	116	-0.74**	-2.36	0.05	115	-0.82	-1.26	0.01
Iron Ore Spot (-1)	116	-0.50***	-3.96	0.12	115	-1.20**	-2.45	0.15
Iron Ore Spot (-2)	116	-0.22	-1.62	0.02	115	-0.82***	-2.85	0.07
BFI (-1)	116	-0.06	-0.20	0.00	115	-1.49*	-1.80	0.04
BFI (-2)	116	-0.65**	-2.07	0.04	115	-0.94	-1.60	0.02
DRI (-1)	116	-0.21	-0.66	0.00	115	-1.01**	-2.01	0.02
Iron Ore Exports (-2)	116	-0.31**	-2.08	0.04	115	-0.26	-1.10	0.01
C. Coal Imports (-1)	116	-0.21*	-1.69	0.02	115	-0.37*	-1.82	0.02
S. Coal Imports (-1)	116	-0.06	-0.44	0.00	115	-0.23	-1.12	0.00
S. Coal Exports (-1)	115	-0.04	-0.23	0.00	115	-0.10	-0.36	0.00
Chinese Imports (-1)	116	-0.17	-0.94	0.01	115	-0.46**	-2.07	0.02
Dry Bulk Exports (-2)	116	-0.38*	-1.91	0.03	115	-0.46*	-1.70	0.01
Spread (-2)	116	0.38*	1.92	0.03	115	0.47*	1.77	0.01
Brent Spot (-1)	116	-0.39**	-2.49	0.05	115	-1.18**	-2.34	0.10
Brent Spot (-2)	116	-0.31*	-1.94	0.03	115	-0.74**	-2.38	0.04
Propane Spot (-1)	116	-0.22*	-1.83	0.03	115	-0.57	-1.56	0.04
Gasoline Spot (-1)	116	-0.21	-1.55	0.02	115	-0.76***	-3.91	0.06
Gasoline Spot (-2)	116	-0.32**	-2.39	0.05	115	-0.46**	-2.08	0.02

Notes: This table reports 1- and 2-month horizon OLS forecasting regressions of Panamax risk premia, $f(t, T) - s(t + T)$, on numerous demand- and trade-related variables. BFI and DRI refer to the Blast Furnace Iron and Directly Reduced Iron indices, respectively. C. Coal and S. Coal refer to coking coal and steaming coal, respectively. The variable spread denotes the spread between the one-month log growth of dry bulk fleet capacity and the one-month log growth of dry bulk exports. To establish the stationarity condition, we incorporate the first differences of these predictors. Note that, by (-1) and (-2) we denote the first and second corresponding lags of the predictor, respectively. The number of observations is denoted by n . To deal with the overlapping nature of the variables, t-statistics are estimated using the Newey-West (1987)

¹⁰⁶ We have also examined several financial variables as predictors – such as exchange, inflation, and interest rates. These variables, however, do not appear to possess any forecasting power over future risk premia.

HAC correction. The slope coefficient, β , is accompanied by *, **, or *** when the absolute t^{NW} statistic indicates significance at the 10%, 5% or 1% level, respectively. Accordingly, a predictor appears in bold whenever it is statistically significant at any of the three conventional levels.

Similar to the “supply” regressions, we regress the 1- and 2-month risk premia on past realisations of those predictors for each contract under consideration. Note that, to satisfy the stationarity condition, we have expressed all trade and demand variables as first (log) differences. Tables 4.9 and 4.10 summarise the results for the Capesize and Panamax contracts, respectively. For reasons of conciseness, we include only the lags of the regressor that possess statistically significant predictive power regarding at least one contract and maturity. Furthermore, for expositional simplicity, wherever we find statistically significant predictability at any of the three conventional levels, the respective regressor appears in bold.

Evidently, the results in Tables 4.9 and 4.10 are in line with the previous subsections. Namely, there is significantly more profound predictability in the Panamax sector compared to the Capesize one in both the 1- and 2-month horizons; specifically, only the two steaming coal variables – out of the fifteen different predictors – do not forecast future Panamax risk premia. In contrast, only five variables possess statistically significant forecasting power in the Capesize case. More important from an economic perspective, however, is the consistency in the signs of the slope coefficients. In particular, we observe that – in the statistically significant cases – past changes in trade and demand variables always negatively forecast future risk premia; in other words, a recent realised improvement in demand conditions – which, *ceteris paribus*, implies an improvement in concurrent physical market conditions – is negatively related to future risk premia. Of course, as expected by the definition of the variable, the sign in the spread regressions coefficients is positive. These results validate the previous finding that recent changes in market conditions are inversely related with future risk premia.

Table 4.11: Regressions of future risk premia on lagged second-hand vessel sales.

Variable	$f(t, 1) - s(t + 1)$				$f(t, 2) - s(t + 2)$			
	n	β	t^{NW}	R^2	T	β	t^{NW}	R^2
Panel A: Capesize Sector (BCI 4TC)								
Capesize Sales	116	-1.68	-0.19	0.00	115	-14.26	-1.03	0.01
Dry Bulk Sales	116	-19.34	-1.65	0.02	115	-52.52**	-2.33	0.05
Panel B: Panamax Sector (BPI 4TC)								
Panamax Sales	116	-12.10**	-2.51	0.05	115	-26.00***	-2.63	0.06
Dry Bulk Sales	116	-15.73***	-2.69	0.06	115	-44.83***	-3.68	0.11

Notes: Panels A-B report 1- and 2-month horizon OLS forecasting regressions of future risk premia, $f(t, T) - s(t + T)$, on lagged second-hand vessel sales scaled by the corresponding fleet size, for the Capesize BCI 4TC and Panamax BPI 4TC contracts, respectively. Namely, in the first and second rows of each panel, the predictors are the first lag of the sector-

specific and total dry bulk fleet sales, respectively. The number of observations is denoted by n . To deal with the overlapping nature of the variables, t-statistics are estimated using the Newey-West (1987) HAC correction. The slope coefficient, β , is accompanied by *, **, or *** when the absolute t^{NW} statistic indicates significance at the 10%, 5% or 1% level, respectively.

Finally, we also test the forecasting power of a variable that has been incorporated as an indicator of liquidity in the physical shipping markets (see Chapter 3 of this thesis) but also as an investor sentiment index (Papapostolou *et al*, 2014). Specifically, we examine whether the number of second-hand vessel transactions within a given month, scaled by the corresponding fleet size, can predict future risk premia. As with the two supply variables, also this indicator is constructed using both the sector-specific and the aggregate dry bulk fleet data. Table 4.11 summarises the results from OLS bivariate regressions of 1- and 2-month risk premia on the first lag of the second-hand vessel transaction variables. Accordingly, we observe that there is evidence of statistically significant predictability of future risk premia. In particular, vessel transactions appear to negatively forecast future risk premia. Similar to all previous forecasting tests, this predictability is substantially stronger in the Panamax sector; namely, all slope coefficients are significant at least at the 5% level.

This finding is interesting since, from an economic point of view, we cannot identify any causality between the two variables. A potential indirect explanation, however, could be the following one. Namely, in line with Chapter 3 of this thesis, second-hand vessel transactions are positively correlated with physical market conditions; in our dataset, the correlation coefficients between Capesize settlement rates and Capesize and total dry bulk fleet transactions are 0.25 and 0.38, respectively, while in the Panamax case they are 0.34 and 0.39, respectively. In addition, as illustrated above, prosperous market conditions negatively forecast future risk premia. Thus, these two facts combined might be able to explain the observed predictability.

The documented strong predictability of FFA risk premia is very interesting from an economic perspective. As analysed in Section V, a natural explanation for these stylised facts could be based on the composition of the investor population. Namely, one should a priori expect that in a derivative market where the investor population consists of both “physical hedgers” and non-hedgers, that is, investors that do not only have different objectives but are potentially also asymmetrically informed. Specifically, it is reasonable to assume that the former, by participating also in the physical market, have “inside” information and are more experienced regarding the physical shipping market conditions compared to speculators. As a result, they are expected to form more accurate forecasts of future spot market conditions than the latter. Since, however, the

prevailing FFA rate is determined in equilibrium, this heterogeneity results in biased FFA prices and, in turn, risk premia predictability.

4.III.D. Trading Activity, Spot Market Conditions, and Risk premia

In order to conclude the analysis of predictable variation in the dry bulk FFA market, this subsection examines the forecasting power of FFA trading activity variables. For this purpose, we incorporate measures of trading volume and open interest as regressors.¹⁰⁷ Namely, trading volume refers to the number of FFA contracts traded over a given period while open interest to the number of contracts outstanding at a given point in time; thus, open interest measures the number of long – or, equivalently, short – positions in the market. Accordingly, our aim is to test whether – changes in – market activity can predict either future spot growth or/and risk premia. In other words, we want to examine whether FFA market liquidity is related to either of these two variables. An additional motivation for this analysis is the finding by Hong and Yogo (2011) that movements in commodity market interest can predict commodity returns.


Table 4.12: Descriptive statistics for the trading volume and open interest variables.

Variable	T	n	Mean	Median	SD	Max	Min	ρ_1
Panel A: Capesize Sector (BCI 4TC)								
Trading Volume Growth	1	43	0.00	-0.01	0.60	1.73	-1.39	-0.37
	2	42	0.00	0.05	0.50	1.21	-1.63	-0.18
Open Interest Growth	1	43	0.02	0.08	0.23	0.28	-0.97	-0.05
	2	42	0.01	0.08	0.24	0.25	-0.92	-0.20
Open Interest MA	1	41	0.03	0.11	0.29	0.30	-0.93	0.54
	2	40	0.02	0.08	0.26	0.33	-0.82	0.44
Panel B: Panamax Sector (BPI 4TC)								
Trading Volume Growth	1	43	0.00	-0.05	0.44	0.92	-0.77	-0.48
	2	42	0.00	0.13	0.36	0.60	-1.18	-0.45
Open Interest Growth	1	43	0.02	0.05	0.16	0.31	-0.57	-0.28
	2	42	0.02	0.11	0.18	0.27	-0.57	-0.34
Open Interest MA	1	41	0.03	0.06	0.15	0.38	-0.45	0.19
	2	40	0.02	0.04	0.15	0.35	-0.44	0.03


¹⁰⁷ “Hedging pressure” is an additional market activity variable representing the unbalance of traders’ hedging positions (Ekeland *et al*, 2016). In the commodity finance literature, it is empirically quantified by the ratio of traders’ net short position to the open interest in the corresponding commodity. Accordingly, academic researchers in commodity markets usually examine the relation between “hedging pressure” and futures risk premia to test, among others, the Keynesian Theory of Normal Backwardation (Gorton *et al*, 2012). In shipping, however, there is no collective data regarding FFA traders’ positions; therefore, we are not able to incorporate this variable.

Notes: This table presents descriptive statistics for the 1-month trading volume growth, 1-month open interest growth, and the logarithm of current open interest scaled by the moving average (MA) of open interest over the previous three months. Panels A and B correspond to the BCI 4TC and BPI 4TC FFA contracts, respectively, for the 1- and 2-month maturities. The maturity of the contract and the number of observations are denoted by T and n , respectively. The included statistics are the mean, median, standard deviation, maximum, minimum, and 1-month autocorrelation coefficients.

Our dataset is obtained from the London Clearing House (LCH) and consists of monthly observations on trading volume and open interest related to the BCI 4TC and BPI 4TC contracts with



Panel A: BCI 4TC 1-month contract.



Panel B: BPI 4TC 1-month contract.

Figure 4.4: Trading Volume and Open Interest.

Panels A-B plot the evolutions of trading volume growth, open interest growth, and current open interest scaled by the moving average (MA) of open interest over the previous three months, for the BCI 4TC and the BPI 4TC 1-month contracts, respectively. All variables correspond to log differences. The sample runs from January 2013 to August 2016.

1- and 2-month maturities. Consistent with the previous analysis, open interest is sampled at the last trading day of each month. Trading volume, however, corresponds to the entire month under consideration while, due to data limitations, the sample runs from January 2013 to August 2016.¹⁰⁸

In the remainder of this subsection, we examine the predictive power of the following set of regressors: 1-month growth in trading volume, 1-month growth in open interest, and the logarithm of current open interest scaled by the moving average of open interest over the previous three months, defined as “open interest MA”. Table 4.12 summarises descriptive statistics related to those variables while Figure 4.4 illustrates their evolution for the 1-month contract case. Evidently, both

Table 4.13: Regressions of future risk premia and spot growth on trading activity variables.

Variable	T	n	Δs			r		
			β	t^{NW}	R^2	β	t^{NW}	R^2
Panel A: Capesize Sector (BCI 4TC)								
Trading Volume Growth	1	43	0.29**	2.14	0.10	-0.13	-1.39	0.05
	2	42	0.24	1.28	0.03	-0.03	-0.22	0.00
Open Interest Growth	1	43	0.26	0.93	0.02	-0.38	-1.58	0.06
	2	42	0.75***	3.72	0.05	-0.52***	-4.28	0.04
Scaled Open Interest MA	1	41	0.26	1.11	0.03	-0.47**	-2.70	0.13
	2	40	1.15***	2.81	0.15	-0.99***	-2.92	0.18
Panel B: Panamax Sector (BPI 4TC)								
Trading Volume Growth	1	43	0.14*	1.80	0.07	-0.06	-1.00	0.02
	2	42	0.12	0.98	0.02	0.02	0.14	0.00
Open Interest Growth	1	43	0.08	0.50	0.01	-0.01	-0.04	0.00
	2	42	-0.25	-0.78	0.02	0.22	0.91	0.02
Scaled Open Interest MA	1	41	0.09	0.66	0.01	-0.10	-0.76	0.01
	2	40	0.03	0.06	0.00	-0.09	-0.32	0.00

Notes: This table reports 1- and 2-month horizon OLS forecasting regressions of future spot growth, Δs , and risk premia, r , on 1-month trading volume growth, 1-month open interest growth, and the logarithm of current open interest scaled by the moving average of open interest over the previous three months. Panels A and B correspond to the Capesize BCI 4TC and Panamax BPI 4TC contracts, respectively. In the trading activity case, spot growth corresponds to settlement growth. The maturity of the contract and the number of observations are denoted by T and n , respectively. t-statistics are estimated using the Newey-West (1987) HAC correction. The slope coefficient, β , is accompanied by *, **, or *** when the absolute t^{NW} statistic indicates significance at the 10%, 5% or 1% level, respectively.

¹⁰⁸ While the Baltic Exchange reports figures related to trading volume and open interest for a longer period, the corresponding dataset has two important limitations. First, it is provided on a weekly basis; hence, there is a time inconsistency with the monthly frequency of spot, settlement, and FFA rates. Second, the Baltic Exchange trading volume and open interest figures correspond to the entire Capesize and Panamax sectors and not to specific contracts and maturities; hence, there is also a contract mismatching. Note that, while we have reproduced the estimation presented in this subsection using the Baltic Exchange data, the corresponding empirical results are not reported since they are neither statistically nor economically significant. However, they can be provided upon request.

trading volume and open interest growth exhibit significantly volatile behaviour characterised by negative first-order autocorrelation coefficients. In contrast, the scaled open interest variable is – by construction – much more persistent and significantly less volatile.

Accordingly, we perform OLS regressions of 1- and 2-month future spot growth and risk premia on this set of predictors. While in the open interest regressions the spot growth variable is estimated using the daily spot rate prevailing at the issuance of the FFA contract as the initial spot price, $s(t)$, when using trading volume growth as a predictor, spot growth is defined using the settlement rate at t instead of the prevailing spot price – since trading volume corresponds to the total number of transactions during each entire calendar month and not to the ones within a single day.¹⁰⁹ Unfortunately, a similar adjustment cannot be achieved in the risk premium case; hence, there is a time inconsistency. Table 4.13 presents the results for this set of predictive regressions.

Table 4.14: Contemporaneous regressions of settlement growth on trading volume growth.

Contract	n	β	t	R^2	ρ
Panel A: Capesize Sector (BCI 4TC)					
1-month Contract	43	0.36**	2.53	0.16	0.40
2-month Contract	42	0.33*	1.99	0.09	0.30
Entire Contract	43	0.42*	2.01	0.09	0.30
Panel B: Panamax Sector (BPI 4TC)					
1-month Contract	43	0.09	1.14	0.03	0.17
2-month Contract	42	0.14	1.42	0.05	0.22
Entire Contract	43	0.15	1.48	0.05	0.23

Notes: Panels A-B report contemporaneous regressions of 1-month settlement growth on 1-month trading activity growth for the Capesize BCI 4TC and Panamax BPI 4TC contracts, respectively. Apart from for the specific contract maturities, we also examine the entire contract. The number of observations and the correlation coefficient are denoted by n , and ρ , respectively. The slope coefficient, β , is accompanied by *, **, or *** when the absolute t -statistic indicates significance at the 10%, 5% or 1% level, respectively.

To begin with, we observe that trading volume growth does not have significant predictive power over future risk premia.¹¹⁰ In contrast, there appears to be a positive relation between trading volume growth and future settlement growth. This relation becomes statistically significant in the 1-month maturity for both contracts. From an economic perspective, the latter feature suggests that an increase in trading volume positively forecasts future physical market conditions. We further

¹⁰⁹ Note that the results are very similar when using the alternative definition of spot growth.

¹¹⁰ One might argue that these results are contaminated because the risk premium variable corresponds to a contract initiated on a single day while trading volume to the total transactions during an entire month. For this reason, we have performed the same set of predictive regressions using the corresponding daily trading volume figures; the obtained results, however, still do not indicate statistically significant predictability of risk premia by trading volume.

examine this relationship by running contemporaneous regressions of 1-month settlement growth on the 1-month trading volume growth, that is, regressions where growth in trading volume does not enter the equation lagged. Note that we perform these regressions not only for the specific contract maturities but also for each entire contract (i.e., total trading volume related to the respective 4TC contract). Furthermore, we quantify the correlation between the levels of trading volume and settlement – and spot – rates. Tables 4.14 and 4.15 summarise these additional results.

Evidently, the first two columns of Table 4.15 indicate that trading volume is positively correlated with current market conditions in both sectors. In addition, Table 4.14 suggests that an improvement in market conditions is accompanied by a contemporaneous increase in trading volume. Thus, market participants appear to trade more aggressively during prosperous market conditions and, in turn, increased trading volume forecasts a further improvement in market conditions. Notice that these results are stronger for the Capesize contracts compared to the Panamax ones.

Table 4.15: Correlation between trading volume, open interest, spot, and settlement rates.

Contract	TV and Spot	TV and Settlement	OI and Spot	OI and Settlement
Panel A: Capesize Sector (BCI 4TC)				
1-month Contract	0.36	0.46	0.11	0.13
2-month Contract	0.48	0.59	0.07	0.11
Entire Contract	0.71	0.70	0.28	0.24
Panel B: Panamax Sector (BPI 4TC)				
1-month Contract	0.17	0.27	-0.18	-0.20
2-month Contract	0.28	0.41	-0.09	-0.09
Entire Contract	0.59	0.65	0.15	0.11

Notes: Panels A-B present correlation coefficients for the following pairs of variables: trading volume (TV) and spot rates, trading volume and settlement rates, open interest (OI) and spot rates, and open interest and settlement rates.

We now turn to the relation between open interest and future spot growth and risk premia. Before analysing the results from the predictive regressions, though, notice in Table 4.15 that open interest appears to be very loosely related to both current spot and settlement rates. Thus, in contrast to trading volume, we cannot argue that open interest strongly depends on current market conditions.¹¹¹ Accordingly, we begin by examining the predictive power of the 1-month growth in open interest. Interestingly, the corresponding figures in Table 4.13 indicate the existence of

¹¹¹ Similar to trading volume (Table 4.14), we have also performed contemporaneous regressions of 1-month spot growth on 1-month open interest growth and the results do not indicate a significant relation between these two variables. The corresponding results can be provided upon request.

significant differences between the two contract types. Specifically, the results in the BCI 4TC contract clearly suggest that growth in open interest forecasts future spot growth and risk premia in a positive and negative manner, respectively. Noticeably, in the 2-month contract, the results are statistically significant at the 1% level. In contrast, in the BPI 4 TC contract open interest growth does not have significant predictive power over either future spot growth or risk premia while the signs of the slope coefficients do not follow a consistent pattern.

When we incorporate the “open interest MA” as a predictor, the obtained results are qualitatively similar to the open interest growth ones, albeit, substantially stronger in the case of the BCI 4TC contract. Specifically, both the magnitudes of the slope coefficients and the R^2 s are significantly higher compared to the previous set of regressions. Furthermore, there is statistically significant predictability of future premia also in the one-year case. In contrast, in the BPI 4TC case, the results once again do not suggest a robust relationship between open interest and future spot growth and risk premia. This time, however, the open interest variable is always accompanied by a negative – positive – sign in the predictive risk premia – spot growth – regressions.

In conclusion, the results from the trading activity regressions are indicative of the following patterns. First, market participants appear to trade more aggressively during prosperous market conditions (Table 4.15). In turn, increased trading volume forecasts a further improvement in market conditions (Table 4.13). These features are more profound in the case of the BCI 4TC contract. However, trading volume does not have statistically significant forecasting power regarding future risk premia. Second, in the case of the BCI 4TC contract, the open interest growth and “open interest MA” variables forecast both future market conditions and risk premia. Namely, both variables are positively and negatively related to future spot growth and future risk premia, respectively. These results are significantly stronger in the 2-month maturity. In the case of the BPI 4TC contract, however, there is no robust evidence of either sort of predictability. Note that, in a cross-sector comparison, the trading activity findings are in contrast to the predictability results in the previous subsections since they appear to be significantly stronger in the case of the BCI 4TC contract compared to the BPI 4TC one. However, given the small sample employed in this subsection’s analysis, our results should be treated with caution.

4.III.E. Unbiased Expectations Hypothesis

The unbiased expectations hypothesis (UEH) states that the rate of a futures – forward – contract before maturity must be equal to the rational expectation of the settlement price at maturity. Thus:

$$f(t, T) = E_t[s(t + T)], \quad (4.4a)$$

where $f(t, T)$ is the log-price of the futures – forward – contract at t , $s(t + T)$ is the log-settlement rate of the contract at $t + T$, and $E_t[\cdot]$ is the rational expectations' operator conditional on the time t information filtration.

This hypothesis is closely related to a definition of the efficient market hypothesis (EMH). Namely, according to the primary definition of the weak-form market efficiency, future asset returns should not be predicted by past returns; mathematically, in weak-form efficient markets, returns should follow a random walk process. Fama (1991) extends this definition by arguing that future returns should not be predicted – not only by past realisations of the variable, but also – by \mathcal{F}_t -measurable variables. Equivalently, future returns should be unpredictable given the current information filtration (Kavussanos and Nomikos, 1999). Thus, if FFA markets are efficient, FFA returns – risk premia – should not be predicted by \mathcal{F}_t -measurable variables such as valuation ratios, lagged risk premia, realised physical market conditions, and economic indicators. In this regard, the documented predictability in the BCI and BPI 4TC contracts suggests both that FFA rates are not unbiased forecasts of the realised settlement rates and that FFA markets are not efficient in the sense of Fama.

4.III.E.i. The Wald Test Approach

Apart from the existence of return predictability, a straightforward way to test the unbiasedness hypothesis is by performing a Wald test on the coefficients of the regression equation 4.3a. Namely,

Table 4.16: Wald test on the coefficients of regression equation 4.3a.

Variable	$H_0: \alpha = 0$		$H_0: \beta = 1$		$H_0: \alpha = 0$ and $\beta = 1$
	Value	t -statistic	Value	t -statistic	Chi-square
Panel A: Capesize Sector (BCI 4TC)					
1-month Contract	-0.048 (0.032)	-1.518 [0.1319]	-0.201 (0.125)	-1.607 [0.1109]	7.570 [0.023]
2-month Contract	-0.132 (0.095)	-1.396 [0.166]	0.037 (0.133)	0.282 [0.778]	1.956 [0.376]
Panel B: Panamax Sector (BPI 4TC)					
1-month Contract	-0.009 (0.016)	-0.566 [0.573]	-0.368 (0.115)	-3.197 [0.002]	13.524 [0.001]
2-month Contract	-0.047 (0.045)	-1.039 [0.301]	-0.410 (0.208)	-1.970 [0.051]	4.000 [0.135]

Notes: Panels A-B report results from Wald tests on the coefficients of regression equation 4.3a of the main text for the Capesize BCI 4TC and Panamax BPI 4TC contracts, respectively. Specifically, we examine – both separately and jointly – whether the intercept, α , and the slope coefficient, β , are statistically equal to 0 and 1, respectively. When we examine these two hypotheses separately, we present the normalised value of the restriction, its standard error in parenthesis (\cdot),

the corresponding t-statistic, and the exact significance level in square brackets [·]. When we examine the joint hypothesis, we present the Chi-square statistic and the exact significance level in square brackets [·].

if the log basis is an unbiased estimator of future spot growth, then $\alpha_{\Delta S, T}$ and $\beta_{\Delta S, T}$ should be – jointly – equal to 0 and 1, respectively. Table 4.16 summarises the results from these Wald tests for both contracts and maturities. Interestingly, we observe that in the 1-month horizon, the null hypothesis of $\alpha_{\Delta S, T} = 0$ and $\beta_{\Delta S, T} = 1$ is rejected – at least at the 5% level – for both contracts. In the 2-month horizon, however, the unbiasedness hypothesis appears to be rejected only for the Panamax contract; namely, the null hypothesis that $\beta_{\Delta S, T} = 1$ is rejected at the 10% level. However, we should bear in mind that a limitation of the Wald test is that inference might not be valid when variances are estimated using the Newey-West method. Therefore, the results in the 2-month horizon should be treated with caution.

4.III.E.ii. The Johansen Cointegration Approach

In addition to the previous arguments, a frequently incorporated method to test the unbiasedness hypothesis in futures and forward markets is the use of cointegration techniques and, in particular, Johansen's (1988 and 1991) approach. Regarding the existing shipping literature, Kavussanos and Nomikos (1999) apply this technique to test the hypothesis in the Baltic International Freight Futures Exchange (BIFFEX) market. However, as in the majority of markets where the unbiasedness hypothesis has been tested using this approach, their results provide mixed evidence. Specifically, as the authors argue, the validity of the hypothesis depends on both the type – that is, the idiosyncrasies of the market under investigation – and the time-to-maturity of the contract. As illustrated in the following, our results are in line with this argument. Since Johansen's approach is extensively analysed in the literature, we simply outline the most important points of this framework before presenting our empirical results.

To begin with, we can empirically test the unbiased expectations hypothesis – that is, equation 4.4a – by the means of the regression equation

$$s(t + T) = \beta_1 + \beta_2 f(t, T) + \varepsilon_t; \varepsilon_{t+1} \sim iid(0, \sigma_\varepsilon^2). \quad (4.4b)$$

Importantly, though, since in most markets under examination the futures and spot rates are non-stationary series, we cannot directly perform this regression.¹¹² However, as Engle and Granger (1987) illustrate, the non-stationarity caveat can be circumvented if futures and spot rates are cointegrated.

¹¹² For a detailed description of the stationarity and cointegration topics, the reader can refer to the related literature.

Accordingly, Johansen's (1988) Vector Error Correction Model (VECM) approach can be applied to test for the unbiasedness hypothesis; namely, the VECM specification is

$$\Delta X_t = \mu + \sum_{i=1}^{p-1} \Gamma_i \Delta X_{t-i} + \Pi X_{t-1} + V_t; V_t \sim iid(0, \Sigma), \quad (4.5)$$

where, in our case, X_t is the 2×1 vector of spot and futures rates $[s(t+T) f(t, T)]'$; μ is a 2×1 vector of deterministic components that may include an intercept term, a linear trend term or both; and V_t is a 2×1 vector of white noise residuals with a 2×2 positive definite covariance matrix, Σ . This VECM specification contains information regarding both the short- and long-term adjustments to changes in vector X_t through the estimates of Γ_i and Π , respectively (Johansen, 1988 and 1991; Johansen and Juselius, 1990). As the Granger Representation Theorem (Engle and Granger, 1987) states, the rank of matrix Π is the crucial parameter for cointegration; therefore, we can distinguish between three cases.

First, if $\text{rank}(\Pi) = 2$ then both futures and spot rates are stationary in levels, $I(0)$; hence, we can examine the unbiasedness hypothesis directly from (4.4b). Second, if $\text{rank}(\Pi) = 0$ then spot and futures rates are both $I(1)$ variables but not cointegrated; in this case, the unbiasedness hypothesis is a priori rejected. Third, if $\text{rank}(\Pi) = 1$ then there is a single cointegrating vector, that is, a single cointegrating relationship between the futures and spot rates. Accordingly, matrix Π can be factorised into two separate (2×1) matrices with full column rank, α and β , such that $\Pi = \alpha\beta'$. In this case, while both futures and spot rates are $I(1)$ variables, the product $\beta'X_t$ is $I(0)$ – where the vector of cointegrating parameters is given by the column of β (Martin *et al*, 2013). Moreover, vector α contains the error correction coefficients which measure the speed of convergence to the long-term steady state.

Notice that it is of utmost importance to specify correctly the deterministic components in the VECM since the asymptotic distributions of the cointegration test statistics depend on this choice. Having specified correctly μ , the vector series becomes $X_{t-1} = [s(t-1+T) f(t-1, T)]'$ with a cointegrating vector $\beta' = [1 \beta_1 \beta_2]$ where the spot rate coefficient is normalised to one and β_1, β_2 correspond to the intercept term and the coefficient of the futures rate, respectively. Finally, note that while – in the case of $I(1)$ spot and futures rates – cointegration is a necessary condition for the unbiasedness hypothesis it is not a sufficient one (Hakkio and Rush, 1989). Namely, for the unbiasedness hypothesis to hold, the system of restrictions $[1 \beta_1 \beta_2] = [1 0 - 1]$ must be satisfied.¹¹³

¹¹³ We can test these restrictions using the likelihood-ratio (LR) statistic proposed by Johansen and Juselius (1990).

Having analysed the main points of this framework, we now apply it to the dry bulk FFA market. Specifically, as illustrated in Table 4.4, all $s(t + T)$ and $f(t, T)$ variables in our case appear to be non-stationary. Therefore, instead of performing regression 4.4b, we test for unbiasedness through the VECM specification 4.5. Accordingly, using a combination of model selection criteria,¹¹⁴ we specify a robust lag structure and the appropriate deterministic components for the VECM corresponding to each contract and maturity (Kavussanos and Nomikos, 1999). In all cases under consideration, futures and spot rates appear to be cointegrated, that is, $\text{rank}(\Pi) = 1$.¹¹⁵

Table 4.17: Johansen Cointegration Test.

Variable	VECM Specification		Coefficients		Hypothesis Tests on β'		
	Lags	Trend	β_1	β_2	$H_0:$ $\beta_1 = 0$	$H_0:$ $\beta_2 = -1$	$H_0:$ $\beta_1 = 0$ and $\beta_2 = -1$
Panel A: Capesize Sector (BCI 4TC)							
1-month Contract	1	I/N.T.	-1.06 (0.02)	0.64 (0.23)	5.198 [0.023]	NA NA	5.198 [0.023]
2-month Contract	1	I/N.T.	-1.09 (0.04)	0.96 (0.37)	5.217 [0.022]	NA NA	5.217 [0.022]
Panel B: Panamax Sector (BPI 4TC)							
1-month Contract	0	I/N.T.	-1.03 (0.01)	0.31 (0.14)	4.347 [0.037]	NA NA	4.347 [0.037]
2-month Contract	3	L.I./T	-0.90 (0.06)	0.003 (0.001)	4.327 [0.038]	NA NA	4.327 [0.038]

Notes: Panels A-B report results from the Johansen cointegration approach for the Capesize BCI 4TC and Panamax BPI 4TC contracts, respectively. Columns 2 and 3 report the lag order and the deterministic components of the applied VECM model (equation 4.5). The abbreviations I/N.T. and L.I./T correspond to the specifications “The level data have no deterministic trends and the cointegrating equations have intercepts” and “The level data and the cointegrating equations have linear trends”, respectively. Columns 4 and 5 report the values of the normalised cointegrating coefficients. The corresponding standard error appears in parenthesis (.). Finally, columns 6, 7, and 8 present the LR statistic related to the validity of the cointegrating vector restrictions $\beta_1 = 0$, $\beta_2 = 1$, and $\beta_1 = 0$ and $\beta_2 = 0$ (jointly). The corresponding exact significance level appears in square brackets [.]. Note that the abbreviation “NA” refers to the case where the restriction is not binding.

Having established this necessary condition, we examine the unbiasedness hypothesis by testing – jointly – the restrictions $\beta_1 = 0$ and $\beta_2 = -1$ of the cointegrating relationship. Table 4.17 summarises the specification for each model, the values of the normalised cointegrating coefficients,

¹¹⁴ Namely, we base our decision on a combination of the Likelihood-Ratio (LR), Final Prediction Error (FPE), Akaike Information Criterion (AIC), Schwartz Information Criterion, and Hannan-Quinn Information Criterion.

¹¹⁵ This result was expected since, as indicated by the Phillips-Perron and Augmented Dickey-Fuller unit root tests (see Table 4.4), all risk premium variables are stationary. Note that the null hypothesis of non-stationarity was rejected at the 1% level for both contracts and maturities.

and the LR statistic related to the restrictions on the cointegrating vector. As it becomes evident, the unbiasedness hypothesis is rejected for all contracts and maturities. Finally, in line with the literature (Nomikos and Kavussanos, 1999), for the 2-month contracts – due to the overlapping nature of risk premia and growth rates variables (Hansen and Hodrick, 1980) – we also incorporate the Philips and Hansen (1990) fully modified ordinary least squares (FMOLS) estimation procedure. The results from the FMOLS regressions also reject the unbiasedness hypothesis.¹¹⁶

In conclusion, the obtained results from those econometric tests unequivocally suggest that there exists a bias in the formation of the 1-month FFA rates in both contracts. Regarding the 2-month contracts, our findings point towards the existence of a bias, especially in the Panamax BPI 4TC case. Note that these results are perfectly aligned with the ones obtained in Section 4.II and Subsections 4.III.A, 4.III.B, and 4.III.C. Consequently, our findings – which suggest the rejection of the unbiased expectations hypothesis and, in turn, of the dry bulk FFA markets' efficiency, especially in the 1-month maturity and the BPI 4TC contract – are robust. Finally, in a cross-sector comparison, there is no clear economic or financial justification for the difference in the obtained results. Namely, in terms of fundamentals, the physical markets related to each sector are highly correlated (see Chapter 2 and Subsection 4.II of this chapter). Regarding the FFA market structure, while we cannot examine the composition of the investor population related to each sector, we can plausibly assume that there are no significant differences in the characteristics of the investors participating in each one. Furthermore, in terms of FFA market liquidity, while the Capesize sector is relatively more liquid than the Panamax one (as analysed in Chapters 1 and 4), we cannot justifiably argue that the observed discrepancy in results can be attributed to that liquidity difference.

4.IV. A Heterogeneous Expectations Model for the FFA Market

It has been well-analysed in the asset pricing literature that there exist two potential explanations for the rejection of the unbiased expectations hypothesis: the formation of irrational expectations and the existence of – time-varying – required risk premia on behalf of investors. As analysed in Chapter 3 of this thesis, most of the potential “rational” explanations incorporate “exotic preferences” rendering them almost indistinguishable from the distorted beliefs ones (Cochrane, 2011). Equivalently, their predictions stem from auxiliary assumptions and not from the rationality assumption per se (Arrow, 1986). The fact, however, that almost any biased beliefs model can be re-expressed as a rational expectations' one with time-varying preferences/discount factors (Cochrane, 2011) does not validate the latter approach or, vice versa, invalidate the former one. Specifically, as

¹¹⁶ The results from the FMOLS estimation can be provided by the author upon request.

Lof (2015) argues, biased beliefs models are very appealing when modelling boom-bust cycles as the ones documented in the shipping industry (Greenwood and Hanson, 2015).

Furthermore, as analysed in Section 4.II, Fama and French (1987) argue that any potential irrational forecasts of future spot prices will appear as a time-varying risk premium. However, while Fama and French do not exclude the existence of irrational forecasts as a potential explanation for time-varying risk premia in futures/forwards markets – in line with most papers in the literature – they do not examine it. Namely, most commodity futures models incorporate the “theory of storage” explanation of – a time-varying “hedging pressure” bias which, in turn, results in – “time-varying” risk premia (Gorton *et al*, 2012; Ekeland *et al*, 2016). However, in our case, since shipping services are a non-storable commodity, this “rational” justification cannot be applied.

Therefore, in order to justify economically and, in turn, reproduce our main empirical findings – namely, the momentum effect and the predictability of future risk premia by recent changes in market conditions – we develop in the remaining of this chapter a theoretical model of FFA price determination that allows us to depart from the rational expectations benchmark. While the proposed framework draws its main features from the last generation of structural economic models in the commodity futures literature (Gorton *et al*, 2012; Acharya *et al*, 2013), apart from the standard “hedging pressure” bias, it can also account for distorted beliefs on behalf of a fraction of the investor population – that is, for heterogeneous expectations. Accordingly, by analysing and simulating several alternative specifications of the model, we show that one must depart from the rational expectations benchmark of the economy in order to reproduce the observed regularities

4.IV.A. Economic Environment and Model Solution

Consider a discrete-time environment where the passage of time is denoted by t . The economy consists of one commodity – a numéraire – which is the freight service and two markets. There is a spot market related to a specific shipping route and a derivative market with a forward contract (FFA) on the freight service corresponding to this route. While in the FFA market both short and long positions are allowed, in the physical market short positions are not. Both markets operate in every period, that is, they clear at each t and, in turn, the respective equilibrium rate is determined. Naturally, the FFA contract at each t is related to the spot rate at $t + 1$.¹¹⁷

Let S_t denote the spot price at t , observed at each period by the entire investor population. In the context of our theoretical model, the spot price is stochastic and exogenously determined. Thus,

¹¹⁷ The implicit assumption is that the spot rate coincides with the settlement rate of the contract at maturity. Furthermore, in the following, we focus on the 1-month contract since the evidence is stronger for this horizon.

we examine the formation of FFA rates in a partial equilibrium framework. In line with the data (Tables 4.1, 4.4, and 4.5), the evolution of spot prices is assumed to be given by:

$$S_{t+1} = S_t + \kappa_{t+1} + \lambda_{t+1}, \quad (4.6)$$

where $\kappa_{t+1} \sim iid N(0, \sigma_\kappa^2)$ and $\lambda_{t+1} \sim iid N(0, \sigma_\lambda^2)$ are two uncorrelated random error terms. More interestingly, κ_{t+1} is realised at t but it cannot, in principle, be observed by all market participants with the same precision. The reader can think of κ_{t+1} as a signal or, equivalently, as private information about future market conditions. In contrast, λ_{t+1} is realised at time $t + 1$ and all market participants at t have the same prior information about its distributional properties.

The FFA market consists of three investor types, i : “ship owners”, “charterers”, and “speculators”, denoted by o , c , and s , respectively. We normalise the investor population related to each type to a unit measure. Ship owners – this group also includes operators of vessels – are the providers of the freight service; therefore, they want to hedge their exposure to freight risk through FFA contracts. Equivalently, they can be thought of as the producers of the freight service commodity. In 2006, ship owners accounted for approximately 60% and 46% of the total dry bulk FFA market participants and traded volume, respectively (Alizadeh and Nomikos, 2009). A ship owner has two incentives to trade in the FFA market. First and most importantly, he is interested in hedging his production risk. Second, he speculates on the difference between the FFA rate and the expected settlement rate, that is, for market-making purposes (Vives, 2008). In equilibrium, as hedgers of future sales, ship owners are expected to take the short position on the FFA contract.

Charterers are the consumers of the commodity since they transport their cargoes through ship owners’ vessels. In practice, this group may correspond to large trading houses, including commodity and energy firms; in 2006, trading houses accounted for 25% and 39% of the dry bulk FFA market participants and traded volume, respectively (Alizadeh and Nomikos, 2009).¹¹⁸ By participating in the FFA market, charterers want to reduce their consumption risk. Like ship owners, however, their demand also consists of a speculative component. In equilibrium, as hedgers of future purchases, charterers are expected to take the long position on the derivative contract. Since ship owners and charterers participate in both markets, they can be defined as “physical hedgers” or, equivalently, “traditional players”.

The third investor type corresponds to speculators or, equivalently, “non-hedgers”; in practice, this group may consist of finance houses such as hedge funds and investment banks but also from individual investors. While in 2006 finance houses accounted for 15% of both the dry bulk FFA

¹¹⁸ Note that, in addition to charterers, this role can also be attributed to “cross hedgers”, that is, diversified investors whose portfolio exposure is negatively correlated with freight rates; hence, they can hedge their exposure by taking the long position in the FFA market.

market participants and traded volume, in the following ten years this percentage has significantly increased (Alizadeh and Nomikos, 2009).¹¹⁹ As is always the case in this type of models, speculators' trade is motivated by purely speculative incentives. Thus, we implicitly assume that speculators' participation in the FFA market is not part of a diversification policy, that is, they are not "cross-hedgers". Equivalently, they can be considered as market makers who participate in the FFA market aiming to profit from absorbing part of the freight risk that ship owners and charterers want to hedge (Vives, 2008).

In line with the literature (Hong and Yogo, 2012; Acharya *et al*, 2013), agents are assumed to have mean-variance objective functions where both the risk aversion parameter, γ_i , and the time t expectations operator, E_t^i , depend on the agent type. Importantly, the only source of uncertainty in the model is the realisation of the future spot price, S_{t+1} . The crucial assumption of our framework is that agents form heterogeneous expectations regarding future market conditions for two potential reasons. Specifically, we assume that in agent i 's mind, spot prices evolve according to

$$S_{t+1} = (1 - \vartheta_i)[S_t + \rho_i \kappa_{t+1} + \lambda_{t+1}] + \vartheta_i[S_t + \psi_i(S_{t-1} - S_t) + \lambda_{t+1}] \quad (4.7a)$$

$$\Rightarrow S_{t+1} = S_t + (1 - \vartheta_i)\rho_i \kappa_{t+1} + \vartheta_i \psi_i(S_{t-1} - S_t) + \lambda_{t+1} \quad (4.7a')$$

in which $\vartheta_i \in [0,1]$, $\rho_i \in [0,1]$, and $\psi_i > 0$.

The right-hand side of (4.7a) consists of two terms or, equivalently, signals. Regarding the first term, the quantity in the square brackets, $S_t + \rho_i \kappa_{t+1} + \lambda_{t+1}$, represents the *fundamental evolution of the spot price* as perceived by investor i ; we call this the "fundamental value signal". As mentioned above, while the value of S_t and the distributional properties of λ_{t+1} are public information, the random term κ_{t+1} is not since it depends on the private information of the investor type. Specifically, for an investor with perfect information about future market conditions the "coefficient of precision", ρ_i , is equal to 1; equivalently, the less informed an investor is the more ρ_i approaches zero.

Regarding the second term, the quantity in the square brackets, $S_t + \psi_i(S_{t-1} - S_t) + \lambda_{t+1}$, represents the *contrarian evolution of the spot price* as perceived by investor i ; we call this the "contrarian value signal". This indicates that spot prices will fall if they have recently risen and vice versa. The coefficient ψ_i measures the "degree of gambler's fallacy" or, equivalently, the "degree of contrarian beliefs" of investor i ; thus, for a totally rational investor $\psi_i = 0$. Therefore, for investor i , the evolution of the spot price variable is given by a weighted average of these two signals; we call

¹¹⁹ Unfortunately, we do not have a more recent measure of the composition of either traded volume or market participants.

the coefficient ϑ_i “degree of wavering”. Equivalently, $(1 - \vartheta_i)$ quantifies the degree of confidence that investor i has about his private information.

Accordingly, we assume that physical hedgers are both perfectly informed and totally rational; thus, we set $\rho_o = \rho_c = 1$, $\psi_o = \psi_c = 0$, and $\vartheta_o = \vartheta_c = 0$. Equivalently, they only trust the fundamental value signal which they receive with perfect precision. In contrast, speculators are both less than perfectly informed and irrational, that is, $\rho_s \in [0,1)$, $\psi_s > 0$, and $\vartheta_s \in (0,1)$. Thus, they waver between the two signals. The assumption regarding asymmetric and imperfect information can be justified by the fact that traditional players operate also in the physical shipping market, potentially for a long period; therefore, they are more experienced and/or better informed – since they have “inside” information regarding the actual future market conditions – than speculators. Hence, they are expected to form more accurate forecasts of future spot market conditions than the latter.

Regarding the behavioural bias assumption, as analysed in Section 4.1, speculators are assumed to suffer from a variation of “the law of small numbers” bias which is also known as “regression – reversion – to the mean” and “gambler’s fallacy”. In line with Shefrin (2000), “the law of small numbers” arises because people think that “...the law of large numbers applies to small as well as to large samples” or, equivalently, “they exaggerate how likely it is that a small sample resembles the parent population from which is drawn” (Tversky and Kahneman, 1971; Terrell, 1994; Rabin, 2002). As a result, individuals that suffer from this misperception inappropriately predict – rapid – reversal of a trend or shock.

In line with equation 4.7a, this behavioural bias is introduced in a rather straightforward manner. Namely, speculators believe that spot price shocks tend to cancel out each other rapidly; thus, they expect that a price shock at t will be followed by one of the opposite sign at $t + 1$. Equivalently, they believe that the spot price variable tends to revert rapidly to its level before the last realised shock. As Rabin (2002) argues, an individual suffering from the “gambler’s fallacy” believes that draws of one signal – a spot price shock in our case – increase the odds of next drawing other signals – that is, a spot price shock of the opposite sign. A natural consequence of this bias is a contrarian investment behaviour on behalf of speculators.

In practice, traders frequently form expectations about future market conditions and, in turn, devise investment strategies following simple technical analysis rules that are based on contrarian strategies – which can be – influenced by behavioural biases such as the “gambler’s fallacy”. In particular, Kaniel *et al* (2008) provide evidence that numerous traders indeed select contrarian strategies while laboratory experiments, conducted by Bloomfield *et al* (2009), suggest that mainly uninformed investors usually adopt contrarian behaviour. What is more, Grinblatt and Keloharju

(2000) show that, in Finnish markets, inexperienced investors frequently act as contrarians while more sophisticated ones tend to follow momentum strategies (Lof, 2015). Those findings are particularly related to our model since speculators correspond to financial investors who, as non-participants in the physical market, are assumed to be less sophisticated and informed regarding future shipping market conditions compared to “traditional players”.

The speculator-specific parameters ρ_s , ψ_s , and ϑ_s characterise completely the information structure of our model. When $\rho_s = 1$ and either ψ_s or ϑ_s equals zero, all agents are totally rational and have perfect and, thus, symmetric information about the economy. We define this case as the benchmark “rational” economy of our model, R . When $\rho_s < \rho_o = \rho_c = 1$, information is both imperfect and asymmetric, irrespective of ψ_s and ϑ_s (Wang, 1993). When $\psi_s, \vartheta_s > 0$, the aggregate –average– investors’ expectations in the market are formed in an irrational manner.

Incorporating in equation 4.7a’ the expectation and variance operators – conditional on both public information available at time t and the specific agent’s private information and beliefs – we obtain

$$E_t^i[S_{t+1}] = S_t + (1 - \vartheta_i)\rho_i\kappa_{t+1} + \vartheta_i\psi_i(S_{t-1} - S_t) \quad (4.7b)$$

and

$$\text{Var}_t^i[S_{t+1}] = \text{Var}_t[S_{t+1}] = \sigma_\lambda^2. \quad (4.7c)$$

Therefore, while the expectation of the future spot price depends on both the agent-specific information and beliefs, the perceived variance is equal to the variance of the random cash flow shock which, in turn, is common knowledge.

The timeline of the model is as follows. At each t , λ_t is realised and, in turn, S_t is observed by the entire investor population. In addition, κ_{t+1} is also realised, however, it is not observed with the same precision by each investor type. Accordingly, agents determine their optimal time t demands for the FFA contracts with the aim of maximising their respective mean-variance objective functions. First, for each ship owner this corresponds to

$$\max_{h_t^o} E_t^o[S_{t+1}Q_{t+1} + h_t^o(S_{t+1} - F_t)] - \frac{\gamma_o}{2} \text{Var}_t[S_{t+1}Q_{t+1} + h_t^o(S_{t+1} - F_t)], \quad (4.8a)$$

where Q_{t+1} are his time $t + 1$ holdings of the physical asset (i.e., ship owner’s fleet capacity) while h_t^o and F_t are his time t demand for and the price of the FFA contract, respectively. The optimisation yields

$$h_t^o = \frac{E_t^o[S_{t+1} - F_t]}{\gamma_o \text{Var}_t[S_{t+1}]} - Q_{t+1}. \quad (4.8b)$$

Second, each charterer maximises

$$\max_{h_t^c} E_t^c[-S_{t+1}D_{t+1} + h_t^c(S_{t+1} - F_t)] - \frac{\gamma_c}{2} \text{Var}_t[-S_{t+1}D_{t+1} + h_t^c(S_{t+1} - F_t)], \quad (4.9a)$$

where D_{t+1} is his time $t + 1$ demand for shipping services while h_t^c is his time t demand for the FFA contract. This yields

$$h_t^c = \frac{E_t^c[S_{t+1} - F_t]}{\gamma_c \text{Var}_t[S_{t+1}]} + D_{t+1}. \quad (4.9b)$$

Following Gorton *et al* (2012) and Hong and Yogo (2012), we assume that ship owners at time t know with certainty – they commit to – the amount of shipping services they will sell at time $t + 1$, Q_{t+1} . This assumption is in line with the nature of the industry where the realisation of newbuilding decisions requires a significant construction lag (Kalouptsidei, 2014; Greenwood and Hanson, 2015) and, more importantly, each firm is perfectly informed about its own delivery schedule. Regarding scrapping and second-hand vessels sale and purchase decisions, we assume that these are taken at time t . In turn, the corresponding investment decisions at time t will affect the respective hedging decision at t . Finally, regarding the operation of vessels, we assume that ship owners at time t know with certainty how many of their vessels will be available in the spot market at $t + 1$, that is, how many vessels will not be engaged in either time-charter or voyage contracts by that time.

We make an analogous assumption for charterers; namely, charterers at time t know with certainty – they commit to – the amount of shipping services they will demand at time $t + 1$, D_{t+1} . As one can imagine, this assumption is plausible for large commodity producers and consumers and established trading houses. Note that the Capesize and Panamax dry bulk sectors, examined in this research, are related to the trades of iron ore, coal, grain, bauxite, and the larger minor bulks trades. All these commodities are transported in large cargoes occupying the entire vessel for a given contract.

Third, speculator's maximisation problem is

$$\max_{h_t^s} E_t^s[h_t^s(S_{t+1} - F_t)] - \frac{\gamma_s}{2} \text{Var}_t[h_t^s(S_{t+1} - F_t)], \quad (4.10a)$$

where h_t^s is his time t demand for the FFA contract. This yields

$$h_t^s = \frac{E_t^s[S_{t+1} - F_t]}{\gamma_s \text{Var}_t[S_{t+1}]} \quad (4.10b)$$

In equilibrium, FFA contracts are in zero net supply. Therefore, the market-clearing condition at each t requires

$$h_t^o + h_t^c + h_t^s = 0. \quad (4.11)$$

Substituting equations 4.8b, 4.9b, and 4.10b in 4.11, we obtain the – endogenously – determined equilibrium FFA rate at t , F_t^* :

$$F_t^* = \frac{\gamma_c \gamma_s E_t^o[S_{t+1}] + \gamma_o \gamma_s E_t^c[S_{t+1}] + \gamma_o \gamma_c E_t^s[S_{t+1}]}{\gamma_c \gamma_s + \gamma_o \gamma_s + \gamma_o \gamma_c} - \frac{\gamma_o \gamma_c \gamma_s}{\gamma_c \gamma_s + \gamma_o \gamma_s + \gamma_o \gamma_c} \sigma_\lambda^2 (Q_{t+1} - D_{t+1}). \quad (4.12)$$

Equation 4.12 indicates that the FFA rate consists of two terms. The first one is a weighted average of market expectations regarding the future spot price – with the weights being determined by the agent-specific coefficients of risk aversion. If all agents in the market held symmetric, perfect information and formed rational expectations this term would reduce to $E_t^R[S_{t+1}]$. This is a standard term in rational expectations models with symmetric information. The second term quantifies the “hedging pressure” bias in the FFA price, the direction of which depends only on the sign of the parenthesis, that is, on the fundamental structure of the economy under consideration. In the literature, “hedging pressure” is defined as the imbalance of traders’ hedging positions (Ekeland *et al*, 2016). Note that, apart from the “hedging pressure”, the magnitude of the second term also depends on the agent-specific coefficients of risk aversion as well as on the volatility of the cash flow shock.

Importantly, note that in most commodity markets structural models, the hedging pressure variable is endogenously determined by incorporating the theory of storage and, specifically, by modelling explicitly the level of inventories. Accordingly, inventories, hedging pressure, and spot rates are interdependent. In the case of shipping, however, the underlying asset is non-storable; thus, hedging pressure cannot be determined endogenously through this mechanism. Therefore, to account for time-varying hedging pressure, we need to impose an assumption – based on plausible economic arguments – that relates it explicitly to the – exogenously – determined spot rate process.¹²⁰

In conclusion, in rational expectations symmetric information models, when “hedging pressure” is equal to zero, the derivative contract’s price is an unbiased predictor of the future spot price. For

¹²⁰ Note that it is out of the scope of this research to model the shipping freight rate mechanism. For more on this topic the reader can refer to Kalouptsi (2014) and Greenwood and Hanson (2015).

conciseness, in the following, we define the hedging pressure variable as $Q_t - D_t = HP_t$; thus, when HP_t is positive, the physical market position of short hedgers (ship owners) exceeds the one of long hedgers (charterers) and vice versa. When HP_t equals zero the two positions exactly offset each other.

The innovative idea proposed by this framework, however, is that even in the absence of hedging pressure the FFA price can be a biased predictor of future spot rates due to the heterogeneity of beliefs among the investor population. Unfortunately, since there is neither data availability regarding FFA traders' positions nor surveys regarding their beliefs and investment strategies – as is the case in the equity markets literature (Greenwood and Shleifer, 2014) – we are not able to test formally the actual source of the documented bias. Thus, in the following, we aim to provide the most plausible explanation by simulating the “rational” and “irrational” versions of our framework and examining which one reproduces more sufficiently the observed regularities.

Without loss of generality and for expositional simplicity, we assume that $\gamma_o = \gamma_c = \gamma_s = \gamma$. Accordingly, equation 4.12 is simplified to

$$F_t^* = \frac{1}{3} \{E_t^o[S_{t+1}] + E_t^c[S_{t+1}] + E_t^s[S_{t+1}]\} - \frac{\gamma\sigma_\lambda^2}{3} HP_{t+1}. \quad (4.13)$$

Incorporating in (4.13) equation 4.7b for $i = o, c, s$ yields

$$F_t^* = S_t + \frac{2 + (1 - \vartheta_s)\rho_s}{3} \kappa_{t+1} + \frac{\vartheta_s\psi_s}{3} (S_{t-1} - S_t) - \frac{\gamma\sigma_\lambda^2}{3} HP_{t+1}. \quad (4.14)$$

It is also useful to examine the benchmark rational economy, R , in which the market solely consists of totally rational and perfectly informed agents. In this case, the expected spot price at $t + 1$ and the time t FFA rate, $F_t^{R,HP}$, are given by

$$E_t^R[S_{t+1}] = S_t + \kappa_{t+1}, \quad (4.15)$$

and

$$F_t^{R,HP} = S_t + \kappa_{t+1} - \frac{\gamma\sigma_\lambda^2}{3} HP_{t+1}, \quad (4.16)$$

respectively.

Comparing (4.14) to (4.16), we observe that $F_t^* = F_t^{R,HP}$ if and only if $S_{t-1} = S_t$ – that is, if there is spot price shock between the two consecutive dates, $t - 1$ and t – and $\kappa_{t+1} = 0$ – that is, if there is no private information/signal about future spot market conditions – or equivalently, if in three consecutive dates, $t - 1$, t , and $t + 1$, the spot rate is (expected to be) the same. Whenever a

shock perturbs the equilibrium, however, the future price deviates from its rational equilibrium analogue. The sign and magnitude of this deviation depend on the values of the shocks κ_t , λ_t , and κ_{t+1} and the speculator-specific coefficients.

The realised bias in the FFA rate at $t + 1$ in the heterogeneous-agent economy can be quantified by subtracting equation 4.6 from (4.14):

$$F_t^* - S_{t+1} = \left[\frac{(1 - \vartheta_s)\rho_s - 1}{3} \kappa_{t+1} + \frac{\vartheta_s \psi_s}{3} (S_{t-1} - S_t) \right] - \frac{\gamma \sigma_\lambda^2}{3} HP_{t+1} - \lambda_{t+1}. \quad (4.17)$$

This bias can be decomposed into three terms. We define the first one as the “heterogeneous expectations bias”; this arises if and only if there is asymmetry of information and/or existence of the “gambler’s fallacy” in the market. The second term is the familiar “hedging pressure bias”; this arises if and only if $HP_{t+1} \neq 0$, that is, if $Q_{t+1} \neq D_{t+1}$. The third one is the “random bias”; this arises if and only if the – unpredictable – error term of the cash flow process corresponding to time $t + 1$, λ_{t+1} , is different than zero.

Thus, in the absence of asymmetric information and gambler’s fallacy, the rationally expected and the realised bias in the FFA rate at $t + 1$ are

$$F_t^{R,HP} - E_t^R[S_{t+1}] = -\frac{\gamma \sigma_\lambda^2}{3} HP_{t+1} \quad (4.18a)$$

and

$$F_t^{R,HP} - S_{t+1} = -\frac{\gamma \sigma_\lambda^2}{3} HP_{t+1} - \lambda_{t+1}, \quad (4.18b)$$

respectively. Moreover, in the absence of the first two biases, the FFA rate is

$$F_t^R = S_t + \kappa_{t+1}, \quad (4.19)$$

and, in turn, the rationally expected risk premium at t is

$$F_t^R - E_t^R[S_{t+1}] = 0, \quad (4.20)$$

while the realised risk premium at $t + 1$ is given by

$$F_t^R - S_{t+1} = -\lambda_{t+1}. \quad (4.21)$$

Hence, even in the absence of the first two biases, the realised risk premium can be significantly different than the rationally expected one – which in this case will always be statistically equal to

zero. Specifically, the realised risk premium in this case would depend only on the distributional properties of the error term. Thus, since $\lambda_{t+1} \sim N(0, \sigma_\varepsilon^2)$, *i. i. d.* over time, the average realised risk premium would be statistically equal to zero and, furthermore, there would be neither statistically significant momentum nor predictability of risk premia – as documented in Section 4.III.

In conclusion, both the fundamental structure of the economy –as quantified by the hedging pressure – and market participants’ beliefs – as quantified by the speculator-specific coefficients – can affect the realised risk premia. In order to illustrate the effect of these two potential sources of bias on realised risk premia, we calibrate our model for several alternative specifications and, accordingly, provide a comparison between the obtained results. Note that the simulation exercise focuses on the Panamax BPI 4TC 1-month contract since the evidence of predictability in this case is more significant.

A final note is that we could have modelled the “gambler’s fallacy” bias through a straightforward contrarian investment strategy indicating to go long (short) on the current FFA contract when the realised risk premium is positive (negative), that is, when the short (long) position on the expired FFA contract realises a profit. This would result in a speculator demand function of the form

$$h_t^s = (1 - \vartheta_s) \frac{S_t + (1 - \vartheta_i) \rho_i \kappa_{t+1} - F_t}{\gamma \sigma_\lambda^2} + \vartheta_s \psi_s \frac{S_t - F_{t-1}}{\gamma \sigma_\lambda^2}. \quad (4.22)$$

From a modelling point of view, however, both mechanisms yield the same result; that is a contrarian investment behaviour on behalf of speculators which, in turn, can create the observed form of predictability and momentum in the market.

4.IV.B. Simulation of the Model

In this subsection, we calibrate the economy described above for several different specifications of the model – defined as scenaria – depending on the characteristics of the population and the fundamental structure of the market. Accordingly, for every scenario, we generate 10,000 sample paths using equation 4.6, each one corresponding to 120 periods or, in other words, 10 years. If somewhere in a simulation either the spot rate variable or the FFA rate attain a negative value we discard this path.¹²¹ Finally, we estimate the average of each statistic under consideration across all valid paths and we compare it to its empirical value (Barberis *et al*, 2015a). In particular, we are interested in (i) the predictive power of the FFA basis regarding future spot growth and future risk premia – that is, the slope coefficients, their p-values, and the R^2 s of the regressions, (ii) the predictive power of lagged spot growth and lagged risk premia regarding future risk premia – that is,

¹²¹ We impose this restriction since neither spot nor FFA rates can be negative.

the slope coefficients, their p-values, and the R^2 s of the regressions, (iii) the mean of the FFA log basis and its p-value, (iv) the mean of the FFA log risk premium and its p-value, and (v) the correlation between spot growth and realised risk premia.

4.IV.B.i. Scenario 1: Rational Benchmark without Hedging Pressure

We begin by examining our model's predictions in the simplest case, that is, when all agents are perfectly informed, totally rational and, furthermore, there is no hedging pressure in the FFA market. Recall that the FFA rate and the realised risk premium in this scenario are given by (4.19) and (4.21), respectively. Therefore, we only need to calibrate parameters S_0 , σ_κ^2 , and σ_λ^2 . We set $S_0 = 20$; that is, the initial spot rate is assigned the value of the mean of the spot rate variable (in thousand US dollars, see Panel B of Table 4.1). We set the standard deviations of the private information, σ_κ^2 , and the unpredictable random shock, σ_λ^2 , both equal to 1 to reduce the number of discarded paths but at the same time ensure a sufficient degree of spot price volatility. Note that, in this case, the values of S_0 , σ_κ^2 , and σ_λ^2 per se have no direct impact on the estimation and the results remain qualitatively the same for different plausible values of the parameters.

As expected, the simulation results (Scenario 1 in Table 4.19) suggest that this scenario can neither generate risk premia predictability – and, thus, nor a momentum effect – nor a positive mean basis nor a positive mean realised risk premium. In line with equation 4.20, the reason is that the rationally expected risk premium is zero in this case. The only two statistics qualitatively matched are the negative correlation between spot growth and risk premia and the positive predictability of future spot growth by the current basis. This can be explained by the fact that the basis is an unbiased and, thus, a very accurate predictor of future spot rates; namely, the basis is perfectly positively correlated with the rationally expected future spot rates. Accordingly, an unexpected random shock in spot rates, λ_{t+1} , will result in a shock of the opposite sign in the risk premium (equation 4.21); this, in turn, generates negative correlation between these two variables.

4.IV.B.ii. Scenario 2: Rational Benchmark with Constant Hedging Pressure

The second scenario describes an economy where all agents are perfectly informed and totally rational, however, there exists constant hedging pressure in the FFA market, that is, there is a constant difference in the positions of physical agents. The FFA rate and the realised risk premium are given by (4.16) and (4.18b) with $HP_0 = HP \neq 0$, respectively. Following Barberis *et al* (2015 and 2016), we set the coefficient of risk aversion, γ , equal to 0.1 while for the constant hedging pressure we choose a value that ensures that the simulated average realised risk premium will be close to the observed one (Table 4.1). Namely, we set $HP_0 = HP = -20$, that is, we assume that

long hedgers' – charterers' – positions in the physical market constantly exceed the ones of short hedgers – ship owners.

The simulation results (Scenario 2 in Table 4.19) suggest that neither this specification can generate risk premia predictability by market conditions nor a momentum effect. This can be explained by the fact that the constant negative hedging pressure implies a constant positive rationally expected risk premium *and not a time-varying one* (as implied by equation 4.18a for $HP = -20$). In turn, however, the constant positive rationally expected risk premium results in both a positive mean basis a positive mean realised risk premium (the latter can be shown by taking unconditional expectations on both sides of equation 4.18b, for $HP = -20$). The positive predictability of future spot growth by the current basis can be explained by the fact that the bias in the FFA rate is constant and, thus, the basis is perfectly positively correlated with the rationally expected future spot rates. Following the same line of reasoning, the realised risk premium is negatively correlated with the realised spot growth.

4.IV.B.iii. Scenario 3: Rational Benchmark with Time-Varying Hedging Pressure

In this scenario, all agents are perfectly informed and totally rational as before, however, there exists time-varying hedging pressure in the FFA market. As analysed above, however, we cannot apply the “theory of storage” in shipping to model explicitly the hedging pressure variable and its interdependence with the spot rate process. Furthermore, we do not have data on the hedging pressure variable to empirically examine and, accordingly, conclude about its relationship with spot rates. Therefore, to account for time-varying hedging pressure, we need to assume a stochastic process for the variable. Equivalently, we can impose an assumption, based on plausible economic arguments, that relates hedging pressure to the – exogenously – determined spot rate process. Since hedging pressure is defined as the difference between demand for short hedging positions – which, in turn, is related to fleet supply – and demand for long ones – which, in turn, is related to demand for seaborne trade – one should expect the former and the latter to be negatively and positively related to the corresponding physical market conditions, respectively. Accordingly, one should expect hedging pressure, $HP_t = Q_t - D_t$, to be negatively related to S_t . Hence, hedging pressure's evolution can be indirectly modelled through the evolution of the exogenous spot rate process.

Following the usual convention in the shipping literature (Kalouptsidei, 2014; Greenwood and Hanson, 2015), we assume that the spot rate is determined through an – linear – inverse demand function:

$$S_t = \alpha T_t - \beta F_t, \quad (4.23)$$

where F_t and T_t correspond to the time t available fleet capacity and demand for seaborne services, respectively. The positive coefficients α and β are positively and negatively related to the elasticity of the demand curve, respectively.

Accordingly, we relate hedging pressure to equation 4.23 in a very straightforward manner. Specifically, recall that at t physical market participants determine their hedging demands related to $t + 1$; this corresponds to Q_{t+1} for ship owners and D_{t+1} for charterers. For simplicity, we assume that these variables are equal to the rationally expected values of F_t and T_t , respectively:

$$\begin{cases} Q_{t+1} = E_t^R[F_{t+1}] \\ D_{t+1} = E_t^R[T_{t+1}] \end{cases} \quad (4.24)$$

Importantly, this assumption can be directly related to the signal κ_{t+1} realised at time t ; that is, why physical market participants at t receive a private signal about the spot rate at $t + 1$ with perfect precision. Furthermore, since fleet supply in the short run is highly inelastic, we set F_t and, in turn, Q_t equal to a constant, Q .¹²² This implies that ship owners have a constant hedging demand for FFA contracts. In turn, the evolution of charterers' hedging demand, D_t , can be quantified through equations 4.6, 4.23, and 4.24:

$$D_{t+1} = \frac{\beta}{\alpha} Q + \frac{S_t + \kappa_{t+1}}{\alpha}. \quad (4.25)$$

Therefore, the hedging pressure variable corresponding to $t + 1$ is given by

$$HP_{t+1} = Q_{t+1} - D_{t+1} = \left(1 - \frac{\beta}{\alpha}\right) Q - \frac{S_t + \kappa_{t+1}}{\alpha}. \quad (4.26)$$

Thus, hedging pressure is a decreasing function of both current market conditions and the signal about future market conditions. Plugging in (4.16) equation 4.26 yields the expression for the rational expectations time-varying hedging pressure FFA rate:

$$F_t^{R,HP} = \left(1 + \frac{\gamma\sigma_\lambda^2}{3\alpha}\right) (S_t + \kappa_{t+1}) - \frac{\gamma\sigma_\lambda^2}{3} \left(1 - \frac{\beta}{\alpha}\right) Q. \quad (4.27)$$

Finally, the rationally expected bias in the FFA rate is given by

¹²² Note that while this simplifying assumption can be easily relaxed it does not have any qualitative or quantitative implication on the model.

$$F_t^{R,HP} - E_t^R[S_{t+1}] = \frac{\gamma\sigma_\lambda^2}{3\alpha}(S_t + \kappa_{t+1}) - \frac{\gamma\sigma_\lambda^2}{3}\left(1 - \frac{\beta}{\alpha}\right)Q, \quad (4.28)$$

while the realised one equals

$$F_t^{R,HP} - S_{t+1} = \frac{\gamma\sigma_\lambda^2}{3\alpha}(S_t + \kappa_{t+1}) - \frac{\gamma\sigma_\lambda^2}{3}\left(1 - \frac{\beta}{\alpha}\right)Q - \lambda_{t+1}. \quad (4.29)$$

Equations 4.27, 4.28, and 4.29 suggest that the FFA rate, the rationally expected bias, and the realised bias are all increasing functions of both current spot rates and the signal about future market conditions.

Accordingly, we calibrate parameters α , β , D_0 , and Q in the following manner. Equations 4.23 and 4.24 imply that $S_0 = \alpha T_0 - \beta Q \Rightarrow S_0 = \alpha D_0 - \beta Q$ while from the constant hedging pressure case we have defined $HP = HP_0 = Q - D_0 = -20$. Thus, assuming $D_0 = 100$ yields $Q = 80$. In turn, since $S_0 = 20$ we can calibrate α and β from $20 = 100\alpha - 80\beta$; therefore, setting $\beta = 0.1$ yields $\alpha = 0.28$. Finally, in order to illustrate how spot rates and hedging pressure are determined and interrelated, assume that at $t = 0$ the signal κ_1 equals 1. Hence, the rationally expected spot price at $t = 1$ is equal to 21 and the long hedging demand related to $t = 1$, D_1 , equals 103.5714. In turn, the corresponding hedging pressure variable, HP_1 , becomes -23.5714, that is, it decreases by 3.5714.

The simulation results (Scenario 3 in Table 4.19) suggest that while this specification provides a better approach for the observed regularities compared to the previous two, it cannot simultaneously match two of the most important stylised facts, that is, the momentum effect and the negative predictability of risk premia by lagged spot market conditions. Regarding the former, we observe that, while the coefficient in the lagged risk premia regression appears to be positive, it remains statistically insignificant at any conventional level. What is more, the coefficient in the lagged spot growth regression is positive (although statistically insignificant). Note that even if we recalibrate the coefficients – namely, the variance of the random shock, σ_λ^2 – to obtain significant slope coefficients in the lagged risk premium regression, there will still be no negative predictability of future risk premia by past market conditions.

This result can be easily justified by examining equation 4.29 at $t + 1$:

$$\begin{aligned} F_{t+1}^{R,HP} - S_{t+2} &= \frac{\gamma\sigma_\lambda^2}{3\alpha}(S_{t+1} + \kappa_{t+2}) - \frac{\gamma\sigma_\lambda^2}{3}\left(1 - \frac{\beta}{\alpha}\right)Q - \lambda_{t+2} \\ &= \frac{\gamma\sigma_\lambda^2}{3\alpha}(S_t + \Delta S_{t+1} + \kappa_{t+2}) - \frac{\gamma\sigma_\lambda^2}{3}\left(1 - \frac{\beta}{\alpha}\right)Q - \lambda_{t+2}, \end{aligned}$$

where by $\Delta S_{t+1} = S_{t+1} - S_t$ we denote the change in the spot rate. Therefore, we observe that, ceteris paribus, the realised risk premium is an increasing function of lagged spot rate changes. In turn, this explains the non-negative predictability of risk premia by lagged spot growth in this scenario.

In a similar manner, the positive sign in the lagged risk premium regression can be explained if we restate equation 4.29 at $t + 1$ in the following manner

$$F_{t+1}^{R,HP} - S_{t+2} = (F_t^{R,HP} - S_{t+1}) + \frac{\gamma\sigma_\lambda^2}{3\alpha} \kappa_{t+2} + \left(1 + \frac{\gamma\sigma_\lambda^2}{3\alpha}\right) \lambda_{t+1} - \lambda_{t+2}.$$

The explanation for the remaining simulation results directly follows from the analysis in the previous two scenarios.

4.IV.B.iv. Scenario 4: Distorted Expectations and Constant Hedging Pressure

The fourth scenario corresponds to the economy with asymmetric information and irrationality of beliefs. In addition, we assume that there is constant hedging pressure in the market as in scenario 2. Note that all predictive regression results remain qualitatively the same if we set hedging pressure equal to zero. In this case, the FFA rate and the realised risk premium are given by equations 4.14 and 4.17, respectively, for $HP_0 = HP = -20$. Accordingly, we examine several parameterisations for the speculator specific parameters, $\{\vartheta_s, \rho_s, \psi_s\}$.

In the following, we present and discuss the results for the set $\{0.9, 0.5, 1\}$; namely, we allow speculators to “worry” about the “fundamental value signal” but weigh more heavily the “contrarian value” one (Barberis *et al*, 2016). In addition, we assume that they receive the private value signal with 50% precision; thus, there is asymmetry of information in the market. Finally, we set the “degree of gambler’s fallacy” equal to 1, implying that speculators believe that the last spot price shock will be immediately cancelled out. For the ease of reference, Table 4.18 summarises all relevant parameter values.

The corresponding simulation results (Scenario 4 in Table 4.19) suggest that this specification can match simultaneously almost all stylised facts. Most importantly, we observe that it can account not only for the momentum effect – the lagged risk premia coefficient is positive and statistically significant – but also for the negative predictability of future risk premia by lagged spot growth – the lagged spot growth coefficient is negative and statistically significant.

Table 4.18: Parameter values.

Parameter	Assigned Value
-----------	----------------

S_0	20
σ_k^2	1
σ_λ^2	{1,2.5}
γ	{0.04,0.1}
HP_0	-20
D_0	100
Q	80
α	0.28
β	0.1
ϑ_s	0.9
ρ_s	0.5
ψ_s	1

Notes: This table summarises the assigned values regarding the initial level of the spot rate variable, S_0 ; the variance of the private signal, σ_k^2 ; the variance of the unexpected error term, σ_λ^2 ; the coefficient of risk aversion, γ ; the initial level of the hedging pressure variable, HP_0 ; the initial level of the long hedging demand variable, D_0 ; the level of the short hedging demand variable, Q ; the two coefficients related to the linear inverse demand function, α and β ; the “degree of wavering”, ϑ_s ; the coefficient of precision, ρ_s ; and the “degree of gambler’s fallacy”, ψ_s .

The latter feature can be explained simply by examining equation 4.17 at $t + 1$:

$$F_{t+1}^* - S_{t+2} = \left[\frac{(1 - \vartheta_s)\rho_s - 1}{3} \kappa_{t+2} + \frac{\vartheta_s \psi_s}{3} (-\Delta S_{t+1}) \right] - \frac{\gamma \sigma_\lambda^2}{3} HP - \lambda_{t+2}. \quad (4.30)$$

Namely, the realised risk premium is, *ceteris paribus*, a decreasing function of lagged spot rate changes, $\Delta S_{t+1} = S_{t+1} - S_t$. Furthermore, equation 4.17 at t can be re-expressed as

$$-\Delta S_{t+1} = F_t^* - S_{t+1} - \left[\frac{(1 - \vartheta_s)\rho_s + 2}{3} \kappa_{t+1} + \frac{\vartheta_s \psi_s}{3} (S_{t-1} - S_t) \right] + \frac{\gamma \sigma_\lambda^2}{3} HP. \quad (4.31)$$

Plugging (4.31) in (4.30) we observe that the realised risk premium at $t + 1$ is an increasing function of the realised risk premium at t .

Note, that the only stylised facts poorly matched in this case are the ones related to the variance decomposition (Scenario 4 in Panel B of Table 4.19) since essentially none of basis variation is attributed to time-varying risk premia. This result can be explained by the fact that the “contrarian value signal” significantly reduces the volatility of the realised risk premia. If we increase, however, either the variance of the unexpected shock, σ_λ^2 , or the “degree of fallacy”, we can match sufficiently well also this regularity. The former adjustment is presented in scenario 4’ of Table 4.19 for $\sigma_\lambda^2 = 2.5^2$ and $\gamma = 0.04$.

Furthermore, note that when there is irrationality of beliefs but no information asymmetry, the results are qualitatively very similar to the ones above. Finally, the case with asymmetric information and rational beliefs closely resembles scenario 2. Hence, we can argue that the “gambler’s fallacy” feature – that is, the behavioural bias component – appears to be the most crucial source of heterogeneity of beliefs. In turn, according to our simulation results, the contrarian investment behaviour on behalf of a population fraction is the main determinant of the observed risk premia predictability.

4.IV.B.v. Scenario 5: Distorted Expectations and Time-Varying Hedging Pressure

This last case combines the features of scenaria 3 and 4; namely, it corresponds to the economy with asymmetric information, irrationality of beliefs, and time-varying hedging pressure. In line with scenario 4, we present and discuss the results for the speculator-parameterisation $\{0.9, 0.5, 1\}$. In this case, the equilibrium FFA rate is obtained by plugging in (4.14) the expression for hedging pressure in (4.26):

$$F_t^* = S_t + \frac{2 + (1 - \vartheta_s)\rho_s}{3} \kappa_{t+1} + \frac{\vartheta_s \psi_s}{3} (S_{t-1} - S_t) - \frac{\gamma \sigma_\lambda^2}{3} \left[\left(1 - \frac{\beta}{\alpha}\right) Q - \frac{S_t + \kappa_{t+1}}{\alpha} \right]. \quad (4.32)$$

Accordingly, the rationally expected bias and the realised one are given by

$$F_t^* - E_t^R[S_{t+1}] = \frac{(1 - \vartheta_s)\rho_s - 1}{3} \kappa_{t+1} + \frac{\vartheta_s \psi_s}{3} (S_{t-1} - S_t) - \frac{\gamma \sigma_\lambda^2}{3} \left[\left(1 - \frac{\beta}{\alpha}\right) Q - \frac{S_t + \kappa_{t+1}}{\alpha} \right] \quad (4.33)$$

and

$$F_t^* - S_{t+1} = \frac{(1 - \vartheta_s)\rho_s - 1}{3} \kappa_{t+1} + \frac{\vartheta_s \psi_s}{3} (S_{t-1} - S_t) - \frac{\gamma \sigma_\lambda^2}{3} \left[\left(1 - \frac{\beta}{\alpha}\right) Q - \frac{S_t + \kappa_{t+1}}{\alpha} \right] - \lambda_{t+1}, \quad (4.34)$$

respectively. The corresponding simulation results (scenario 5 in Table 4.19) suggest that this specification can simultaneously match most observed regularities in a sufficient manner, albeit, worse than scenario 4. This result was expected since this scenario combines the features of the previous two economies.

Of course, one can obtain values closer to the actual ones either through finer adjustment of the set of parameters or by using exact closed-form expressions for the moments of interest. However, the obtained results will be very similar to the ones described in scenario 4 and, in turn, the economic intuition will be the same.

In conclusion, both the theoretical predictions and the simulation of our model suggest that, in order to simultaneously match all observed regularities sufficiently well, one has to depart from the rational benchmark of the economy since the – time-varying – hedging pressure dimension alone cannot capture the negative predictability of risk premia by lagged market conditions. While the predictions are not particularly sensitive to the degree of information asymmetry this is not true for the behavioural bias feature; namely, a fraction of investors must suffer from the “gambler’s fallacy” and, in turn, follow a contrarian investment strategy.

4.V. Conclusion

This chapter examines the formation of FFA rates in the dry bulk shipping industry. Specifically, our empirical analysis concentrates upon the Capesize BCI 4TC and Panamax BPI 4TC monthly contracts. We illustrate that the bulk of volatility in the FFA basis can be attributed to expectations

about future physical market conditions rather than expectations about future risk premia, as is commonly suggested in the commodity finance literature. Furthermore, we provide both an economic interpretation of this result and a comparison to the ones obtained from other futures and forward markets. Our results validate and extend the economic arguments presented in the seminal commodity market papers that examine the forecasting power of derivative contracts. Namely, predictability of spot rates appears to be an increasing function of the commodity cost of storage. In shipping, where the commodity is a service – hence, non-storable – and the industry is subject to significant supply and demand shocks which cannot be attenuated through adjustments of the short-term supply, we observe predictable variation of spot rates and, in turn, substantial forecasting ability on behalf of the FFA rates.

Despite this finding, though, there appears to be a bias in the FFA rates in the form of both a strong momentum effect and significant predictability of risk premia by lagged price-based signals and economic variables that reflect recent changes in the physical market conditions. An additional interesting finding of this chapter is the evidence of “contango” in the FFA market. Furthermore, we examine whether future market conditions and risk premia can be predicted by market activity variables that incorporate the FFA trading volume and open interest figures related to the corresponding contracts. While there appears to exist some sort of predictability, especially in the Capesize sector, the results cannot yet be generalised given the small size of the incorporated trading activity dataset.

Importantly, the existence of statistically significant predictability of future risk premia contradicts the unbiased expectations hypothesis and, in turn, the efficiency of the FFA market. We further examine the validity of the unbiasedness hypothesis by performing three frequently incorporated econometric tests. The obtained results unequivocally suggest that there exists a bias in the formation of the 1-month FFA rates in both contracts. Regarding the 2-month contracts, our findings point towards the existence of a bias, especially in the Panamax BPI 4TC case.

We further contribute to the literature by developing a behavioural asset pricing framework that, among other features, can explain both the existence of momentum and the documented predictability of future risk premia by lagged physical market conditions. Since vessels are non-storable commodities, the proposed framework departs from the “theory of storage” and the “cost-of-carry” model. Importantly, our dynamic framework can simultaneously account for both the familiar “hedging pressure” feature – the rational dimension – and a heterogeneous beliefs explanation – the irrational dimension. The proposed model incorporates three types of agent: ship owners, charterers, and speculators.

The distinct feature of our framework is that, apart from having – as is standard in the literature – different objective functions, agents might also differ in the way they form expectations about future market conditions. Specifically, we develop an asymmetric information environment where speculators suffer from a behavioural bias known as “the law of small numbers” – or, equivalently, “reversion to the mean” or “gambler’s fallacy”. Accordingly, it is illustrated formally that, to simultaneously match the observed regularities, one must depart from the rational expectations benchmark of the model since the – time-varying – hedging pressure dimension alone cannot capture the negative predictability of risk premia by lagged market conditions. While the predictions are not particularly sensitive to the degree of information asymmetry this is not true for the behavioural bias feature; namely, a fraction of investors must suffer from the “gambler’s fallacy” and, in turn, follow a contrarian investment strategy.

To the best of our knowledge, the FFA market had never been examined from the perspective of a structural behavioural economic model before. In addition, we contribute to the generic commodity finance literature by incorporating explicitly the behavioural dimension in the formation of derivative contracts rates.

Appendix 4

A.4.A. Complementary Results for the 3- and 4-Month Maturity Contracts

Table 4.A1: Descriptive statistics for the variables of interest.

x	T	n	Start	End	\bar{x}	MD	SD	CV	Max	Min	ρ_1	ρ_2	ρ_{12}
Panel A: Variables in Levels (in '000 \$) for the Capesize Sector (BCI 4TC)													
FFA3	3	76	3.10	6.16	14.3	12.0	8.2	0.57	37.1	3.2	0.91	0.82	0.11
FFA4	4	44	10.12	5.16	12.1	9.3	6.4	0.53	26.2	4.2	0.84	0.71	-0.06
Panel B: Variables in Levels (in '000 \$) for the Panamax Sector (BPI 4TC)													
FFA3	3	76	3.10	6.16	10.8	9.2	5.9	0.54	29.0	4.0	0.96	0.92	0.55
FFA4	4	44	10.12	5.16	7.9	7.6	2.5	0.32	14.1	4.3	0.92	0.78	-0.12
Panel C: Variables in Log Differences for the Capesize Sector (BCI 4TC)													
b	3	76	3.10	6.16	0.22	0.31	0.54	-	1.39	-0.93	0.49	0.13	0.36
Δs	3	76	6.10	9.16	-0.03	0.07	0.83	-	1.88	-1.76	0.58	0.05	0.44
r	3	76	6.10	9.16	0.25	0.32	0.63	-	1.85	-1.03	0.65	0.30	0.10
b	4	44	10.12	5.16	0.26	0.38	0.63	-	1.61	-1.02	0.47	0.15	0.48
Δs	4	44	2.13	9.16	-0.02	-0.05	1.03	-	1.78	-2.33	0.56	0.30	0.40
r	4	44	2.13	9.16	0.28	0.32	0.68	-	2.08	-1.26	0.62	0.41	-0.04
Panel D: Variables in Log Differences for the Panamax Sector (BPI 4TC)													
b	3	76	3.10	6.16	0.11	0.08	0.28	-	1.00	-0.41	0.52	0.09	0.15
Δs	3	76	6.10	9.16	-0.03	-0.06	0.38	-	0.85	-0.83	0.52	-0.01	0.11
r	3	76	6.10	9.16	0.15	0.14	0.33	-	0.89	-0.58	0.69	0.35	-0.15
b	4	44	10.12	5.16	0.16	0.13	0.31	-	1.12	-0.48	0.62	0.24	0.03
Δs	4	44	2.13	9.16	0.01	0.01	0.48	-	0.85	-0.93	0.68	0.30	0.18
r	4	44	2.13	9.16	0.15	0.17	0.44	-	1.02	-0.67	0.81	0.54	-0.07

Notes: Panels A-B present descriptive statistics for the levels of the FFA rates corresponding to the 3- and 4-month BCI 4TC and BPI 4TC FFA contracts. These variables (x) are expressed in thousand U.S. dollars. Panels C-D present descriptive statistics for the basis, b , the spot growth, Δs , and the risk premium, r , corresponding to the 3- and 4-month BCI 4TC and BPI 4TC FFA contracts. These variables (x) are expressed in log differences. The maturity of the contract and the number of observations are denoted by T and n , respectively. The first and last months of the variable in our sample analysis are indicated by columns 4 and 5 (labelled "Start" and "End"), respectively (e.g., 3.10 refers to March 2010). The included statistics are the mean (\bar{x}), median (MD), standard deviation (SD), coefficient of variation (CV), maximum (max), minimum (min), and 1-month (ρ_1), 2-month (ρ_2), and 12-month (ρ_{12}) autocorrelation coefficients.

Table 4.A2: Correlation matrix.

Variable	Levels						Log Differences
	Spot	Settlement	FFA1	FFA2	FFA3	FFA4	$Corr(\Delta s, r)$
Panel A: Capesize Sector (BCI 4TC)							
FFA3	0.75	0.59	0.87	0.96	1.00		-0.76
FFA4	0.44	0.41	0.64	0.82	0.96	1.00	-0.81
Panel B: Panamax Sector (BPI 4TC)							
FFA3	0.92	0.84	0.97	0.99	1.00		-0.71
FFA4	0.65	0.18	0.78	0.90	0.98	1.00	-0.78

Notes: Panels A and B of this table correspond to the BCI 4TC and BPI 4TC FFA contracts, respectively. Columns 2-7 present the correlation coefficients for spot, settlement, and FFA rates. All these variables are in levels. The last column presents the corresponding correlation coefficients for the log spot growth, Δs , and the log risk premium, r , for the 3- and 4-month BCI 4TC and BPI 4TC FFA contracts. The latter two variables are expressed in log differences.

Table 4.A3: Significance of FFA bases and risk premia.

T	n	Mean Basis	t of Basis	An. Mean Premium	An. SD Premium	t^{NW} of Premium
Panel A: Capesize Sector (BCI 4TC)						
3	76	22.43%	3.62	101.43%	125.48%	1.96
4	44	25.93%	2.75	82.55%	118.57%	1.45
Panel B: Panamax Sector (BPI 4TC)						
3	76	11.48%	3.63	88.61%	66.68%	2.11
4	44	16.32%	3.50	92.93%	76.14%	1.43

Notes: This table presents descriptive statistics related to FFA bases and risk premia for the Capesize BCI 4TC and Panamax BPI 4TC contracts with maturities equal to 3 and 4 months. The maturity of the contract and the number of observations are denoted by T and n , respectively. The included statistics are the mean and t-statistic of the basis, the annualised mean and standard deviation of risk premium, and the t-statistic, t^{NW} of the risk premium. To deal with the overlapping nature of risk premia, the corresponding t-statistics are estimated using the Newey-West (1987) HAC correction. When the t-statistic indicates significance at least at the 10% level, the respective mean statistic appears in bold.

Table 4.A4: Regressions of future risk premia and spot growth on current FFA basis.

Variable	T	n	α	t^{NW}	β	t^{NW}	R^2	Wald Test
Panel A: Capesize Sector (BCI 4TC)								
Δs	3	76	-0.26	-1.66	1.02***	5.57	0.43	4.61*
r	3	76	0.26	1.66	-0.02	-0.10	0.00	
Δs	4	44	-0.34	-1.47	1.25***	5.84	0.58	2.17
r	4	44	0.34	1.47	-0.25	-1.18	0.05	
Panel B: Panamax Sector (BPI 4TC)								
Δs	3	76	-0.12*	-1.84	0.74***	4.50	0.28	9.72***
r	3	76	0.12*	1.84	0.26	1.58	0.05	
Δs	4	44	-0.10	-0.93	0.69***	3.27	0.20	11.68***
r	4	44	0.10	0.93	0.31	1.46	0.05	

Notes: Panels A-B report the results from 3- and 4-month horizon OLS forecasting regressions of future spot growth, Δs , and risk premia, r , on the current basis for the Capesize BCI 4TC and Panamax BPI 4TC contracts, respectively. To deal with the overlapping nature of returns and growth rates, t-statistics are estimated using the Newey-West (1987) HAC correction. The maturity of the contract and the number of observations are denoted by T and n , respectively. The intercept, α , and the slope coefficient, β , are accompanied by *, **, or *** when the absolute t^{NW} statistic indicates significance at the 10%, 5% or 1% level, respectively. In addition, the last column reports the Chi-square statistic associated with a Wald Coefficient Test on the restrictions $\alpha = 0$ and $\beta = 1$ in regression 4.3a of the main text.

Table 4.A5: Regressions of future spot growth on lagged spot growth and current FFA basis.

T	n	Lagged Spot Growth			FFA Basis		
		β	t^{NW}	R^2	β	t^{NW}	R^2
Panel A: Capesize Sector (BCI 4TC)							
3	76	-0.28**	-2.22	0.02	1.02***	5.57	0.43
4	44	-0.52***	-4.39	0.05	1.25***	5.84	0.58
Panel B: Panamax Sector (BPI 4TC)							
3	76	-0.36**	-2.56	0.03	0.74***	4.50	0.28
4	44	-0.26**	-2.24	0.01	0.69***	3.27	0.20

Notes: Panels A-B report 3- and 4-month horizon OLS forecasting regressions of future spot growth, Δs , on one period lagged 1-month spot growth and the current basis for the Capesize BCI 4TC and Panamax BPI 4TC contracts, respectively. Spot growth is defined as the log of the ratio of the settlement rate to the spot price at the end of the previous month. To deal with the overlapping nature of returns and growth rates, t-statistics are estimated using the Newey-West (1987) HAC correction. The maturity of the contract and the number of observations are denoted by T and n , respectively. The slope coefficient, β , is accompanied by *, **, or *** when the absolute t^{NW} statistic indicates significance at the 10%, 5% or 1% level, respectively.

Table 4.A6: Regressions of future risk premia on lagged risk premia.

Variable	$f(t, 3) - s(t + 3)$				$f(t, 4) - s(t + 4)$			
	T	β	t^{NW}	R^2	T	β	t^{NW}	R^2
Panel A: Capesize Sector (BCI 4TC)								
$f(t - 1, 1) - s(t)$	76	0.36	1.51	0.04	44	0.02	0.06	0.00
$f(t - 2, 1) - s(t - 1)$	76	-0.08	-0.41	0.00	44	-0.04	-0.17	0.00
$f(t - 3, 1) - s(t - 2)$	76	-0.26	-1.10	0.02	44	-0.53	-1.31	0.07
$f(t - 4, 1) - s(t - 3)$	76	-0.47**	-2.01	0.06	44	-0.55***	-3.13	0.06
$f(t - T, T) - s(t)$	73	0.09	0.62	0.01	40	-0.04	-0.20	0.00
Panel B: Panamax Sector (BPI 4TC)								
$f(t - 1, 1) - s(t)$	76	0.55*	1.93	0.07	44	0.61**	2.15	0.03
$f(t - 2, 1) - s(t - 1)$	76	0.34	1.49	0.03	44	0.12	0.34	0.00
$f(t - 3, 1) - s(t - 2)$	76	-0.05	-0.16	0.00	44	-0.41	-0.92	0.03
$f(t - 4, 1) - s(t - 3)$	76	-0.33	-1.09	0.02	44	-0.47	-1.04	0.06
$f(t - T, T) - s(t)$	73	0.18	1.36	0.02	40	-0.05	-0.34	0.00

Notes: Panels A-B report 3- and 4-month horizon OLS forecasting regressions of future risk premia, $f(t, T) - s(t + T)$, on lagged risk premia, for the Capesize BCI 4TC and Panamax BPI 4TC contracts, respectively. Namely, in the first four rows of each panel, the predictor is the lagged one-period risk premium, $f(t - l, 1) - s(t - l + 1)$; that is, the lagged risk premium related to the one-month contract where the number of lags, l , varies from 1 to 4. In the fifth row, the predictor is the corresponding previous risk premium for each contract, $f(t - T, T) - s(t)$; e.g., for the 3-month contract expiring in $t + 3$ months, the predictor is the realised risk premium related to the three-month contract that expired at t . The maturity of the contract and the number of observations are denoted by T and n , respectively. To deal with the overlapping nature of the variables, t-statistics are estimated using the Newey-West (1987) HAC correction. The slope coefficient, β , is accompanied by *, **, or *** when the absolute t^{NW} statistic indicates significance at the 10%, 5% or 1% level, respectively.

Table 4.A7: Regressions of future risk premia on lagged spot growth.

Variable	$f(t, 3) - s(t + 3)$				$f(t, 4) - s(t + 4)$			
	T	β	t^{NW}	R^2	T	β	t^{NW}	R^2
Panel A: Capesize Sector (BCI 4TC)								
$s(t) - s(t - 1)$	76	-0.28*	-1.97	0.03	44	-0.08	-0.42	0.00
$s(t - 1) - s(t - 2)$	76	-0.10	-0.69	0.00	44	-0.08	-0.84	0.00
$s(t - 2) - s(t - 3)$	76	0.03	0.27	0.00	44	0.17	0.44	0.01
$s(t - 3) - s(t - 4)$	76	0.09	0.47	0.00	44	0.27*	1.75	0.02
$s(t) - s(t - T)$	73	-0.12	-1.07	0.02	40	0.06	0.39	0.01
$sp(t) - s(t - 1)$	76	-0.19**	-2.51	0.04	44	-0.05	-0.32	0.00
$sp(t - 1) - s(t - 2)$	76	-0.10	-0.96	0.01	44	-0.02	-0.45	0.00
$sp(t - 2) - s(t - 3)$	76	0.01	0.11	0.00	44	0.07	0.46	0.00
$sp(t - 3) - s(t - 4)$	76	0.04	0.42	0.00	44	0.15	1.46	0.02
$sp(t) - s(t - T)$	73	-0.13	-1.50	0.04	40	0.05	0.31	0.00
Panel B: Panamax Sector (BPI 4TC)								
$s(t) - s(t - 1)$	76	-0.22	-1.27	0.01	44	-0.41***	-2.71	0.03
$s(t - 1) - s(t - 2)$	76	-0.30**	-2.45	0.03	44	-0.25	-1.35	0.01
$s(t - 2) - s(t - 3)$	76	-0.15	-1.05	0.01	44	0.22	1.27	0.01
$s(t - 3) - s(t - 4)$	76	0.29*	1.81	0.02	44	0.62**	2.37	0.07
$s(t) - s(t - T)$	73	-0.23**	-2.37	0.07	40	0.08	0.55	0.01
$sp(t) - s(t - 1)$	76	-0.09	-0.84	0.01	44	-0.29**	-2.10	0.04
$sp(t - 1) - s(t - 2)$	76	-0.22**	-2.50	0.04	44	-0.19	-1.46	0.02
$sp(t - 2) - s(t - 3)$	76	-0.08	-0.95	0.01	44	0.13	0.99	0.01
$sp(t - 3) - s(t - 4)$	76	0.22**	2.42	0.04	44	0.41**	2.43	0.08
$sp(t) - s(t - T)$	73	-0.19*	-1.90	0.06	40	0.02	0.18	0.00

Notes: Panels A-B report 3- and 4-month horizon OLS forecasting regressions of future risk premia, $f(t, T) - s(t + T)$, on lagged spot growth, for the Capesize BCI 4TC and Panamax BPI 4TC contracts, respectively. Namely, in the first four rows of each panel the predictor is the lagged one-period spot growth $s(t - l) - s(t - l - 1)$; that is, the one-month lagged spot growth, where the number of lags, l , varies from 1 to 4. In the fifth row, the predictor is the corresponding previous spot growth for each contract, $s(t) - s(t - T)$; e.g., for the 3-month contract expiring in $t + 3$ months, the predictor is the realised spot growth related to the three-month contract that expired at t , that is, the one corresponding to period $t - 3$ to t . In rows 6-10 of each panel we perform the same set of regressions as in the first five rows, with the only difference being that spot growth is estimated using the respective daily spot rate, $sp(t)$, as the final spot price instead of the current settlement rate. The maturity of the contract and the number of observations are denoted by T and n , respectively. To deal with the overlapping nature of the variables, t-statistics are estimated using the Newey-West (1987) HAC correction. The slope coefficient, β , is accompanied by *, **, or *** when the absolute t^{NW} statistic indicates significance at the 10%, 5% or 1% level, respectively.

Table 4.A8: Descriptive statistics for trading volume and open interest variables.

Variable	T	n	Mean	Median	SD	Max	Min	ρ_1
Panel A: Capesize Sector (BCI 4TC)								
Trading Volume Growth	3	41	0.00	0.05	0.51	1.25	-1.51	-0.32
	4	40	0.00	0.02	0.51	0.74	-1.62	-0.36
Open Interest Growth	3	41	0.02	0.09	0.21	0.23	-0.87	-0.20
	4	40	0.02	0.09	0.23	0.24	-0.93	-0.28
Open Interest MA	3	39	0.03	0.09	0.23	0.27	-0.76	0.41
	4	38	0.02	0.08	0.23	0.26	-0.80	0.32
Panel B: Panamax Sector (BPI 4TC)								
Trading Volume Growth	3	41	0.00	0.01	0.41	0.78	-0.91	-0.45
	4	40	-0.01	-0.01	0.48	0.74	-1.43	-0.46
Open Interest Growth	3	41	0.02	0.09	0.19	0.29	-0.54	-0.30
	4	40	0.02	0.10	0.20	0.29	-0.63	-0.37
Open Interest MA	3	39	0.02	0.06	0.15	0.22	-0.46	0.16
	4	38	0.02	0.06	0.17	0.21	-0.52	0.12

Notes: This table presents descriptive statistics for the 1-month trading volume growth, 1-month open interest growth, and the logarithm of current open interest scaled by the moving average (MA) of open interest over the previous three months. Panels A and B correspond to the BCI 4TC and BPI 4TC FFA contracts, respectively, for the 3- and 4-month maturities. The maturity of the contract and the number of observations are denoted by T and n , respectively. The included statistics are the mean, median, standard deviation, maximum, minimum, and 1-month autocorrelation coefficients.

Table 4.A9: Regressions of future risk premia and spot growth on trading activity variables.

Variable	T	n	Δs			r		
			β	t^{NW}	R^2	β	t^{NW}	R^2
Panel A: Capesize Sector (BCI 4TC)								
Trading Volume Growth	3	41	0.50***	4.07	0.08	-0.25***	-3.16	0.04
	4	40	0.27*	1.93	0.02	-0.05	-0.49	0.00
Open Interest Growth	3	41	0.95**	2.46	0.06	-0.44***	-3.11	0.02
	4	40	1.13***	3.58	0.06	-0.33**	-2.25	0.01
Scaled Open Interest MA	3	39	1.91***	4.97	0.26	-1.18***	-2.98	0.17
	4	38	2.42***	6.11	0.28	-1.07***	-3.13	0.12
Panel B: Panamax Sector (BPI 4TC)								
Trading Volume Growth	3	41	0.17*	1.84	0.03	-0.02	-0.24	0.00
	4	40	0.01	0.21	0.00	0.00	-0.04	0.00
Open Interest Growth	3	41	0.05	0.18	0.00	0.10	0.54	0.00
	4	40	0.10	0.70	0.00	-0.06	-0.48	0.00
Scaled Open Interest MA	3	39	0.62	1.27	0.05	-0.40	-0.96	0.02
	4	38	0.80***	3.02	0.07	-0.56*	-1.98	0.04

Notes: This table reports 3- and 4-month horizon OLS forecasting regressions of future spot growth, Δs , and risk premia, r , on 1-month trading volume growth, 1-month open interest growth, and the logarithm of current open interest scaled by the moving average of open interest over the previous three months. Panels A and B correspond to the Capesize BCI 4TC and Panamax BPI 4TC contracts, respectively. In the trading activity case, spot growth corresponds to settlement growth. The maturity of the contract and the number of observations are denoted by T and n , respectively. t-statistics are estimated using the Newey-West (1987) HAC correction. The slope coefficient, β , is accompanied by *, **, or *** when the absolute t^{NW} statistic indicates significance at the 10%, 5% or 1% level, respectively.

Table 4.A10: Contemporaneous regressions of settlement growth on trading volume growth.

Contract	n	β	t	R^2	ρ
Panel A: Capesize Sector (BCI 4TC)					
3-month Contract	41	0.28	1.66	0.07	0.26
4-month Contract	40	0.18	1.00	0.03	0.16
Panel B: Panamax Sector (BPI 4TC)					
3-month Contract	41	0.03	0.28	0.00	0.04
4-month Contract	40	0.03	0.35	0.00	0.06

Notes: Panels A-B report contemporaneous regressions of 1-month settlement growth on 1-month trading activity growth for the Capesize BCI 4TC and Panamax BPI 4TC 3- and 4-month contracts, respectively. The number of observations and the correlation coefficient are denoted by n , and ρ , respectively. The slope coefficient, β , is accompanied by *, **, or *** when the absolute t -statistic indicates significance at the 10%, 5% or 1% level, respectively.

Table 4.A11: Correlation between trading volume, open interest, spot, and settlement rates.

Contract	TV and Spot	TV and Settlement	OI and Spot	OI and Settlement
Panel A: Capesize Sector (BCI 4TC)				
3-month Contract	0.64	0.68	0.02	0.04
4-month Contract	0.50	0.47	-0.16	-0.17
Panel B: Panamax Sector (BPI 4TC)				
3-month Contract	0.33	0.44	-0.14	-0.13
4-month Contract	0.41	0.45	-0.12	-0.16

Notes: Panels A-B present correlation coefficients for the following pairs of variables: trading volume (TV) and spot rates, trading volume and settlement rates, open interest (OI) and spot rates, and open interest and settlement rates.

Chapter 5: Conclusion

This thesis examined the dry bulk sector of the shipping industry. We began by analysing the relation between second-hand vessel prices, net earnings, and holding period returns. Specifically, we provided strong statistical evidence that almost the entire volatility of shipping earnings yields can be attributed to variation in expected net earnings growth; almost none to expected returns variation and almost none to varying expectations about the terminal earnings yield. According to our results, earnings yields are negatively and significantly related to future net earnings growth. Furthermore, we found no consistent, strong statistical evidence supporting the existence of time-varying risk premia in the valuation of dry bulk vessels.

From an economic point of view, our analysis suggested that in order for valuation ratios to significantly predict future cash flows, current cash flows must have a profound second-order effect on the current price of the asset through the future cash flow stream. From a statistical perspective, the significant predictability of earnings growth by the earnings yield is driven by the extreme volatility of shipping net earnings.

Accordingly, we integrated the examination of the second-hand market by incorporating in the analysis the trading activity related to dry bulk vessels. For this purpose, we developed a heterogeneous expectations asset pricing model that – among other stylised facts – can account for the actual behaviour of vessel prices and the positive correlation between net earnings, vessel prices, and second-hand vessel transactions. The proposed economy consists of two agent types, conservatives and extrapolators, who form heterogeneous expectations about future net earnings and at the same time under (over) estimate the future demand responses of their competitors. Formal estimation of the model suggested that, to simultaneously match the empirical regularities, the average investor expectations in the second-hand market for ships must be “near-rational”. In particular, the investor population must consist of a very large fraction of agents with totally – or very close to – rational beliefs while the remaining ones must hold highly extrapolative beliefs; thus, there must exist significant heterogeneity of beliefs in the market.

From an economic perspective, this finding is in accordance with the nature of the shipping industry; namely, the large fraction of conservative investors corresponds to the large number of established shipping companies that operate in the industry, having strong prior experience and expertise about the freight rate mechanism. In turn, their superior knowledge translates into more accurate forecasts about future market conditions compared to relatively new investors. Extrapolators, on the other hand, reflect new entrants such as private equity firms with little or no previous experience of the market. It is well-documented that during prosperous periods, new

entrants, impressed by the high prevailing earnings and short-term returns, are eager to buy vessels which, subsequently, are more than keen to sell as conditions begin to deteriorate. In contrast, there are many cases where traditional, established owners have realised significant returns by selling vessels at the peak of the market and buying at the trough.

Having concluded the analysis of the physical shipping market for second-hand vessels – that is, real assets – we turned to the derivative market for Forward Freight Agreements (FFAs) – that is, financial instruments – related to the dry bulk shipping sector. Accordingly, we illustrated formally that the bulk of volatility in the FFA basis can be attributed to expectations about future physical market conditions rather than expectations about future risk premia – a result perfectly aligned with the respective finding regarding the physical market for ships. Despite this finding, though, we documented the existence of a bias in the FFA rates in the form of “contango” but also of both a strong momentum effect and significant predictability of risk premia by price-based signals and economic variables reflecting physical market conditions. The evidence of bias was further supported by the results of three econometric tests which suggested rejection of the unbiased expectations hypothesis.

In order to justify these findings, we developed a dynamic asset pricing framework that can incorporate both the “hedging pressure” feature and a heterogeneous beliefs explanation. In the proposed model, apart from having different objective functions, agents – that is, ship owners, charterers, and speculators – also differ in the way they form expectations about future market conditions. Specifically, speculators form biased expectations due to asymmetric-imperfect information but mainly due to a behavioural bias known as “gambler’s fallacy”. Empirical estimation of the model suggested that, to simultaneously match the observed regularities, one must depart from the rational expectations benchmark. While the predictions were not particularly sensitive to the degree of information asymmetry, this was not true for the behavioural bias feature; namely, a fraction of investors must suffer from the “gambler’s fallacy” and, in turn, follow a contrarian investment strategy.

From an economic perspective, the heterogeneous expectations feature of our model can be justified by the fact that ship owners and charterers – who participate also in the physical market and, thus, have “inside” information regarding the actual future market conditions – are expected to be able to form more accurate forecasts about future spot rates than speculators – who correspond to financial investors that participate only in the FFA market. It is well-documented that, in practice, traders frequently follow contrarian strategies which can be influenced or motivated by such behavioural biases. Specifically, there is market evidence that mainly uninformed and inexperienced investors usually adopt contrarian behaviour.

References

Quoted Statistical Sources

Baltic Exchange: www.balticexchange.com

Clarksons Shipping Intelligence Network (SIN): www.clarksons.net

London Clearing House (LCH Group – LCH.Clearnet): www.lch.com

Thomson Reuters Datastream Professional: www.thomsonone.com

Books, Articles, and Reports

Abraham, J.M. and Hendershott, P.H. (1996) Bubbles in metropolitan housing markets, *Journal of Housing Research* **7**, 191–207.

Acharya V.V., Lochstoer, L.A., and Ramadorai, T. (2013) Limits to arbitrage and hedging: Evidence from commodity markets, *Journal of Financial Economics* **109**, 441-465.

Adland R., Jia, H., and Strandenes, S. (2006) Asset bubbles in shipping? An analysis of recent history in the dry bulk market, *Maritime Economics and Logistics* **8**, 223-233.

Akaike, H. (1973) Information theory and extension of the maximum likelihood principle, in: Petrov, B. and Csake, F.(eds.), *Second international symposium on information theory*, Akademiai Kiado, Budapest.

Alizadeh, A.H. and Nomikos, N.K. (2007) Investment timing and trading strategies in the sale and purchase market for ships, *Transportation Research Part B: Methodological* **41**, 126-143.

Alizadeh, A.H. and Nomikos, N.K. (2009) *Shipping Derivatives and Risk Management*, Palgrave Macmillan, London.

Alizadeh, A.H. and Nomikos, N.K. (2010a) An overview of the dry bulk shipping industry, in: Grammenos, C.T. (ed.), *The Handbook of Maritime Economics and Business* 2nd edition, Lloyd's List, London, 319-353.

Alizadeh, A.H. and Nomikos, N.K. (2010b) Managing Freight Rate Risk Using Freight Derivatives: An Overview of the Empirical Evidence, in: Grammenos, C.T. (ed.), *The Handbook of Maritime Economics and Business* 2nd edition, Lloyd's List, London, 745-776.

Allen, F. and Michaely, R. (2003) Payout policy, in: Constantinides, G., Harris, M., and Stulz, R. (eds.) *Handbook of the economics of finance: Corporate Finance* Vol 1A, Elsevier, Amsterdam, 337-429.

Alti, A. and Tetlock, P. (2014) Biased Beliefs, asset prices, and investment: a structural approach, *Journal of Finance* **69**, 325-361.

Ang, A. and Bekaert, G. (2007) Stock return predictability: Is it there?, *Review of Financial Studies* **20**, 651–707.

Arrow, K.J. (1986) Rationality of self and others in an economic system, *Journal of Business* **59(4)**, 385-399.

Bansal, R. and A. Yaron (2004) Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* **59**, 1481–1509.

Barberis, N., Greenwood, R., Lawrence, J., and Shleifer, A. (2015a) X-CAPM: An extrapolative capital asset pricing model, *Journal of Financial Economics* **115**, 1-24.

Barberis, N., Greenwood, R., Lawrence, J., and Shleifer, A. (2015b) Extrapolation and Bubbles, Unpublished working paper, California Institute of Technology, Harvard business School, and Yale School of Management.

Barberis, N., Shleifer, A., and Vishny, R. (1998) A model of investor sentiment, *Journal of Financial Economics* **49**, 307-343.

Beenstock, M. (1985) A theory of ship prices, *Maritime Policy and Management* **12(3)**, 215-225.

Beenstock M. and Vergottis, A. (1989) An econometric model of the world market for dry cargo freight and shipping, *Applied Economics* **21(3)**, 339-356.

Bloomfield, R., O' Hara, M., and Saar, G. (2009) How noise trading affects markets: An experimental analysis, *Review of Financial Studies* **22(6)**, 2275-2302.

Brock, W.A. and Hommes, C.H. (1997) A rational route to randomness, *Econometrica* **69**, 1059-1095.

Brock, W.A. and Hommes, C.H. (1998) Heterogeneous beliefs and routes to chaos in a simple asset pricing model, *Journal of Economic Dynamics and Control* **22**, 1235-1274.

Camerer, C. and Lovo, D. (1999) Overconfidence and excess entry: An experimental approach, *American Economic Review* **89**, 306-318.

- Campbell, S.D., Davis, M.A., Gallin, J. and Martin, R.F. (2009) What moves housing markets: A variance decomposition of the rent-price ratio, *Journal of Urban Economics* **66**, 90-102.
- Campbell, J.Y. and Shiller, R.J. (1988a) The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* **1**, 195-228.
- Campbell, J.Y. and Shiller, R.J. (1988b) Stock prices, earnings, and expected dividends, *Journal of Finance* **43(3)**, 661-676.
- Capozza, D.R., Hendershott, P.H., and Mack, C. (2004) An anatomy of price dynamics in illiquid markets: Analysis and evidence from local housing markets, *Real Estate Economics* **32**, 1-32.
- Chen, L. (2009) On the reversal of return and dividend growth predictability: A tale of two periods, *Journal of Financial Economics* **92**, 128-151.
- Chen, L., Da, Z., and Priestley, R. (2012) Dividend smoothing and predictability, *Management Science* **58**, 1834-1855.
- Choi, J. and Mertens, T. (2013) Extrapolative expectations and the equity premium. Unpublished working paper, New York University and Yale University.
- Cochrane, J.H. (1992) Explaining the variance of price-dividend ratios, *Review of Financial Studies* **5**, 243-280.
- Cochrane, J.H. (2005) *Asset Pricing* Revised Edition, Princeton University Press, Princeton.
- Cochrane, J.H. (2011) Presidential address: Discount rates, *Journal of Finance* **66(4)**, 1046-1108.
- Collin-Dufresne, P., Johannes, M., and Lochstoer, L. A. (2016) Asset Pricing when 'This Time is Different', Unpublished working paper, EPFL and Columbia Business School.
- Davis, M.A., Lenhart, A., and Martin, R.F. (2008) The rent-price ratio for the aggregate stock of owner-occupied housing, *Review of Income and Wealth* **54(2)**, 279-284.
- Ekeland, I., Lautier, D., and Villeneuve, B. (2016) Speculation in commodity futures markets: A simple equilibrium model, Unpublished working paper, Université Paris-Dauphine.
- Engle, R.F. and Granger, C.W. (1987) Cointegration and error correction: Representation, estimation, and testing, *Econometrica* (**55**), 251-276.
- Engsted, T. and Pedersen, T.K. (2010) The dividend-price ratio does predict dividend growth: International Evidence, *Journal of Empirical Finance* **17**, 585-605.
- Fama, E.F. (1970) Efficient capital markets: A review of theory and empirical work, *Journal of Finance* **25**, 383-417.
- Fama, E.F. (1984a) Forward and spot exchange rates, *Journal of Monetary Economics* (**14**), 319-338.
- Fama, E.F. (1984b) The information in the term structure, *Journal of Financial Economics* (**13**), 509-528.
- Fama, E.F. (2014) Two pillars of asset pricing, *American Economic Review* **104**, 1467-1485.
- Fama, E.F. and French, K.R. (1987) Commodity Futures Prices: Some Evidence on Forecast Power, Premiums, and the Theory of Storage, *The Journal of Business* **60(1)**, 55-73.
- Fama, E.F. and French K.R. (1988a), Permanent and temporary components of stock prices, *The Journal of Political Economy* **96**, 246-273.
- Fama, E.F. and French, K.R. (1988b) Dividend yields and expected stock returns, *Journal of Financial Economics* **2**, 3-25.
- French, K. R. (1986) Detecting Spot Price Forecasts in Futures Prices, *The Journal of Business* **59(2, Part 2)**, S39-S54.
- Fuster, A., Hebert, B., and Laibson, D. (2011) Natural expectations, macroeconomic dynamics, and asset pricing, *NBER Macroeconomics Annual* **26**, 1-48.
- Gallin, J., (2008) The long-run relationship between house prices and rents, *Real Estate Economics* **36**, 635-68.
- Geltner, D.M and Miller, N.G. (2006) *Commercial Real Estate Analysis and Investments*, South-Western Educational Publishing, Cincinnati.
- Gilles, C. and LeRoy S.F (1991) Econometric aspects of the variance-bounds tests: A Survey, *The Review of Financial Studies* **4(4)**, 753-791.
- Ghysels, E., Plazzi A., Torous W.N., and Valkanov R.I. (2012) Forecasting real estate prices, in: Elliot, G. and Timmermann, A. (eds.), *Handbook of economic forecasting* Vol 2, Elsevier, Amsterdam.
- Glaeser, E.L. (2013), A nation of gamblers: Real Estate speculation and American history, *The American Economic Review* **103(3)**, 1-42.
- Gorton, G.B., Hayashi, F., and Rouwenhorst, G.K. (2012) The Fundamentals of Commodity Futures Returns, *Review of Finance* **17**, 35-105.
- Greenwood, R. and Hanson, S.G. (2015) Waves in ship prices and investment, *Quarterly Journal of Economics* **130**, 44-109.
- Greenwood, R. and Shleifer, A. (2014) Expectations of returns and expected returns, *Review of Financial Studies* **27**, 714-746.

- Grinblatt, M. and Keloharju, M. (2000) The investment behaviour and performance of various investor types: A study of Finland's unique data set, *Journal of Financial Economics* **55(1)**, 43-67.
- Hakkio, C. and Rush, M. (1989) Market efficiency and cointegration: An application to the Sterling and Deutschmark exchange rates, *Journal of International Money and Finance* **8**, 75-88.
- Hamilton, B.W. and Schwab, R.M. (1985) Expected appreciation in urban housing markets, *Journal of Urban Economics* **18**, 103-118.
- Hannan, E.J. and Quinn, B.G. (1979) The determination of the order of an autoregression, *Journal of the Royal Statistical Society (Series B)* **41**, 190-195.
- Hansen, L. and Hodrick, R. (1980) Forward exchange rates as optimal predictors of futures spot rates: An econometric analysis, *Journal of political Economy* **88**, 829-853.
- Hazuka, T.B. (1984) Consumption betas and backwardation in commodity markets, *Journal of Finance* **39**, 647-655.
- Hirshleifer, D. and Wu, J. (2013) Asset pricing in production economies with extrapolative expectations. Unpublished working paper, University of California, Irvine and University of Minnesota.
- Hodrick, R. J. (1992) Dividend yields and expected stock returns: Alternative procedures for inference and measurement, *Review of Financial Studies* **5**, 357-386.
- Hong, H., Yogo, M. (2012) What does futures market interest tell us about the macroeconomy and asset prices?, *Journal of Financial Economics* **105**, 473-490.
- Jegadeesh, N. and Titman, S. (1995) Overreaction, Delayed Reaction, and Contrarian Profits, *The Review of Financial Studies* **8(4)**, 973-993.
- Johansen, S. (1988) Statistical analysis of cointegrating vectors, *Journal of Economic Dynamics and Control* **12**, 231-254.
- Johansen, S. (1991) Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models, *Econometrica* **59**, 1551-1580.
- Johansen, S. and Juselius, K. (1990) The full information maximum likelihood procedure for inference on cointegration – with applications to the demand for money, *Oxford Bulletin of Economics and Statistics* **52**, 169-210.
- Kahneman, D. (2011) *Thinking Fast and Slow*, Farrar, Straus and Giroux, New York.
- Kahneman, D. and Lovallo, D. (1993) Timid choices and bold forecasts: A cognitive perspective on risk taking, *Management Science* **39**, 17-31.
- Kalouptzidi, M. (2014) Time to build and fluctuations in bulk shipping, *American Economic Review* **104(2)**, 564-608.
- Kaniel, R., Saar, G., and Titman, S. (2008) Individual investor trading and stock returns, *The Journal of Finance* **63(1)**, 273-310.
- Kavussanos, M.G. and Nomikos N.K. (1999) The Forward Pricing Function of the Shipping Freight Futures Market, *The Journal of Futures Markets* **19(3)**, 353-376.
- Killian, L. (2009) Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market, *American Economic Review* **99**, 1053-1069.
- Kindleberger, C. (1978) *Manias, Panics, and Crashes: A History of Financial Crises*, MacMillan, New York.
- Lakonishok, J., Shleifer, A., and Vishny, R.W. (1994), Contrarian Investment, Extrapolation, and Risk, *The Journal of Finance* **49(5)**, 1541-1578.
- Leary, M.T. and Michaely R. (2011) Determinants of dividend smoothing: Empirical evidence, *Review of Financial Studies* **24**, 3179-3249.
- Leroy, S.F. and Porter, R.D. (1981) The present-value relation: Tests on implied variance bounds, *Econometrica* **49(3)**, 555-574.
- Lintner, J. (1956) Distribution of incomes of corporations among dividends, retained earnings and taxes, *American Economic Review* **46(2)**, 97-113.
- Lof M. (2015) Rational speculators, contrarians, and excess volatility, *Management Science* **61(8)**, 1889-1901.
- Marsh, T.A. and Merton R.C. (1987) Dividend behavior for the aggregate stock market, *Journal of Business* **60(1)**, 1-40.
- Martin, V., Hurn, S., and Harris, D. (2013) *Econometric Modelling with Time Series*, Cambridge University Press, New York.
- Mehra, R. and Prescott, E. (1985) The equity premium: a puzzle, *Journal of Monetary Economics* **15**, 145-161.
- Mozes, H.A. and Rapaccioli, D.C. (1998) The link between dividend changes and future earnings, *Journal of Financial Statement Analysis* **3**, 29-39.
- Newey, W.K. and West, K.D. (1987) A simple, positive-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* **55**, 703-708.

- Papapostolou, N.C., Nomikos N.K., Poulialis P.K., and Kyriakoy, I. (2014) Investor sentiment for real assets: The case of dry bulk shipping market, *Review of Finance* **18**, 1507-1539.
- Park, A. and Sabourian, H. (2011) Herding and Contrarian Behavior in Financial Markets, *Econometrica* **79(4)**, 973-1026.
- Phillips, P. and Hansen, B. (1990) Statistical Inference in instrumental variables regression with I(1) processes, *Review of Economic Studies* **57**, 99-125.
- Phillips, P. and Perron, P. (1988) Testing for a unit root in time series regressions, *Biometrika* (**75**), 335-346.
- Plazzi, A., Torous, W., and Valkanov, R. (2010) Expected returns and expected growth in rents of commercial real estate, *Review of Financial Studies* **23**, 3469–3519.
- Rabin, M. (2002) Inference by believers in the law of small numbers, *Quarterly Journal of Economics* **117**, 775-816.
- Rangvid, J., Schmeling, M., and Schrimpf, A. (2014) Dividend predictability around the world, *Journal of Financial and Quantitative Analysis* **49**, 1255-1277.
- Reinhart, C. and Rogoff, K. (2009) *This time is different: Eight Centuries of Financial Folly*, Princeton University Press, Princeton.
- Sargent, T.J. (1979) A note on the estimation of the rational expectations model of the term structure, *Journal of Monetary Economics* **5**, 133-143.
- Schwarz, G. (1978) Estimating the dimension of a model, *Annals of Statistics* (**6**), 461-464.
- Shefrin, H. (2000) *Beyond Greed and Fear: Understanding Behavioural Finance and the Psychology of Investing*, Harvard Business School Press, Boston.
- Shiller, R. (1981) Do stock prices move too much to be justified by subsequent changes in dividends?, *American Economic Review* **71**, 421–436.
- Shiller, R. (2000) *Irrational Exuberance*, Princeton University Press, Princeton, New Jersey.
- Shiller, R. (2003) From Efficient Markets Theory to Behavioral Finance, *The Journal of Economic Perspectives* **17(1)**, 83-104.
- Stafford, E., Chao, A., and Luchs, K. (2002) *Ocean Carriers*, Harvard Business School Case: 202-027.
- Stopford, M. (2009) *Maritime Economics* 3rd edition, Routledge, New York.
- Szymanowska, M., De Roon, F., Nijman, T., and Van Den Goorbergh, R. (2014) An Anatomy of Commodity Futures Risk Premia, *The Journal of Finance* **69(1)**, 453-482.
- Ter Ellen, S. and Zwinkels, R.C.J. (2010) Oil price dynamics: A behavioural finance approach with heterogeneous-agents, *Energy Economics* **32**, 1427-1434.
- Terrell, D. (1994) A Test of the Gambler's Fallacy – Evidence from Pari-Mutuel Games, *Journal of Risk and Uncertainty* **8**, 309-317.
- Tirole, J. (2006) *The theory of corporate finance*, Princeton University Press, Oxfordshire.
- Tversky, A. and Kahneman D. (1971) Belief in the Law of Small Numbers, *Psychological Bulletin* **76**, 105-110.
- Tversky, A. and Kahneman D. (1974) Judgment under uncertainty: Heuristics and biases, *Science* **185**, 1124-1131.
- UNCTAD (2015) *Trade and Development Report*, E.15.II.D.4, New York and Geneva, United Nations.
- Vives, X. (2008) *Information and Learning in Markets: The Impact of Market Microstructure*, Princeton University Press, Princeton, New Jersey.
- Vuolteenaho T. (2002) What drives firm-level stock returns?, *Journal of Finance* **57**, 233–264.
- Wang, J. (1993) A model of intertemporal asset prices under asymmetric information, *The Review of Economic Studies* **60(2)**, 249-282.
- Wheaton, W.C. and Torto, R. (1988) Vacancy rates and the future of office rents, *AREUEA Journal* **19**, 430–436.
- White, H. (1980) A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroscedasticity, *Econometrica* **48**, 817–838.