2

THE THEORETICAL FOUNDATIONS OF THE FREIGHT MARKET

EXECUTIVE SUMMARY

In the traditional model of the freight market, which dates back to the 1930s, freight rates are determined in a perfectly competitive market, where the stock of fleet is predetermined at any point in time. This implies that freight rates adjust instantly to clear the demand–supply balance, where supply is fixed. Accordingly, freight rates respond exclusively to fluctuations in demand, rising when demand increases and falling when demand falls. This assumption introduces an element of irrationality on the part of both owners and charterers because the supply is constantly changing – new ships arrive in the market all the time. A cursory look at deliveries of new vessels shows that there are significant changes in supply from month to month. Hence, the assumption that supply is fixed in the short run is not appropriate. Instead, both charterers and owners form expectations of demand and supply and this requires a dynamic analysis rather then a static one in which the fleet is fixed. In this chapter we suggest an alternative theorising of the freight market, which captures this dynamic analysis of freight rates. This new framework consists of a bargaining process over freight rates in which charterers and owners form expectations of demand and supply over an investment horizon relevant to their decisions. In this framework, freight rates are viewed as asset prices, which are determined by discounting future economic fundamentals.

In the traditional model the demand for and supply of shipping services are functions of freight rates in a perfectly competitive market. But freight rates are not prices determined in auction markets, where many owners bid for the same cargo and the one with the lower bid wins the contract (Dutch auctions). The characteristics of the freight market do not accord with those of perfect competition. Intuitively, perfect competition is a market system where the actions of individual buyers and sellers have a negligible impact on the market and where both are price takers. The assumptions of perfect competition are not satisfied in the freight market. In particular, the product is not homogeneous;[1] the assumption of a very large (in theory infinite) number of buyers and sellers is not applicable, transaction costs are not zero and there is no freedom of entry and exit. Although the product is seemingly homogeneous (the capacity to transport particular categories of products or commodities), the demand for shipping services is restricted by volume, time and route – a given cargo over a particular route that meets a well-specified time schedule. Although there are many ships in the market only a few are available to satisfy the given demand specifications in time and place. These characteristics violate the homogeneous product assumption, the condition of large (infinite) number of buyers and sellers and the hypothesis of zero transaction costs. The latter should be interpreted as the penalties (legal or reputational) that a charterer would incur for waiting for a better deal (lower freight rates). Moreover, there are large barriers to entry both in capital and the operation and management of the fleet.

Rather, the freight rate over a particular cargo is the outcome of a bargaining process that happens at the same time – or approximately the same time – in different places and where information about freight rates agreed is almost instantaneously available to all other participants. Thus, the agreed freight rates do not balance demand and supply in a particular place at a particular point in time, but rather expectations of overall demand and supply in a particular segment of the market or the entire market. Accordingly, freight rates are equilibrium rates in a bargaining game where players form rational expectations about economic conditions.

The negotiations between an owner and a charterer over a contract for a particular cargo can best be viewed as a zero-sum game between the two players.[2] Both players know the freight rates that have been agreed so far. This information permeates to the rates that were agreed on the same or similar routes with ships of the same or different capacity. Players also know not just the latest rates, but their entire history. This enables them to assess and form expectations of the dynamic evolution of future freight rates. Thus, when the charterer and the owner enter the negotiations they would bargain over the deviation of the expected future freight rate from the latest or equilibrium rate. The final outcome will be influenced by the bargaining power of each party. In this context, it is better to formulate the problem as bargaining over the discounted present value of future freight rates. If the bargaining power of the charterer is stronger than that of the owner, the deviation of the agreed freight rate from the latest or equilibrium rate will be negative, implying a lower rate than the latest one. If, on the other hand, the bargaining power of the owner is stronger, the deviation of the agreed freight rate from the latest or equilibrium rate will be positive, implying a higher rate than the latest one. In some negotiations one of the parties is a big player and has the upper hand. The agreed freight rate would be to the advantage of the big player in the context of the market average. But such freight rates are outliers (they belong to the tails of the distribution). In the median negotiation the bargaining power of the charterer and the owner depends on economic conditions. In 'good' or improving economic conditions the owner has stronger bargaining power, whereas in 'bad' or worsening economic conditions the charterer's bargaining power prevails. Hence, in 'good' or improving economic conditions, freight rates would be on an uptrend; and vice versa. It is shown in this chapter that expectations about key shipping variables are formed by expectations of how policymakers (mainly central banks) would respond to current and future economic conditions.

1 A FRAMEWORK FOR MARITIME ECONOMICS

The first part of the book deals with the microfoundations of maritime economics (or shipping markets). Shipping is organised in the form of four markets: the freight market; the shipyard (or newbuilding) market; the scrap market and the secondhand market. The freight market is subdivided into the spot market and the time charter (or period) market. These markets pertain to all ship types (dry bulk, oil tankers, containers and specialised ships, like LNG) and all ship sizes. In the dry sector there are four major ship sizes: Capesize, Panamax/Kamsarmax, Handymax/Supramax and Handysize. In the oil tanker (or wet) market there are five major ship sizes: VLCC, Suezmax, Aframax, Panamax and Handysize. In this taxonomy there are four markets (freight, newbuilding, scrap and secondhand) for each type of ship and for each size.

The microfoundations developed here are common to all types and ship sizes. The microfoundations attempt to derive the general form of the underlying demand and supply functions in all four markets based on the principles of rationality, which is the basis of a scientific approach to shipping. Economic agents are assumed to be utility or profit maximisers subject to well-defined economic and technical constraints. This maximising behaviour gives rise not just to the functional form of the underlying demand and supply functions, but also to their exact determinants and in most cases to the qualitative influence (positive or negative) of each determinant on the demand or supply. These restrictions are important for drawing inferences and in empirical work on maritime economics.

This framework enables one to analyse the impact of exogenous shocks (anticipated or unanticipated) on the equilibrium of the shipping system as a whole and a single market (for example, Capes). In the first part of the book we deal with the microfoundations of each single market (freight, newbuilding, scrap and secondhand). In the second part of the book we integrate the shipping markets and examine the properties of the entire system. Moreover, we analyse the interrelationship of shipping with the macroeconomy. The economy helps to explain the demand for shipping services, which in traditional analysis is treated as exogenous to shipping markets. The integration of shipping markets with macroeconomics sheds lights on how expectations in the shipping markets are formed. In this context the assumptions of rational expectations become more palatable to swallow. For it is one thing to assume that expectations are, on average, correct for freight rates and ship prices (newbuilding and second hand) and another to assume that expectations for short-term interest rates are, on average, correct. The second assumption is more palatable to swallow, given the emphasis of major central banks on shaping and influencing interest rate expectations through announcing the targets of economic policy, the extent to which they would tolerate deviations from conflicting targets and how long they would stick with current policies (like low interest rates). In the macroeconomic approach to maritime economics expectations about policy drive expectations in shipping markets.

In this framework, developments in the major regions of the world economy shape the major forces of demand and supply in the overall shipping market (whether dry, wet or containership). These developments in the overall market infiltrate in time to the various sectors in a manner that takes into account the disequilibrium of each sector from the overall market. For example, if freight rates, the fleet size or the prices of a particular sector, say Capes, are higher than the overall market by more than justified by economic fundamentals, then Capes would adjust through time so that equilibrium is attained once more.

Hence, every market (dry, wet or containership) consists of five variables, which are determined simultaneously. These are the demand for shipping services (the cargo being transported), the stock of the net fleet, freight rates, newbuilding (NB) prices and secondhand (SH) prices. Each variable is shaped in one or more markets, but all markets are interacting with each other. In the freight market, the demand for shipping services by charterers and the supply of shipping services by owners determine the amount of cargo transported and freight rates. In the shipyard market, the demand for vessels by owners and the supply of vessels by shipyards determine NB prices and the deliveries of new vessels, which are added to the existing stock of fleet. In the secondhand market, the demand for vessels by owners and the supply of vessels by other owners, determines the SH prices and the volume of sales/purchases. In the scrap market, the demand for scrap metal by scrapyards and the supply of vessels for demolition by owners determine the price of scrap metal and the volume of ships which are demolished. The stock of the net fleet at the end of each period, say quarter, is simply the stock of old vessels at the end of the previous period, augmented by the deliveries in the current period, less the vessels for demolition in this same period.

This hierarchical structure of the dry market is presented in Box 2.1 for the dry market. A similar structure exists for the wet and containership markets. Macroeconomic developments affect the conditions in the overall dry market, which are then transmitted to each sector of the market. The four interacting markets are presented in Box 2.2. Demand and supply in each market determine the equilibrium price and quantity. With the help of these boxes it is easy to appreciate the interactions of the various markets. A booming world economy spurs world trade and the demand for shipping services. With a fixed stock of net fleet, but assuming some spare fleet capacity, the cargo being transported increases and freight rates go up. If freight rates cover the operational costs, then owners increase the supply of shipping services to meet the higher demand by increasing the speed of the vessels. A sustained increase in demand that cannot be met by higher speed induces owners to buy in the secondhand market and order new ships. This increases SH prices, lowers the fleet demolition and increases the demand for new vessels. Shipyards respond by increasing NB prices, as in the short run they cannot meet the higher demand for vessels; it takes time, approximately two years, to deliver a new vessel.

Box 2.1 A macroeconomic (hierarchical) approach to shipping

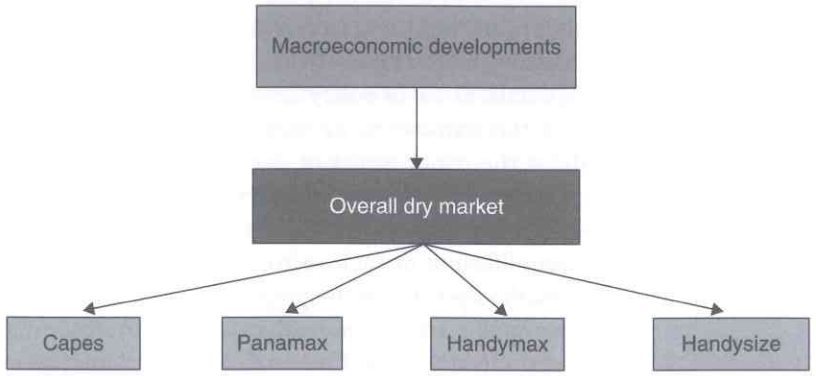

Box 2.2 The structure of the dry market

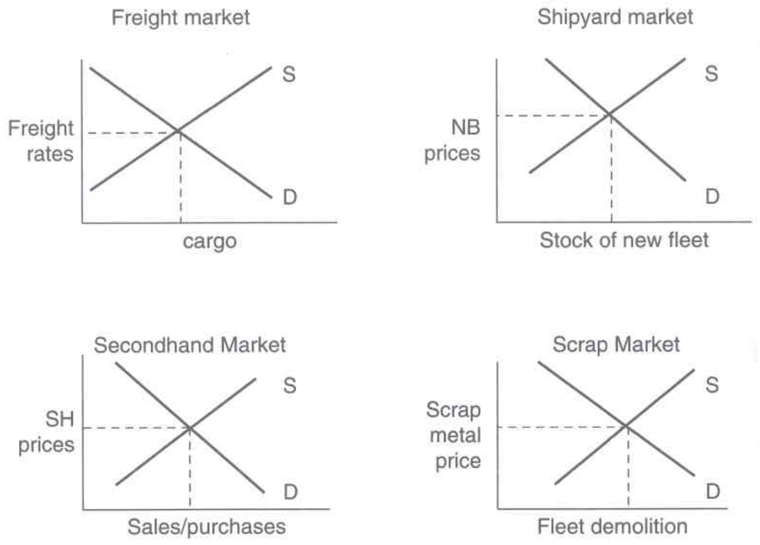

When the world economy slows down or falls into recession the demand for shipping services falls. The cargo being transported is reduced and freight rates slide. The demand for secondhand ships diminishes, SH prices fall and demolition picks up steam. Owners cancel orders and NB prices fall. This means that shipping cycles are primarily caused by economic (or business cycles).

The macroeconomic approach to maritime economics differs from the traditional approach, where the freight market is isolated from the rest of the system. So, in the traditional approach the system is not simultaneous; it is post-recursive, namely it can be arranged in a particular order to be solved. First the freight market is solved and then the equilibrium level of freight rates and cargo enter the other three markets, which are then solved simultaneously. It should be borne in mind that the Beenstock and Vergottis (1993) model, which is regarded as a high-water mark in shipping systems, is entirely post-recursive. This destroys the simultaneity of the shipping system and gives the impression that shipping decisions are easy to take and depend entirely on developments in the shipping market, which are governed exclusively by exogenous developments in the demand for shipping services. This approach has given rise to the relative isolation of maritime economics from other branches of economics and a tendency for treating ship sizes as segmented markets. Instead, the macroeconomic approach gives priority to the overall market and postulates a hierarchical approach to shipping. Thus, the market of each ship size co-moves (or, in the jargon of econometrics, is co-integrated) with the overall market. Therefore, shocks to the overall market are transmitted to each ship market. In time, each ship-size market moves to equilibrium with the overall market.

Part I of the book consists of two chapters. Chapter 2 analyses the theoretical foundations of the freight market, which is split into two markets: the spot market and the time charter market. In this chapter we offer a new framework for analysing spot freight rates. This framework has some implications for the nature of the risk premium in the time charter market.

Chapter 3 examines the theoretical foundations of the shipyard, scrap and secondhand markets. It starts with a single owner's decision problem of the optimal fleet and explains how the demand for newbuilding vessels is derived at the individual and aggregate level. It then goes on to consider the shipyard market and examines the influence of supply in determining the price of new vessels and vessel deliveries. It then analyses the secondhand market and shows how the demand and supply functions are obtained. Finally, it considers the scrap market and explains how the net fleet is determined.

2 THE TRADITIONAL MODEL

The traditional model of freight rates goes back to Tinbergen (1931, 1934), Koopmans (1939), Hawdon (1978), Strandenes (1984, 1986) and Beenstock and Vergottis (1993). The demand for and supply of shipping services are functions of freight rates in a perfectly competitive market. The cargo is measured in tonne-miles in recognition of the fact that both the volume of the cargo to be transported and the distance covered matter. The demand for shipping services is assumed to be a negative function of freight rates, while the supply of shipping services is a positive function. The demand for shipping services is assumed to be very inelastic with respect to freight rates, as charterers have a lot to lose if the entire cargo that is earmarked for transport is not shipped and does not arrive on time at the destination port. The supply of shipping services, on the other hand, is supposed to be a non-linear function of freight rates. At low freight rates the supply of shipping services is very elastic, as there is a glut of vessels. A small increase in freight rates attracts many shipowners willing to take the existing cargo. Alternatively, an increase in the demand for shipping services is met largely by an increase in the volume to be transported at unchanged or slightly higher freight rates. But as the demand for shipping services keeps on increasing a smaller proportion of extra volume is transported, while freight rates increase at a bigger proportion. As the demand for cargo rises to the point where all ships are fully utilised, it becomes impossible to meet the extra demand. Charterers are bidding for higher freight rates to see that their cargo is transported. In the limit the same cargo is transported in aggregate but at much higher freight rates.

The non-linear supply function is thus the result of a fixed supply in the short run, as it takes approximately two years for shipyards to respond to a higher demand for vessels by the owners. At some low level of freight rates the supply curve becomes perfectly elastic, as below that level some owners do not cover the average variable cost and go bust. But as long as they cover the average variable cost, it is worthwhile remaining in business in the short run. In the long run, though, owners must cover the average total cost, which includes fixed costs and the cost of debt service, in order to remain in business. As bankruptcies rise, the total fleet in the market diminishes and the minimum freight rate goes up.

In a perfectly competitive freight market the volume to be transported and the freight rates are those that equilibrate the demand for and supply of shipping

services. In this framework, the equilibrium condition that demand must equal supply provides an equation for determining freight rates, while either the demand for shipping services or the supply of shipping services is used to determine the equilibrium cargo transported. In empirical work, the demand and supply are transformed into functions of capacity utilisation measured in millions of dead weight (dwt). In this framework, the equilibrium freight rate is a positive non-linear function of the fleet capacity utilisation and a negative linear function of bunker costs, while the demand for shipping services is used to determine the equilibrium cargo being transported.

The mechanics of this theory are illustrated in Figure 2.1. The supply of shipping services, labelled S, is plotted as a curve, which is relatively flat at low levels of demand and becomes very steep at high levels. The demand for shipping services, labelled D, is negatively sloped but very steep, implying that it is highly inelastic. At low levels of fleet utilisation equilibrium is attained at A, whereas at high levels of utilisation at C. Because of the curvature of the supply curve, an increase in demand, reflected as a parallel shift from D1 to D2, results in a new equilibrium at B (low level of fleet utilisation) or at D (high level of fleet utilisation). The increase in freight rates from A to B is small compared to that from C to D. Similarly, the equilibrium cargo transported is larger from A to B compared with that from C to D. In the limiting case of perfectly inelastic supply (a vertical segment), no extra cargo is transported; the entire increase in demand is met with higher freight rates.

The supply of shipping services is also a negative function of bunker costs. As the price of oil rises, owners are willing to reduce the supply of shipping services at the same freight rate. Higher bunker costs shift the supply curve to the left, as to transport any given cargo owners would demand higher freight rates to cover the dearer bunker costs. As a result, the equilibrium cargo is lower and the equilibrium freight rate is higher. This is illustrated in Figure 2.2. The supply curve shifts to the left from S_1 to S_2 in response to higher bunker costs. Initial equilibrium is at A and final equilibrium is at B, which implies higher freight rates and slightly smaller cargo, because the demand curve is very inelastic.

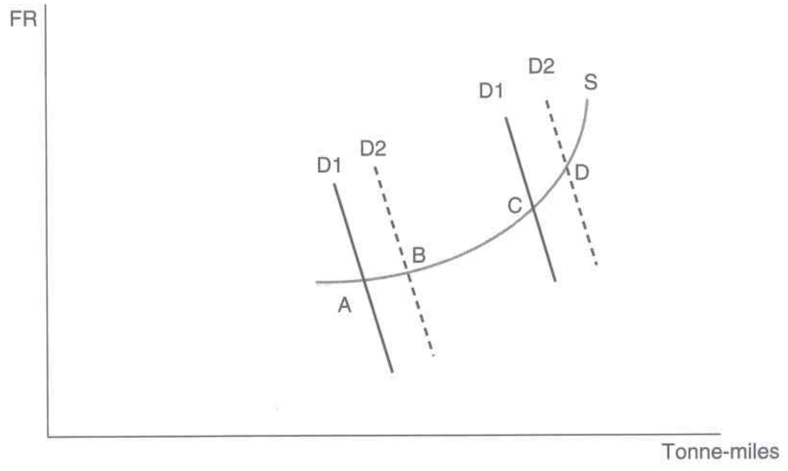

Figure 2.1 Demand and supply of shipping services

4

THE EFFICIENCY OF SHIPPING MARKETS

EXECUTIVE SUMMARY

In this chapter we explore the issue of whether freight rates and ship prices (newbuilding and secondhand) are 'efficient'. This is an issue to which academic economists in the field of maritime economics, following similar lines of research in the field of finance, have devoted a great deal of time and effort. The issue is interesting from not only a theoretical point of view, but also a practical one. For the practitioner owner a key question is whether to employ the vessels in the time charter (period market) or the spot market. If freight markets are efficient, the decision that is taken will make no difference to profitability. If markets are inefficient, profit opportunities arise. The question is under what circumstances and when to switch from one market to the other to maximise profits. Similarly, if ship prices are inefficient, an owner can devise strategies to maximise profits. When ship prices are lower than their fundamental value, defined under the hypothesis of market efficiency, excess profits can be made by buying and operating these ships. When ship prices are higher than their fundamental value, it might be profitable to charter vessels rather than buying them.

These micro issues also have macro implications. Market inefficiency gives rise to misallocation of resources. For example, if expectations of rising freight rates do not reflect economic fundamentals, buying a ship on such expectations will lead to a misallocation of resources.

With the term 'market efficiency' economists are asking whether the difference in profit between two alternative strategies (the so-called excess profit or excess return) is zero. If the excess profit is zero, the underlying market where the strategies are developed, is said to be efficient. If the excess profit is non-zero (positive or negative), the market is inefficient. This sounds very simple, but the Efficient Market Hypothesis (EMH) is posed *ex ante* (the profit opportunity at the time the decision had to be made) and not *ex post* (with hindsight). A simple example, analysed in the main text, highlights the difference. Consider the recent event of the Greek bailout by the troika in May 2010. The event marked the reversal of economic policy in the major economies from restoring growth and eliminating unemployment to the pre-crisis levels, to reining in public finances. It is the major cause of the distressed shipping markets from the spring of 2010 to the first half of 2013, as demand fell behind supply, predetermined from projections of rosy demand conditions before the crisis and in the aftermath following

the swift recovery of demand until the Greek crisis. If markets were efficient in the short run, then time charter rates would have adjusted instantly to the low level of demand from this change of policy in such a way that in terms of expected profit or loss it would have made no difference in choosing to stay in the spot or the time charter market. So, why markets may be inefficient?

Many economic agents simply do not appreciate the implications of the new information. Continuing with our example, one would not expect many owners to have appreciated the full implications of the Greek bailout for the locking in of time charter rates. Only the very well-informed owners (with good consultants) would have been able to grasp the significance of this event for shipping. Even from those owners that did, only a few might have been prepared to act upon it. The reason may well be that owners are not always prepared to take the risk. What was the *ex post* (with hindsight) profit? The average three-year time charter rate in the first half of 2010 for a 52,000 tonne Supramax was $17,300, but the average in the spot market over the same period was $18,800. An owner may have viewed the cost of moving from the spot to the time charter market as giving up $1,500 per day for an event that might have come to nothing at the end. *Ex post*, this would have been a very good decision as an owner operating in the spot market in the three-year period from the second half of 2010 to the end of the first half of 2013 would have made for the same ship only $9,600 per day (assuming no risk of default on the charterer). Therefore, *ex post* the strategy of moving to a three-year time charter contract would have resulted in a net gain of more than $7,600 per day. But the issue of market efficiency is posed not *ex post*, but *ex ante*. In other words, the issue of market efficiency is whether this profit opportunity could have been recognised at the time with information available then, so that freight rates would have adjusted to eliminate the profit opportunity. For market efficiency to hold economic agents must learn instantly all new information, absorb its implications for profit opportunity or loss and react instantly to take advantage of it. These conditions are unlikely to be valid in the real world. This would mean that markets might be inefficient in the short run (that is, in every period of time), but not necessarily in the long run.

In the long run, *ex ante* excessive profits are zero, but in the short run some economic agents may be able to exploit them. But as more and more agents learn about these profit opportunities the *ex ante* excess profitability is eliminated. In the context of our earlier example the adjustment of the time charter rates would have been gradual rather than instantaneous to the new long-run equilibrium of lower demand for shipping services. Accordingly, owners that would have acted swiftly and locked in three-year contracts in the second half of 2010 would have been better off than staying in the spot market. The condition of market efficiency in the long run implies that these profits would, in time, decline and ultimately be eliminated. Hence, in the real world there may be evidence of a weak form of market efficiency. The example also illustrates that the failure of market efficiency in the short run may be due to how risk affects shipping decisions.

Although these conclusions may seem uncontroversial and common sense, at least to the layman, they are much more difficult to prove in lab conditions. This chapter explains the different tests that have been devised to examine the validity of the Efficient Market Hypothesis in the context of freight rates and ship prices.

In describing these tests the aim is not simply to show whether shipping markets are efficient or not, but how some of these elaborate models can help to improve decision making in shipping.

This chapter begins by explaining the Efficient Market Hypothesis and rational expectations, which is an indispensible constituent component of market efficiency. Then the chapter explains the two models that support the Efficient Market Hypothesis, namely the random walk model and the martingale model. Two tests of the validity of the Efficient Market Hypothesis emerge naturally out of this literature: tests of unpredictability of excess profitability (or excess returns) and tests of informational efficiency or orthogonality of past information with future prices. In an efficient market, where economic agents are rational, actual prices should reflect the fundamental value of the asset. In freight markets this amounts to the time charter rate being equal to a weighted average of rolling spot rates for the duration of the time charter contract. For ship prices, the fundamental value is the present value of expected future profitability. In an efficient market excess returns over those implied by the fundamental value should be independent (or uncorrelated) of historical information available at time t or earlier, which implies the unpredictability of excess returns and the orthogonality of past information with future prices.

A second set of battery tests can be devised by comparing the actual price of the asset with that implied by its fundamental value. Regression tests can be conducted to test whether the two prices (actual and fundamental) are equal. If they are, this is evidence of market efficiency. If they are not, it does not necessarily follow that markets are inefficient. The reason is that expected profits require the specification of expectations-generating mechanisms. Two broad categories can be distinguished: forward-looking and backward-looking expectations. Rational expectations fall squarely in the first category, but backward-looking ones are not inconsistent with the Efficient Market Hypothesis, as the information set available at time t includes all past values of all relevant variables. From this angle, a test of market efficiency is a test of the joint hypothesis of market efficiency and the specific expectations scheme. Thus, if the joint hypothesis is rejected by the data, it is not clear which part of the joint hypothesis is responsible for it: is the market inefficient or the assumed expectations scheme is wrong?

The development of the VAR methodology (see the main text and the Statistical Appendix) has helped to shed light on this issue because this methodology ensures that all past information is used efficiently in forming expectations. From this point of view, a third battery of tests emerges. The validity of the Efficient Market Hypothesis implies some (non-linear) restrictions on the VAR. This means that a test of market efficiency is to compare the statistical fit of the restricted with the unrestricted model. Market efficiency requires that the two models should give the same fit. If the unrestricted model performs better than the restricted one, then the test indicates the rejection of market efficiency because an unjustified imposition of restrictions reduces the explanatory power of the model. Such a test may be computationally cumbersome to perform because it requires the estimation of two models: the restricted and the unrestricted one. An alternative test is to work with the unrestricted model and test whether the restrictions are met.

What are the implications of the restrictions on the coefficients of the VAR for the freight market? First, there are no excess profits to be made by owners in choosing a time charter contract over a series of rolling spot contracts that span the duration of the time charter contract. The difference is zero and hence either scheme gives the same profit. Second, no other information is needed to forecast the spread between the time charter and the spot rates. In other words, the forecast error in predicting future changes in spot rates is independent of information available at time *t* or earlier. Therefore, the orthogonality principle of market efficiency is valid. Third, the validity of the non-linear restrictions is a test of the joint hypothesis of the expectations theory of the term structure of freight rates and the VAR system of generating expectations. Therefore, the interpretation of the non-linear restrictions taken all together is that freight markets are efficient.

The VAR methodology gives rise to a third battery set of testing for market efficiency. This method is based on volatility tests in the form of variance inequality or variance bounds tests. The basic idea of volatility tests is that spot freight rates are too volatile to reflect changes in economic fundamentals, namely that the time charter rate is equal to a weighted average of expected changes in spot rates. Therefore, evidence of excess volatility is not consistent with economic fundamentals, thereby rejecting the Efficient Market Hypothesis.

The variance inequality test of market efficiency is an alternative to the regression tests explained earlier. Both are derived on the principle of informational efficiency or orthogonality of rational expectations. This means that the regression test and the variance inequality test (in theory, though not in practice) are equivalent and follow from each other.

On balance, these tests do not support the hypothesis of market efficiency for freight rates and ship prices. But, as has been mentioned, this does not mean that that these markets are not efficient in the long run. In the context of our earlier example a key question is: how long does it take for these profit opportunities to be eliminated? Would owners that locked in after the first six months following the Greek crisis have failed to capitalise on relatively robust charter rates? In other words, how long does it take for disequilibrium to be corrected before the new long-run equilibrium (no profit opportunity) is attained? These questions can be tackled using the framework of cointegration. In the Statistical Appendix to this chapter some effort is made to explain this framework intuitively and mathematically. Broadly speaking, two variables are cointegrated if they are driven together not because of a common trend but because they are moving in such a way as to restore a new long-run equilibrium. In the freight markets this means that the time charter and spot rates are moving in such a way that their spread is eliminated in the long run. This would be true if the spread between the time charter and the spot rates is equal, in the long run, to a weighted average of expected changes in spot rates over the lifetime of the time charter contract. Therefore, the efficiency of freight rates is viewed as the long-term condition of zero excess profitability of staying in the time charter market over a strategy of opting for rolling spot contracts for the duration of the time charter contract. The existence of these long-term relations (cointegrating vectors) is only a necessary condition for market efficiency. The validity of the restrictions on the coefficients of the cointegrating would provide a sufficient condition for long-term market efficiency. But in some cases, such as in ship prices, the theory does not impose numerical restrictions on the coefficients of the cointegrating vectors, other than that these coefficients should be either positive or negative. Although this drawback makes the tests based on cointegration inconclusive on the issue of market efficiency, they are extremely useful in developing models that help to improve decision making in shipping.

The evidence of excess profitability in the short run is not necessarily against market efficiency if risk is allowed to affect decisions in shipping. If the excess profitability is viewed as compensation for risk taking, the hypothesis of market efficiency cannot be rejected both in theory and in empirical tests. This is the literature on time-varying risk premia in the theory of the term structure of freight rates and vessel prices. The contribution of this literature, therefore, is not in testing for market efficiency, but in defining and modelling risk and explaining how shipping decisions are affected by risk. The starting point is that risk can be approximated with the conditional variance of past forecast errors. This concept is intuitively appealing. Consider an owner that takes seriously the advice of an economic consultant and assume that in the recent past the *ex post* variance of the forecast errors is small (the consultant's forecasts are relatively accurate). In this case the owner's perception of risk is low (or is reduced), as the owner can rely more on the consultant's advice to formulate a strategy. Moreover, risk depends on the degree of the owner's risk aversion. The second point of this methodology is that high risk is rewarded with high return. The excess profitability is a function of risk, where the risk is measured by the degree of the owner's risk aversion and the variance of the forecast errors. The more risk averse is an owner and the less accurate the previous forecasts are, the greater is the required profit (the risk premium) needed to compensate the owner for moving from the time charter to the spot market. This interpretation follows from viewing the equation for excess profitability as the demand for spot contracts (a risky asset) against the demand for time charter contracts (the safe asset) in a 'mean-variance' model of asset demands with two assets. Similarly, the excess return of investing in ships over the return on money is a function of risk, where risk is measured again by the conditional variance of the forecast errors. According to the Capital Asset Pricing Model (CAPM) of asset pricing, the higher the risk is, the higher the excess return.

The principles of cointegration and time-varying risk premia help greatly in specifying models of freight rates and ship prices that enable owners to predict them. Therefore, these models help to improve shipping decision making.

1 INTRODUCTION

As we have seen in Chapter 2, freight rates have become asset prices and vessel prices (newbuilding and secondhand) are undisputedly asset prices. Thus, it is plausible to ask whether freight rates and ship prices can be predicted, as improved shipping decision making requires that both are predictable. By nature, asset prices are widely thought to be unpredictable because they discount the implications of 'news' on economic fundamentals, which are extremely volatile. But at a deeper level, the extent to which asset prices are unpredictable is due to the belief that asset markets are 'efficient'. Hence, it is important to review the Efficient

Market Hypothesis and the two widely used models that support it: the martingale and the random walk models. We then examine models of freight rates and vessel prices that are consistent with market efficiency and discuss the tests that can be conducted to test for efficiency in shipping markets. Finally, we present the empirical evidence on the efficiency of shipping markets. This chapter is organised as follows. After this short introduction, we review the Efficient Market Hypothesis and in sections 3 and 4 the martingale model and the random walk models that support it. In section 5 we analyse the tests for the efficiency of freight markets and present the empirical evidence. Section 6 deals with the same issues of ship prices, while the last section concludes.

2 THE EFFICIENT MARKET HYPOTHESIS

The *Efficient Market Hypothesis* (EMH) is the simple statement that asset prices fully reflect all available information (Fama, 1970). More precisely, a market is said to be efficient with respect to given information set, if its prices remain unaffected by revealing that information to all market participants (Malkiel, 1992). A market that would obey this property must be a perfectly competitive one in which market participants must be fully rational. These rational traders rapidly assimilate all available information and prices are accordingly adjusted. If all available information is immediately reflected into current asset prices, then only new information or 'news' on economic fundamentals, such as dividends, can cause changes in prices. This implies that prices are unforecastable because news, by definition, cannot be predicted and hence it is impossible to make economic profits by trading on the basis of that information set.

These ideas of market efficiency can be more formally presented by using the concepts of mathematical expectations and rational expectations.[1] As a first approximation, investors form rational expectations when they forecast an asset price as the mathematical expectation and they apply the principles of mathematical expectations to make such a forecast. What are these principles? Mathematical expectations have three properties: unbiasedness, orthogonality and the law of iterated expectations. Let us examine these properties separately. Under rational expectations the mathematical conditional expectation of an asset price is equal to its actual price, P_t, plus a random error, ε_t

$$E(P_{t+1}|I_t) = P_{t+1} + \varepsilon_{t+1} \quad \varepsilon_t \sim IID(0,\sigma^2) \tag{4.1}$$

where $IID(0,\sigma^2)$ denotes that the random variable ε_t is independently and identically distributed through time with mean 0 and variance σ^2 and E is the mathematical conditional expectation of P_{t+1}. The mathematical expectation of the price next period is conditional on all information available in period t, contained in the information set I_t. An alternative simpler notation of $E(P_{t+1}|I_t)$ is $E_t(P_{t+1})$. Equation (4.1) states that forecasts are on average correct. Any particular forecast may be wrong by the value of ε_t and this may be large or small, but in repeated exercises forecast errors cancel out. From a statistical point of view equation (4.1) implies that the conditional expectation is an ***unbiased*** forecast,[2] which is the first principle of mathematical expectations.

The forecast error (or innovation) is uncorrelated with all information available at time t. This is stated mathematically as

$$E(\varepsilon'_{t+1} I_t) = 0 \tag{4.2}$$

and is known as the *informational efficiency or orthogonality* property of conditional expectations. The justification of this property is simply that if they were correlated then that information could be used to improve the forecast. Since the information set I_t contains also realisations of the error term ε (that is, past forecast errors, also called innovations), then it is obvious that the orthogonality property implies also that the error term is not serially correlated. If the error term is serially correlated, then past innovations can be used to improve the forecast of the asset price. This can easily be seen if it is assumed that the error term follows a first-order autoregressive scheme

$$\varepsilon_{t+1} = \rho\varepsilon_t + u_{t+1} \tag{4.3}$$

where u is a white noise process. Rewrite equation (4.1) as

$$P_{t+1} = E_t(P_{t+1}) + \varepsilon_{t+1} \tag{4.4}$$

Lagging this equation and multiplying by ρ and then subtracting it from the original equation using also (4.3) gives:

$$P_{t+1} = \rho P_t + [E_t(P_{t+1}) - E_{t-1}(P_t)] + u_{t+1} \tag{4.5}$$

The first term on the right-hand side shows that prices are forecastable. Tomorrow's price depends on the price today. The other terms contain no information which can be used to improve the price forecast. Hence, the orthogonality property rules out serial correlation in the residuals.

Consider now using equation (4.4) to make a forecast into the distant future, say two periods ahead. This is written mathematically as

$$E_t[E_{t+1}(P_{t+2})] = E_t(P_{t+2}) \tag{4.6}$$

In deriving this result the ***rule of iterated expectations*** has been used which states that

$$E_t E_{t+1} E_{t+2} \ldots = E_t \tag{4.7}$$

The meaning of the rule of iterated expectations is that with the information available today the investor does not know how the forecast will be revised next period when more information becomes available. Hence, the forecast two periods ahead is based only on the information available today.

It is obvious that investors are rational when they treat their forecast as the mathematical expectation and they apply the principles of mathematical expectations.

However, this is not sufficient to claim that they form rational expectations in the sense of Muth (1961) because their subjective expectations may differ. Expectations are formed rationally, in the sense of Muth, when investors, in the aggregate, do not make systematic errors. This requires that their *subjective* expectations are equal to each other (homogeneous expectations) and, in turn, are equal to the conditional mathematical expectations which are based on the ***true*** economic model. Thus investors are forming rational expectations when they use all available information to form their 'best' forecast of asset prices and in processing this information they are making use not only of the ***true*** economic model but also of the three properties of mathematical conditional expectations, namely, unbiasedness, informational efficiency or orthogonality and the law of iterated expectations.

The extreme version of the EMH can now simply be stated as the requirement that all investors know the true information set I_t and that this is available at no cost to all of them. The extreme version of EMH, namely that asset prices fully reflect all available information, is obviously false since a precondition is that information and transactions costs are always zero. Even in theory, if there are costs in gathering and processing information some profits can still be earned (see Grossman and Stiglitz, 1980).

However, these profits are obviously a reward for the information and processing costs and thus cannot be abnormally high. From this perspective a market is efficient if prices reflect information to the point where the marginal benefits simply match the marginal cost and hence a normal profit is earned. In the real world, though, it is difficult to define what would constitute normal profits since costs cannot be measured precisely. Hence, the issue is not so much whether a market is efficient in absolute terms, because the answer then is that it is not efficient. The issue is whether a market is efficient in relative terms. The advantage of the EMH is that it provides a benchmark against which the ***relative*** efficiency of a market can be assessed. For example, one can compare the efficiency of the futures vs spot markets, the efficiency of an emerging market relative to the New York stock exchange, the efficiency of the freight forward agreements (FFAs) vs spot freight rates.

All empirical tests of the EMH attempt to measure whether profits can be made on trading or whether prices are forecastable for a given information set. Some tests concentrate on whether fund managers in the real world can earn excess or abnormal returns or whether such returns can be earned on a hypothetical trading rule. The advantage of the former is that it concentrates on real trading, but suffers from the drawback that the information set used by portfolio managers is not observable. A hypothetical trading rule is therefore superior but it requires an explicit definition of the information set, the normal and excess return, and the trading costs. Following Roberts (1967), Fama (1970) defined three different information tests and hence three forms of efficiency.

Weak-form Efficiency: The information set includes only the history of prices or returns.

Semi-strong Efficiency: The information set includes all publicly available information.

Strong-form Efficiency: The information set includes in addition private information available to just few participants (insiders' information).

Fama (1991) introduced a slightly different taxonomy of efficiency with the hindsight of twenty years of research conducted between his first and second review articles on market efficiency. Instead of weak-form tests he proposed *tests for return predictability*. These include not only the history of prices or returns but also other variables like dividend yields, price–earning ratios and interest rates. These tests include not only time series tests but also cross-sectional predictability of returns, that is, tests of asset-pricing models and the anomalies (like the size effect) as well as seasonal patterns (like the January effect). In this category are also included tests of excess volatility. Instead of semi-strong efficiency he proposed *event studies* and instead of strong efficiency *tests for private information*.

The second element that needs to be specified for an empirical test of the EMH is a model of 'normal' returns. For a long time the standard assumption was that the normal return is constant over time. However, in recent years the emphasis has been on equilibrium models with time-varying expected returns. Finally, abnormal or excess returns must be defined. These are easily defined as the difference between the actual return on a security and its normal return.

From the above framework it can easily be seen that the issue of whether a market is efficient in an absolute sense is impossible to settle. Laboratory conditions for testing whether a hypothetical trading rule can lead to excess profits are difficult to set. If a trading rule fails to generate excess returns this does not mean that there is no trading rule which can generate excess return. Even worse, assume that someone discovers such a rule and makes excess profits. Would she still be able to make excess returns with the same rule after a while? The information set would soon include this rule and as more people are using it any excess profits will be eliminated. Thus a better test may involve the predictability of asset prices. But even this route is not promising. Assume that prices are forecastable, but the explained variance is small (less than 10 per cent), as is indeed the case with short horizon returns. The question immediately arises as to whether this is significant and that in turn depends on whether profits can be made. Hence, we are back to square one!

But the worst problem in any form of test of market efficiency is the *joint hypothesis* problem. Any test of market efficiency is conditional on the equilibrium model of normal returns. If efficiency is rejected, this could be due to the market being truly inefficient or to the incorrect model being used. This *joint hypothesis* problem means that the market efficiency as such can never be rejected. However, this may not be important, if one accepts that absolute efficiency is an ideal world and that the advantage of the EMH lies in specifying a benchmark against which relative efficiency can be assessed. We now turn to models, which can explain the time series properties of asset prices from the point of view of the EMH. Two such models are explored, the martingale and random walk models. We then examine the issue of whether the EMH implies models that reflect economic fundamentals. We show that the EMH is compatible with models that reflect economic fundamentals. Then we proceed to analyse models that contradict the EMH based on some form of irrationality, like fads. Finally, we examine the empirical evidence of which side may be right.

the methodology developed by Campbell and Shiller (1988) in financial markets. Campbell and Shiller apply equation (4.63) to a stock price, where dividends appear instead of profits and where there is no resale value in period n so that the present value is computed for an infinite horizon. The authors compute a linearized version of (4.63)) around the geometric mean of P, and Π, denoted by $\overline{P}$ and, $\overline{\Pi}$ respectively, using a first order Taylor series expansion:[11]

$$p_t^* = \sum_{i=0}^{n-1} \rho^i (1-\rho) E_t \pi_{t+1+i} - \sum_{i=0}^{n-1} \rho^i E_t r_{t+1+i} + \rho^n E_t ps_{t+n} + c \qquad (4.64)$$

where lower-case letters denote natural logarithms; $E_t\, ps_{t+1} = \ln(E_t\, PS_{t+1})$ and similarly for Π; $E_t\, r_{t+1} = \ln(1 + E_t R_{t+1})$; $\rho = \overline{P}/(\overline{P} + \overline{\Pi})$; and c is a constant.

Note that (4.64) consists of three terms: the present value of expected profits; the present value of the expected interest rates; and the present value of the capital gains. Each present value though is computed with a constant discount rate, ρ, rather than time-varying discount rates. Therefore, the linearisation of (4.63) into (4.64) enables the separation of the impact of expected profits (and capital gains) from the time varying interest rates. This is very useful in evaluating the importance of risk premia as an independent factor of the theoretical price.

In analogy with the spread in the term structure Campbell and Shiller suggest subtracting π from both sides of (4.64). This transformation is necessary to make the underlying series stationary. As freight rates are non-stationary,[12] but can rendered stationary by taking their first difference, similarly vessel prices (new and secondhand) and profits are non-stationary, but can become stationary by taking their first difference. By taking the difference between two such non-stationary variables, the resulting variable is stationary.[13] The transformation enables (4.64) to be written as:

$$S_t^* = \sum_{i=0}^{n-1} \rho^i E_t e\pi_{t+1+i} + \rho^n E_t eps_{t+n} + d \qquad (4.65)$$

where $S_t^* = p_t - \pi_t$, $e\pi_t = \Delta\pi_t - r_t$, $eps = ps_t - \pi_t$. In (5.54) all variables are stationary enabling valid statistical inferences.

Thus the theoretical spread of the (log) price from the (log) profit is equal to the present value of the spread of profits from the interest rate and the present value of the spread of the capital gains.

6.2 TESTS OF MARKET EFFICIENCY

Once the expectations of the variables are substituted out, the theoretical spread can be compared with the actual spread. If vessel prices are efficient, the theoretical spread should be equal with the actual spread. A test like (4.63) can be conducted to examine the EMH of vessel prices.

Similarly with the term structure of freight rates, three different expectations schemes can be considered: rational expectations, autoregressive and VAR.

The VAR approach opens up a battery of tests similar in nature to the efficiency of freight markets. The VAR model should explain the three state variables in (4.65):

$$S_t = k_1 + a(L)S_t + \beta(L)e\pi_t + c(L)eps_t + \varepsilon_{1t} \qquad (4.66a)$$

$$e\pi_t = k_2 + d(L)S_t + e(L)e\pi_t + f(L)eps_t + \varepsilon_{2t} \qquad (4.66b)$$

$$eps_t = k_3 + g(L)S_t + h(L)e\pi_t + m(L)eps_t + \varepsilon_{3t} \qquad (4.66c)$$

where each coefficient is a polynomial in the lag operator of order p. For example:

$$a(L) = a_1 L + a_2 L^2 + \cdots + a_p L^p.$$

The VAR system can be expressed in the matrix notation of (4.19) and forecasts can be computed in the form of (4.30). These can be entered into (4.65) to compute the theoretical spread. If the EMH of vessel prices is valid, the non-linear restrictions (4.32) must hold. Thus, one test of the EMH is to compare the statistical fit of the constrained with the unconstrained VAR through log-likelihood statistic, as in equations (4.37)–(4.39). The second approach is to estimate the unrestricted VAR and test whether the restrictions (4.32) hold. The null hypothesis of efficient prices is that the restrictions are valid, which can be tested through a Wald statistic. As an alternative to a Wald-test block exogeneity tests can be conducted. Finally, the null hypothesis of efficient pricing can be tested through a variance ratio test of the form of (4.44).

6.3 WEAK FORM OF MARKET EFFICIENCY – COINTEGRATING TESTS

Efficient ship pricing may not hold in every period of time, but it may be true in the long run. In other words, there may be a weak form of market efficiency. As in the freight markets, a test of the weak form of efficient pricing is through cointegration tests. A possible cointegrating relation is between prices (newbuilding, secondhand and scrap) and profits.[14] The theory imposes no numerical restrictions on the coefficients of the cointegrating relation, other than a positive sign on the coefficient of profits. In other words, an increase in profitability should increase vessel prices. The cointegration test within a VAR framework takes the form:

$$\Delta p_t = \sum_{i=1}^{p} a_i \Delta p_{t-i} + \sum_{i=1}^{p} \beta_i \Delta\pi_{t-i} + \gamma_1 [p_{t-1} + \phi_0 + \phi_1 \pi_{t-1}] \qquad (4.67a)$$

$$\Delta\pi_t = \sum_{i=1}^{p} c_i \Delta p_{t-i} + \sum_{i=1}^{p} d_i \Delta\pi_{t-i} + \gamma_{21} [p_{t-1} + \phi_0 + \phi_1 \pi_{t-1}] \qquad (4.67b)$$

The VAR is specified with respect to the two possible variables that might form a cointegrating relation, namely the vessel price p_t and profit π_t. There is an equation for each one of the state variables. The order of the VAR is p; there are p-lags in each of the state variables. The order of the VAR is chosen with the Akaike Information Criterion or the Schwarz Information Criterion (see the Statistical Appendix). The existence of a cointegrating relation is tested through the Johansen procedure (see the Statistical Appendix).

6.4 TIME-VARYING RISK PREMIA

According to the EMH the profits from shipping should not exceed those from alternative investments. Efficient ship pricing implies that the ship price reflects the profits from operation and the capital gains, which is the essence of (4.63). But the variables in (4.63) are unobserved. An alternative way to postulating generation expectations mechanisms is to consider one-period holding yields. Let $H\pi_t$ denote the one-period holding shipping yield, which is defined by:

$$H\pi_{t+1} = \ln(P_{t+1} + \Pi_{t+1}) - \ln P_t \tag{4.68}$$

Recall that $r_{t+1} = \ln(1 + R_{t+1})$ denotes the return on money. Efficient ship pricing implies that the one-period holding shipping yield should be equal to the return on money for the same period:

$$H\pi_{t+1} = r_{t+1} \tag{4.69}$$

An alternative way of expressing efficient ship pricing is that the excess return from shipping over the return on money should be zero. Hence a test of efficient pricing is to examine whether the excess return is independent of the information available at time t. This gives rise to two complementary tests: testing whether the excess return is correlated with previous information and testing the predictability of the excess return. These procedures are testing the joint hypothesis of efficient pricing and rational expectations. But failure of the joint hypothesis may be due, as in the case of freight markets, to time varying risk premia. As a result by accounting for risk premia the excess return is zero. But as has been argued in the case of freight rates, one can always specify a function that renders the excess return equal to zero. Accordingly, the contribution of this literature lies in show showing how risk affects ship prices. As with freight rates one can specify a time invariant risk premium and examine whether the hypothesis is consistent with empirical evidence. If it is not, then a function explaining risk has to be postulated. The capital asset pricing model (CAPM) offers a theory of relating return with risk. According to CAPM, there is a positive relationship between risk and return: as risk rises, the excess return required by an owner to buy the ship is increased. Let XR denote the excess return of shipping over the return on money and define the *mean function* in the ARCH/GARCH-M class of models as an AR(m) process augmented by a time varying risk premium that depends on the conditional variance (or standard deviation) of forecast errors in predicting the excess return, σ^2.

$$XR_t = a_0 + \sum_{i=1}^{m} \beta_i XR_{t-i} + \phi\sigma_t^2 + \eta_t \tag{4.70}$$

In the mean function the second term involving the sum represents the AR(m) process; the coefficient φ measures the importance of risk, which according to CAPM should be positive; and η are the residuals, which should be a white noise process.

The forecast errors upon which risk is measured are obtained from an ARMA(p, q) process:

$$XR_t = c_0 + \sum_{i=1}^{p} b_i XR_{t-i} + \sum_{i=1}^{q} d_i \varepsilon_{t-i} + \varepsilon_t, \quad \varepsilon_t \approx IID(0, \sigma_t^2) \tag{4.71}$$

The first sum captures the AR(p) process and the second sum the MA(q) process. The ARMA specification is necessary if there is autocorrelation in the residuals.

The conditional variance of the forecast errors, which measure the time varying risk, can be modelled, for simplicity, as a GARCH process of order (1, 1):

$$\sigma_{t+1}^2 = f_0 + f_1\varepsilon_t^2 + f_2\sigma_t^2 \tag{4.72}$$

6.5 THE EMPIRICAL EVIDENCE

Early empirical work (for example, Beenstock, 1985) simply used the EMH to model vessel prices. Strandenes (1984, 1986) investigates the price formation in the dry bulk and tanker markets using the present value model. The empirical evidence shows that long term profitability is more important than current profits in explaining ship prices, which she interprets as support for the semi-rational expectations. Vergottis (1988) is one of the authors to test the efficiency of ship pricing by using the principles of rational expectations of unbiasedness and informational efficiency or orthogonality. Hale and Vanags (1992) test for market efficiency through the use of block exogeneity tests (Granger causality). But the procedure, as noted above, only tests for the weak form of market efficiency. These authors examine also the existence of cointegration relations between various vessel prices using the Engle–Granger two-stage procedure, but with mixed results. Glen (1997) employs the more powerful Johansen approach to test for cointegration in a multivariate setting. Veenstra (1999b) examines the existence of cointegration between secondhand prices, a time charter rate, newbuilding prices and scrap prices.

Kavussanos and Alizadeh (2002b) provide a thorough and exhaustive examination of efficient pricing using the elaborate testing procedures outlined above. Their empirical findings reject the EMH for newbuilding and secondhand prices. The authors attribute this inefficiency to time varying risk premia, showing that there is a positive relationship between risk and return in shipping in line with CAPM.

A common theme in all the abovementioned studies is that there is no theoretical basis for the cointegrating vector. Rather, the cointegrating vector is an empirical result based on intuition. For example, in the exemplary study of Kavussanos and Alizadeh (2002b) the cointegrating vector is based on the notion that prices are cointegrated with profits, where the profit is modelled as the spread between a time charter equivalent rate and operating costs. The latter are modelled as an exponential growth rate regression. A notable exception to this rule is the study by Tsolakis, Cridland and Haralambides (2003), where the cointegrating vector is the reduced form of a demand–supply framework in the secondhand market. The cointegrating vector to be tested empirically includes secondhand prices, a time charter rate, newbuilding prices, Libor and the orderbook to fleet ratio. Their empirical results suggest that newbuilding prices and time charter rates form a cointegrating vector with secondhand prices. Libor is only significant in the dry bulk market but not in the tanker market, a strange result. Haralambides et al. (2004) extend the above results in the newbuilding market. The possible cointegrating vector for newbuilding prices includes, in addition, cost variables, such as steel prices roll-plates in Japan. These theoretical approaches to the determinants of the cointegrating vector have their foundations on structural models of newbuilding and secondhand prices, such as Koopmans (1939), Hawdon (1978) and Jin (1993), among others. This is a trend in the right direction in establishing the importance of structural models in maritime economics. A notable study in this new trend is that of Jiang and Lauridsen (2012), which analyses the price formation of Chinese dry bulk carriers. The empirical evidence suggests that the time charter rate has the most significant positive impact on new prices followed by the cost of shipbuilding, the profit margin and the shipyard capacity utilisation.

7 CONCLUSIONS

This chapter has explained the EMH and the statistical tests of the efficiency of freight rates and ship prices. The methodology has been borrowed from the financial markets, but has been adapted to shipping mainly by correcting for the finite life of ships and of time charter contracts. The empirical evidence suggests that both freight rates and vessel prices are, on balance, inefficient. The literature on time-varying risk premia has shed light on the nature of this inefficiency and has highlighted models that explain risk and how it affects shipping decisions. The empirical evidence also suggests the advantage of structural models over purely statistical models of shipping as a guideline to the specification of the cointegrating vectors.

STATISTICAL APPENDIX

The literature on the efficiency of shipping markets involves many statistical and econometric concepts. This Appendix provides a summary of concepts and statistical tests which are used in this and other chapters. The objective of this Appendix is simply to make the book self-contained. Therefore the Appendix is not meant to replace econometrics textbooks, such as Hamilton (1994) and Greene (1997), which cover all the material covered here.

STATIONARY AND NON-STATIONARY UNIVARIATE TIME SERIES

All shipping variables, such as freight rates, can be thought of as statistical time series, as they assume a value for each period, be that a day, week, month, quarter or a year.

A very common statistical model of shipping and economic time series is the univariate (single variable) autoregressive model of order 1, denoted by AR(1). Thus

$$y_t = a + \beta \cdot y_{t-1} + \varepsilon_t \qquad |\beta| < 1 \qquad \text{(A.1)}$$

The concept of autoregression means that the variable y depends on past values of itself through an explicit linear function, such as (A.1). The function is linear because the terms are added and α and β are constant coefficients. The order of the autoregression is characterised by the maximum lag of y. Thus, an AR(2) process is written as

$$y_t = a + \beta_1 \cdot y_{t-1} + \beta_2 \cdot y_{t-2} + \varepsilon_t \qquad |\beta_i| < 1 \quad \text{for } i = 1, 2. \qquad \text{(A.2)}$$

The variable y depends linearly, through the constant coefficients α, β_1 and β_2, on two lagged values of itself. In statistical and econometric analysis the aim is to identify the exact functional form (A.1), (A.2) or another more complicated but unknown form) and the precise value of their coefficients.

In each autoregressive scheme a disturbance (error or residual) term, ε, is added. This means that the exact value of y in every period t is determined by a deterministic component captured by $\alpha + \beta \cdot y_{t-1}$ in (A.1) and a stochastic (or random) component, ε. The disturbance term assumes random values in each period and therefore y differs in each period because of the stochastic nature of the disturbance term. The values of ε depend on events such as strikes, weather conditions, political events or the influence of other important variables which were unintentionally omitted from the specification of (A.1) or (A.2). Because of the stochastic nature of ε, the variable y is also stochastic following an autoregressive stochastic process.

The Identification of the 'true' (or population) values of the coefficients through statistical or econometric analysis depends on the properties of the disturbance term. From a sample of data on y the estimates of the coefficients (α and β) would approach the 'true' (population) values as the sample size tends to infinity if the disturbance term is purely random. In statistical analysis the properties of the estimates of the coefficients that must be satisfied so that they approximate the 'true' values are best-linear-unbiased estimators (BLUE). The 'best' property means that the estimator has a minimum variance, while the unbiasedness property means that the sample mean of the estimator is equal to the population value. If the random variable ε satisfies some properties then the estimate of the coefficients are BLUE. These properties of the ε can be summarised as follows.

$$E\varepsilon_t = 0 \qquad \text{for all } t \qquad \text{(A.3a)}$$

$$\operatorname{var}(\varepsilon_t) = E(\varepsilon_t^2) = \sigma^2 \qquad \text{for all } t \qquad \text{(A.3b)}$$

$$\operatorname{cov}(\varepsilon_t, \varepsilon_{t-j}) = 0 \qquad \text{for all } j \neq 0 \text{ and all } t \qquad \text{(A.3c)}$$

Although in the real world we observe just one value of ε in each period, the stochastic nature of ε implies that there is an infinite number of observations that could have been observed forming an entire distribution. For each distribution in every period t the first property (A.3a) implies that the mean of the distribution is zero. The second property[15] (A.3b) implies that the variance of each distribution assumes the same constant value, σ^2. This property is usually referred to as *homoscedasticity*. When the variance is not constant through time the disturbance term is *heteroscedastic*. The third property (A.3c) is that the covariance[16] of any pair of ε, such as ε_t and ε_{t-1} is zero. For the covariance to be zero the correlation of ε with any past or future value of ε must be zero. These three properties of the disturbance term are usually referred to as identically and independently distributed and they are denoted as

$$\varepsilon \approx IID(0,\sigma^2) \tag{A.4a}$$

When the disturbance term satisfies these properties it is called a *white noise process*. If the distribution of the disturbance term is also normally distributed, then it is called a *Gaussian white noise process*, denoted as

$$\varepsilon \approx NIID(0,\sigma^2) \tag{A.4b}$$

When the error term follows a Gaussian white noise process the estimates of the coefficients of (A.1) or (A.2) obtained from a relatively large sample ($T > 30$) are BLUE. When the error term follows a white noise process, (A.4b), the estimates are asymptotically BLUE (i.e. they become BLUE as the sample size tends to infinity).

What is the meaning of an autoregressive process? The variable y oscillates in a random way around the population mean. The population mean can be calculated by assigning to ε its mean value of zero and assuming that y converges to its mean value (that is, when $y_t = y_{t-1} = y_{t-2}$). For convergence[17] the absolute value of β should be less than 1. Hence, the population (or unconditional) mean of y, Ey, is

$$Ey_t = \mu = \frac{a}{1-\beta} \tag{A.5}$$

It is clear from (A.5) that the absolute value of $\beta < 1$. If $\beta = 1$, then y does not converge to its mean value but diverges to infinity.

The population (or unconditional) variance of y can also be calculated as follows. First compute for (A.1) the deviation from the mean, which enters into the definition of the variance.

$$(y_t - \mu) = a + \beta \cdot y_{t-1} + \varepsilon_t - \mu = \beta \cdot (y_{t-1} - \mu) + \varepsilon_t$$

This relationship is obtained by substituting for $a = \mu\,(1-\beta)$ from the definition of the mean (A.5). Substituting this value into the definition of the variance of y we have

$$\begin{aligned}\mathrm{var}(y_t) = E(y_t - \mu)^2 = \beta^2 \cdot E(y_{t-1} - \mu)^2 + \mathrm{var}(\varepsilon_t)\\ + 2\mathrm{cov}(y_{t-1} - \mu, \varepsilon_t) = \sigma^2/(1-\beta^2)\end{aligned} \tag{A.6a}$$

In deriving (A.6) we have made use of (i) the $\mathrm{var}(y_t) = \mathrm{var}(y_{t-1})$ and (ii) the covariance of $(y_{t-1} - \mu)$ and ε_t is zero and (iii) the (absolute) value of $\beta < 1$.[18]

We can now define a *stationary* time series y, given by any stochastic process such as (A.1) or (A.2), as a series whose mean, variance and auto-covariance converge to a finite value.[19] The convergence means that for a series y the mean, the variance and the auto-covariance (the latter for a given lag length) are independent of time and are finite (that is, they do not tend to infinity). For a given lag length j, the auto-covariance of y_t and y_{t-j} is constant. The auto-covariance changes with the lag j. A series y is a stationary if the absolute value of *all* $\beta < 1$.

More formally a series is 'covariance' stationary if:

$$Ey_t = \mu, \quad \mathrm{var}(y_t) = \sigma_y^2, \quad \mathrm{cov}(y_t, y_{t-j}) = \gamma_j \tag{A.8}$$

In plain English a stationary variable is one that is trendless. It is not upward sloping, nor downward sloping against time. The process is mean reverting. A shock causes a deviation from the mean (or the equilibrium value) for a while, but the deviation peters out in the long run. In a graph the series frequently crosses the mean. The variability of the series around the mean is, on average, constant.

As the condition for stationarity is that *all* $\beta < 1$ (in absolute value), it follows that a series is *non-stationary* if at least one β is equal to or greater than 1. There are two widely used non-stationary models: a *random walk with drift* and a *random walk without drift*. Considering (A.1) and setting $\beta = 1$, we obtain a random walk with drift.

$$\Delta y_t = y_t - y_{t-1} = a + \varepsilon_t \tag{A.9}$$

The coefficient a is the drift. The interpretation of (A.9) is appealing if y is the natural logarithm of Y and therefore Δy is the rate of growth of Y. According to (A.9) the rate of growth is equal to a constant plus a white noise process. Therefore Y is non-stationary as its mean is increasing through time. A time varying mean is meaningless as a measure of central location of the distribution of Y.

The statistical properties of non-stationary variables are different from those that are stationary. To illustrate some of these differences it is convenient to define the lag operator L. For any variable y, $Ly_t = y_{t-1}$; $L^2y_t = y_{t-2}$; and in general $L^k y_t = y_{t-k}$. Using the lag operator, a random walk without drift (that is, $a = 0$ in (A.9)) can be written as

$$y_t = (1-L)^{-1}\varepsilon_t = \varepsilon_t + \varepsilon_{t+1} + \cdots \tag{A.10}$$

The (unconditional) mean and (unconditional) variance of (A.10) are:

$$Ey_t = 0, \quad \mathrm{var}(y_t) = E\left[\sum_{i=0}^{n-1}\varepsilon_{t+i}^2 + \sum_{t\neq s}\varepsilon_t\varepsilon_s\right] = n\sigma^2 \tag{A.11}$$

Therefore, the mean of a random walk without drift is zero, but the variance tends to infinity as n increases (that is, it is not independent of time). The difference between a random walk with and without drift is that with drift the mean is time varying, whereas without drift the variance is time varying.

Unlike the unconditional mean, the conditional mean uses the information in the time series to predict future values of y. For a random walk with and without drift the forecast of y 'm' periods ahead is

$$\text{With drift:} \quad E_t y_{t+m} = y_t + a \cdot m \qquad \text{Without drift:} \quad E_t y_{t+m} = y_t \tag{A.12}$$

Another important difference between stationary and non-stationary variables is that the former may give rise to a *deterministic trend*, while the latter to a *stochastic trend*. To illustrate this difference consider again the random walk model with drift, equation (A.9), and assume that the initial value of y is y_0. Then (A.9) can be rewritten as:

$$y_t = y_0 + a \cdot t + \sum_{i=1}^{t} \varepsilon_i \tag{A.13}$$

It can be seen from (A.13) that y_t does not return to the deterministic trend, defined by $(y_0 + a t)$, because of the accumulation of past random error terms. The variable y follows a 'stochastic-trend' because y will drift up or down depending on the sign of a. Note that the first difference of y_t (Δy_t) is stationary. For this reason y_t is referred to as 'difference-stationary'.

In contrast consider the following model:

$$x_t = \delta + a \cdot t + \varepsilon_t \tag{A.14}$$

The variable x moves around a deterministic trend $(\delta + a\, t)$ by the disturbance term ε, which, by assumption, is stationary. Accordingly, x is said to be 'trend-stationary' because although it follows a trend the deviations from the trend are stationary. A comparison of (A.13) and (A.14) shows that both variables follow a linear trend, but the disturbance term in (A.13) is non-stationary, whereas the disturbance term in (A.14) is stationary. The variable y follows a stochastic trend and is difference-stationary, while the variable x follows a deterministic trend and is trend-stationary. The difference between difference-stationary and trend-stationary variables complicates the testing of unit roots, as we shall see later.

The different statistical properties between stationary and non-stationary variables have huge implications for statistical estimation and statistical inference. Estimation of non-stationary variables may give rise to 'spurious'[20] regression results. This means that the estimated regression gives the impression of good 'fit' (high R^2 and adjusted R^2) and statistically significant coefficients (high t-statistics), when there is no 'true' relationship between the variables. This would be the case if the variables for which a model is built are related through a common trend or a third variable that is omitted from the model. There are also huge problems with statistical inference, as the hypothesis underlying testing procedures (such as t-statistics χ^2- and F-statistics) assume that the variables are stationary.

STATIONARITY (UNIT ROOT) TESTS

Luckily, these problems can be resolved because any non-stationary variable can be transformed into a stationary one by differencing it a number of times. If a

variable y becomes stationary when it is differenced once (i.e. Δy is stationary), then it is said that it is integrated of order 1, denoted by $I(1)$. If $\Delta^2 y$ is stationary, then y is integrated of order two, $I(2)$. In general, if $\Delta^k y$ is stationary then y is integrated of order k, $I(k)$.

Granger (1966) shows that most economic variables follow a stochastic process of the form

$$\Delta y_t = a + \varepsilon_t + \lambda \varepsilon_{t+1} \tag{A.15}$$

If y is again the natural logarithm of a variable Y, then (A.15) implies that Y grows at the constant rate a plus a moving average error term. This implies that all economic and shipping variables need to be differenced once or twice to become stationary, namely that they are $I(1)$ or $I(2)$.

A simple approach in detecting whether a variable is stationary or non-stationary, which is frequently referred to in the literature of the efficiency of shipping markets, is the *autocorrelation function* and the *correlogram*. The autocorrelation between y_t and y_{t-j}, denoted by $\rho(j)$, is defined as

$$\rho(j) = \frac{\operatorname{cov}(y_t, y_{t-j})}{\operatorname{var}(y_t)} \tag{A.16}$$

When $j = 0$, $\rho(0) = 1$. For a stationary variable $\rho(j)$ tends to zero as j increases, whereas for a non-stationary variable it remains significantly above zero. The correlogram simply plots the autocorrelation coefficient against j and provides a visual guidance to the stationarity or not of a variable. A quick drop to zero indicates a stationary variable, whereas a flat line above zero suggests a non-stationary variable.

But there are also formal tests of stationarity and we discuss the Dickey–Fuller (DF), the augmented Dickey–Fuller (ADF) tests and the Phillips–Perron test, as they are extensively used in the literature of the efficiency of shipping markets. Consider again the AR(1) process described by equation (A.1) and repeated here for convenience

$$y_t = a + \beta \cdot y_{t-1} + \varepsilon_t \tag{A.1}$$

We have established that if $\beta < 1$ (in absolute terms), y_t is stationary, $I(0)$, provided that ε_t is stationary, $I(0)$, as well.

The random walk model with drift is obtained by setting $\beta = 1$ in (A.1). Thus

$$\Delta y_t = a + \varepsilon_t \tag{A.17}$$

The right hand side of (A.17) is stationary, provided the error term is stationary. Hence, for $\beta = 1$, Δy_t is stationary and therefore y_t is $I(1)$. This reasoning suggests the following test of stationarity. Subtract y_{t-1} from both sides of (A.1)

$$\Delta y_t = a + \phi \cdot y_{t-1} + \varepsilon_t \qquad \phi = \beta - 1 \tag{A.18}$$