



City Research Online

City, University of London Institutional Repository

Citation: Moutzouris, I. C. & Nomikos, N. (2019). The Formation of FFA Rates in Dry Bulk Shipping: Spot Rates, Risk Premia and Heterogeneous Expectations. *Journal of Futures Markets*, 39(8), pp. 1008-1031. doi: 10.1002/fut.21980

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/20713/>

Link to published version: <https://doi.org/10.1002/fut.21980>

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.

The formation of FFA Rates in dry bulk shipping: Spot rates, risk premia and heterogeneous expectations^{*}

IOANNIS C. MOUTZOURIS ^a and NIKOS K. NOMIKOS ^b

^b Cass Business School, Faculty of Finance, City, University of London, 106 Bunhill Row, London EC1Y 8TZ, UK (contact number: +44 (0)7554017706, email: ioannis.Moutzouris.1@cass.city.ac.uk).

^b Corresponding Author. Cass Business School, Faculty of Finance, City, University of London, 106 Bunhill Row, London EC1Y 8TZ, UK (contact number: +44 (0)20 7040 0104, email: N.Nomikos@city.ac.uk).

September 2018

Abstract

This article examines the formation of forward rates in the dry bulk shipping industry. We illustrate that the bulk of basis volatility can be attributed to expectations about future physical market conditions rather than expectations about future risk premia. However, there exists significant predictability of risk premia by both price-based signals and economic variables. To explain this finding, we develop a dynamic asset pricing framework where, apart from having different objective functions, agents might also differ in the way they form expectations about future market conditions. Accordingly, we argue that the average investor should hold “near-rational” but slightly contrarian beliefs.

Keywords: Asset Pricing; Behavioural Finance; Shipping Industry; Biased Beliefs; Law of Small Numbers; Heterogeneous Agents; Contrarian Strategy;

JEL Codes: C13, G12, G13, G40.

^{*} We would like to thank for their useful feedback and suggestions Brian Wright, James Hamilton, Nicolás Merener, Ana-Maria Fuertes, Ioannis Kyriakou, Nikos Papapostolou, Christian Urom, Sjur Westgaard as well as participants at the: 2017 Commodity and Energy Markets Conference in Oxford; 2017 11th International Conference on Computational and Financial Econometrics in London; 2018 Commodity and Energy Markets Annual Meeting in Rome; 2018 New Directions in Commodity Research Conference in Denver.

1 INTRODUCTION

In this paper, we extend the application of heterogeneous beliefs models to commodity derivatives. In particular, we focus on the forward market for shipping freight rates, that is the market for Forward Freight Agreements (FFAs). 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). During the last decade, in addition to traditional shipping investors, the market for freight derivatives has also attracted the interest and participation of investors from other sectors of the economy. Due to the distinct features of the shipping industry and its highly volatile character, trading volume in shipping derivatives markets is expected to increase significantly over the next years. Hence, understanding the pricing and trading dynamics of this market is important.

To the best of our knowledge, this is the first time that a structural, heterogeneous-beliefs asset pricing model is applied to a futures or forward market.¹ Thus, we provide a framework that can be adapted and evaluated empirically in other commodity derivative markets. In doing so, our contribution to the literature is threefold. First, we are the first to apply the variance decomposition framework in a derivatives market where the underlying asset is a non-storable service. Second, we document for the first time several noticeable empirical regularities related to FFA rates and risk premia. Third, we propose a theoretical behavioural asset pricing model that can account for these stylised facts.

We begin by analysing the formation of the basis in the freight forward market. Fama (1984a and 1984b) and Fama and French (1987) show that the variance of the basis of any futures and forward contract can be decomposed into the sum of the covariance between the basis and the expected change in the spot price and the covariance between the basis and the expected premium over the spot price at maturity; this premium reflects the excess return for an investor who goes short on the derivative contract. We illustrate formally that volatility in the FFA basis can be attributed primarily to expectations about future physical market conditions rather than expectations about future risk premia, as is generally the case in commodity markets (Fama and French, 1987). This is justified on the basis that freight rates are subject to supply and demand shocks which cannot be smoothed through short-term adjustments in supply; the reader can parallelise this to lack of commodity storability. This results in predictable variation of spot rates which, consequently, increases the forecasting ability of the FFA basis. This is consistent with previous empirical evidence that

¹ For instance, Ellen and Zwinkels (2010) and Bredin *et al* (2108) also use behavioural models with heterogeneous speculators to explain commodity price dynamics. Despite some similarities to our model, commodity demand in their frameworks is not derived explicitly through a structural economic model as in our case.

predictability of future spot rates is a decreasing function of commodity storability (Hazuka, 1984; French, 1986; Fama and French, 1987).

While volatility in FFA basis is primarily attributed to changes in expected spot rates, we cannot rule-out the existence of (possibly, time-varying) risk premia. Accordingly, we provide evidence of several stylised features that might be of interest to academic and practitioners alike. Specifically, in contrast to most commodity futures markets, we find strong statistical evidence of contango; that is, realised risk premia are, on average, positive. In addition, there exists significant predictability of future risk premia, consistent with the existence of a momentum effect: lagged risk premia strongly and positively forecast future risk premia. Finally, FFA risk premia can be strongly negatively forecasted by both spot market signals and economic indicators related to commodity trade and shipping demand. 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 which unequivocally suggest the existence of a bias in the dry bulk FFA market. From a market participant's perspective, those stylised features can be used to develop potentially profitable trading strategies.

We develop a theoretical model of FFA price determination to reproduce our main empirical findings. The proposed framework draws its main features from the latest generation of structural economic models in the commodity futures literature (e.g. Gorton *et al*, 2012; Acharya *et al*, 2013) and has been modified and extended in two, quantitatively simple but conceptually important, manners. First, our framework departs from the "theory of storage" explanation of "time-varying" risk premia (e.g. Gorton *et al*, 2012; Ekeland *et al*, 2018) since shipping freight is a non-storable service. An immediate consequence of this is the extension of the (widely used) two-period economic environment to an infinite horizon model, which simplifies the empirical evaluation of the generated framework; accordingly, we validate the theoretical predictions of our model through numerical simulations. Second, we incorporate the existence of distorted beliefs on a fraction of the investor population; heterogeneous beliefs models provide an alternative way for researchers to explain empirical regularities in asset prices that cannot be explained by traditional rational expectations models. This way, we contribute to the generic commodity finance literature by incorporating explicitly the behavioural dimension in the formation of derivative contracts rates.

Our discrete-time economy consists of three types of agents; ship-owners, charterers, and speculators. Apart from having different objective functions, agents also differ in the way they form expectations about future market conditions. While ship-owners and charterers are fully rational investors, speculators are characterised by bounded rationality and suffer from a form of

“representativeness heuristic” which means that they “exaggerate how likely it is that a small sample resembles the parent population from which is drawn” (Tversky and Kahneman, 1971; Shefrin, 2000). As a result, following a shock in freight rates, speculators believe that rates will revert more rapidly to their previous level than is the case in reality which results in a contrarian investment behaviour on their behalf.

The use of behavioural models in equity markets is often justified by survey data which confirms the theoretical predictions (Greenwood and Shleifer, 2014). Since there are no comparable detailed surveys regarding shipping industry participants’ beliefs and investment strategies, we justify the use of the proposed model by contradiction using both theoretical predictions and model simulations. Namely, we show that a rational expectations model with a hedging pressure bias, cannot explain the documented empirical regularities. Similarly, simulation tests suggest that to simultaneously match all observed regularities sufficiently well, the average investor should hold “near-rational” but slightly contrarian beliefs.

The remainder of this article is organised as follows. Section II describes the shipping industry and discusses our dataset. Section III performs the empirical analysis. Section IV presents the environment of our economy and the theoretical model. Section V concludes.

2 DESCRIPTION OF THE INDUSTRY AND DATASET

The market for FFA contracts was established in 1992 as a hedging instrument for participants in the physical shipping market. 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” (Alizadeh and Nomikos, 2009).

In the context of this research, we focus on the Capesize and Panamax dry bulk FFA contracts since they constitute by far the most liquid segments: on average, approximately 98,000 dry bulk FFA contracts are traded each month of which 45% and 43%, respectively, are for Capesize and Panamax vessels. The total paper activity corresponds to about 90% of the trading activity in the underlying physical market. Our dataset consists of monthly observations of spot prices, settlement rates, and FFA rates for the BCI 4TC and BPI 4TC contracts with 1- and 2-month maturities from January 2007 to September 2016, obtained from The Baltic Exchange. Those contracts correspond to the equally weighted average of the four trip-charter routes of the Baltic Capesize Index (BCI 4TC) and the Baltic Panamax Index (BPI 4TC), respectively. Incorporating the industry convention, settlement rates are

calculated as the arithmetic average of the respective underlying spot rates over all trading days of the settlement month.²

Following Fama and French (1987) and Cochrane (2011), we decompose the difference between the log FFA rate at time t for a contract expiring in T periods, $f(t, T)$, and the current log spot rate, $s(t)$, into the sum of the expected change in the log spot rate and an expected premium over the log settlement rate at maturity of the contract, $s(t + T)$:

$$f(t, T) - s(t) = E_t[s(t + T) - s(t)] + E_t[f(t, T) - s(t + T)]. \quad (1)$$

The quantity on the left-hand side of (1) is defined as the basis of the FFA contract which is also a frequently used valuation ratio. The two terms on the right-hand side of (1) are the expected spot growth rate and the expected risk premium, respectively; the latter 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. Equation (1) shows that a high basis may be due to higher expected spot growth rate and/or higher expectations of future log risk premia.

Panel A of Figure 1 presents the Forward, Spot, and Settlement rates for the 1-month Panamax case. While all series are very volatile, reflecting the growth in freight rates and the subsequent crisis in 2008, they are also strongly correlated. From an economic perspective, the high volatility and uncertainty regarding freight market conditions justifies the existence of the FFA market as a hedging instrument for participants in the physical market but also attracts the trading interest of investors outside of the shipping markets such as hedge funds and investment banks (Alizadeh and Nomikos, 2009). Panel B of Figure 1 illustrates the evolution of the spot growth, premium and basis for the 1-month Panamax case. Evidently, while all three variables appear to fluctuate significantly over time, spot growth rates exhibit the highest volatility. This finding is also confirmed by the estimated standard deviations presented in Table I and has important implications for the predictive regression results, as discussed in the following section.

Looking at the remaining descriptive statistics in Table I, we note that the basis is, on average, significantly positive and strictly increases with maturity for both contracts, consistent with the forward market being in contango. There is also evidence of positive mean risk premia in both

² Depending on the cargo-carrying capacity of the vessel (measured in dead-weight tonnage, dwt), the dry bulk fleet is subdivided into the Capesize, Panamax, Handymax, and Handysize sectors. At the largest end of the range, Capesize carriers (above 150,000 dwt) are used in the transportation of iron ore and coal while Panamax carriers (around 74,000 dwt) are associated with the trade of a large variety of commodities such as coal, grains, bauxite, and the larger minor bulks. The 4TC routes reflect the main trading routes on which Capesize and Panamax vessels operate: the Atlantic Trade, the Pacific Trade, the Continent-to-Far East Trade, and the Far East-to-Continent Trade. Finally, 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.

contracts and across all horizons suggesting that, on average, FFA rates are higher than the corresponding realised settlement rates at the maturity of the contract. It is interesting to note that the risk premia are high in absolute terms, implying annualised mean returns of more than 30% in all cases, which is much higher than those reported in the literature for other commodity markets. For instance, Szymanowska *et al* (2014) examine eight commodity indices and report annualised mean returns ranging from -7% to 11%. Note though that the premia are only statistically significant in the 1-month case. From a market participant's perspective and in terms of potential trading strategies, those results suggest that an investor should rather take a short position in the FFA market as this seems to be consistently generating positive returns. In addition, it might be preferable for a short hedger (or speculator) to go short on two consecutive 1-month contracts instead of taking a short position on the respective 2-month one. The inverse is true for a long hedger.

A further noticeable stylised fact is that the 1- and 2-month risk premia are positively autocorrelated and their autocorrelation increases with contract maturity which indicates the existence of a momentum effect, as analysed in the following section. Finally, risk premia and spot growth rates are negatively correlated, as is also shown in Panel B of Figure 1, which suggests that an unexpected positive (negative) shock in spot rates will result in a negative (positive) realised risk premium.

Concluding, in the dry bulk FFA market there is no evidence supporting the Theory of Normal Backwardation (Keynes, 1930; Hicks, 1939). This contrasts with the evidence from many other commodity futures markets where, in line with the theory, futures prices are at a premium to expected spot prices to compensate the long side of the futures settlement for providing price insurance (Gorton *et al*, 2012). Following this argument, our results suggest that the short side of the FFA position is – on average – rewarded for providing price insurance. Furthermore, these stylised facts combined appear to verify the common view of practitioners that a positive basis or, equivalently, a negative “roll yield” is a requirement for the existence of a positive risk premium to a short position in futures markets (Gorton *et al*, 2012).³ From an economic perspective, there exist three potential explanations for this finding. First, there is on average net long hedging pressure in the market; that is, more long hedgers than short ones. Second, short physical hedgers – shipowners – require a ‘settlement risk premium’ and/or a ‘basis risk premium’ from the long side; the first arises due to the difference between the average rate used for the settlement of the FFA contract and the freight rate at the which the vessel is fixed in the physical market while the second from the mismatch between the specification of the FFA contract and the exposure of the hedger in the physical market (Alizadeh

³ 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 article.

and Nomikos, 2009). Third, short physical hedgers demand a “liquidity premium” from the long side; Kang *et al* (2018) argue that this is a consequence of the fact that short physical hedgers accommodate the liquidity demands of speculators who readjust their positions much more frequently than physical hedgers. While any (combination) of those explanations appears plausible, we cannot confirm that empirically since detailed commitment of traders reports are not available for this market.

3 PREDICTABILITY OF MARKET CONDITIONS AND RISK PREMIA USING VALUE INDICATORS

We begin by applying the variance decomposition framework (Fama, 1984a and 1984b; Fama and French, 1987) to assess the forecasting ability of the FFA market for ships. While shipping researchers have already examined the question of FFA predictability using cointegration techniques (e.g. Kavussanos and Nomikos, 1999), our estimation framework is more versatile and aims to fill certain gaps in the literature.

First, and most importantly, this 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, since shipping services are considered a commodity, we explain our findings by performing a comparison with other commodity futures and forward markets. Specifically, Hazuka (1984), French (1986), and Fama and French (1987) show that the forecasting ability of futures contracts is directly related to seasonality 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 non-storable service. In addition, this decomposition allows us to quantify precisely the variation in the FFA basis that can be attributed to expectations about future market conditions and time-varying risk premia.

Second, the results from the variance decomposition framework are robust even in the presence of overlapping observations. What is more, the variance decomposition framework examines a question of relative predictability without imposing any restrictions on either the spot rate process, the rationality of expectations or the existence of time-varying risk premia. Third, the incorporated sample corresponds to the most recent available data – including the extreme shipping cycle of the period 2007 to 2009 – regarding the forward shipping markets.

Following Cochrane (2011), it is straightforward to decompose the variance of the basis into two parts. Namely, multiplying both sides of (1) by $f(t, T) - s(t) - E[f(t, T) - s(t)]$ and taking expectations yields

$$\begin{aligned} var[f(t, T) - s(t)] &= cov[f(t, T) - s(t), s(t + T) - s(t)] \\ &\quad + cov[f(t, T) - s(t), f(t, T) - s(t + T)]. \end{aligned} \tag{2a}$$

Therefore, the variance of the basis is equal to the covariance between the basis and the expected spot growth and the covariance between the basis and the expected risk premium. Dividing both sides of (2a) by the variance of the basis yields

$$\frac{cov[f(t, T) - s(t), s(t + T) - s(t)]}{var[f(t, T) - s(t)]} + \frac{cov[f(t, T) - s(t), f(t, T) - s(t + T)]}{var[f(t, T) - s(t)]} = 1 \quad (2b)$$

$$\Rightarrow \beta_{\Delta s, T} + \beta_{r, T} = 1,$$

where $\beta_{i, T}$ is the T -period contract coefficient corresponding to the i^{th} element of the decomposition. Incorporating (1) in (2b), these two coefficients are further analysed into

$$\beta_{\Delta s, T} = \frac{var[s(t + T) - s(t)] + cov[s(t + T) - s(t), f(t, T) - s(t + T)]}{var[s(t + T) - s(t)] + var[f(t, T) - s(t + T)] + 2cov[s(t + T) - s(t), f(t, T) - s(t + T)]} \quad (2c)$$

and

$$\beta_{r, T} = \frac{var[f(t, T) - s(t + T)] + cov[s(t + T) - s(t), f(t, T) - s(t + T)]}{var[s(t + T) - s(t)] + var[f(t, T) - s(t + T)] + 2cov[s(t + T) - s(t), f(t, T) - s(t + T)]} \quad (2d)$$

The expressions above suggest that the variance of the basis depends on the variances of the spot growth and the risk premium as well as the covariance between those two components. We can examine which of those two sources is the major determinant of the observed variability in basis by running forecasting OLS regressions in the spirit of Fama (1984a and 1984b), Fama and French (1987), and Cochrane (2011). Namely, we regress realised log spot growth and realised log risk premia on the current log basis:

$$s(t + T) - s(t) = \alpha_{\Delta s, T} + \beta_{\Delta s, T} \cdot [f(t, T) - s(t)] + \varepsilon_{\Delta s, t+T}, \quad (3a)$$

$$f(t, T) - s(t + T) = \alpha_{r, T} + \beta_{r, T} \cdot [f(t, T) - s(t)] + \varepsilon_{r, t+T}, \quad (3b)$$

In line with Fama and French (1987), statistical evidence that $\beta_{\Delta s, T}$ is positive means that the basis has forecasting power on 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 (1) and (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 restriction implies that regressions (3a) and (3b) will always allocate all basis variation to either expected spot growth or expected risk premia or some combination of the two; thus, those regressions examine a question of relative predictability through the magnitudes of the two slope coefficients, $\beta_{\Delta s,T}$ and $\beta_{r,T}$.

We run the predictive regressions (3a) and (3b) for the Capesize BCI 4TC and Panamax BPI 4TC contracts for both maturities.⁴ 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 II, all spot growth coefficients are significantly positive at the 1% level in every sector and horizon and the forecasting power of the log basis appears to be strong with the R^2 s of growth regressions being at least 14%. Turning next into the regressions for risk premia, the slope coefficients and the respective t-statistics are much smaller in magnitude and the R^2 s are below 8% in all cases. Overall, the 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 are good predictors of future spot rates.

There is a long-standing debate in asset pricing regarding the forecasting ability of futures markets. In many markets, futures prices 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 markets, we now show that FFA rates are better predictors of future market conditions compared to the current spot rates. Table II presents the results from regressions of future spot growth on the first lag of the 1-month spot growth. We note that, irrespective of the maturity and the sector under consideration, the coefficients for lagged spot growth are not statistically significant which is consistent with the view that FFA rates contain superior information compared to contemporaneous spot and settlement rates.⁵

In order to interpret our findings, it is worth noting that if the current spot price equals the expected spot price, futures prices cannot provide a better forecast of the future spot price than current spot prices (French, 1986). In other 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 profound and important for the interpretation of the forecasting results.

⁴ We also examine formally whether the variables of interest satisfy the necessary stationarity condition. Evidently, the null hypothesis of a unit root is rejected for the three variables for both contracts and maturities. Results are available from the authors.

⁵ This 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 dependent one. The results can be provided by the authors on request.

One should expect that in markets where realised spot rates are volatile there will be strong predictability of future spot rates from futures prices. Specifically, if investors know – up to a certain degree – the underlying data generating process of the spot rate, they will be able to predict the future spot rate; in the presence of futures markets, this expectation is reflected – at least partially – in futures rates and, in such cases, spot rate volatility results in futures rate volatility. Consequently, in line with French (1986) and Fama and French (1987), for futures prices to provide reliable forecasts of future spot rates, the volatility of both spot rate changes and futures basis must be high.

Hazuka (1984), French (1986), 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 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 variation in spot prices. As a result, for commodities that are non-storable (e.g. electricity) or require high storage costs relative to value (e.g. broilers and eggs), the respective futures prices exhibit significant forecasting power. In contrast, for commodities with low storage costs relative to value, such as precious metals, prices are not informative regarding future market conditions (French, 1986; Fama and French, 1987).

Following this discussion, the results 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, in line with Table I and Panel B of Figure 1, log-changes in both the basis and spot rates changes are highly volatile.⁶

In addition, it is well-documented that freight rates are very volatile and subject to demand shocks along with significant construction lags on the supply side, the combination of which results in very volatile yet mean-reverting (in longer horizons) rates. As illustrated in the recent shipping literature (Greenwood and Hanson, 2015; Nomikos and Moutzouris, 2018a and 2018b) however, due to the nature of the industry,⁷ future spot rates can be predicted – up to a certain degree – based on the time- t public information filtration and/or investors’ private information. Accordingly, FFA rates are

⁶ The fact that spot rates changes are relatively more volatile than the basis – the ratio of the standard deviation of the basis to the spot growth ones ranges from 0.54 to 0.72 – is attributed to the existence of time-varying risk premia, as illustrated in the following section.

⁷ There is a large number of established private 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); consequently, they have strong prior experience and expertise about the key supply and demand drivers of the shipping industry which translate into better forecasts about future market conditions.

expected to reflect, to some extent, the economic predictions of market participants. In contrast, if expected 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. Note that in the applied variance decomposition methodology, the FFA basis is the sole state variable and thus, is assumed to summarise the historical and prevailing market conditions (Fama and French, 1988). Equivalently, the “more storable” a commodity is, the lower the predictability of future spot rates is expected to be. In shipping, however, the commodity is a non-storable service which results in predictable variation of spot rates and, in turn, substantial forecasting ability of FFA rates.⁸

Notice that the arguments presented above apply not only to derivatives contracts but also to financial (e.g. stocks and bonds) and physical (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 (such as dividend yields) is positively related to cash flow volatility and 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 commodity markets. Closely related to these arguments is the analysis in Nomikos and Moutzouris (2018a) who examine the formation of vessel prices in a framework similar to the one discussed here. Their findings indicate that vessel valuation ratios (namely, the earnings yield which is defined as the ratio of net earnings to the current vessel price) strongly and negatively predict future market conditions.

Finally, recall that we examine a question of relative predictability; since $\beta_{\Delta s,T} + \beta_{r,T} = 1$, basis variation must be either due to predictability of future risk premia and/or predictability of future spot growth. Accordingly, the fact that spot growth changes are more volatile than the respective risk premia predisposes us, through equations (2c) and (2d), for the variance decomposition results. In conclusion, we argue that FFA basis moves *mainly* due to expectations about future changes in the spot prices because the latter can be predicted – up to a certain degree – by market agents at time t through the shipping supply and demand mechanism. For the *residual* proportion of basis variability that is attributed to time-varying risk premia, there can be two plausible economic justifications; a “rational” and an “irrational” one.

⁸ One might argue that shipping services are storable as well in the sense that shipowners can affect the supply of shipping services by ordering new vessels and scrapping or laying-up existing ones. While this argument is valid, it only affects the supply in the longer term and not in the 1- and 2-month horizons which is the focus in this paper. In line with this argument, in the following we illustrate that shipping supply variables have no predictive power regarding future risk premia. Furthermore, 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 (1987), 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 are statistically insignificant and thus, the theory is rejected.

Regarding the former, there exist two, usually interconnected (Gorton *et al*, 2012; Ekeland *et al*, 2018), “rational” theories to explain 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 IV, we illustrate formally why the latter explanation appears to be more plausible in the FFA market.⁹

We should note that regressions (3a) and (3b) are designed to detect variation in expected premia; hence, failure to identify time-varying expected premia does not imply that expected premia are zero (Fama and French, 1987). Indeed, as reported in Table I, there appears to be statistical evidence of positive mean risk premia for both contracts. We examine formally the predictive power of additional factors that may explain the existence of risk premia. Table III 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 1-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 the second lag, and so on. 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 .

Results in Table III 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. Notably, in the Panamax contracts the slope coefficients are significant at the 1% level and the second lag of the 1-month risk premium strongly positively predicts both the 1- and 2-month risk premia. We also note that when we use higher lags as regressors the values of the slope coefficients strictly decrease and become less significant. Thus, a high realised risk premium forecasts high future premia and vice versa which indicates the existence of a momentum effect. Importantly, the fact that predictability is attenuated as the lag of the regressor increases reinforces this argument. Therefore, from a trading strategy perspective, this suggests that taking the short (long) position on the FFA contract after a positive (negative) risk premium is realised might be a profitable investment strategy.

⁹ According to this theory, the degree of irrationality should be analogous to the relative variability of the risk premia component, as measured by the magnitude of the $\beta_{r,T}$ coefficient. Thus, the fact that the bulk of volatility is attributed to spot growth changes implies that, while distorted expectations can justify the observed bias, the average degree of expectations’ irrationality is not extreme.

We examine next whether future risk premia can be forecasted by changes in lagged spot growth and by lags of the Baltic Dry Index (BDI). The BDI is a composite freight index widely used by practitioners as a general market indicator in the dry-bulk market. It is in other words the ‘barometer’ of dry-bulk shipping (Alizadeh and Nomikos, 2009). Results presented in Table III indicate that there exists statistically significant predictability of future risk premia from realised physical market conditions. Specifically, in both contracts, the first lag of monthly changes in spot rates is statistically significant and negatively predicts 1-month future risk premia. Therefore, a recent increase in the spot market strongly predicts a decrease in future risk premia. These results are consistent across maturities and sectors, although it appears that predictability is stronger for the Panamax contract and the 1-month horizon. This heterogeneity can be explained by the fact that Capesizes are less diversified than Panamaxes in terms of both the nature of cargo transported and the trading routes they operate on. As a result, it might be easier to forecast future market conditions – and, in turn, more difficult to predict future premia – in the Capesize sector compared to the Panamax one.

The finding that spot market indicators and lagged risk premia have significant predictive power regarding future risk premia becomes more interesting if we recall that the same indicators have very little explanatory power over future market conditions (Table II). Therefore, this implies that these variables may affect in an inefficient or biased manner the formation of current FFA rates and, in turn, future risk premia. Moreover, comparing the risk premia and spot market indicators results, we observe that when using the first lag of the regressors, the corresponding risk premia and physical market indicators slope coefficients have opposite signs and the magnitudes of the former coefficients are higher than the latter ones. This feature suggests that lagged risk premia may contain more information regarding future risk premia compared to physical market conditions.¹⁰

Finally, these findings may have useful implications for devising profitable investment strategies. Namely, current FFA basis, lagged risk premia and lagged changes in physical market conditions can be incorporated as signals/indicators for taking a position in the FFA market. It is thus interesting to examine the potential drivers of this sort of predictability and momentum in the FFA market. To this end, in Section IV we develop a theoretical model that can justify and reproduce these findings.

4 PREDICTABILITY OF RISK PREMIA USING ECONOMIC VARIABLES

As illustrated in the previous subsection, there is evidence of risk premia predictability by realised physical market conditions. Since spot rates are determined in equilibrium through the freight rate mechanism, we further examine the predictability of FFA risk premia by economic variables related to

¹⁰ With the exception of the 1-month BCI 4TC contract, this finding is verified using bivariate regressions where the realised risk premium is regressed on the first lags of the risk premium and the spot growth.

the supply of and demand for shipping services. Namely, we focus on predictors that reflect current and recent short-term changes in supply and demand conditions. This new dataset is retrieved from Clarksons Shipping Intelligence Network and runs at a monthly frequency for the period corresponding to the previous analysis.

We first examine several shipping supply variables related to the capacity and availability of the fleet: 1-, 2-, and 3-month changes in fleet capacity (both for the specific sector and for the total dry bulk fleet) and congestion in main dry bulk ports scaled by the corresponding fleet capacity. None of these variables appears to be significant in explaining variations in risk premia at conventional levels. An explanation for this is that fleet supply is slow-moving and inelastic in the short term.

We now turn to the demand variables which consist of trade and demand indicators that, in line with the existing literature (e.g. Kalouptsi, 2014), aim to capture the prevailing conditions in the dry bulk sector (i.e. global iron ore seaborne exports, global coking coal and steam coal seaborne imports and exports, and global seaborne dry bulk trade), the shipping industry as a whole (i.e. aggregate Chinese imports, crude oil, gasoline, and propane prices) but also the general global macroeconomic environment (i.e. world steel production, the trade-weighted steel production index, the Blast Furnace Iron (BFI) and Directly Reduced Iron (DRI) indices, iron ore spot prices but also OECD inflation and LIBOR rates), expressed in monthly log-differences. In addition, we incorporate as a regressor the spread between the one-month growth rates of dry bulk fleet supply and commodity demand (as quantified by total dry bulk seaborne trade). This variable, defined in Nomikos and Moutzouris (2018a), aims to capture imbalances between shipping supply and demand

Results from these regressions are consistent with previous evidence. To begin with, predictability is stronger in the Panamax sector compared to the Capesize for both the 1- and 2-month horizons; namely, only steaming coal imports and exports – out of fifteen different predictors – do not forecast future Panamax risk premia. In contrast, only five variables (namely, iron ore spot prices, global coking coal imports, global steaming coal imports and exports, and aggregate Chinese imports) are significant for Capesizes. It is important to note that the signs of the regression coefficients that are statistically significant are always negative. This implies that past changes in trade and demand variables always negatively forecast future risk premia; in other words, futures risk premia are negatively affected by a recent improvement in demand conditions.

Finally, we also examine whether the number of second-hand vessel transactions within a given month, scaled by the corresponding fleet size, can predict future risk premia. This can be used both as a measure of liquidity (Nomikos and Moutzouris, 2018b) and as a proxy for investor sentiment (Papapostolou *et al*, 2014). The results (not presented here) suggest that there is evidence of statistically significant predictability of future risk premia. In particular, vessel transactions negatively

forecast future risk premia. Similar to all previous forecasting tests, predictability is stronger in the Panamax sector. This finding is interesting since there two variables are only indirectly linked with each other. Specifically, second-hand vessel transactions are positively correlated with physical market conditions, the correlation coefficients between freight rates and dry bulk fleet transactions being 0.38 and 0.39 for the Capesize and Panamax sectors, respectively. In turn, as illustrated above, prosperous market conditions negatively forecast future risk premia.

In conclusion, the results in this section are in line with economic theory and reinforce our previous findings. Namely, increased demand for shipping services implies, *ceteris paribus*, an improvement in physical market conditions or, equivalently, an increase in spot rates both of which negatively predict future risk premia.¹¹

5 A DYNAMIC BEHAVIOURAL ASSET PRICING MODEL OF FFA RATES

The documented predictability in the BCI and BPI 4TC contracts suggests that FFA rates are not unbiased forecasts of the realised settlement rates and that FFA markets are not efficient in the sense of Fama (1970). Namely, the unbiased expectations hypothesis states that futures prices before maturity must be equal to the rational expectation of the settlement price at maturity:

$$f(t, T) = E_t[s(t + T)], \quad (4)$$

where $E_t[\cdot]$ is the rational expectations' operator conditional on the time t information filtration. This hypothesis is closely related to weak-form market efficiency, i.e., that future asset returns cannot be predicted by past returns. This framework can be extended to include not only past realisations of the variable, but also any other \mathcal{F}_t -measurable variable (Fama, 1991). Equivalently, future returns should be unpredictable by \mathcal{F}_t -measurable variables, such as valuation ratios, lagged risk premia, realised physical market conditions and economic indicators (Kavussanos and Nomikos, 1999).

Apart from the existence of return predictability, a straightforward way to test for unbiasedness is by performing a Wald test on the coefficients of the regression equation (3a). Namely, 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. The unbiasedness hypothesis is also examined by testing parameter restrictions in the cointegrating vector (Johansen, 1991) as well as using the Philips and Hansen (1990) fully modified

¹¹ We have also incorporated trading activity variables to examine whether liquidity in the FFA market can forecast either future market conditions and/or future risk premia. Our dataset is obtained from the London Clearing House (LCH) and consists of monthly observations on trading volume related to the BCI 4TC and BPI 4TC contracts with 1- and 2-month maturities. Our results indicate that trading volume is positively correlated with current market conditions and, furthermore, an improvement in market conditions is accompanied by a contemporaneous increase in trading volume; however, it does not have statistically significant forecasting power regarding future risk premia.

ordinary least squares (FMOLS) procedure. Results from these tests indicate the rejection of the unbiasedness hypothesis across all contracts and maturities, as expected.¹²

There exist two possible explanations for the rejection of the unbiasedness hypothesis: namely, the formation of irrational expectations and the existence of – either constant or time-varying – risk premia. To capture the stylised features of this market, we develop a dynamic asset pricing model of Forward Freight Agreements rates that can account both for “hedging pressure” bias and irrational, heterogeneous expectations. Accordingly, by analysing and simulating several alternative specifications of the model, we show that to reproduce these findings in a sufficient manner we need to depart from the rational expectations benchmark of the economy.

5.1 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 derivatives market with a forward contract (FFA) on the freight service corresponding to this route. Both markets operate in every period, that is, they clear at each t and, in turn, the respective equilibrium rate is determined. Naturally, the forward contract at each t is related to the spot rate at $t + 1$.

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, we examine the formation of FFA rates in a partial equilibrium framework. In line with the data (Tables I and II), the evolution of spot prices follows a random walk with zero drift process:

$$S_{t+1} = S_t + \varepsilon_{t+1} = S_t + \kappa_{t+1} + \lambda_{t+1}, \quad (5)$$

where $\varepsilon_{t+1} \sim iid N(0, \sigma_\varepsilon^2)$. We further assume that the random error term, ε_{t+1} , consists of two uncorrelated parts, $\kappa_{t+1} \sim iid N(0, \sigma_\kappa^2)$ and $\lambda_{t+1} \sim iid N(0, \sigma_\lambda^2)$; κ_{t+1} is realised a priori, that is at t , but is not observed by all market participants with the same precision. The reader can think of κ_{t+1} as a private information signal about future market conditions; 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 types of investors, 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-owners¹³ are the providers of the freight service who will hedge

¹² These results are available from the authors.

¹³ This group may also include vessel operators; the distinction between the two is that a shipowner owns a controlling interest in the ship while an operator is a management company that deals with the day-to-day operations of the ship (Stopford, 2009).

their freight income by taking a short position in the FFA market. 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 spot-settlement rate.

Charterers (cargo owners) are the consumers of the commodity since they transport their cargoes using ships. In practice, this group may correspond to large trading houses, commodity producers, mining companies and energy firms. By participating in the FFA market, charterers want to reduce their consumption risk and, like ship-owners, their demand also consists of a speculative component. In equilibrium, charterers are expected to take long positions on the derivative contract. Since ship-owners and charterers participate in both markets, they can be defined as “physical hedgers” or, equivalently, “traditional players”.

Finally, the third investor type corresponds to speculators; in practice, this group may consist of finance houses such as hedge funds and investment banks but also from individual investors. Speculators are motivated by purely speculative incentives thus their participation in the FFA market is not part of a diversification policy and their aim is to profit from absorbing part of the freight risk that ship-owners and charterers wish to hedge (Vives, 2008).

In line with the literature (e.g. 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 innovation of the proposed framework is that agents form heterogeneous expectations regarding future market conditions and they may be asymmetrically informed. Specifically, we assume that in agent i 's mind, spot prices evolve according to

$$S_{t+1} = (1 - \vartheta_i)[S_t + \rho_i \kappa_{t+1} + \lambda_{t+1}] + \vartheta_i[S_t + \psi_i(S_{t-1} - S_t) + \lambda_{t+1}] \quad (6a)$$

in which $\vartheta_i \in [0,1]$, $\rho_i \in [0,1]$, and $\psi_i > 0$.

The specification in (6a) consists of two terms or 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 distribution of λ_{t+1} are public information, the random term κ_{t+1} is not since it depends on the private information of each 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 closer ρ_i is to 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 (Shefrin, 2000); thus, for a totally rational investor $\psi_i = 0$. Overall, for investor i , the evolution of the spot price variable is given by a weighted average of these two signals; we call the weight coefficient ϑ_i “degree of wavering”. Equivalently, $(1 - \vartheta_i)$ quantifies the degree of confidence that investor i has about his private information.

We assume that physical hedgers are both perfectly informed and totally rational; thus, they only trust the fundamental value signal which they receive with perfect precision, so we set $\rho_o = \rho_c = 1$, $\psi_o = \psi_c = 0$, and $\vartheta_o = \vartheta_c = 0$. In contrast, speculators believe that spot price shocks tend to cancel out each other and spot rates tend to revert rapidly to their level before the last realised shock; that is, a price shock at t is followed by one of the opposite sign at $t + 1$. Therefore, speculators are both less than perfectly informed and irrational, that is, $\rho_s \in [0,1)$, $\psi_s > 0$, and $\vartheta_s \in (0,1)$ and 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 and have been doing so potentially for a long period; therefore, they are more experienced and/or better informed than speculators since they have “inside” information regarding physical market conditions. Hence, they are expected to form “more accurate” forecasts of future spot market conditions than the latter.¹⁴ However, it is important to note that the assumption of which group – or combination of groups – is the least sophisticated does not matter as the bias in FFA rates depends on the aggregate views of the individuals or, equivalently, the average market view. In other words, it does not matter who is irrational as long as the aggregate view is irrational.

The speculator-specific parameters ρ_s , ψ_s , and ϑ_s characterise completely the information structure of this model. When $\rho_s = 1$ and either ψ_s or ϑ_s equals zero, all agents are totally rational and have perfect and symmetric information about the economy. We define this case as the benchmark “rational” economy of the 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 $\psi_s, \vartheta_s > 0$, then, on average, aggregate investors’ expectations in the market are formed in an irrational manner. In concise form, equation (6a) can be written as:

¹⁴ We assume that speculators are less sophisticated traders in the sense that they are less informed and follow contrarian strategies compared to physical market participants. Although this assumption is in line with the structure and composition of participants in the shipping industry, one might argue that in practice the opposite is true since non-commercial investors, such as hedge funds, may use advanced analytics that are not available to traditional players who are more conservative. In fact, it may be the case that shipowners and/or cargo owners are the least sophisticated players. In support of this argument, Kang *et al* (2018) suggest that commercial hedgers may hedge based also on their market views and thus, their position may also include a speculative component. Namely, their empirical results suggest that in many cases commercials trade as contrarians.

$$S_{t+1} = S_t + (1 - \vartheta_i)\rho_i\kappa_{t+1} + \vartheta_i\psi_i(S_{t-1} - S_t) + \lambda_{t+1} \quad (6a')$$

Incorporating in equation (6a') 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

$$E_t^i[S_{t+1}] = S_t + (1 - \vartheta_i)\rho_i\kappa_{t+1} + \vartheta_i\psi_i(S_{t-1} - S_t) \quad (6b)$$

and

$$\text{Var}_t^i[S_{t+1}] = \text{Var}_t[S_{t+1}] = \sigma_\lambda^2. \quad (6c)$$

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 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 FFA contracts with the aim of maximising their respective mean-variance objective functions. First, for each ship-owner this corresponds to

$$\max_{h_t^o} E_t^o[S_{t+1}Q_{t+1} + h_t^o(S_{t+1} - F_t)] - \frac{\gamma_o}{2} \text{Var}_t[S_{t+1}Q_{t+1} + h_t^o(S_{t+1} - F_t)], \quad (7a)$$

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. Following Gorton *et al* (2012) and Hong and Yogo (2012), we assume that ship-owners at time t know with certainty the amount of shipping services they will sell at time $t + 1$, Q_{t+1} , which is plausible for large shipping companies. The optimisation yields

$$h_t^o = \frac{E_t^o[S_{t+1} - F_t]}{\gamma_o \text{Var}_t[S_{t+1}]} - Q_{t+1}. \quad (7b)$$

Second, each charterer maximises

$$\max_{h_t^c} E_t^c[-S_{t+1}D_{t+1} + h_t^c(S_{t+1} - F_t)] - \frac{\gamma_c}{2} \text{Var}_t[-S_{t+1}D_{t+1} + h_t^c(S_{t+1} - F_t)], \quad (8a)$$

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. As in the case of shipowners, we assume that charterers at time t know with certainty the amount of shipping services they will demand at time $t + 1$. The optimisation yields

$$h_t^c = \frac{E_t^c[S_{t+1} - F_t]}{\gamma_c \text{Var}_t[S_{t+1}]} + D_{t+1}. \quad (8b)$$

Third, speculator's maximisation problem is

$$\max_{h_t^s} E_t^s[h_t^s(S_{t+1} - F_t)] - \frac{\gamma_s}{2} \text{Var}_t[h_t^s(S_{t+1} - F_t)], \quad (9a)$$

where h_t^s is his time t demand for the FFA contract. This yields

$$h_t^s = \frac{E_t^s[S_{t+1} - F_t]}{\gamma_s \text{Var}_t[S_{t+1}]} \quad (9b)$$

In equilibrium, FFA contracts are in zero net supply. Therefore, the market clearing condition at each t requires

$$h_t^o + h_t^c + h_t^s = 0. \quad (10)$$

Substituting equations (7b), (8b), and (9b) in (10), the equilibrium FFA rate at t , F_t^* , is endogenously determined and equals

$$F_t^* = \frac{\gamma_c \gamma_s E_t^o[S_{t+1}] + \gamma_o \gamma_s E_t^c[S_{t+1}] + \gamma_o \gamma_c E_t^s[S_{t+1}]}{\gamma_c \gamma_s + \gamma_o \gamma_s + \gamma_o \gamma_c} - \frac{\gamma_o \gamma_c \gamma_s}{\gamma_c \gamma_s + \gamma_o \gamma_s + \gamma_o \gamma_c} \sigma_\lambda^2 (Q_{t+1} - D_{t+1}). \quad (11)$$

Equation (11) 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 $E_t^R[S_{t+1}]$. The second term measures the balance between supply and demand for freight services and thus quantifies the “hedging pressure” bias in the FFA rate, the direction of which depends only on the sign of the term in brackets.

In structural models for commodity markets, hedging pressure is determined endogenously by incorporating the theory of storage and 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.

For conciseness, in the following we define the hedging pressure variable as $Q_t - D_t = HP_t$. When HP_t equals zero the two positions exactly offset each other; on the other hand, when HP_t is positive, the physical market position of short hedgers (ship-owners) exceeds the one of long hedgers (charterers) and vice versa. In conclusion, in rational expectations symmetric information models, when “hedging pressure” is equal to zero, hedgers’ positions in the derivative market exactly offset each other and the derivative contract’s price is an unbiased predictor of the future spot price.

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. Thus, in the following, we aim to provide the most plausible explanation for the existence of biases 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 (11) is simplified to

$$F_t^* = \frac{1}{3} \{E_t^o[S_{t+1}] + E_t^c[S_{t+1}] + E_t^s[S_{t+1}]\} - \frac{\gamma\sigma_\lambda^2}{3} HP_{t+1}. \quad (12)$$

Plugging equation (6b) in (12) for $i = o, c, s$ subject to $\rho_o = \rho_c = 1$, $\psi_o = \psi_c = 0$, and $\vartheta_o = \vartheta_c = 0$ yields

$$F_t^* = S_t + \frac{2 + (1 - \vartheta_s)\rho_s}{3} \kappa_{t+1} + \frac{\vartheta_s\psi_s}{3} (S_{t-1} - S_t) - \frac{\gamma\sigma_\lambda^2}{3} HP_{t+1}. \quad (13)$$

It is also useful to examine the benchmark rational economy, R , in which the market solely consists of fully rational agents who know precisely the actual stochastic process that governs the evolution of spot prices. In this case, the expected spot price at $t + 1$ and the time t FFA rate, $F_t^{R,HP}$, are given by

$$E_t^R[S_{t+1}] = S_t + \kappa_{t+1}, \quad (14)$$

and

$$F_t^{R,HP} = S_t + \kappa_{t+1} - \frac{\gamma\sigma_\lambda^2}{3} HP_{t+1}, \quad (15)$$

respectively.

Comparing (13) to (15), we observe that $F_t^* = F_t^{R,HP}$ if and only if $S_{t-1} = S_t$ and $\kappa_{t+1} = 0$; that is, if there is no random shock, ε_t , between $t - 1$ and t and $\kappa_{t+1} = 0$ and there is no private information/signal about future spot market conditions. Whenever a shock perturbs the equilibrium,

however, the futures price deviates from its rational equilibrium counterfactual. The sign and magnitude of this deviation depend on the values of the shocks $S_t - S_{t-1} = \varepsilon_t$ 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 (5) from (13):

$$F_t^* - S_{t+1} = \left[\frac{(1 - \vartheta_s)\rho_s - 1}{3} \kappa_{t+1} + \frac{\vartheta_s \psi_s}{3} (S_{t-1} - S_t) \right] - \frac{\gamma \sigma_\lambda^2}{3} HP_{t+1} - \lambda_{t+1}. \quad (16)$$

The latter bias can be decomposed into three terms. We define the first term in brackets 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 error term of the cash flow process, λ_{t+1} , is different than zero.

Similarly, using equations (14), (15), and (5), 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

$$F_t^{R,HP} - E_t^R[S_{t+1}] = -\frac{\gamma \sigma_\lambda^2}{3} HP_{t+1} \quad (17a)$$

and

$$F_t^{R,HP} - S_{t+1} = -\frac{\gamma \sigma_\lambda^2}{3} HP_{t+1} - \lambda_{t+1}, \quad (17b)$$

respectively. Moreover, in the absence of hedging pressure, this becomes

$$F_t^R = S_t + \kappa_{t+1}, \quad (18)$$

and, in turn, the rationally expected risk premium at t is

$$F_t^R - E_t^R[S_{t+1}] = 0, \quad (19)$$

while the realised risk premium at $t + 1$ is given by

$$F_t^R - S_{t+1} = -\lambda_{t+1} \quad (20)$$

Therefore, even if there is no hedging pressure and all investors are rational and have access to the same information, the realised risk premium can be different than the rationally expected one, which

in this case will always be equal to zero. Specifically, from equation (20), the realised risk premium depends on the realisation and the distributional properties of the error term. Thus, since $\lambda_{t+1} \sim i.i.d. N(0, \sigma_\varepsilon^2)$ 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 in general – as documented in Section III.

In conclusion, both the fundamental structure of the economy, as quantified by 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 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 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

$$h_t^s = (1 - \vartheta_s) \frac{S_t + (1 - \vartheta_i) \rho_i \kappa_{t+1} - F_t}{\gamma \sigma_\lambda^2} + \vartheta_s \psi_s \frac{S_t - F_{t-1}}{\gamma \sigma_\lambda^2}. \quad (21)$$

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, would create the observed form of predictability and momentum in the market.

5.2 Calibration of the Model

We calibrate the economy described above for several different specifications of the model. For each scenario, we generate 10,000 sample paths using equation (5) each one corresponding to 120 monthly periods. If somewhere in a simulation either the spot rate variable or the FFA rate attain a negative value, we discard that path. Finally, we estimate the average of each statistic and regression estimate under consideration across all valid paths and we compare it to its empirical value (Barberis *et al*, 2015). In particular, we are interested in (i) the predictive power of the FFA basis regarding future spot growth and future risk premia, (ii) the predictive power of lagged spot growth and lagged risk premia regarding future risk premia, (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. The initial parameters used in the simulations are presented in Table IV and are described in greater detail below, for each simulation experiment.

We begin by examining our model's predictions in the simplest case, that is, when all agents are perfectly informed and totally rational and, furthermore, there is no hedging pressure in the FFA market. The FFA rate and the realised risk premium in this scenario are given by (18) and (20), respectively and 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 our sample (in thousand US dollars). 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 V) suggest that there is no predictability in risk premia and, thus, no momentum effect in prices; similarly, the mean basis and mean realised risk premium are both zero which contrasts sharply with the actual data presented in the last column of the same Table. In line with equation 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 [20]) which induces negative correlation between these two variables.

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. The FFA rate and the realised risk premium are given by (15) and (17b), respectively and $HP_0 = HP \neq 0$. Following Barberis *et al* (2015 and 2018), we set the coefficient of risk aversion, γ , equal to 0.1 while for the constant hedging pressure we select a value that ensures that the simulated average realised risk premium will be close to the observed one (Table I). We thus set $HP_0 = HP = -20$, that is, we assume that long hedgers' (charterers') position in the physical market constantly exceeds that of short hedgers (ship-owners).

As in the previous case, the simulation results (Scenario 2 of Table V) suggest that this specification is not consistent with risk premia predictability or a momentum effect. This is expected since (negative) constant hedging pressure implies a rationally expected (positive) constant risk premium *and not a time-varying one* (since $HP_{t+1} = -20$ in equation 17a). In turn, this results in both a positive mean basis and a positive mean realised risk premium (the latter can be shown by taking unconditional

expectations of both sides of equation [17b], for $HP = -20$). The positive predictability of future spot growth by the current basis can also be explained as the bias in the FFA rate is constant and, thus, 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.

In the third scenario, all agents are perfectly informed and totally rational, as before, but there exists time-varying hedging pressure. As analysed above, we cannot apply the “theory of storage” 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 for freight to empirically examine its relationship with spot rates. As such, to account for time-varying hedging pressure, we assume a stochastic process based on plausible economic arguments, that relates it to the exogenously determined spot rate process. Since hedging pressure is defined as the difference between demand for short hedging positions, which is related to fleet supply, and demand for long hedging positions, which is related to demand for seaborne trade, one should expect the former and the latter to be negatively and positively related to conditions in the freight market, respectively; specifically, $Q_t - D_t = HP_t$ should be negatively related to S_t . Hence, the evolution of hedging pressure 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 a linear inverse demand function:

$$S_t = \alpha T_t - \beta F_t, \quad (22)$$

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 (22) in a straightforward manner. Specifically, recall that at t , 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:

$$\begin{cases} Q_{t+1} = E_t^R[F_{t+1}] \\ D_{t+1} = E_t^R[T_{t+1}] \end{cases} \quad (23)$$

Importantly, this assumption can be directly related to the private signal about the spot rate, κ_{t+1} , realised at time t that market participants receive with perfect precision. Furthermore, since fleet

supply in the short run is inelastic, we set F_t and, in turn, Q_t equal to a constant, Q .¹⁵ 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 5, 22, and 23:

$$D_{t+1} = \frac{\beta}{\alpha}Q + \frac{S_t + \kappa_{t+1}}{\alpha}. \quad (24)$$

Accordingly, the hedging pressure variable corresponding to $t + 1$ is given by

$$HP_{t+1} = Q_{t+1} - D_{t+1} = \left(1 - \frac{\beta}{\alpha}\right)Q - \frac{S_t + \kappa_{t+1}}{\alpha}. \quad (25)$$

Thus, it is a decreasing function of both current market conditions and the signal about future market conditions. Plugging in (15) equation (25) yields the expression for the rational expectations, time-varying hedging pressure FFA rate:

$$F_t^{R,HP} = \left(1 + \frac{\gamma\sigma_\lambda^2}{3\alpha}\right)(S_t + \kappa_{t+1}) - \frac{\gamma\sigma_\lambda^2}{3}\left(1 - \frac{\beta}{\alpha}\right)Q. \quad (26)$$

Finally, using equation (14), we calculate the rationally expected bias in the FFA rate as

$$F_t^{R,HP} - E_t^R[S_{t+1}] = \frac{\gamma\sigma_\lambda^2}{3\alpha}(S_t + \kappa_{t+1}) - \frac{\gamma\sigma_\lambda^2}{3}\left(1 - \frac{\beta}{\alpha}\right)Q, \quad (27)$$

while the realised one equals

$$F_t^{R,HP} - S_{t+1} = \frac{\gamma\sigma_\lambda^2}{3\alpha}(S_t + \kappa_{t+1}) - \frac{\gamma\sigma_\lambda^2}{3}\left(1 - \frac{\beta}{\alpha}\right)Q - \lambda_{t+1}. \quad (28)$$

Equations 26, 27, and 28 suggest that the FFA rate, the rationally expected bias, and the realised bias are all increasing functions of both current market conditions and the signal about future market conditions.

We calibrate parameters α , β , D_0 , and Q in the following manner. Equations 22 and 23 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 $HP = HP_0 = Q - D_0 = -20$. Thus, assuming $D_0 = 100$ yields $Q = 80$. Accordingly, since $S_0 = 20$, we can calibrate α and β from $20 = 100\alpha - 80\beta$; thus, assuming $\beta = 0.1$ implies $\alpha = 0.28$. To

¹⁵ In principle, ship-owners can affect short-run elasticity by slow steaming and reducing available utilisation. While this simplifying assumption can be easily relaxed, it does not have any qualitative or quantitative implication on the model's predictions and results.

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 of Table V) suggest that while this specification provides a better match for the observed regularities, it cannot simultaneously match two of the most important stylised features: the momentum effect and the negative predictability of risk premia by lagged spot market conditions since the coefficients in the respective regressions are statistically insignificant. Note that even if we were to 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 would still be no negative predictability of future risk premia by past market conditions.

This can be easily justified by examining equation (28) 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 the realised risk premium is an increasing function of lagged spot rate changes. In turn, this explains the positive 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 (28) at $t + 1$ as

$$F_{t+1}^{R,HP} - S_{t+2} = (F_t^{R,HP} - S_{t+1}) + \frac{\gamma\sigma_\lambda^2}{3\alpha}\kappa_{t+2} + \left(1 + \frac{\gamma\sigma_\lambda^2}{3\alpha}\right)\lambda_{t+1} - \lambda_{t+2}.$$

The remaining results under this scenario can be explained in a similar way.

The fourth case corresponds to the economy with asymmetric information and irrationality of beliefs. Furthermore, we assume that there is constant hedging pressure in the market as in scenario 2.¹⁶ Therefore, the FFA rate and the realised risk premium are given by equations 13 and 16 , respectively (setting $HP_0 = HP = -20$). 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

¹⁶ Note that the predictive regression results remain the same if we set hedging pressure equal to zero.

value” one (Barberis *et al*, 2018). 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.

The simulation results (Scenario 4 of Table V) suggest that this specification can match simultaneously almost all stylized facts. Most importantly, we observe that it can account not only for the momentum effect – the lagged risk premium coefficient being positive and statistically significant – but also for the negative predictability of future risk premia by lagged spot growth, since the lagged spot growth coefficient is negative and statistically significant.

The latter feature can be explained by examining equation 16 at $t + 1$:

$$F_{t+1}^* - S_{t+2} = \left[\frac{(1 - \vartheta_s)\rho_s - 1}{3} \kappa_{t+2} + \frac{\vartheta_s\psi_s}{3} (-\Delta S_{t+1}) \right] - \frac{\gamma\sigma_\lambda^2}{3} HP - \lambda_{t+2}. \quad (29)$$

Specifically, the realised risk premium is a decreasing function of the one-period lagged spot rate changes, $\Delta S_{t+1} = S_{t+1} - S_t$. Furthermore, equation 16 at t can be re-expressed as

$$-\Delta S_{t+1} = F_t^* - S_{t+1} - \left[\frac{(1 - \vartheta_s)\rho_s + 2}{3} \kappa_{t+1} + \frac{\vartheta_s\psi_s}{3} (S_{t-1} - S_t) \right] + \frac{\gamma\sigma_\lambda^2}{3} HP. \quad (30)$$

Plugging (30) in (29), we observe that the realised risk premium at $t + 1$ is an increasing function of the realised risk premium at t .

The only stylised facts poorly matched in this case are the ones related to the variance decomposition (Scenario 4 in Panel B of Table V); essentially, none of the variation in basis is attributed to time-varying risk premia in this case, because the “contrarian value signal” significantly reduces the volatility of realised risk premia. However, if we increase 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 for $\sigma_\lambda^2 = 2.5^2$ and $\gamma = 0.04$ is presented in Scenario 4’ in Table V. Finally, when there is irrationality of beliefs but information is symmetric the results are very similar to the ones above.

The last scenario combines some features from 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 the expression for hedging pressure in (25) in equation (13):

$$F_t^* = S_t + \frac{2 + (1 - \vartheta_s)\rho_s}{3} \kappa_{t+1} + \frac{\vartheta_s \psi_s}{3} (S_{t-1} - S_t) - \frac{\gamma \sigma_\lambda^2}{3} \left[\left(1 - \frac{\beta}{\alpha}\right) Q - \frac{S_t + \kappa_{t+1}}{\alpha} \right]. \quad (31)$$

Accordingly, the rationally expected bias and the realised one are given by

$$F_t^* - E_t^R[S_{t+1}] = \frac{(1 - \vartheta_s)\rho_s - 1}{3} \kappa_{t+1} + \frac{\vartheta_s \psi_s}{3} (S_{t-1} - S_t) - \frac{\gamma \sigma_\lambda^2}{3} \left[\left(1 - \frac{\beta}{\alpha}\right) Q - \frac{S_t + \kappa_{t+1}}{\alpha} \right] \quad (32)$$

and

$$F_t^* - S_{t+1} = \frac{(1 - \vartheta_s)\rho_s - 1}{3} \kappa_{t+1} + \frac{\vartheta_s \psi_s}{3} (S_{t-1} - S_t) - \frac{\gamma \sigma_\lambda^2}{3} \left[\left(1 - \frac{\beta}{\alpha}\right) Q - \frac{S_t + \kappa_{t+1}}{\alpha} \right] - \lambda_{t+1}, \quad (33)$$

respectively. The corresponding simulation results (Scenario 5 of Table V) suggest that this specification can simultaneously match all observed regularities in a sufficient manner. This result was expected since this scenario combines the features of the previous two economies.¹⁷

In conclusion, both the theoretical predictions and the simulation results of the proposed 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 distorted expectations feature; namely, a fraction of investors must suffer from the “gambler’s fallacy” or, equivalently, follow a contrarian investment strategy. From a market composition perspective, our results suggest that the average FFA investor should hold “near-rational” but slightly contrarian beliefs to match the observed risk premia predictability.

6 CONCLUSION

This article examines the formation of forward prices 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 suggested in the commodity

¹⁷ One may obtain values closer to the actual moments either through finer adjustment of the set of parameters or, by using exact closed-form expressions for the moments of interest. Nevertheless, we believe that the results and the realised patterns will be qualitatively similar.

finance literature. Our results validate and extend the economic arguments presented in the seminal commodity market papers that examined the forecasting power of derivative contracts. Namely, predictability of spot rates is an increasing function of the commodity cost of storage. In shipping, where the commodity is a non-storable service and the industry is subject to significant supply and demand shocks, we evidence 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. 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 FFA rates in both contracts.

To capture these stylised features we develop a dynamic behavioural asset pricing framework that can explain both the existence of momentum and the documented predictability of future risk premia. The proposed framework departs from the “theory of storage” and the “cost-of-carry” model and incorporates both the “hedging pressure” feature – the rational dimension – and a behavioural finance explanation – the irrational dimension. Accordingly, we show that the average FFA investor should hold “near-rational” but slightly contrarian beliefs to match the observed risk premia predictability in the market.


REFERENCES

- Acharya V.V., Lochstoer, L.A., & Ramadorai, T. 2013. Limits to arbitrage and hedging: Evidence from commodity markets. *Journal of Financial Economics* 109, 441-465. doi: [10.1016/j.jfineco.2013.03.003](https://doi.org/10.1016/j.jfineco.2013.03.003)
- Akaike, H. 1973. Information theory and extension of the maximum likelihood principle, in: Petrov, B., & Csake, F. (eds.), *Second international symposium on information theory*. Budapest, HU: Akademiai Kiado.
- Alizadeh, A.H., & Nomikos, N.K. 2009. *Shipping Derivatives and Risk Management*. London, UK: Palgrave Macmillan.
- Barberis, N., Greenwood, R., Lawrence, J., & Shleifer, A. 2015. X-CAPM: An extrapolative capital asset pricing model. *Journal of Financial Economics* 115, 1-24. doi: [10.1016/j.physa.2018.06.004](https://doi.org/10.1016/j.physa.2018.06.004)
- Barberis, N., Greenwood, R., Lawrence, J., & Shleifer, A. 2018. Extrapolation and Bubbles. *Journal of Financial Economics* 129, 203-227. doi: [10.1016/j.jfineco.2018.04.007](https://doi.org/10.1016/j.jfineco.2018.04.007)
- Bredin, D., Potì, V., & Salvador, E. 2018. Commodity Pricing: Evidence from Rational and Behavioral Models, Unpublished working paper. University College Dublin and Universitat Jaume I.
- Chen, L., Da, Z., & Priestley, R. 2012. Dividend smoothing and predictability. *Management Science* 58, 1834–1855. doi: [10.1287/mnsc.1120.1528](https://doi.org/10.1287/mnsc.1120.1528)
- Cochrane, J.H. 2011. Presidential address: Discount rates. *Journal of Finance* 66(4), 1046-1108. doi: [10.1111/j.1540-6261.2011.01671.x](https://doi.org/10.1111/j.1540-6261.2011.01671.x)
- Ekeland, I., Lautier, D., & Villeneuve, B. 2018. Speculation in commodity futures markets: A simple equilibrium model. *Economic Theory*, Forthcoming.


- Ellen, S.T., & Zwinkels, R.C.J. 2010. Oil price dynamics: A behavioural finance approach with heterogeneous agents. *Energy Economics* 32, 1427-1434. doi: [10.1016/j.eneco.2010.03.003](https://doi.org/10.1016/j.eneco.2010.03.003)
- Engle, R.F., & Granger, C.W. 1987. Cointegration and error correction: Representation, estimation, and testing. *Econometrica* 55, 251-276.
- Fama, E.F. 1970. Efficient capital markets: A review of theory and empirical work. *Journal of Finance* 25, 383-417. doi: [10.1111/j.1540-6261.1970.tb00518.x](https://doi.org/10.1111/j.1540-6261.1970.tb00518.x)
- Fama, E.F. 1984a. Forward and spot exchange rates. *Journal of Monetary Economics* 14, 319-338. doi: [10.1016/0304-3932\(84\)90046-1](https://doi.org/10.1016/0304-3932(84)90046-1)
- Fama, E.F. 1984b. The information in the term structure. *Journal of Financial Economics* 13, 509-528. doi: [10.1016/0304-3932\(84\)90046-1](https://doi.org/10.1016/0304-3932(84)90046-1)
- Fama, E.F., & 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., & French K.R. 1988. Permanent and temporary components of stock prices. *The Journal of Political Economy* 96, 246-273. doi: [10.1086/261535](https://doi.org/10.1086/261535)
- Fama, E.F. 1991. Efficient Capital Markets: II. *The Journal of Finance* 46(5), 1575-1617. doi: [10.1111/j.1540-6261.1991.tb04636.x](https://doi.org/10.1111/j.1540-6261.1991.tb04636.x)
- French, K. R. 1986. Detecting Spot Price Forecasts in Futures Prices. *The Journal of Business* 59(2, Part 2), S39-S54.
- Gorton, G.B., Hayashi, F., & Rouwenhorst, G.K. 2012. The Fundamentals of Commodity Futures Returns. *Review of Finance* 17, 35-105. doi: [10.1093/rof/rfs019](https://doi.org/10.1093/rof/rfs019)
- Greenwood, R., & Hanson, S.G. 2015. Waves in ship prices and investment. *Quarterly Journal of Economics* 130, 44-109. doi: [10.1093/qje/qju035](https://doi.org/10.1093/qje/qju035)
- Greenwood, R., & Shleifer, A. 2014. Expectations of returns and expected returns. *Review of Financial Studies* 27, 714-746. doi: [10.1093/rfs/hht082](https://doi.org/10.1093/rfs/hht082)
- Hazuka, T.B. 1984. Consumption betas and backwardation in commodity markets. *Journal of Finance* 39, 647-655. doi: [10.1111/j.1540-6261.1984.tb03653.x](https://doi.org/10.1111/j.1540-6261.1984.tb03653.x)
- Hicks, J. R. 1939. *Value and Capital*. Cambridge, UK: Oxford University Press.
- 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. doi: [10.1016/j.jfineco.2012.04.005](https://doi.org/10.1016/j.jfineco.2012.04.005)
- 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. Doi: [10.2307/2938278](https://doi.org/10.2307/2938278)
- Kalouptzidi, M. 2014. Time to build and fluctuations in bulk shipping. *American Economic Review* 104(2), 564-608. doi: [10.1257/aer.104.2.564](https://doi.org/10.1257/aer.104.2.564)
- Kang, W., Rouwenhorst, K.G., & Tang, K. 2018. A Tale of Two Premiums: The Role of Hedgers and Speculators in Commodity Futures Markets. Unpublished working paper, Shanghai University of Finance and Economics, Yale School of Management, and Tsinghua University.
- Kavussanos, M.G., & Nomikos N.K. 1999. The Forward Pricing Function of the Shipping Freight Futures Market. *The Journal of Futures Markets* 19(3), 353-376. doi: [10.1002/\(SICI\)1096-9934\(199905\)19:3<353::AID-FUT6>3.0.CO;2-6](https://doi.org/10.1002/(SICI)1096-9934(199905)19:3<353::AID-FUT6>3.0.CO;2-6)
- Keynes, J. M. 1930. *Treatise on Money*. New York, NY: Macmillan.
- 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. doi: [10.1257/aer.99.3.1053](https://doi.org/10.1257/aer.99.3.1053)
- Newey, W.K., & West, K.D. 1987. A simple, positive-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55, 703-708.
- Nomikos, N.K. Moutzouris, I.K. 2018a. The earnings yield and predictability of earnings in the dry bulk shipping industry. Unpublished working paper, Cass Business School.
- Nomikos, N.K., & Moutzouris, I.K. 2018b. Extrapolative Expectations and the Second-Hand market for Ships. Unpublished working paper, Cass Business School.
- Papapostolou, N.C., Nomikos, N.K., Poulialis, P.K., & Kyriakoy, I. 2014. Investor sentiment for real assets: The case of dry bulk shipping market. *Review of Finance* 18, 1507-1539. doi: [10.1093/rof/rft037](https://doi.org/10.1093/rof/rft037)
- Rangvid, J., Schmeling, M., & Schrimpf, A. 2014. Dividend predictability around the world. *Journal of Financial and Quantitative Analysis* 49, 1255-1277. doi: [10.1017/S0022109014000477](https://doi.org/10.1017/S0022109014000477)
- Shefrin, H. 2000. *Beyond Greed and Fear: Understanding Behavioural Finance and the Psychology of Investing*. Boston, MA: Harvard Business School Press.
- Stopford, M. 2009. *Maritime Economics* (3rd ed.). New York: Routledge.

- Szymanowska, M., De Roon, F., Nijman, T., & Van Den Goorbergh, R. 2014. An Anatomy of Commodity Futures Risk Premia. *The Journal of Finance* 69(1), 453-482. doi: [10.1111/jofi.12096](https://doi.org/10.1111/jofi.12096)
- Tversky, A., & Kahneman D. 1971. Belief in the Law of Small Numbers. *Psychological Bulletin* 76, 105-110. doi: [10.1037/h0031322](https://doi.org/10.1037/h0031322)
- Vives, X. 2008. *Information and Learning in Markets: The Impact of Market Microstructure*. Princeton, NJ: Princeton University Press.

FIGURES AND TABLES



Panel A: FFA, Spot, and Settlement Rates.



Panel B: Basis, Spot Growth, and Risk Premium.

Figure 1: Variables of Interest.

Panel A plots the evolution of spot, settlement, and FFA rates from January 2007 to August 2016, for the 1-month BPI 4TC contract; figures are in thousand US\$/day. The spot and settlement rates are the prices observed at the issuance and maturity of the corresponding FFA 1-month contract, respectively. Panel B plots the evolution of basis, spot growth, and risk premium for the 1-month BPI 4TC contract. All variables correspond to log differences. Note that spot rates, FFA rates, and the basis are reported at time t while settlement rates, spot growth, risk premia at time $t + 1$.




Figure 2: Risk Premia, Lagged Risk Premia, and Lagged Spot Growth.

This figure plots the evolution of risk premia, lagged risk premia, and lagged spot growth, for the 1-month BPI 4TC contract. All variables are in 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.

Table I: Descriptive statistics for the variables of interest.

Statistics	BCI 4TC		BPI 4TC	
	1-month	2-month	1-month	2-month
n	116	115	116	115
Mean of $f - s$	9.02%	12.90%	4.83%	8.88%
t of $f - s$	4.07	3.58	4.07	4.20
Annualised Mean of r	79.63%	76.58%	31.98%	49.95%
t^{NW} of r	2.22	1.43	1.75	1.25
SD of $f - s$	0.24	0.39	0.13	0.23
SD of Δs	0.37	0.72	0.18	0.36
SD of r	0.32	0.60	0.16	0.35
$Corr(\Delta s, r)$	-0.77	-0.84	-0.72	-0.80
ρ_1 of r	0.21	0.48	0.32	0.60
ρ_2 of r	0.10	0.10	0.22	0.34

Notes: This table presents descriptive statistics related to the FFA basis, $f - s$; spot growth, Δs , and risk premium, r , for the 1- and 2-month Capesize BCI 4TC and Panamax BPI 4TC contracts. All variables are expressed in log differences. The sample runs from January 2007 to September 2016. The number of observations is denoted by n . The included statistics are: the mean and t-statistic of the basis; the annualised mean and t-statistic, t^{NW} , of the risk premium, estimated using the Newey-West (1987) HAC correction; the standard deviations of the three variables, SD; and the correlation coefficient between spot growth and risk premium, $Corr(\Delta s, r)$. Note that when the t-statistic indicates significance at least at the 10% level, the respective mean statistic appears in bold.

Table II: Variance Decomposition and regressions of spot growth on lagged spot growth.

Regression	T	n	β	t^{NW}	R^2
Panel A: Capesize Sector (BCI 4TC)					
Δs on $f - s$	1	116	0.80***	6.39	0.26
r on $f - s$	1	116	0.20	1.61	0.02
Δs on $\Delta s(-1)$	1	115	-0.08	-0.83	0.01
Δs on $f - s$	2	115	1.04***	7.83	0.31
r on $f - s$	2	115	-0.04	-0.28	0.00
Δs on $\Delta s(-1)$	2	114	-0.24	-1.31	0.01
Panel B: Panamax Sector (BPI 4TC)					
Δs on $f - s$	1	116	0.63***	5.48	0.21
r on $f - s$	1	116	0.37***	3.20	0.08
Δs on $\Delta s(-1)$	1	115	-0.06	-0.62	0.00
Δs on $f - s$	2	115	0.59***	2.84	0.14
r on $f - s$	2	115	0.41*	1.97	0.07
Δs on $\Delta s(-1)$	2	114	-0.10	-0.46	0.00

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, $f - s$, and future spot growth, Δs , on lagged 1-month spot growth, $\Delta s(-1)$, 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 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 slope coefficient, β , is accompanied by *, **, or *** when the absolute t^{NW} statistic indicates significance at the 10%, 5% or 1% level, respectively.

Table III: Regressions of future risk premia on lagged risk premia and spot 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)								
$f(t - 1, 1) - s(t)$	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
$s(t) - s(t - 1)$	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)								
$f(t - 1, 1) - s(t)$	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
$s(t) - s(t - 1)$	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 lagged risk premia and 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 1-period risk premium, $f(t - l, 1) - s(t - l + 1)$ where the number of lags, l , varies from 1 to 3. In the fourth row, the predictor is the corresponding previous 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. In rows five to seven of each panel, the predictor is the lagged 1-period spot growth $s(t - l) - s(t - l - 1)$ where the number of lags, l , varies from 1 to 3. In the eighth 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 . Finally, in rows nine and ten of each panel, the predictor is the first and second lag of the log growth of the Baltic Dry Index (BDI), 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.

Table IV: Parameter values.

Parameter	Assigned Value
S_0	20
σ_κ^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_κ^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 .

Table V: Model predictions for the quantities of interest.

Quantity	Rational Benchmark			Heterogeneous-Agent Economy			Actual Data
	Scenario 1	Scenario 2	Scenario 3	Scenario 4	Scenario 4'	Scenario 5	
Panel A: Basis, Risk Premium, and Spot Growth							
Mean of Basis	0.00	0.04	0.04	0.03	0.08	0.05	0.05
p-value of Basis	0.44	0.00	0.17	0.00	0.00	0.17	0.00
Mean of Risk Premium	0.00	0.05	0.04	0.02	0.07	0.04	0.03
p-value of Risk Premium	0.56	0.00	0.16	0.00	0.00	0.16	0.08
Risk Premium and Spot Growth Correlation	-0.70	-0.71	-0.60	-0.81	-0.87	-0.75	-0.72
Panel B: Predictive Power of the FFA Basis							
Future Spot Growth Coefficient	0.99	1.00	0.75	1.02	0.56	0.80	0.63
Future Spot Growth p-value	0.00	0.00	0.00	0.00	0.03	0.00	0.00
Future Spot Growth R^2	0.50	0.50	0.38	0.36	0.08	0.28	0.21
Future Risk Premium Coefficient	0.01	0.00	0.25	-0.02	0.44	0.20	0.37
Future Risk Premium p-value	0.44	0.45	0.13	0.50	0.10	0.28	0.00
Future Risk Premium R^2	0.01	0.02	0.09	0.01	0.07	0.06	0.08
Panel C: Predictability of Risk Premia							
Lagged Spot Growth Coefficient	0.01	-0.01	0.07	-0.30	-0.28	0.35	-0.15
Lagged Spot Growth p-value	0.45	0.46	0.39	0.01	0.05	0.05	0.00
Lagged Spot Growth R^2	0.01	0.01	0.02	0.14	0.09	0.08	0.08
Lagged Risk Premium Coefficient	-0.01	0.02	0.14	0.31	0.30	0.23	0.32
Lagged Risk Premium p-value	0.45	0.44	0.28	0.02	0.03	0.01	0.00
Lagged Risk Premium R^2	0.01	0.01	0.04	0.10	0.10	0.14	0.10

Notes: This table summarises the theoretical model's predictions for the quantities of interest presented in the left column. The right column presents the empirical values of these quantities as illustrated in Sections II and III of the paper. Columns 2-6 report the average value of each quantity across 10,000 simulated paths, for a given market scenario as analysed in the main text. The basic model parameters are presented in Table IV. Scenario 4' corresponds to scenario 4 for $\sigma_\lambda^2 = 2.5^2$ and $\gamma = 0.04$.