CAND.MERC. APPLIED ECONOMICS & FINANCE

MASTER'S THESIS

Quantifying Freight Rate Risk using Stochastic Modelling

Author:
Sindre SJØBAKK

Supervisor:
Prof. Bent Jesper CHRISTENSEN

Number of pages: 80
Number of characters: 160 509

May 11, 2016

Essentially, all models are wrong,
but some are useful.

– George E. P. Box

ABSTRACT

This thesis is devoted to management of freight rate risk in the dry bulk sector of shipping. The freight rate risk is modelled in a multivariate setting including Panamax, Supramax and Handysize vessels. Freight rates are proposed to follow a geometric Brownian motion including stochastic volatility as modelled by Heston. In order to maintain interdependency between vessel classes, a *Student's t* copula is used to couple the Wiener processes of the freight rate model. The freight rate model captures the fat tails and volatility clustering of historical freight rate data as well as the historical correlation between the freight rate indices.

The freight rate risk is quantified by estimating the potential Cash Flow-at-Risk and Expected Shortfall. The cash flow effects of operating in the exposed spot rate market are compared to hedging by operations in the time-charter market. Three different lengths of time-charter contracts are used in the comparison; three months, six months, and one year. Simulations show that this type of hedging reduces the potential cash flow losses by 6%, 19% and 37% for these contracts, respectively.

Keywords. *Risk Management, Shipping, Cash Flow-at-Risk, Expected Shortfall, Stochastic Modelling of Freight Rates, Time-Charter.*

CONTENTS

List of Figures

List of Tables

ABBREVIATIONS

ACF	Autocorrelation function
ADF	Augmented Dickey-Fuller test
AR(1)	Autoregressive model of order one
BHSI TCA	Baltic Handysize Index Time-Charter Average
BPI TCA	Baltic Panamax Index Time-Charter Average
BSI TCA	Baltic Supramax Index Time-Charter Average
$CFaR$	Cash Flow-at-Risk
DoF	Degrees of freedom
dwt	Deadweight tonnage
ES	Expected Shortfall
EWMA	Exponentially weighted moving average
GBM	Geometric Brownian motion
iid	Independent and identically distributed (variable)
mt	Metric tonne
O-U	Ornstein-Uhlenbeck process
PACF	Partial autocorrelation function
TC	Time-charter contract
TCE	Time-charter equivalent
VaR	Value-at-Risk

ACKNOWLEDGMENTS

Writing this thesis has been a project for the past half year and during this time I have received valuable help and inspiration from several people. First and foremost, I would like to thank my supervisor Prof. Bent Jesper Christensen for continuous help throughout the project. Furthermore, I am grateful for the help I got from Steen Koekebakker a Professor at Universitetet i Agder. Even though not being a part of this project, he found time to help me with the freight rate modelling.

Getting access to freight rate data is not easy for an individual researcher, but thanks to Kristoffer Røstad and Bernhard Baardson at the research department of Fearnleys Shipping, I got enough data to complete the thesis.

Moreover, I would like to thank Egil Husby the Chief Risk Officer and Andrew Redfern a Business Analyst at Bulk Invest, for giving me access to company data and answering many company specific questions.

Finally, I am thankful for all valuable discussions and support received from dear friends and family during the research process.

1 INTRODUCTION

The demand for shipping is closely linked to the world economy and as such, shipping serves as a key indicator on its growth. Few will argue that shipping is not a risky industry. Over the last decade, we have seen historical highs and lows in shipping freight rates, in line with the financial turbulence of the world economy in this period. Historically, freight rates have always fluctuated, giving the rise of mega wealthy shipowners like Onassis and Fredriksen, and ruined many more. Fluctuations are the nature of freight rates, and as far as shipowners are concerned these cycles are like the dealer in a poker game, dangling the prospects of riches on the turn of each card. This keeps them struggling through dismal recessions and upping the stakes as the cash rolls in during booms. Market agents in shipping are players in the world's biggest poker game, in which the chips are valued in tens of millions of dollars, betting on ships which may or may not be needed (Stopford, 2009, p. 94). Just as in a poker game, succeeding in the shipping industry depends on a blend of skill, luck and psychology.

This thesis is devoted to managing the substantial freight rate risk involved in shipping.[1] Any poker player can win a hand on pure luck, but good players win frequently by using risk managing tools such as calculating the probability of having the best hand or reading the opponent's play. One way of reducing the risk in shipping is to model the future freight rates and estimate the probability of incurring losses greater than manageable. Quantitative risk measures used in this thesis are Cash Flow-at-Risk and Expected Shortfall, which are quantile-based (predicted) fre-

[1] It should be noted that there are several types of risks in shipping, but this thesis focuses only on the freight rate risk, which is a type of price risk (Alizadeh and Nomikos, 2009).

quency models.

In order to model freight rates, it is common to view them as stochastic processes. The word *stochastic* is derived from the Greek "stochazestai", the art of guessing, or "stochastikos" meaning skilled at aiming ("stochos" being a target) (McNeil et al., 2015, p. 5). By using stochastic models for risk management, this thesis will hopefully emphasize the skill aspect rather than the guesswork, of risk management in shipping.

Originally, the aim of this thesis was to examine how freight rate risk could be managed in the Norwegian shipping company Western Bulk AS. However, during the research process, Western Bulk AS sold one of its two divisions and the remaining one (which changed name to Bulk Invest ASA) went bankrupt. Consequently, the aim of the thesis turned to a more general perspective on how shipping companies can manage freight rate risk. Bulk Invest ASA is then used as an illustrative example of how the risk can be modelled and measured, as well as an excellent example of how risky and unpredictable this industry is.

1.1 Problem and Objective

In shipping, there are numerous types of risk. Broadly speaking, shipping companies face business risks which can be categorized into price risk, liquidity risk, credit risk and pure risk (Alizadeh and Nomikos, 2009).[2] Freight rate risk is a kind of price risk and refers to the variability in the earnings of a shipping company due to changes in freight rates. Because volatility in the freight market has a direct impact on the profitability of the company, this risk is perhaps the most important of all risks encountered by a shipping company (Alizadeh and Nomikos, 2009, p. 3). In addition, the freight rate risk has shown in the past to be extensive. Due to this, the problem statement of the thesis is the following.

Problem statement. *How can shipping companies in the dry bulk sector manage their freight rate risk?*

This is a broad and comprehensive statement and as such, certain research questions are needed to limit the scope of the thesis.

[2] Along with business risk, companies also face legal and political risk, strategic and operational risk, and environmental risk. Pure risk is the risk of accidents and the like.

Research questions.

1. What are the distributional properties of the return on freight rates?
2. How can freight rate risk be modelled?
3. How to maintain interdependency between several freight rate indices?
4. How sensitive are the freight rate model variables to changes?
5. Which are the most appropriate and effective quantitative risk measures available?
6. What is the cash flow effect of actively managing the freight rate risk as opposed to operating purely in the exposed spot freight rate market?

In the next section, some limitations and assumptions of the thesis are presented and discussed.

1.2 Limitations

Bulk Invest ASA operates only in the dry bulk sector of shipping and therefore – as the problem statement disclose – this thesis is limited to freight rate risk in that particular segment. Dry bulk is the largest sector in shipping (cf. fig. A.1) and the profundity of empirical research within dry bulk is equally large. However, to the best of the author's knowledge, few (if any) articles have the same structure and combination of research questions as this thesis. Thus, there is a possibility to add to the insight of the already vast empirical research library of dry bulk shipping.

During the process of collecting relevant data for this thesis, I learned that shipping is not the most open and publicly available industry. After several weeks trying to get the correct freight rate data, I finally managed to meet some kind people at the research department of Fearnley Shipping, which provided me with spot freight rate data. However, at that time, I did not know that I would also need data on long-term time-charter contract. Therefore, a shortcut had to be made to be able to finish the thesis within deadline, although this reduced the precision of the analysis.

Consequently, the time-charter business of Bulk Invest ASA is modelled through fixing the simulated spot rates at continuous intervals of three months, six months and one year (i.e. mimicking TC contracts of respective lengths).[3]

There are several ways to *mitigate* freight rate risk, but Bulk Invest ASA mainly uses time-charter contracts as they operate within both the spot- and time-charter market. Hence, other ways to hedge freight rate risk, as for instance Forward Freight Agreements, are not assessed in this thesis.

To limit the scope of the thesis, only market risk factors affecting the revenues (i.e. freight rates) are modelled and assessed. As such, the cost side is assumed proportionally static or assumed to follow a predetermined trajectory (e.g. bunker costs).

The remainder of this thesis is organized as follows. In chapter two, the most prominent relevant literatures are reviewed. Next, a presentation of the Risk Management Process as a framework for the empirical research and a theoretical review of relevant risk measures are provided in chapter three. Chapter four surveys the dry bulk shipping market and the empirical and theoretical dynamics and formation of freight rates. In the last part of this chapter, Bulk Invest ASA is presented as an illustrative case and its risk factors are assessed. The building blocks of the stochastic model are presented in chapter five, starting with a presentation and analysis of the dataset. The following subsections present stochastic processes and volatility, and lastly copulas as a remedy for simulating correlations between processes. In the final subsection of this chapter, the stochastic freight rate model and the cash flow model are proposed. Chapter six includes the empirical results from simulations and discusses the cash flow effect for Bulk Invest ASA. Validation of the risk measures is done by backtesting and sensitivity analysis in chapter seven and eight, respectively. Chapter nine concludes.

[3] The decision whether or not to enter into a time-charter contract depends on several factors, as for instance the prevailing market and its prospects. Assuming constant time-charter fixing is thus not a complete mirroring of the real world.

2 Literature Review

2.1 Freight Rate Modelling

Most of previous literature involving freight rate modelling has been devoted to ship valuation or to price implied real options embedded in various types of ship- and freight contracts. In their paper addressed to price freight rate options, Koekebakker, Sødal and Aadland (2007) use the classical geometric Brownian motion to model freight rate dynamics. This is a stochastic process which is well used in financial modelling.

In order to account for mean reversion, Bjerksund and Ekern (1995) were one of the first to postulate that freight rates follow an Ornstein-Uhlenbeck (O-U) process.[1] Tvedt (1997) extended this modelling framework by proposing a geometric O-U process to correct the lack of downward restriction in ordinary O-U processes (i.e. eliminate the possibility of negative prices). However, few articles have followed up on this downward restriction because freight rates modelled by the original O-U process do not result in frequently negative prices. Furthermore, when they occur, the rates are often slightly negative and hence negligible (Jørgensen and De Giovanni, 2010). In addition, modelling the log of spot freight rates ensures positive freight rates no matter what.

Jørgensen and Giovanni (2010) model freight rates by the original O-U process to value time-charter contracts with purchase options.[2] In their article they note that in addition to the lack of downward restriction, the O-U process also has some other less desirable properties. For instance, this process implies that future freight rates follow a Gaussian distribution,

[1] See section 4.3.5 for the description of mean reversion in freight rates.

[2] A purchase option gives the charterer the right to buy the vessel at a specified date (usually at the end of the time-charter period) and for a specified price.

even though actual freight rates often appear to be skewed and exhibit high kurtosis. The process also has a constant rate of volatility, which does not match the empirical facts of freight rate volatility (see section 5.1.3).

It should be noted that there is some dissension within the shipping literature about the stationarity of spot freight rates. Koekebakker, Sødal and Aadland review this topic in their article from 2006. Although sound theoretical arguments propose mean reversion (i.e. stationarity),[3] almost any empirical test of stationarity conclude that spot freight rates are non-stationary processes (i.e. following random walks; see for instance Tvedt, 2003).[4] Thus, modelling spot freight rates as an O-U process may not be the exact remedy for capturing the freight rate dynamics after all.

Koekebakker, Sødal and Aadland implement the O-U process in two other articles where they value the real option of switching between the dry bulk market and the wet bulk market for a combination carrier (2008), and switching between the dry bulk- and tanker market (2009). However, in these articles they look at the spread between two markets and argue that it must converge to a long-run mean. The reason is that if the differential between the freight rates in two shipping markets becomes too large, the market agents will switch to the high yielding market. This switching (i.e. interaction between supply and demand) ensures that the spread will revert to a long-term mean.

In some of the more recent papers involving freight rate modelling, more complex models are utilized in order to capture more of the empirical properties of freight rate dynamics.[5] Nomikos et al. (2013) suggest jump-diffusion models that can capture the fat tails of the logarithmic returns of freight rates. Benth, Koekebakker and Taib (2015) propose modelling freight rates based on exponential Lévy processes, including stochastic volatility by the work of Bandorff-Nielsen and Shepard (2001).[6] Benth and Koekebakker (2015) use the same model to price Forward Freight Agreements for Supramax vessels.

---

[3]See section 4.3.5.

[4]Freight rates have commonly been tested by the classical linear ADF-test. Koekebakker et al. find that spot freight rates are stationary, but only by using a non-linear framework for stationarity testing.

[5]Some empirical properties besides mean reversion are: fat-tailed returns, stochastic volatility (i.e. volatility clustering) and short term persistence.

[6]Bandorff-Nielsen and Shepard construct continuous-time stochastic volatility models based on non-gaussian processes of O-U type.

All of the literature presented above model univariate freight rate distributions only. As one of few, Merikas et al. (2013), examine joint distribution decomposition in the dry bulk sector with the use of different copulas, and validate them by Value-at-Risk (VaR). For risk management purposes, they find that using no-time dependent (i.e. static) copulas deliver VaR estimates with breach numbers within the confidence level (as opposed to dynamic copulas).

2.2 Risk Measuring in Shipping

Measuring risk has since the introduction of Value-at-Risk in the late 1980s been a manageable practice for most firms. The simplicity and flexibility of the *at-Risk* measures have made them common in use and in empirical research. The literature on Value-at-Risk and other similar risk measures is extensive. Philipe Jorion – by many seen as the number one expert on Value-at-Risk – has written an excellent book on the general theory and use of VaR.[7] However, there is not the same extent of articles and empirical research regarding practical use of quantitative risk measure (i.e. *at-Risk* models) within the shipping literature.

In their article from 2008, Angelidis and Skiadopoulos use a Value-at-Risk approach to measure the market risk of freight rates. With data from both the dry and wet bulk market, they measure the risk by several different types of parametric and non-parametric models, as well as the extreme value method. They find that the simplest non-parametric models (i.e. filtered historical simulations) provide the most reliable and robust risk measures.

Alesii (2005) explains total variability in net present value for an industrial project in the presence of real options, and derives the VaR of the expanded net present value of the project. This research is illustrated by a numerical example from the shipping industry. He also derives the Cash Flow-at-Risk for each epoch in the life of the project.

The previous mentioned article by Benth et al. (2015) applies VaR models to analyze the efficiency of different stochastic processes meant to capture freight rate dynamics.

[7] Jorion, 2007.

3 Methodology

3.1 Risk Management

Running a shipping company (or any company in general) can be illustrated by sailing a ship through a strait. When there is high tide, it is easy to sail and almost no decision or wrong turn could make you run aground. This is the equivalent of a market state where freight rates are high and revenues are good. However, when the market cycle turns and the low tide sets in, reefs and rocks awash make every bad decision a potential ship wrecker. Sailing a company aground may sink the ship or at least make it costly to recover for sailing again.

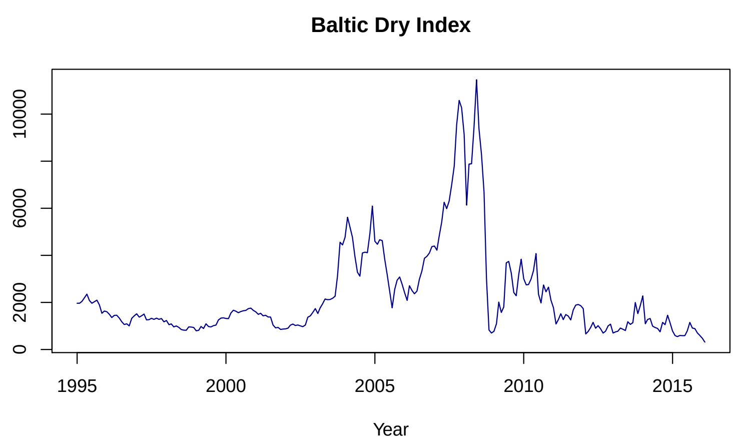

Figure 3.1: The Baltic Dry Index reflecting the daily movements in the dry bulk sector (averagely weighted across vessel types)

After a couple of years with a very high tide, cf. fig. 3.1, the shipping market has turned and it may never quite reach the same peak again. The dry bulk market experiences freight rate levels which are the lowest in several decades. Against this backdrop, risk assessment, management and mitigation are critical components if companies are to survive and thrive in the challenging new world.

But what is risk? In finance theory, risk is defined as the dispersion of unexpected outcomes owing to movements in financial variables (Jorion, 2007, p. 75). This includes both positive and negative deviations. Hence, some risks are good and some are bad. For instance, building on the previous metaphor; taking the risk of steering through a shallow strait may incur some extra risk compared to sailing around in open waters. But, passing through the strait will reduce the total voyage time and therefore have a positive effect on expected revenues. Conversely, navigating through the strait when the weather is bad, the tide is low and the piracy activity is high incurs risks which the company will be better off sailing around in open waters. Thus, a need for active risk management is prevalent.

Table 3.1: Expected returns and volatility by different dry bulk vessels classes

Vessel class	**Mean $/day**	**Standard deviation**		**Mean return**
		$/day	**% of mean**	
Capesize	36 448	38 867	107%	-0.04%
Panamax	19 060	17 194	90%	-0.03%
Supramax	20 045	15 282	76%	-0.07%
Handysize	14 450	10 845	75%	-0.08%

Calculations are based on daily time-charter rates in the horizon: 1998 – 2016.

In this thesis, risk management is defined as a set of financial or operational activities that maximize the value of a company or a portfolio by reducing the costs associated with cash flow volatility.[1] As shipping being an industry with very high volatility in earnings and relatively low expected returns, cf. table 3.1, managing risk is essential for any shipping company pursuing profitable operations.

[1]This definition is in accordance with Stulz (2003).

Managing risk is a dynamic and everlasting process which can be divided into the following stages:

- Identification
- Measurement
- Mitigation
- Review and monitoring

The first stage of the Risk Management Process is to identify the potential risk factors for the company.[2] This may be factors affecting the financial, operational or strategic performance, and it could be both internal and external sources of risk. After having identified the potential risk factors, one needs to measure and quantify the potential impact on the business, and the likelihood that the events will happen.

The next step is to minimize unwanted effects of negative impacts on the business from the risk taken. In shipping, this can be done in numerous ways depending on the risk mitigated. For instance, freight rate risk can be hedged by using time-charter contracts, Forward Freight Agreements or options. It should be noted that some risk is desirable for the company and some risk is not. In order to earn more, certain risks must be taken. The choice of either hedging or operating exposed to the market must be taken at the top level management of the shipping company.

The final stage of the Risk Management Process is to monitor and control the actions taken in the process so far. This review makes the foundation for the next risk assessment. After reviewing the steps, the process starts over again. The different stages of the Risk Management Process is illustrated in fig. 3.2.

[2] Also known as mapping of risk factors.

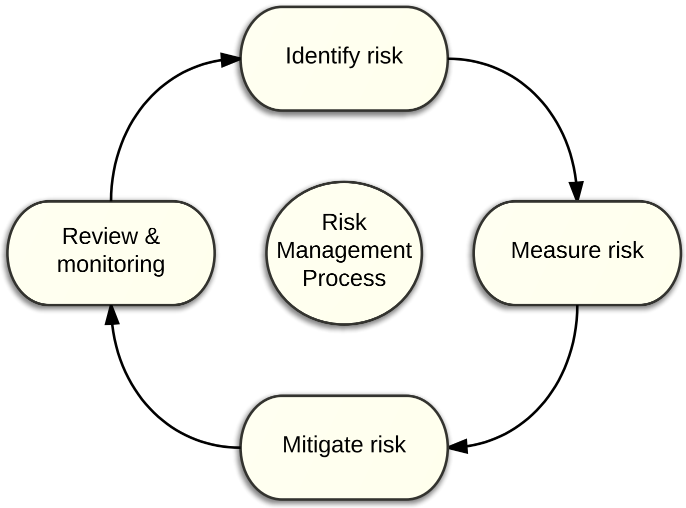

Figure 3.2: The Risk Management Process. *Source*: Own making

The Risk Management Process will work as a framework for this thesis. Firstly, risk factors for Bulk Invest ASA are identified in order to simulate potential risk and effects on the business. Then, the simulated risk is measured by quantitative methods described in the next subsections. Time-charter contracts are examined as a hedging tool for mitigating spot freight rate risk. This is reviewed by comparing it to operating in the (unhedged) spot market. Lastly, the reliability of the quantitative measures are reviewed and verified.

3.2 Measuring Risk

Measuring the risk taken is an important part of the Risk Management Process. There are several ways to measure the risk depending on the type of risk factor assessed. For instance, some risk factors, like customer or employee satisfaction, are easier to measure and disclose by qualitative methods, e.g. surveys and interviews. Other, more numerical factors (e.g. cash flow fluctuations) can be measured by quantitative methods. In this thesis the focus is on the latter type of risk factors, making the use of quantitative methods appropriate.

In the following subsections relevant quantitative measures of risk are presented. In addition, the use of backtesting and sensitivity analysis are presented as remedies for reviewing the risk measuring stage.

3.2.1 Value-at-Risk

Value-at-Risk (VaR) is a risk measure commonly used among banks and other financial firms. VaR is a statistical measure of downside risk based on the current market positions (Jorion, 2007, p. 105). It expresses the maximum loss over a target horizon that will not be exceeded at a given level of confidence. Let R be the final return of a position and r be all possible values the return can take over the holding period. Furthermore, let the target horizon span from time t to time T. VaR is then a predetermined quantile of the predicted distribution of the returns on the current positions (e.g. returns on freight rates). A quantile is defined as the value such that the probability mass to the right or left is equal to α (Jorion, 2007, p. 89). As we are interested in the downside risk of the position, the probability mass to the left of the quantile is the relevant one. Assuming that the underlying distribution of the returns is continuous, the quantile can be expressed as

$$\alpha = P(R \leq -VaR_{t,T}^{\alpha}) = \int_{-\infty}^{-VaR_{t,T}^{\alpha}} f_{t,T}(r)\,dr \tag{3.1}$$

The predicted Value-at-Risk will, amongst other factors, depend on the probability distribution of the returns.[3] There are two branches of methods to estimate this distribution: A non-parametric approach like the so called historical method estimates the shape of the distribution based on the historical data. Value-at-Risk is then estimated by scaling and ranking the observations and finding the α% quantile of the distribution. For example, if we want to find VaR at a 95% confidence level, based on 1 000 historical observations, we simply find the (1 - 95%) x 1 000 = 50th lowest return. This is the 5% quantile of the past 1 000 observations.

The other branch assumes, in contrast to the first one, that there is a known underlying shape of the distribution of returns (e.g. the Gaussian

[3] Other factor include: volatility of returns, total exposure, length of the horizon and confidence level.

distribution). This branch is called the parametric approach to Value-at-Risk. Assuming the probability distribution of the returns follows a Gaussian one, with mean μ and standard deviation σ, we can rewrite eq. (3.1) in terms of the distribution function (Φ)

$$\alpha = \Phi([-VaR^{\alpha}_{t,T} - \mu_{t,T}]/\sigma_{t,T})$$

and hence, VaR can be expressed by the inverse of the distribution function as

$$VaR^{\alpha}_{t,T} = -[\mu + \sigma\Phi^{-1}(\alpha)] \tag{3.2}$$

However, the Gaussian probability distribution (or other distributions) does not always fit the historical returns perfectly. For instance, the skewness or kurtosis of the observed returns may be higher than the theoretical values. Using a Monte Carlo method to predict the Value-at-Risk can account for this as anything in the simulation process can be tweaked to fit the observed features of the historical distribution.

In cases of more than one asset (i.e. several freight indices), the *portfolio* VaR can be estimated by matrix algebra and using the correlation between the assets. The matrix formula for VaR is:

$$VaR_p = (\mathbf{V}'\mathbf{\Sigma}\mathbf{V})^{1/2} \tag{3.3}$$

where $\mathbf{\Sigma}$ is a correlation matrix, $\mathbf{V}$ is a matrix of individual VaRs and $\mathbf{V}'$ is its transposed equivalent.

3.2.2 Cash Flow-at-Risk

Estimating the Value-at-Risk is especially suitable for financial firms like banks as they often have positions in the market related to trading purposes. This means they would only consider a potential change in fair value to be the main risk. Non-financial firms, on the other hand, have a different approach to risk management because these companies aim to stabilize the prospective cash flows from their positions in the market. Thus, the main risk factor for such a firm is the operational business. Furthermore, as banks and other financial firms can easily cover or liquidate their positions in the short-term, VaR with a short time horizon (i.e. daily or weekly) is appropriate. For a non-financial firm, its portfolio of assets may not have the same liquidity as a financial firm. A longer horizon in the risk measure is then more consistent with the quarterly or yearly profit and loss measure of the non-financial firm (whereas banks for instance, often have a daily profit and loss measure).

Cash Flow-at-Risk ($CFaR$) is the cash flow equivalent of Value-at-Risk. It possesses the same nice features as it sums up all the company's risk exposures in a single number that can be used to guide corporate risk management decisions. It is this single number – the maximum predicted shortfall (in cash flows) given the targeted probability level – and the fact that it can be directly compared to the firm's risk tolerance that are the uniquely attractive features of both VaR and $CFaR$ (Andrén et al., 2005). Cash Flow-at-Risk is estimated on longer horizons (e.g. quarterly or yearly) than VaR because it focuses on revenue flows from the operational business, making it a natural risk measure for non-financial firms. $CFaR$ is analogous to VaR and it is calculated the same way, but on cash flow rather than value, which is the key here: Applying VaR to a non-financial firm's portfolio of financial instruments and/or assets would only capture a small part of the company's overall exposure since VaR ignores the risk of the company's underlying commercial cash flows. $CFaR$ then represents a transfer of the concept underlying VaR to a setting where cash flows are the targeted variable (Andrén et al., 2005).

The calculation of $CFaR$ requires an estimate of the probability distribution for future levels of cash flow. In this thesis, the cash flows are estimated based on a stochastic model of freight rates. Fluctuations in freight rates will directly affect the fleet cash flow of any (unhedged) shipping company, and consequently, cash flow performance is one of the topmost concerns in shipping. In capital intensive industries like shipping, where there is a lot of project financing, one important principle is that repayment must come from the operating cash flows of the financed asset (Gatti, 2013). In that sense, what really matters in measuring freight market risk is the impact of freight rate variability on cash flow performance. In this thesis, the simulated cash flow paths from the stochastic model will work as a foundation for the cash flow risk estimation.

3.2.3 Extreme Value Theory

Value-at-Risk has been criticized as a risk measure since it is not coherent when the underlying distribution is non-normally distributed (see for instance Artzner et al., 1999 and Acerbi and Tasche, 2002b; Acerbi and Tasche, 2002a). The problem is that VaR does not fulfill the axiom of sub-additivity.[4] Acerbi and Tasche explain this problem: For a sub-additive measure, portfolio diversification always leads to risk reduction, while for measures which violate this axiom diversification may produce an increase in their value even when partial risks are triggered by mutually exclusive events (Acerbi and Tasche, 2002a, p. 3).

In addition, when measuring risk by *at-Risk* measures (e.g. $CFaR$ or VaR), we only compute a specific quantile of a desired significance level. This means that these measures do not account for how the probability mass is distributed in the tail (i.e. values exceeding the quantile). If the distribution of returns is perfectly bell-shaped (i.e. normally distributed), then there are not many extreme observations in the tail as it decays faster towards zero than the tail of distributions with higher kurtosis. But, as touched upon in the VaR section 3.2.1, distributions of returns often exhibit higher kurtosis than the normal distribution. This makes the question: "What is the expected loss if we experience a worst case scenario (i.e. exceeding the $CFaR$ number)?" interesting.

One answer to this question is provided by Expected Shortfall (ES), which in fact is a coherent risk measure. ES estimates the expected loss in the α% worst cases of a specified period. Keeping notations constant, let R be the final return of a specific position in the market over a specified time horizon T. Furthermore, let $\alpha = A \in (0, 1)$ be some percentage representing a sample of "worst cases" for this position (i.e. some specified probability level). The A% Expected Shortfall is then defined as (Acerbi and Tasche, 2002a)

$$ES^{\alpha}(R) = -\alpha^{-1}\left(E\left[R\,\mathbf{1}_{\{R\leq r^{\alpha}\}}\right] - r^{\alpha}(P\left[R \leq r^{\alpha}\right] - \alpha)\right) \tag{3.4}$$

The term $r^{\alpha}\left(P\left[R \leq r^{\alpha}\right] - \alpha\right)$ is interpreted as the exceeding part to be subtracted from the expected value $E\left[R\,\mathbf{1}_{\{R\leq r^{\alpha}\}}\right]$ when $\{R \leq r^{\alpha}\}$ has probability larger than $\alpha = A$%.[5] On the contrary, when the probability distribution is continuous we always have that $P\left[R \leq r^{\alpha}\right] = a$. Thus, the

[4]Let Θ be the set of risk, X and Y be random variables representing the final net position of an asset, and ρ the risk measure. The sub-additivity axiom then states that: $For\ all\ X\ and\ Y \in \Theta \Rightarrow \rho(X+Y) \leq \rho(X) + \rho(Y)$

[5]Explanation of notation: $\mathbf{1}_{\{Relation\}} = \begin{cases} 1, \text{if Relation is true} \\ 0, \text{if Relation is false} \end{cases}$

term vanishes and equation (3.4) is reduced to

$$ES^{\alpha} = -E\left[R|R \leq r^{\alpha}\right] \tag{3.5}$$

3.3 Validation of Risk Measures

An important part of the Risk Management Process is to review the steps taken in order to verify that the process is reliable and robust. By validating the risk measures, we check that the models are adequate and that they predict the risk reasonably well. Unrealistic assumptions and flawed risk models could be more damaging to the business than helpful because they may underestimate risks and thus lead to unjustified investment decisions (Alizadeh and Nomikos, 2009, p. 333).

3.3.1 Validating *at-Risk* Models

Backtesting

Back testing is a formal statistical framework that consists of verifying that actual losses are in line with the projected ones (Jorion, 2007, p. 139). This involves systematically comparing the history of the forecasts provided by the risk measure with their associated portfolio returns. The most commonly used framework in backtesting VaR models was developed by Christoffersen (1998).[6] As $CFaR$ is analogical with VaR, some of the tests intended for VaR are applicable for this risk measure as well. A sequence of out-of-sample VaR estimates is said to be efficient with respect to the information set available at $t-1$; Ω_{t-1}, if the following condition holds (Alizadeh and Nomikos, 2009, p. 333):

$$E\left[\Phi_t|\Omega_{t-1}\right] = \alpha \qquad \text{with } \Phi_t = \begin{cases} 1, & R_t < VaR_t^{\alpha} \\ 0, & R_t \geq VaR_t^{\alpha} \end{cases} \tag{3.6}$$

The above equation implies that the expected failures, Φ_t, of the *at-Risk* measure should be $i)$ equal to the nominal significance level on average

[6] Colletaz et al. (2013) expanded this framework to look at how severe each violation of the VaR is, not just the number of violations as in Christoffersen's framework. Colletaz et al.'s framework will not be utilized in this thesis.

and ii) uncorrelated with any variable in the information set available at $t-1$. Building on the work of Kupiec (1995), Christoffersen provided a likelihood ratio (LR) test for this property (eq. (3.6)) in order to state the accuracy of a VaR model.

An *at-Risk* model must have an unconditional coverage in the sense that the number of actual violations (Φ) of the modelled risk level (i.e. greater cash flow losses than the $CFaR$) must not be statistically different from the expected number of violations at the given confidence level. In addition, the actual violations must be independent from each other, meaning that the previous violation should not contain information about the presence of future violations.

The unconditional coverage can be tested by the following LR test:

$$LR_{uc} = 2\ \ln\left[\left(1-\frac{\Phi}{\Gamma}\right)^{\Gamma-\Phi}\left(\frac{\Phi}{\Gamma}\right)^{\Phi}\right] - 2\ \ln\left[(1-\tau)^{\Gamma-\Phi}\tau^{\Phi}\right] \sim \chi_1^2 \qquad (3.7)$$

Where Γ is the number of *at-Risk* measures estimated over the period and τ is the theoretical proportion of violations in that period for a given confidence level (ϱ):

$$\tau = \frac{\Gamma\,(1-\varrho)}{\Gamma}$$

The likelihood ratio test of unconditional coverage is asymptotically chi square distributed with one degree of freedom. The null hypothesis of this test is that the average number of *at-Risk* violations is correct.

Christoffersen also provided a likelihood ratio to test for the independence between recorded violations. This test is applicable for VaR where the value of a position is directly related to the return on it. However, when estimating $CFaR$, we also need to consider the prevailing costs of the company. Due to these costs (which are not present in VaR estimations), the cash flows are likely to violate the estimated quantile loss several days in a row if revenues (e.g. freight rates) are low. The independence test will therefore not provide any reasonable test results for $CFaR$ and will not be considered in the backtesting chapter.

Sensitivity Analysis

Sensitivity analysis (also known as stress testing or scenario analysis) involves examining the risk model's responsiveness and its sensitivity to some extreme events (hypothetical or real). This is not a direct test of the *at-Risk* model, but more a test of the solidity of the business. It should be noted that this type of model testing has been criticized (see for instance Alizadeh and Nomikos, 2009, p. 335). First of all, the test is biased because the scenarios tested are subjectively chosen by the individual performing the test. Secondly, it is not easy to make a valid inference from the test as probabilities are difficult to assign to extreme events. Finally, it is impossible to combine the test in a quantitative manner with the actual risk measure (i.e. $CFaR$). Berkowitz (1999) expresses this nicely: 'stress testing is a statistical purgatory. We have some loss numbers but who is to say whether we should be concerned about them?'

Even so, a sensitivity analysis could give the user of the *at-Risk* measure a pointer at which factors have greatest impact on the risk of the portfolio, and in that manner be a useful tool.

3.3.2 Validating Expected Shortfall

Backtesting *at-Risk* measures is very easy and robust due to the simple framework described in the previous subsection. Unfortunately, the same is not true for Expected Shortfall. The literature is debating whether 1) it is possible to backtest ES at all and 2) if it is possible; what is the best way to do it.

One of the reasons why the backtesting ability of ES is discussed is highlighted by Kerkhof and Melenberg (2004). They argue that it is difficult to backtest ES because there is a misalignment between estimated and actual losses. In other words; ES is essentially the average loss of the α% worst cases, but in real life this number is just a single scenario (i.e. not a combination calculated through mathematical formulas).

Gneiting (2011) put further contribution to the distrust of the backtesting ability by proving that ES lacks a mathematical property called elicitability. If we define ES as a forecast y given the realization x (e.g. the realized loss of that period), lack of elicitability means it is not possible to find a scoring function $S(y, x)$ that is to be minimized by a corresponding forecast evaluation. An example of minimizing a scoring function could be to assess the best of several weather forecasts: Using the sum of squared forecasting errors (SSE), the one with the lowest SSE is chosen. In practice, this is equivalent to minimizing the scoring function of the weather forecasts. However, this is not possible to do for ES because such a scoring function does not exists, as proven by Gneiting.

Other researchers do believe Expected Shortfall can be backtested and several tests have been proposed (see e.g. Acerbi and Szekely, 2014; Emmer et al., 2013; Righi and Ceretta, 2013). However, as opposed to the case of *at-Risk* measures, there is no *one* framework for backtesting which is commonly reckoned as adequate. Since the research on this topic is so fragmented and in order to limit the scope of this thesis (the backtesting ability of ES could be a thesis in its own), backtesting Expected Shortfall is omitted in the empirical research.

4 The Shipping Market

4.1 The Dry Bulk Market

From the decades following the Second World War, bulk shipping (including liquid and dry bulk) has developed into the major sector of the shipping industry and bulk tonnage now accounts for about two-third of the world merchant fleet. The dry bulk sector which includes shipping of iron ore, grain, coal and other minor bulks accounts for about one-third of the world merchant fleet (Danish Shipowner's Association, 2015). The seaborne trade of main bulk cargo has experienced a linear growth since the 1970s as depicted in figure 4.1.

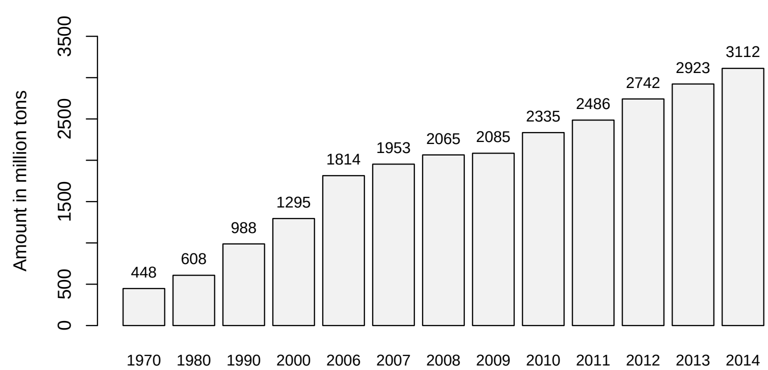

Figure 4.1: The growth in shipping of main bulk cargo. *Source:* UNCTAD; Clarkson Research Services © Statista 2016

Table 4.1: Vessels sizes and commodities shipped

Vessel class	Deadweight tonnage	Main commodity shipped
Capesize	80 000 + dwt	Iron ore, coal
Panamax	60 000 – 80 000 dwt	Coal, grain
Supramax	35 000 – 60 000 dwt	Grain, bauxite, phosphate
Handysize	10 000 – 35 000 dwt	Steel products, cement, logs

Source: Alizadeh and Nomikos (2009). Note that there is not *one* standard classification on vessel types due to the everlasting change in size and specification of the vessels. Consequently, vessel names and size intervals may vary from source to source.

By January 2015 there were approximately 16 900 dry bulk vessels in the world merchant fleet (fig. A.1). Dry bulk vessels are mainly divided into four groups: Capesize, Panamax, Supramax (also known as Handymax) and Handysize. These ship types are differentiated by their size, i.e. tonnage, with Capesize being the largest vessel, and Handysize the smallest. Capesize, which are vessels above 80 000 dwt are mainly used for iron ore and coal. Panamax vessels have their name from the Panamax canal which limits the vessel to be up to 80 000 dwt. Panamax vessels are mainly used for coal and grain. Supramax and Handysize vessels are smaller ships between 35 000 – 60 000 dwt and 10 000 – 35 000 dwt respectively. These vessels are mainly used for grain, bauxite, phosphate, steel products, sugar and other minor bulks. Table 4.1 sums up the different dry bulk vessel classes.

4.2 Freight Contracts

A shipping freight contract is an agreement between a shipowner and a charterer where the shipowner provides a service to the charterer for a specified amount of money per day for ship hire, or per ton cargo transported between two ports. This specified amount is called a freight rate. Depending on the needs of the charterer in terms of the type and duration of the service, several different types of freight contracts have been developed and used in international shipping. In the following, the two most relevant freight contracts are listed and explained:

- Voyage charter contracts – A contract where the shipowner agrees to transport a specified amount of cargo between a designated loading and discharging port. These are also known as *spot contracts* in the shipping industry, and they are quoted in US dollars per metric ton (US $/mt) or as a lump sum.

- Time-charter (TC) contracts – Under a TC contract, the charterer agrees to hire the vessel from the shipowner for a specified period of time and under certain conditions. These conditions include the vessel's performance specifications (speed, consumption, etc.), the condition and location of the vessel during delivery and redelivery, fuel on board, and trading areas, as well as several other terms.

These two contracts are the most common types used in the dry bulk market. Single-voyage contracts can be classified as short-term or spot-charter shipping, while time-charter (TC) contracts are often long-term, ranging in length from a few weeks or months to several years. TC contracts can be seen as a method of fixing the spot rate in order to secure a fixed revenue stream over a certain period. This means the shipping company hedges its freight rate risk for a period equalling the length of the time-charter contract.

There is also a difference in the two contracts regarding the cost allocation between shipowner and charterer. When entering a voyage contract, the shipowner is responsible for all costs related to the shipping; i.e. cargo-handling costs, voyage costs, capital cost, and operating and maintenance costs.[1] In a TC contract however, the shipowner only has to pay the operating and maintenance cost as well as the capital cost of owning the ship.

4.3 Freight Rates

In the following subsections, the behavior, formation and implications of freight rates are thoroughly explained.

[1]Voyage costs include: fuel costs, port charges, pilotage, and canal dues. Operating and maintenance costs include: crew wages, stores and provisions, periodic maintenance and insurance costs.

4.3.1 Shipping Cycles

Market cycles pervade the shipping industry. Martin Stopford elegantly describes them as waves hitting the beach: From a distance they look harmless, but once you are in the surf it's a different story. No sooner has one finished than another starts and, like surfers waiting for a wave, shipowners cluster in the through, paddling to keep afloat and anxiously scanning the horizon for the next big roller (Stopford, 2009, p. 93).

These shipping cycles are a crucial part of the market mechanism of freight rates and they have both long and short components as well as seasonal ones. The longer cycles are the heart of the mechanism and they are driven by technical, economical or regional changes. These cycles are more difficult to detect as they can span from century-wide to half a century. For instance, from 1869 to 1914 the shipping industry experienced a downward spiral in freight rates which was driven by the increasing efficiency of steamships and the phasing out of the much less efficient sailing ships (Stopford, 2009).

The shorter cycles are more easily detected as they are created by the supply and demand for seaborne trade. These cycles works as a mechanism devoted to removing imbalances in the supply and demand for ships. If there is a shortage of supply, the market grants investors with high freight rates until more ships are ordered. On the other side, an excess of ships results in lower freight rates and squeezes the cash flow until the owners of the oldest ships give up the struggle and ships are scrapped. Although these short cycles are more observable than the longer ones and have a periodically behavior (i.e. a market through, followed by a recovery, a market peak and then a collapse), regularity is not a part of the process. Thus, there is no simple formula for predicting the start or end of the next cycle. If, for instance, investors classified the cycles by length and used them as a forecasting aid, they would just postpone the ongoing cycle as the excess or surplus of ships would not be altered (Stopford, 2009).

The last component of shipping cycles is seasonality which is fluctuations in freight rates occurring within the year. These fluctuations happen usually at specific seasons in response to seasonal patterns of demand for sea transport. For instance, in the agriculture trades, there is a noticeable cycle in freight rates for ships carrying grain, caused by the timing of harvest. Another example is the stocking up of oil for periods of peak demand in the winter (Stopford, 2009, p. 97). It should be noted however, that a study of the dry bulk market by Kavussanos and Alizadeh (2001) has rejected the existence of stochastic seasonality and found that freight rates exhibit deterministic seasonality only at a very low level (Benth et al., 2015, p. 275).

4.3.2 Spot Rate Formation

The spot market for freight rates is characterized by the classical supply and demand functions, each of which depends on multiple factors that interact constantly so that the equilibrium freight rate can be determined (Alizadeh and Nomikos, 2009, p. 44).

The demand for shipping services is influenced by several factors including the world economy (drivers of shipping cycles), seaborne commodity trades (e.g. iron ore, coal, grain), the average distance the cargo is shipped, random shocks (e.g. financial crisis) and transport costs (driven up and down by the market). On the other hand, the supply of shipping services is the number of ships available (i.e. gross tonnage) in the world merchant fleet at any particular time and the rate of newbuilding/scrapping of ships. The logistical efficiency – for instance speed and waiting time – of the operating fleet is also influencing the supply of shipping services.

The adjustment mechanism linking supply and demand of shipping services is the freight market. The freight market can be analyzed on the basis of the perfect competition model. The demand function for ocean shipping is inelastic, mainly because there is no competing transport mode for most bulk commodities. The supply function however, is convex in shape due to the limitation of supply at any given point in time. The equilibrium price-level and any changes to it is thus highly dependent on the current level of demand and supply. See figure 4.2.

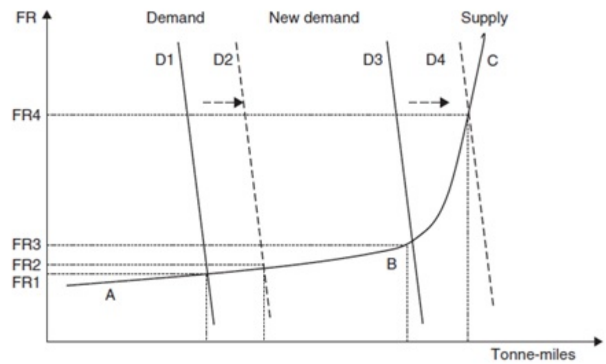

Figure 4.2: Supply and Demand in the Spot Freight Market. *Source:* Alizadeh and Nomikos (2009)

When freight rates are low (between point A and B), the supply curve is highly elastic, and any shift in demand will have a small impact on the freight rate (i.e. $FR_1 \rightarrow FR_2$). However, when freight rates are high (between point B and C), the supply curve becomes inelastic. Assuming initial freight rates at level FR_3, supply and demand schedules are very tight, and the fleet is fully utilized. In this case, if the demand shifts positively from D_3 to D_4, assuming supply in the short term is constant, the new equilibrium freight rate shoots up from FR_3 to FR_4, which is a relatively large increase. This means that in times of high freight rates, the volatility in the freight market is high. Similarly, when freight rates are low, the supply market absorbs most changes in the demand, implying lower volatility.

4.3.3 Time-Charter Formation

While the spot freight rate is determined by the interaction between supply and demand, time-charter rates are believed to be determined through the market's expectations about future spot rates, i.e. time-charter rates are the conditional expectation of the spot rates. The relation between spot- and time-charter rates is based on the expectations hypothesis and the theory on arbitrage-free term structures commonly used in finance (see for instance Kavussanos and Alizadeh, 2002). According to this theory, a shipowner or a charterer should be indifferent between entering into a time-charter contract or a series of spot contracts equaling the length of the TC contract. Assuming the buyer of the shipping contract has some expectation of the future freight rates, notated as $E_t FR^m_{t+im}$, then the discounted cash flow received under each type of contract must be equal or else there would be an arbitrage opportunity in the market. Mathematically, this relationship can be expressed as (Alizadeh and Nomikos, 2009, p. 49)

$$TC^n_t = \sum_{i=1}^{k} \frac{(E_t FR^m_{t+im} - E_t VC^m_{t+im})}{(1+r)^i} \Bigg/ \sum_{i=1}^{k} \frac{1}{(1+r)^i} \qquad k = n/m \tag{4.1}$$

where, r is the discount rate, TC^n_t is the time-charter rate for a n period contract at time t, $E_t FR^m_{t+im}$ is the expected spot charter rate at time t of a contract over m periods from $t + im$ to $t + (i+1)m$, and $k = n/m$ is a positive integer indicating the number of spot contracts corresponding to the length of the time-charter contract. $E_t VC^m_{t+im}$ is the expected voyage cost. Note that as opposed to the spot rate, the time-charter rate is quoted net of voyage costs due to the difference in cost allocation between the two contract types.

However, one important difference between spot- and TC operations is the fact that the latter is a sort of hedge against underlying spot rate fluctuations. That is, a TC contract obtains a certain security in revenues by fixing the freight rate over the period, while the earnings of a corresponding series of spot contracts may vary depending on the market conditions of that period. Hence, a risk element should be included in the relationship between spot- and TC contracts. This risk premium, φ, can be interpreted as the price the shipowner is willing to pay to pass the uncertainty of the spot market to the charterer. The final equation expressing the spot and TC relationship is then

$$TC^n_t = \sum_{i=1}^{k} \frac{(E_t FR^m_{t+im} - E_t VC^m_{t+im})}{(1+r)^i} \Bigg/ \sum_{i=1}^{k} \frac{1}{(1+r)^i} - \varphi_t \qquad k = n/m \tag{4.2}$$

Kavussanos et al. (2002) investigates the expectations hypothesis of the term structure and finds evidence pointing on a time varying element in the risk premium, meaning the size of it might differ according to the prevailing market uncertainty.

4.3.4 Time-Charter Equivalent

The market quotation for spot- and time-charter contracts differ as mentioned in section 4.2. Spot (voyage) contracts are quoted in US $/mt, while TC contracts are quoted as US $/day and net of voyage costs. As such, the two rates are not directly comparable. To overcome this, market agents calculate the time-charter equivalent (TCE) of spot rates. The TCE is calculated in two steps. First, multiplying the spot rate (US $/mt) by the amount of cargo, and deducting the total voyage costs from the total freight payment. Then, dividing this by the number of voyage days. The result is a (spot) freight rate quoted in US $/day comparable with time-charter rates (Alizadeh and Nomikos, 2009, p. 52).

4.3.5 Mean Reversion

Fundamental maritime economics propose theoretical arguments for mean reversion in spot freight rates (Koekebakker et al., 2006). Building on the supply and demand framework explained in section 4.3.2, mean reversion is a property of the spot rates driving them to a long-term mean (i.e. the marginal cost of seaborne trade). Shipping is a highly competitive market and the spot freight rate is normally determined by the marginal cost of the marginal vessel required to satisfy the demand for transportation. Spot rates fluctuate around a long-term mean because there is a theoretical ceiling to the level of them (Tvedt, 1995). If freight rates for instance reach a level which is too high, other vessel types or modes of transportation (i.e. containers or pipelines) can economically substitute bulk vessel shipping. Newbuilding of ships when freight rates are high will also drive spot rates down as the new ships are set afloat. Conversely, when freight rates are low shipowners eventually scrap their old and unprofitable vessels, creating a theoretical lower bound on the freight rates (Koopmans, 1939).

Due to the theoretical arguments of dynamic interaction between supply and demand and the resulting lower and upper bound, freight rates cannot exhibit the same explosive behavior as a non-stationary process in the long run.[2] This is enforced by the long run cycles presented in section 4.3.1. In the short run however, spot rates can fluctuate as a seemingly

[2]This reasoning is also applicable for commodity prices in general.

pure random walk. This is empowered by the fact that most empirical research fail to detect stationarity in spot freight rates (see Koekebakker et al., 2006 for a review on this subject). One reason for this is the time lag between changes in demand and corresponding changes in supply. This time lag stems from a long delivery time of new vessels and a limited ability to increase supply in the short run (Koekebakker et al., 2006). Note that in shipping, the short run may be as long as several years because of the aforementioned time lag.

4.4 Bulk Invest ASA

Bulk Invest ASA (formerly known as Western Bulk ASA) is a dry bulk shipping company originating from the dry bulk operator Western Bulk, established under the names Western Bulk Shipping and Western Bulk Carriers in 1982. The headquarter is located in Oslo, but offices are spread worldwide including Singapore, Seattle, Miami and Santiago. Bulk Invest ASA is used as a numerical example of how a dry bulk shipping company can manage its freight rate risk.

4.4.1 Company Structure

As of February 2016, Bulk Invest ASA (from now on referred to as Bulk Invest) has sold its division called Western Bulk Chartering AS, and now only has one division left; called Bulk Shipholding AS.[3] Through this entity, Bulk Invest operates a fleet of dry bulk vessels on long-term charter contracts at fixed charter rates. The contracts include options to extend the charters and purchase options for the vessels. In addition, the Bulk Invest Group has some direct investments in four of their vessels.

[3] http://www.westernbulk.com/on-wb/news-archive/wester-bulk-chartering-transfer-of-ownership

4.4.2 Identifying the Risk

Revenue

The fleet consists of seven Panamax vessels, nine Supramax vessels and three Handysize (see table 4.2).[4] The Capesize vessel class will therefore not be considered in the rest of the thesis. Bulk Invest has an order book of seven Panamax vessels, which are expected to be delivered during 2016. I will thus include these seven ships in the operating fleet, even though not all of them are operable the whole year.

Table 4.2: Fleet per 2016

Bulk Carrier	Number
Panamax	14
Supramax	9
Handysize	3
Total fleet	26

By the risk management framework presented in section 3.1, it is important to identify the main risk factors of Bulk Invest. The main factor affecting the income is the freight rate. Bulk Invest combines both spot chartering (voyage contracts) and time-chartering in their business model. This thesis seeks to estimate the effects of operating in both of these markets.

The spot chartering business is modelled by assuming that all of Bulk Invest's vessels are continuously employed by voyage contracts. This implies that all vessels enter into a new contract as soon as the current one is finished. Assuming 26 vessels in the fleet and an average voyage length of 40 days, each vessel makes approximately 9 trips per year.[5] This equals 237 voyages and a total of 9 490 ship days.

The time-charter business is modelled by assuming that all vessels are continuously fixed into TC contracts of three different lengths throughout the horizon. The contract lengths are three months, six months and one year. Figure 4.3 illustrates the realized freight rate through time-chartering.

[4] Bulk Invest names the Panamax class *Ultramax*.

[5] The shipping market is operative year-round: $365/40 \approx 9$ trips per year .

From this illustration it can be seen that fixing the vessels into long-term TC contracts makes the realized freight rates (i.e. the revenues) smoother, and thus act as a hedging tool against more extreme fluctuations. A longer contract will eliminate more fluctuations and thus be smoother and less risky (in terms of spot freight rate risk).

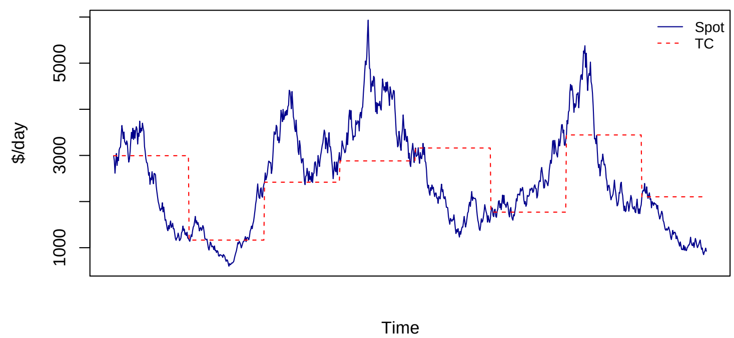

Figure 4.3: The smoothing effect of time-charter (TC) contracts

Costs

Even though Bulk Invest have some direct investment in four of the vessels (they own one Handysize 100% and three Supramax 20%), I will assume that all vessels incur both voyage- and cargo-handling costs when operating. This is because Bulk Invest charters in its vessels and thus, it become the *charterer* of the ship. The costs are allocated between Bulk Invest and the shipowner according to a standard TC contract (cf. section 4.2). In addition, Bulk Invest has to pay TC expenses, which is the cost of chartering a vessel from an external owner. The main costs related to each time-charter contract are presented in table 4.3.

In order to model the future fuel costs, the forward curve of bunker is used as an approximation.[6] According to Geman (2009), forward prices provide information about the views of market participants anticipated price trends and expectations about future supply and demand. This means they can be used as a qualified guess about the future spot price. However,

[6] The collected bunker forward curve ranges from March 2016 to December 2017.

Table 4.3: The different costs included in the cash flow analysis

Cost	Description	Unit	Amount
Commisions	Paid to brokers	of TC exp.	1.25%
	Address commissions	of TC exp.	3.75%
Bunkers	Fuel used whilst steaming	22 mt	70%
	Fuel used in port	2 mt	30%
TC expenses	Panamax	$ per day	13 400
	Supramax	$ per day	12 500
	Handysize	$ per day	9 800
Port costs	Dockage, port fees etc.	$ per port call	48 000

TC expenses are the charter hire paid for the vessels.

it should be stressed that they act merely as a guess and that the actual forecasting ability of forward prices is rather poor.

Bunker fuel is essentially a residual product of refined oil (Alizadeh and Nomikos, 2009, p. 338). Thus, the price of bunker is closely linked to the world crude oil price. As the price of crude oil is at a historical low level the market is in contango, which means that it expects the spot price to increase (see fig. A.2).[7] This implies a growth in bunker costs throughout the horizon of the simulated cash flows.

According to Bulk Invest, vessels under charter are steaming 70% of the time and the rest of the time they are in ports. While steaming, vessels consume approximately 22 mt bunker per day and approximately 2 mt per day while in ports.[8]

Whenever Bulk Invest enters into a voyage contract, it has to pay a commission to the broker, amounting to approximately 1.25% of the contracted freight rate. In addition, an address commission of 3.75% is paid to the cargo owner, making the net income for Bulk Invest 95% of the revenue (i.e. the freight rate).

The TC expenses incurred through chartering the vessels depend heavily on the prevailing market, type of ship, region and the like. However, to maintain certain simplicity in the simulations, the TC expenses are as-

[7]The lowest level of the last decade except for one "dip" right after the Financial Crisis of 2008. `http://www.tradingeconomics.com/commodity/brent-crude-oil`

[8]The actual consumption depends on the freight rate level (see e.g. Stopford, 2009) and the vessel's fuel efficiency (i.e. new vs. old ships).

sumed to be fixed at the amounts presented in the table 4.3. The address commission is applicable in these "ingoing" TC contracts as well, reducing the net TC expense by 3.75%. For a Panamax vessel for instance, the net TC expense is thus 12 900 USD/day.[9]

The port costs are assumed to be fairly market independent and amount to $48 000 per port call. Based on historical data from Bulk Invest, the average number of port calls is assumed to be 1.8 per voyage contract. This results in a total port cost of $20 476 800 per year.[10]

Having identified the main risk factors affecting the operating cash flow in Bulk Invest, the next stage of the Risk Management Process is to quantify the risk. The first step in this stage is to stochastically model the freight rates, which is the subject of the next chapter.

---

[9] 13 400 USD/day – 3.75% $\approx$ 12 900 USD/day.

[10] \$48 000 $\cdot$ 237 voyages $\cdot$ 1.8 calls = \$20 476 800.

5 Model Development

Modelling freight rates is the core of this thesis as it makes the foundation for the cash flow analysis. There are several ways to do the modelling, varying from a simple geometric Brownian motion (i.e. a random walk) and a mean reversion model (i.e. the Ornstein-Uhlenbeck process) to more sophisticated models including jumps (e.g based on an exponential Lévy process) and stochastic volatility. In any case, to develop a model for time series data, we first have to do a preliminary analysis of the data at hand.

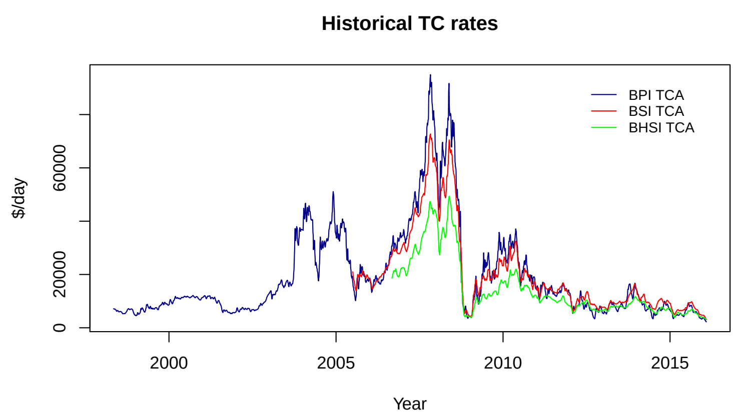

Figure 5.1: The historical daily time-charter rates for the three different types of vessels

5.1 Data Sources

The data used in this thesis was kindly provided by the research department at Fearnleys Shipping. The data consists of three indices of the daily average time-charter freight rates, which correspond to the time-charter equivalents of voyage charter rates.[1] This is basically the cost of hiring a vessel for one day, and hence, the data sets allow for modelling of spot freight rates. The indices are called BPI TCA, BSI TCA, BHSI TCA, which is the index for Panamax, Supramax and Handysize vessels, respectively. Figure 5.1 shows the historical daily time-charter rates for each index. It should be noted that the indices start at different dates, but they all end at February 2, 2016, which will act as the starting date of the simulated freight rates. In the following subsections, an econometric analysis is conducted on observations in the data set up to 2012. This part of the data is used as the *in-sample* part to fit the freight rate model and the rest of the data set is used for backtesting. See chapter 7 for more on the choice of sample parts.

Table 5.1: Start and end date, and the number of observations per index

Index	**Start date**	**End date**	**Observations**
BPI TCA	06-05-98	02-02-16	4 439
BSI TCA	01-07-05	02-02-16	2 644
BHSI TCA	04-09-06	02-02-16	2 350

5.1.1 Test of Stationarity

A formal test of stationarity in the time series is necessary before estimating the freight rate model. A process X_t is stationary if it is in a particular state of statistical equilibrium. When modelling any kind of econometric time series, stationarity is an essential part of the process. Non-stationarity in the time series used for modelling means that exogenous shocks to the data will not fade away with time. Furthermore, non-stationarity can lead to spurious regressions and wrongful conclusions, due to violations of the standard assumptions of an asymptotic analysis (Brooks, 2008, p. 318-320). On the other hand, stationary processes revert to a long-run equilibrium after being shocked, which implies that they have a finite memory. This fluctuation around the long-run equilibrium of stationary series

[1]Weekends and holidays are excluded, which corresponds to 250 trading days a year.

is also known as mean reversion (cf. section 4.3.5). A process is weakly stationary if it satisfies the following equations for $t = 1, 2, \ldots, \infty$:[2]

$$\begin{aligned} E(y_t) &= \mu \\ E(y_t - \mu)(y_t - \mu) &= \sigma^2 < \infty \\ E(y_{t_1} - \mu)(y_{t_2} - \mu) &= y_{t_2 - t_1} \quad \forall\, t_1, t_2 \end{aligned}$$

The first two equations state that the mean and variance are constant over time. Lastly, the covariance between two observations should only depend on the lag (l) between them, and not on the actual time at which it is computed. In other words; the covariance between any two points in time is the same, i.e. $Cov(y_t, y_{t-l}) = \gamma_l$ (Brooks, 2008).

To test for stationarity, the well-known test of a unit root proposed by Dickey and Fuller (1979) and later augmented by Said and Dickey (1984), is used. The Augmented Dickey-Fuller test (ADF) corrects the original Dickey-Fuller test, which only assumes the residual term u_t not to be autocorrelated and further assumes it to be white noise. The ADF test is formally written as

$$\Delta y_t = \psi y_{t-1} + \sum_{i=1}^{p} \alpha_i \Delta y_{t-i} + u_t \tag{5.1}$$

and have the same test statistics and critical values as the original Dickey-Fuller test.[3]

The p lags of Δy_t in the ADF will capture any dynamic structure in the dependent variable ensuring that the residual term is not autocorrelated. The statistical analysis is conducted in R, and the number of lags is chosen based on the plot of the partial autocorrelation function (PACF) for each data series. The plot of the autocorrelation function (ACF) and PACF of the log returns of BPI TCA is presented in section 5.1.4 (fig. 5.3) for illustration.[4]

[2] A *strictly* stationary series have all its moments of the probability distribution invariant over time (Brooks, 2008, p. 208).

[3] ADF test statistics: $\tau = \hat{\psi} / \hat{SE}(\widehat{\psi})$

[4] ACF is the coefficient of correlation between two values in a time series; $Corr(y_t, y_{t-l})$ (Brooks, 2008).

Table 5.2: Results from the Augmented Dickey-Fuller test

Unit root testing	**BPI TCA**	**BSI TCA**	**BHSI TCA**
		Series in log	
Test statistic (τ)	–0.45	–0.74	–0.86
P-value	0.52	0.40	0.35
Lags included	3	1	1
		Series in first difference	
Test statistic (τ)	–24.52	–12.36	–10.24
P-value	0.00	0.00	0.00
Lags included	2	1	1

The basic objective of the ADF test is to examine the null hypothesis that $\psi = 1$, which implies that the series has a unit root, against the one-sided alternative $\psi < 1$, meaning the series is stationary. The results from the ADF test are presented in table 5.2. All of the test statistics correspond to zero mean and no trend since none of the historical means of the indices are significantly different from zero (cf. table 5.3 in the next subsection), implying that there is no trend in the log returns of the indices. This is also consistent with a graphical examination of the plotted freight rate paths in fig. 5.1.

The ACF plots (figs. 5.3 and A.3 to A.5) and the test results show that none of the data series are stationary in log values (i.e. they contain a unit root). This coincides with the fact that a conditional expectation (i.e. time-charter rates, cf. section 4.3.3) is always non-stationary by the martingale property (Koekebakker et al., 2006, p. 23). However, all indices are first-difference stationary, which is shown in the bottom half of table 5.2. Essentially, this means that the log freight rates are non-stationary, but the log returns of them are stationary.

5.1.2 Historical Freight Rate Analysis

As explained in the previous subsection, the moments of a non-stationary time series are not stable over time and an analysis of them will therefore not be valid as time goes. This subsection describes an econometric analysis conducted on the log returns of each index (see table 5.3), as they were found to be stationary.

The means ($\bar{x}$) of BPI TCA, BSI TCA and BHSI TCA are -0.03%, -0.07% and -0.08%, which corresponds to an annual log return of -6.5%, -18.3% and -18.9%, respectively.[5] The means are statistically tested if they are different from zero (H$_0$: $\bar{x} = 0,$ H$_1$: $\bar{x} \neq 0$) and the corresponding p-values are presented in the second row.[6] The p-values clearly show that we fail to reject the H$_0$ hypothesis for all of the indices.

Table 5.3: Summary statistics of daily log returns

Summary statistics	**BPI TCA**	**BSI TCA**	**BHSI TCA**
Mean	–0.03%	–0.07%	–0.08%
P-value mean test	0.50	0.48	0.48
Maximum	13.8%	20.3%	9.2%
Minimum	–21.7%	–11.6%	–13.1%
Standard deviation	0.02	0.02	0.01
Skewness	–0.04	0.41	–1.22
Kurtosis	6.72	18.86	13.55
Jarque-Bera	11 048	39 311	21 206
P-value	0.000	0.000	0.000
Observations	4 438	2 643	2 349

The standard deviation presented in row six can be interpreted as the volatility of the returns. Annualized standard deviation yields 38%, 25% and 20% respectively for the indices. Thus, it seems that there is a positive relation between vessel size and volatility in freight rates. This is consistent with the findings of Alizadeh and Talley (2011) and is argued to be so because larger vessels operate in narrower markets.[7]

The third and fourth moment of each series are presented in row seven and eight of the table. The BHSI TCA is left-skewed, meaning this vessel class experiences a higher frequency of large losses than for instance the

[5]Assuming 250 trading days per year

[6]The p-value is calculated in Excel by the following formula: $NORM.DIST(\bar{x}; 0; s; TRUE)$, where s is the sample standard error.

[7]See e.g. the commodities shipped per vessel type in table 4.1.

BSI TCA (which is slightly right-skewed). The BPI TCA is more symmetric around the mean. All of the indices exhibit a high kurtosis which indicates fatter tails and a possibly higher peak than a normal distribution (Brooks, 2008, p. 162). Note that these historical moments are not constant over time. Figures A.13 and A.14 (in the appendix) show how the skewness and kurtosis have evolved over time for the three indices.

Both of the higher moments of the log returns indicate that the series might deviate from normality.[8] To formally test for deviation from normality the Jarque-Bera (1980) test is used.[9] The p-values show that we reject the null hypothesis of normality for all indices. This is consistent with empirical studies on freight rates (e.g. Kavussanos and Visvikis, 2004; Angelidis and Skiadopoulos, 2008). Quantile-by-quantile plots which support this finding can be found in section 5.1.4. The fact that the freight rates appear to deviate from normality implies that the freight risk cannot be captured by the standard deviation alone.[10] An *at-Risk* measure is more applicable since it can account for the tail risk (Angelidis and Skiadopoulos, 2008, p. 8).

5.1.3 Volatility Clustering

McNeil et al. (2015), list some stylized facts of financial time series based on empirical observations and inferences drawn from these observations. These facts apply to many time series of risk-factor changes such as log returns on equities, indices, exchange rates and commodity prices. Some of these facts are noticeable for freight rates as well: 1) Return series are not independent and identically distributed (*iid*) 2) Volatility appears to vary over time and 3) Extreme returns appear in clusters. In fig. 5.2 the returns of the BPI TCA index are plotted for illustration.[11]

[8] A normal distribution has skewness of zero and a kurtosis of three. Note that the table presents the *excess* kurtosis as calculated in Excel, which means that the actual kurtosis is +3 higher (i.e. BPI TCA's actual kurtosis is 9.72).

[9] This test follows χ^2_2. The 5% critical value is 5.99.

[10] The standard deviation can be seen as a measure of risk in accordance with the portfolio theory proposed by Markowitz (1952).

[11] A plot of log returns of BSI TCA and BHSI TCA can be found in fig. A.6.

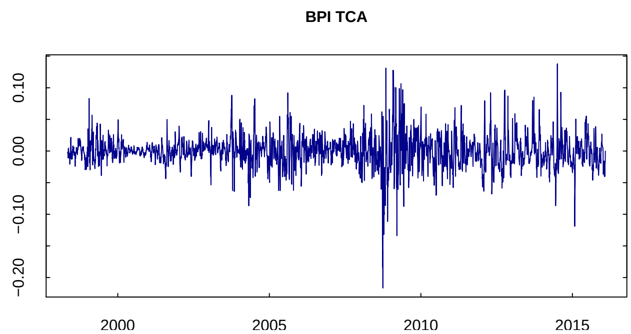

Figure 5.2: The log returns of BPI TCA

The figure clearly shows the existence of volatility clustering, which is the tendency for extreme returns to be followed by other extreme returns, although not necessarily with the same sign (McNeil et al., 2015, p. 80). This phenomenon is supported by empirical research (e.g. Ådland, 2000; Kavussanos, 1996) which have found that the volatility increases with the level of the freight rates. The reason for this is the nonlinear supply curve described in section 4.3.2, as explained by Alizadeh and Nomikos (2009). When the freight rate level is low, the market easily absorbs any changes and shocks to demand because there is an oversupply of tonnage. However, when the market is tight and freight rates are high (i.e. due to a shortage of tonnage and excessive demand), any changes or shocks to the market will cause sharp changes in freight rates and high volatility.

If we compare the return series with the blue line of fig. 5.1 (the historical time-charter average of BPI), we see that high volatility (i.e. large spikes) in the returns coincides with peaks in the freight rate level, confirming the theory of volatility clustering in shipping.[12]

A series which is *iid* and normally distributed will have no extreme spikes in the returns (cf. fig. A.7). Thus, the tails of the returns of the indices are fatter than of the normal distribution. One should be aware of this when simulating the freight rates; any stochastic model is as good (or bad) as its assumptions.[13]

[12]Note that the question whether freight rate levels are low or high depends mainly on the prevailing market equilibrium and marginal costs. The historical price levels are thus not a totally correct measure on the extremity of the freight rate level.

[13]See for instance Jorion, 2007, ch. 21, on the effect of model risk.

5.1.4 Examining the Returns

As opposed to market forecasting where one tries to predict the future in order to make a profit, market risk measurement seeks to prepare for the future and to prevent unexpected losses. When forecasting the market, the means of the distribution is the relevant factor. However, when estimating the potential risk by $CFaR$ and ES, we are interested in the (left) tail of the distributed returns, i.e. the downside risk. Therefore, it is necessary to conduct a closer examination of the properties of the return distribution before choosing a stochastic freight rate model. If for instance normality is (wrongfully) assumed when simulating, there will not be many extreme observations in the tail and the risk would be underestimated.

The starting point of the stochastic freight rate model is to investigate how the returns can be statistically represented. A common starting point is to check whether current return values depend upon the return values in previous periods. In short, this is done by fitting an autoregressive (AR) model to the data based on the number of lags in the PACF of the time series. For spacing purposes, only BPI TCA is discussed and illustrated in this section. Corresponding plots for the other two indices will be written in parentheses beside the respective figure reference and the figures can be found in the appendix. If more info is needed about the appendix plots, it is put in a footnote. The ACF and PACF plot of BPI TCA log returns (first differences) are presented in fig. 5.3.[14]

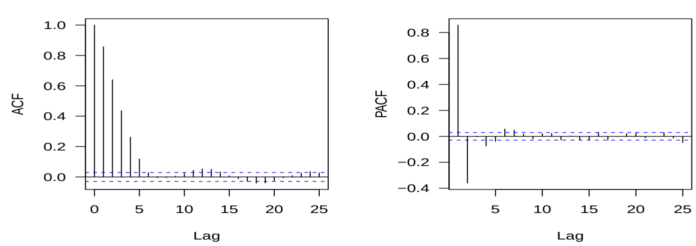

Figure 5.3: The ACF and PACF plot of the log returns (first differences) of BPI TCA

Note that there are two spikes outside the confidence band in the PACF

[14]Plots of BSI TCA and BHSI TCA can be found in figs. A.4 and A.5, together with the ACF/PACF plot of log BPI TCA values in fig. A.3.

for the first difference of BPI TCA. Strictly speaking, this resembles an autoregressive model of order two (an AR(2) model). However, to maintain the complexity of the simulations at a certain level and to keep the focus on other important aspects, I assume that the log returns of the BPI TCA can be represented by an AR(1) model. The first spike is also greater than the second, meaning that the first lag effect is much stronger than the second lag. For the other two indices, the PACFs clearly suggest an AR(1) model. Matematically, this model can be written as

$$AR(1)\ :\ Y_t = \beta Y_{t-1} + \varepsilon_t \tag{5.2}$$

When assuming any statistical model, it is important to check whether the errors from the model have any predictive information. That is, all of the explanatory information should be included in the predictor variable (i.e. β) and not in the error term ϵ_t.[15] Plotting the residuals of the model against the fitted values will reveal any pattern or dependency inherent in predictive error terms. As can be seen in fig. 5.4, the error terms of the AR(1) model are fairly stochastic and symmetrically spread around the dotted line (marking $\hat{\epsilon}_t = 0$), and thus fulfill the assumption of no prediction power in $\hat{\epsilon_t}$. Similar result yields for the two other indices as well (cf. fig. A.8).

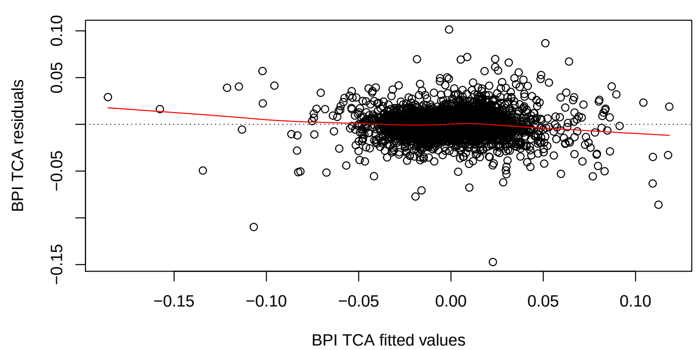

Figure 5.4: The residuals of the AR(1) model plotted against the fitted values

[15]This is essentially two assumptions of the ordinary least squares regression and can be mathematically written as $cov(Y_{t-1}, \epsilon_t) = 0$ and $E(\epsilon_t|Y_{t-1}) = 0$.

In the historical analysis in section 5.1.1, I found that the distributions of returns for all of the indices exhibit high kurtosis (i.e. fat tails). To illustrate this (tail) deviation from normality, a quantile-by-quantile plot of BPI TCA returns is plotted in the left pane of fig. 5.5 (fig. A.9). The returns are plotted as blue circles, and the red line shows the theoretical returns from a normal distribution. The deviations from the red line in both ends illustrate the fat tails of the empirical distribution. For comparison, the same plot is reproduced only fitted against the quantiles of the theoretical *Student's t* distribution (right pane), which is a distribution that exhibits fatter tails than the normal one.[16]

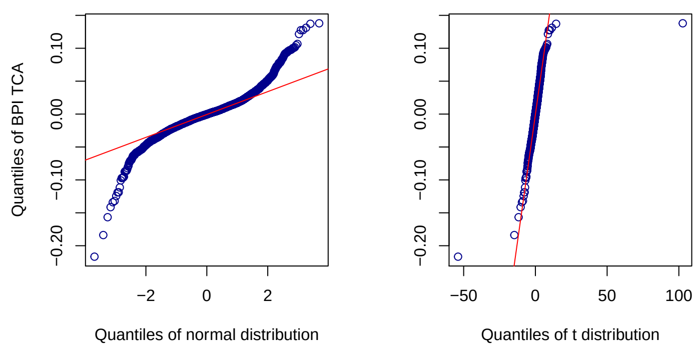

Figure 5.5: Quantile-by-quantile plots showing the returns of BPI TCA versus the theoretical returns of the normal distribution (left) and the *Student's t* distribution with 3.29 degrees of freedom (right)

The graphical inspection of the log returns in section 5.1.3 indicated that the freight rate volatility might not be constant. Therefore, it would be interesting to examine whether there is any long-term memory in the time series of the log returns in order to validate this indication. This is done by looking at the ACF plot of the squared log returns. Figure 5.6 (fig. A.10) plots the autocorrelation for BPI TCA up to 200 lags and shows how the squared returns exhibit a long-range dependency. This is consistent with volatility clustering and confirms the visual inspection in section 5.1.3.

[16]The returns of BPI TCA are fitted against the *Student's t* distributions using the fitdistr function in R, resulting in 3.29 degrees of freedom.

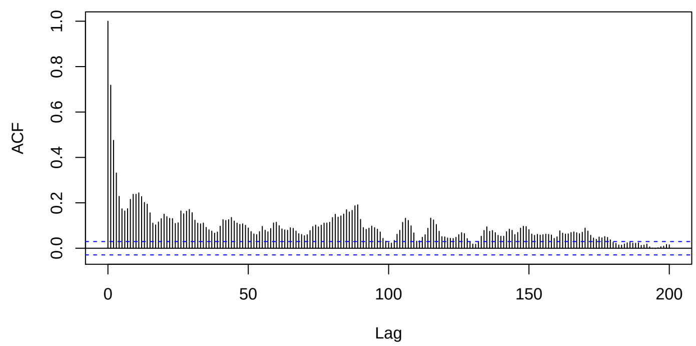

Figure 5.6: Autocorrelation function for the squared log returns for BPI TCA.

Lastly, the empirical density of the log returns is plotted along a normal distribution fitted to the data.This is done to further support the use of stochastic volatility in the freight rate model, and to illustrate the high kurtosis of the index. In fig. 5.7, the empirical density of BPI TCA (black line) is plotted against a standard frequency axis. The high kurtosis is clearly present in the returns of the time-charter index, taking form as a leptokurtic distribution with much higher concentration of probability in the center versus the normal distribution (red line).

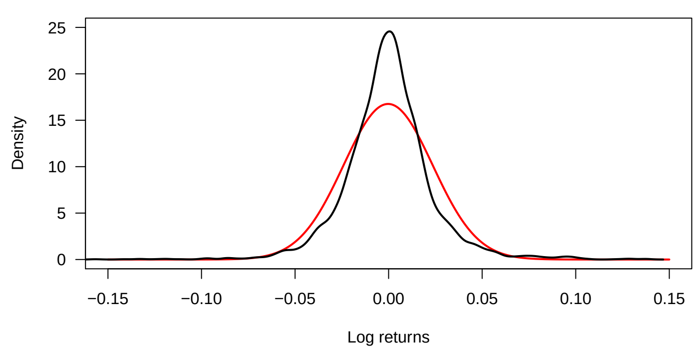

Figure 5.7: Empirical density of the BPI TCA log returns (black line) along with a fitted normal distribution (red line)

In order to zoom in on the tails of the empirical distribution, the same plot is reproduced only this time with a logarithmic frequency axis (bottom panel).[17] Looking at the tails, it is clear that the index has bigger deviations than explained by the normal distribution. This result also points at the presence of stochastic volatility in the data (Benth and Koekebakker, 2015). The results for the two other indices are consistent with this conclusion as well (see appendix A.9).

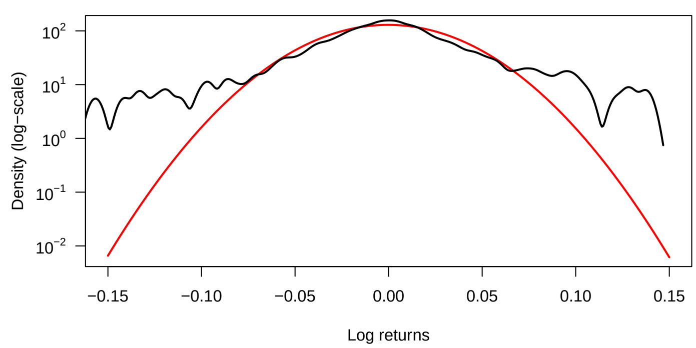

Figure 5.8: Log-scaled empirical density of the BPI TCA log returns (black line) and a fitted normal distribution (red line)

After examining the distribution of the returns thoroughly, the conclusion is that it is necessary to include stochastic volatility in the freight rate model in order simulate the leptokurtic distributions of the returns. In the next section, the building blocks of the freight rate model are presented.

[17]The empirical densities are based on kernel smoothing and the fitted normal distributions are based on the means and standard deviations from table 5.3.

5.2 Stochastic Processes

A formal defintion of a stochastic process is provided by Castañeda et al. (2014):

Definition 5.2.1. *A real stochastic process is a collection of random variables* $\{S_t;\, t \in T\}$ *defined on a common probability space* $(\Omega, \mathcal{F}, \mathcal{P})$ *with values in* $\mathbb{R}$.[18] *T is called the parametric space of the process and is usually a subset of* $\mathbb{R}$. *The set of values that the random variable* S_t *can take is called the state space of the process, denoted by* C.

The mapping defined for each fixed element of the state of nature $(\omega \in \Omega)$,

$$\begin{aligned} S_{(\omega)} \quad : \quad & T \to C \\ & t \mapsto S_t(\omega) \end{aligned}$$

is called the sample path of the process over time. In the following subsections an important stochastic process called the geometric Brownian motion is presented, as well as a popular model for stochastic volatility, namely Heston's model from 1993. Lastly, copulas are introduced as a tool for simulating interdependent stochastic processes, which is the case of the dry bulk indices.

5.2.1 Brownian Motions

Originally Brownian motions were the random motion of pollen grains floating in water discovered by Robert Brown in the early nineteenth century (Ibe, 2013). The Brownian motion lends itself to several other fields, and the geometric version is extensively used in finance (e.g. in the Black-Scholes model for options).[19] The change in an underlying asset (S) assumed to follow a geometric Brownian motion is written as

$$\mathrm{d}S_t = \mu S_t \mathrm{d}t + \sigma S_t \mathrm{d}W_t \tag{5.3}$$

[18]Where Ω is the state of nature with generic elements ω, $\mathcal{F}$ is a set of all events, $\mathcal{P}$ is the probability measure and $\mathbb{R}$ is a set of real numbers.

[19]In the geometric Brownian motion, the logarithm of the underlying quantity follows a Brownian motion with drift.

where μ is the drift parameter (e.g. the asset return) and $\mathrm{d}W_t$ is a standard Brownian motion (also known as a Wiener process[20]) on the probability space $(\Omega, \mathcal{F}, \mathcal{P})$, driven by the (constant) volatility σ. For $0 \leq i < t$, the increments of the Wiener process $\mathrm{d}W = W_t - W_i$ is assumed to be normally distributed with mean 0 and variance $\sigma_W^2 = \sigma_{t-i}^2$. That is, $\mathrm{d}W \sim N\left(0, \sigma_{t-i}^2\right)$.

A geometric Brownian motion (GBM) is the scaling limit of the discrete-time *random walk* (Ibe, 2013). A random walk is a non-stationary process which states that the asset value of tomorrow is best predicted as today's value plus a random shock or error term. This coincides with the efficient market hypothesis assumed to prevail financial markets (Fama and Malkiel, 1970). This hypothesis states that all information regarding a financial asset is incorporated in the current price. Consequently, the asset is assumed to be perfectly priced by the market and any changes to the price are due to new (unexpected) information. This information (i.e. the error term of the random walk) is assumed to be random as we have no prescient knowledge about it.

As opposed to the (stationary) mean reverting Ornstein-Uhlenbeck process, the non-stationary random walk can wander freely during simulations as it is not drawn towards a long-term mean.

5.2.2 Stochastic Volatility

Most financial returns series (including freight rates as shown in table 5.3) are skewed and show a higher degree of kurtosis than assumed in the normal distribution. This results in fat tails and is caused by time varying (i.e. stochastic) volatility and volatility clustering (as described in section 5.1.3). Standard models based on Brownian motions (e.g. the Black-Scholes model) assume constant volatility and thus neglecting the higher kurtosis. One solution to this problem is to let the volatility vary stochastically through stochastic volatility models (SVM).

In general, stochastic volatility models can be analyzed within the following framework (Cerrato, 2012):

[20]Norbert Wiener established the modern mathematical framework of what today is known as the Brownian motion random process (Ibe, 2013).

$$\mathrm{d}S_t = \mu S_t \mathrm{d}t + \sigma_t S_t \mathrm{d}W_{1,t} \tag{5.4}$$
$$\sigma_t = f(V_t) \tag{5.5}$$
$$\mathrm{d}V_t = a\,(t, V_t)\,\mathrm{d}t + b\,(t, V_t)\,\mathrm{d}W_{2,t} \tag{5.6}$$
$$\mathrm{d}W_{1,t}\mathrm{d}W_{2,t} = \rho \mathrm{d}t \tag{5.7}$$

where $W_{1,t}$ and $W_{2,t}$ are two correlated Brownian motions (satisfying eq. (5.7)) with correlation coefficient ρ, the variance process V and time varying parameters a, b. Using Cholesky decomposition, equation eq. (5.7) can be rewritten as

$$W_{1,t} = \rho W_{2,t} + \sqrt{1-\rho^2} Z_t \tag{5.8}$$

where Z_t is a standard normal variable. The correlation coefficient ρ controls the skewness of the density of the underlying variable. A positive correlation implies a rise in variance when the price of the underlying rises. The resulting effect is a fatter right tail and a thinner left tail. The opposite is true for a negative correlation (Rouah, 2013).

There are several different types of stochastic volatility models, but in the following, the popular Heston model from 1993 is presented. This variance process follows a square root diffusion process, given by

$$\mathrm{d}S_t = \mu S_t \mathrm{d}t + \sqrt{V_t} S_t \mathrm{d}W_{1,t} \tag{5.9}$$
$$\mathrm{d}V_t = \kappa\left(\bar{V} - V_t\right)\mathrm{d}t + \xi\sqrt{V_t}\mathrm{d}W_{2,t} \tag{5.10}$$
$$\mathrm{d}W_{1,t}\mathrm{d}W_{2,t} = \rho \mathrm{d}t \tag{5.11}$$

The asset (i.e. freight rate) dynamics in eq. (5.9) follows a geometric Brownian motion, but the volatility dynamics in eq. (5.10) is described by a mean reversion process where κ is the speed of mean reversion and $\bar{V}$ is the long-term variance at which the variance reverts to. ξ is a coefficient which determines the fluctuation in the variance (the standard deviation of the variance). High values of ξ disperse the variance process and thus increase the kurtosis and the fatness of the tails. Low values provide opposite results.

5.2.3 Multivariate Distributions

The market cycles and spot rate formation described in sections 4.3.1 and 4.3.2 are valid for all vessel classes in shipping. By and large, all vessels in the dry bulk sector are influenced by the same exogenous factors (e.g. world economic growth) as they ship homogeneous products (i.e. dry bulk cargo). As such, any shocks to the demand of dry bulk will have much of the same effect on each freight rate index. This implies a certain correlation between the indices. Note that there are some differences in these simultaneous effects, due to the difference in the flexibility of larger versus smaller ships (e.g. Panamax vs. Handysize), cf. table 4.1. This means that the indices are highly, but not perfectly correlated (Kavussanos, 1996). Figure 5.9 shows that the Panamax class is slightly more correlated with the Supramax than the Handysize, and that the two smallest vessel classes (Supramax and Handysize) exhibit the highest historical correlation.

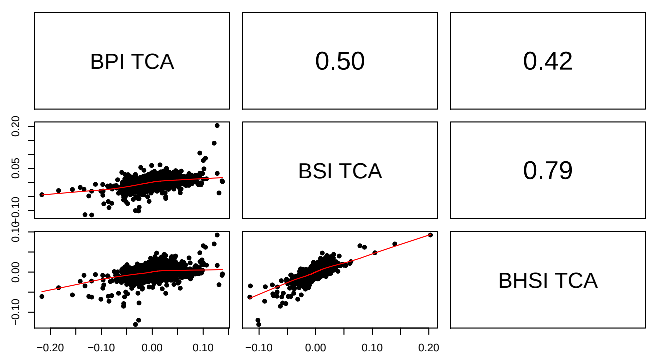

Figure 5.9: Upper right panel shows the correlation matrix of the indices. Lower left panel visualizes the correlation by corresponding pair plots of the returns

A shipping company engaged in multiple segments will thus be jointly exposed to the freight rates. To correctly incorporate the dependency between freight rates in the simulations, the model should include the probability distribution of all segments together. This is done by the use of copulas. A copula provides a way of isolating the description of the dependence structure in a random vector of risk factors.[21] A d-dimensional copula can be defined as a distribution function on $[0,1]^d$ with standard uniform marginal distributions (McNeil et al., 2015, p. 221). By the theorem of Sklar (1959), the importance of copulas in the study of multivariate distribution functions is summarized:

Theorem 5.2.1. (Sklar 1959) *Let F be a joint distribution function with margins $F_1, \ldots, F_d$. Then there exists a d-dimensional copula C: $[0,1]^d \to [0,1]$ such that, for all $x_1, \ldots, x_d$ in $\mathbb{R}^d = [-\infty, \infty]$,*

$$F(x_1, \ldots, x_d) = C\left(F_1\left(x_1\right), \ldots, F_d\left(x_d\right)\right). \tag{5.12}$$

This theorem shows, firstly, that all multivariate distribution functions contain copulas and, secondly, that copulas may be used in conjunction with univariate distribution functions to construct multivariate ones. This is important because it allows simulating multiple freight rate indices while maintaining the interdependency.

There exist several types of copulas that can be used to model interdependency. One major set of copulas is the class of elliptical copulas. In this class we have the Gaussian and *Student's t* copula, based on the well-known distributions of the same names. As the *Student's t* distribution exhibit fatter tails and greater dependences in the tails, this copula is utilized on the Wiener processes driving the freight rates in the simulations, to model the interdependency between the freight rate indices.[22,23] More precisely; the *static* version of the *Student's t* copula is used in the simulations, in line with the findings of Merikas et al. (2013) presented in section 2.1, and for the sake of computational simplicity.

Define the joint *Student's t* distribution as $\mathbf{T}_n\left(\epsilon_1, \ldots, \epsilon_n; \mathbf{R}, \nu\right)$ linking together a set of n variables with correlation represented by the n-dimensional matrix $\mathbf{R}$, and denote by $\mathbf{T}_\nu(x)$ the univariate *Student's t* distributions. The multivariate *Student's t* copula is then defined as (Cherubini et al., 2011):

[21] In a sense, every joint distribution function for a random vector of risk factors implicitly contains both a description of the marginal behavior of individual risk factors and a description of their dependence structure (McNeil et al., 2015, p. 220).

[22] *Fat tails* is a phenomenon in a univariate setting, where the probability of extreme events is larger than expected according to the normal distribution. *Tail dependence* is the concept of extreme events occurring jointly in two different markets, i.e. in a multidimensional setting (Cherubini et al., 2011).

[23] The choice of copula is also backed by the VineCopula package in R, which bases the choice on the Akaike information criterion (AIC).

$$C(u_1, \ldots, u_n) \equiv \mathbf{T}_n \left(\mathbf{T}_\nu^{-1}(u_1), \ldots, \mathbf{T}_\nu^{-1}(u_n); \mathbf{R}, \nu \right) \tag{5.13}$$

Cherubini et al. (2004) provide an algorithm which can be used to generate random variates u_n from the n-copula $T_{\mathbf{R},\nu}$:

- Find the Cholesky decomposition A of $\mathbf{R}$
- Simulate n *iid* variates $\mathbf{z} = (z_1, z_2, \ldots, z_n)'$ from $N(0,1)$
- Simulate a random variate s from χ^2_ν independent of $\mathbf{z}$
- Set $\mathbf{y} = A\mathbf{z}$
- Set $\mathbf{x} = \sqrt{(\nu/s)}\mathbf{y}$
- Set $u_i = T_\nu(x_i)$ with $i = 1, 2, \ldots, n$
- Return $(u_1, \ldots, u_n)' = (F_1(t_1), \ldots, F_n(t_n))'$ where F_i denotes the ith margin

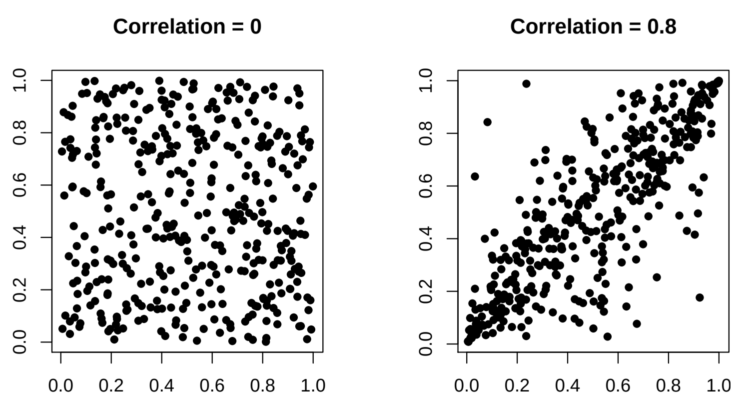

Figure 5.10: Simulated variates by the *Student's t* copula. Left panel shows 400 variates with zero correlation. Right panel shows 400 variates with 0.8 correlation. Both simulations are done with four degrees of freedom

The effect of using a copula when simulating is illustrated in fig. 5.10. The left panel shows how the variates lack interdependency when the correlation between them is set to zero. The right panel clearly shows how the correlation between the variates are maintained when simulating (i.e. low/high values in one of the series correspond to low/high values in the other series). The correlated case also illustrates how the *Student's t* copula includes some extreme observations (the points in the top left and bottom right corner).

5.2.4 Correlated Brownian Motions

In this subsection, the correlation between the two Wiener processes in Heston's stochastic volatility model is analyzed with regards to the three indices. As mentioned section 5.2.2, the correlation ρ in eq. (5.11) affects the skewness of the time series. But how to determine what the value of ρ is? There is several ways to do this and in the following, three methods are described.

The ρ is essentially the correlation between the dynamics of the returns of the underlying (eq. (5.9)) and its variance dynamics (eq. (5.10)). To estimate ρ, it is hence possible to look at the correlation between the historical squared returns and their corresponding variance. The latter part is suggested as:

- Weekly rolling variance
- Change in weekly rolling variance
- Weekly non-overlapping rolling variance
- Change in non-overlapping variance

These methods are used on the three freight rate indices, and the resulting correlations between the variances and the corresponding squared (log) returns are presented in table 5.4.

Table 5.4: Empirical correlations between returns and variance

Method	**BPI TCA**	**BSI TCA**	**BHSI TCA**
Weekly rolling	0.48	0.69	0.70
Change in weekly rolling	–0.12	0.03	0.04
Non-overlapping	0.39	0.64	0.77
Change in non-overlapping	–0.03	0.11	0.35

Both the rolling window methods (overlapping and non-overlapping) yield a relatively high positive correlation between returns and variance, while the "changing" methods yield correlations closer to zero. The implication of a high correlation is that when the freight rate increases, the variance increases as well. This is consistent with the convex supply curve of freight rates and the relation between high freight rates and volatility, as described in sections 4.3.2 and 5.1.3. In that sense, the weekly rolling and non-overlapping methods exhibit correlations which fit the theory better. As they both give relatively similar results, the average correlations between them are used as inputs to the freight rate model. The average correlations are 0.44, 0.67 and 0.74 for BPI TCA, BSI TCA and BHSI TCA, respectively.

5.3 Modelling

In this section, the final stochastic freight rate model is proposed and explained, as well as a simple model for calculating the cash flows corresponding to the simulated freight rate paths. In addition, an explanation of the estimation of the respective parameters is provided. Finally, the assumptions and limitations of the two models are listed and discussed.

5.3.1 Stochastic Model for Freight Rates

The historical data analysis found all the indices to only be first-difference stationary (recall that the log freight rates contained unit roots), cf. section 5.1.1. In addition, since mean reversion is most often not empirically observable (cf. section 4.3.5), and at least not in the short term, the spot freight rate dynamics are proposed to follow a geometric Brownian motion. Let the spot freight rate S be a function of its logarithmic value X_t.

$$S_t = \exp(X_t) \tag{5.14}$$

Then

$$\mathrm{d}X_t = \mu \mathrm{d}t + \sqrt{V_t}\mathrm{d}W_{1,t} \tag{5.15}$$

Here, $\mathrm{d}t$ is an infinitesimal fraction of time equal to 1/250 = 0.004. The driving force of the freight rate paths is $\mathrm{d}W_{1,t}$ which is a Wiener process, and parameter μ is a drift term. In order to incorporate fat tails and volatility clustering in the log returns (cf. section 5.1.4), the variance V is in itself assumed to be a stochastic process

$$\mathrm{d}V_t = \kappa\left(\bar{V} - V_t\right)\mathrm{d}t + \xi\sqrt{V_t}\mathrm{d}W_{2,t} \tag{5.16}$$

where $E^{\mathbb{P}}\left[\mathrm{d}W_{1,t}, \mathrm{d}W_{2,t}\right] = \rho\mathrm{d}t$.[24] Equation 5.16 is a stochastic volatility model by Heston as described in section 5.2.2, meaning that the time varying variance V is explained through a mean reverting process.

To ensure interdependency between the simulated paths of the freight rate indices, the Wiener processes driving the log freight rate dynamics of each index ($\mathrm{d}W_{1,t}$) are coupled together using a *Student's t* copula as explained in section 5.2.3. The Wiener processes driving the variance ($\mathrm{d}W_{2,t}$) are calculated based on Cholesky decomposition as in eq. (5.8).

The starting variance V_1 is set by estimating the exponentially weighted moving average (EWMA) on a rolling window of weekly variance.[25] This is a method which, unlike the simple historical moving average, allows more recent observations to have stronger impact on the estimation than older data points (Brooks, 2008). The weighting of recent versus older data points is set by a decay factor λ. According to the recommendations of RiskMetrics, this decay factor is set to 0.94 in the estimation of V_1 (Zumbach, 2007).

In the proposed freight rate model there are a number of coefficients needed to be estimated. To estimate the drift term in eq. (5.15) the model is discretized (i.e. as a random walk) into

$$X_t - X_{t-1} = \mu + \epsilon_t \tag{5.17}$$

The drift term then follows as

$$\hat{\mu} = \frac{1}{T}\sum_{t=1}^{T}\left(X_t - X_{t-1}\right) \tag{5.18}$$

The volatility in eq. (5.15) is modelled through eq. (5.16), where several coefficients need estimation. There are multiple ways to estimate the parameters of Heston's model, depending on the level of precision pursued.

[24] $\mathbb{P}$ is a historical probability measure (also called real world or physical measure). Thus, this is just another way of writing eq. (5.11) from section 5.2.2.

[25] EWMA: $\hat{\sigma}_t^2 = (1-\lambda)r^2 + \lambda\hat{\sigma}_{t-1}^2$, where σ_t^2 is the estimate of the variance for period t, r^2 is the squared return at time t and λ is a decay factor.

In order to maintain certain simplicity in this thesis a somewhat rough but yet reasonable estimation method is used. The correlation between $\mathrm{d}W_{1,t}$ and $\mathrm{d}W_{2,t}$ is estimated as described in section 5.2.4. The variance in eq. (5.16) can be written as

$$V_t = (X_t - X_{t-1} - \hat{\mu})^2 \tag{5.19}$$

Changes in the variance from time t to $t+1$ can then be written on the form

$$\Delta V_t = q + \beta V_{t-1} + u_t \tag{5.20}$$

and the coefficients $\hat{q}$ and $\hat{\beta}$ are estimated by ordinary least squares regression. These coefficients are then used to estimate the first two parameters of the stochastic volatility model in eq. (5.16):

$$\hat{\kappa} = -\hat{\beta} \tag{5.21}$$

$$\hat{\kappa}\bar{V} = \hat{q} \Rightarrow \bar{V} = \frac{\hat{q}}{\hat{\kappa}} \tag{5.22}$$

The coefficient representing the standard deviation of the variance in the stochastic volatility model (ξ) must be estimated through yet another least squares regression. The squared error terms u_t in eq. (5.20) can be written as

$$\hat{u}_t^2 = \delta + \gamma V_t + w_t \tag{5.23}$$

Coefficients $\hat{\delta}$ and $\hat{\gamma}$ are then estimated, and we finally have that

$$\hat{\xi} = \sqrt{\hat{\gamma}} \tag{5.24}$$

5.3.2 Cash Flow Model

Assuming freight rates to be the source of income for the spot business of Bulk Invest, the cash flow π at time t can be calculated as

$$\pi_t^{spot} = \mathbb{F}\left(S_t - \mathbb{C}_t\right) \tag{5.25}$$

where $\mathbb{F}$ is the fleet presented in table 4.2, S_t is the spot freight rate simulated by the stochastic model presented in section 5.3.1, and $\mathbb{C}_t$ is the total costs presented in table 4.3. For the time-charter business the spot freight rate is fixed every three month, six month, and one year (depending on the contract length) all else being equal:

$$\pi_t^{TC} = \mathbb{F}\left(S_{TC} - \mathbb{C}_t\right) \tag{5.26}$$

5.3.3 Assumptions and Limitations

The presented freight rate- and cash flow models are by no means the most detailed and precise models nor are they the only way to model the respective dynamics. However, for the scope of this thesis they provide a reasonable and informative basis to answer the research questions. In these models lay certain assumptions and limitations listed and discussed below.

Freight rate model:

- Type of model – The geometric Brownian motion may not be the theoretically optimal process for the spot freight rate dynamics, but it coincides with empirical research and it is fully adequate at the horizon of the simulations as explained in section 4.3.5.
- Dependency – The correlation between each index is modelled by a *Student's t* copula on the Wiener processes driving the log freight rate paths, cf. section 5.2.3. A Gaussian copula could just as easily have been used, but the *Student's t* copula fits better with the three indices.
- Static copula - The copula used in simulations is assumed to be static. A dynamic copula would capture any time varying change of dependency and market co-movement, but for the sake of simplicity, a static version of the copula is chosen.

- Trading days vs. Regular days – Even though there are, on average, 250 listed freight rates per year (i.e. 250 trading days) on the indices, the simulations assume 365 days per year. This is done in order to simplify the simulations and it coincides with the number of shipping days per year.

Cash flow model:

- Horizon – The horizon is assumed to be approximately two years; ranging from February 2016 to year-end 2017. This coincides with the available bunker forward curve.
- Generalization of costs – Each cost type can be split into more detailed outlines (e.g. depending on region or type of vessel), but they are generalized in order to limit the computational complexity.
- Canal costs – It is not accounted for canal costs because it is hard to model them appropriately without focusing on specific routes. The proposed model focuses on a more aggregated cash flow (cf. the indices for time-charter averages are used for freight rate/revenue modelling).
- Time-charter – Due to the difficulty in getting data, the TC business is modelled through the spot freight rate even though there exists a separate TC index for longer contracts. This index is highly correlated with the spot rate but it exhibits less volatility because it is the conditional expectation of the future spot rates.
- Capital costs – Because Bulk Invest charters in (almost) all their vessels, the capital costs as well as operating and periodic maintenance costs are allocated to the original owner of the ship. These costs are therefore not included in the cash flow model.

6 Empirical Results

6.1 Simulated Freight Rates

In the simulations, the last day of the data set: February 2, 2016, is set as the starting day. The simulated freight rate paths stretch throughout 2017 when the available bunker forward curve ends (December 31, 2017), making each path 698 days. Ten-thousand iterations are performed with the following data inputs.[1]

Table 6.1: Data inputs of the simulations

Input	BPI TCA	BSI TCA	BHSI TCA
X_1	7.7231	8.0143	8.0513
$\sqrt{V_1}$	0.0130	0.0074	0.0033
ρ	0.4377	0.6646	0.7357
μ	−0.0003	−0.0007	−0.0008
κ	0.2525	0.3987	0.2444
$\bar{V}$	0.0006	0.0003	0.0002
ξ	0.0448	0.0401	0.0349

X_1 is the log value of the respective indices at February 2, 2016 and $\sqrt{V_1}$ is the corresponding weekly rolling standard deviation estimated by EWMA.

The simulations are done by first fitting the *Student's t* copula to the log returns, and then generating a set of 698 Wiener processes which have

[1]The number of iteration is arbitrary, but a higher number increase the accuracy of the simulation. However, the computational power needed increases with the number of iterations as well, creating a natural limit to the total number of iterations.

approximately the same correlation as the three indices. Then, a vector of the corresponding Wiener processes in the variance dynamics (i.e. $\mathrm{d}W_2$ in eq. (5.16)) is calculated using eq. (5.8).

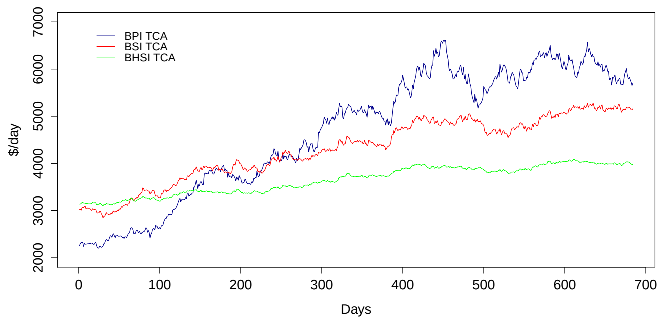

Figure 6.1: One sample freight rate path of each index showing the correlating effect of the copula used in simulations

Having generated the necessary Wiener processes $\mathrm{d}W_1$ and $\mathrm{d}W_2$ with the appropriate interdependencies, the next stage is to simulate the variance dynamics in eq. (5.16) and then the corresponding log freight rate dynamics in eq. (5.15). Lastly, truncated price paths for the time-charter business are estimated. Figure 6.1 shows one sample spot rate path simulated for each index. Due to the copula, the freight rate paths are correlated with each other, just as we expect to see in real life (cf. section 5.2.3).

The average correlations between all simulated paths are presented in the bottom left cells under the 1-diagonal of table 6.2. For comparison, the corresponding historical correlations (cf. fig. 5.9) are presented in the top right cells.

Table 6.2: Comparison of simulated and historical correlations between the indices.

	Historical		
	1	0.50	0.42
Simulated	0.47	1	0.79
	0.38	0.72	1

From left (top) to right (bottom), the numbers represent BPI TCA, BSI TCA, and BHSI TCA.

6.2 Simulated Returns

It would be interesting to analyze the log returns of the simulated freight rates in order to check whether they match the empirical properties found in chapter 5. The simulated properties are presented in table 6.3.

Table 6.3: Statistical summary of simulated returns

Moment	**BPI TCA**	**BSI TCA**	**BHSI TCA**
Mean	0.00%	0.00%	0.00%
Annualized	–0.53%	–0.11%	–0.12%
Std.deviation	2.03%	1.64%	1.36%
Annualized	32.13%	25.93%	21.42%
Skewness	0.14	0.19	0.23
Kurtosis	5.53	6.26	6.65

The annualized average means of the simulated log returns are -0.53%, -0.11% and -0.12% for BPI TCA, BSI TCA and BHSI TCA, respectively. This does not coincide with the empirical means from table 5.3. The reason is that the expected value of eq. (5.15) is just the drift term, because the increments of the Wiener process are assumed to be normally distributed; $\mathrm{d}W \sim N\left(0, \sigma_{t-i}^2\right)$, cf. section 5.2.1. As the drift terms μ in the freight rate model is essentially 0 for each index, the simulated means converge to zero with the number of iterations. It is important to remember that the simulations are by no means a forecast of the future freight rates. In that manner, close to zero drift in the model is an appropriate estimate of μ. Setting the drift parameter greater or less than zero would yield an expected return greater or less than zero (since $E\left[r\right] = \mu \mathrm{d}t$) and this is juxtaposed with forecasting the future freight rates to either increase or decrease.

The average standard deviations of the simulated returns on the TC indices are more in line with the empirical ones. It is good to see that the model returns a positive correlation between vessel size and volatility just as theory and empirical research suggest (see Kavussanos, 1996).[2] In

[2]The correlation between vessel size and volatility is due to the differences in flexibility of larger vs. smaller vessels, as explained in section 5.2.3.

addition, the model is capable of incorporating volatility clustering in the daily returns. This is illustrated in fig. 6.2 where one sample path of the simulated returns of the BPI TCA is plotted.

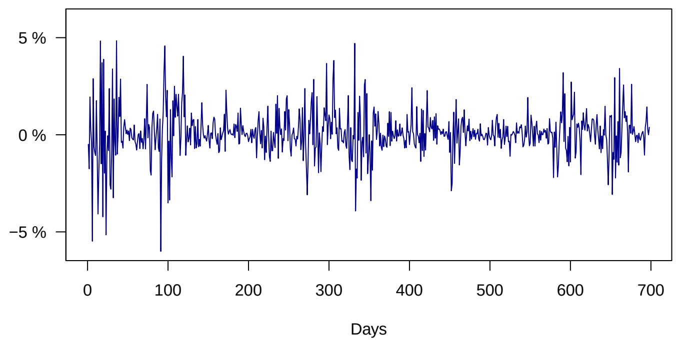

Figure 6.2: One sample path of the simulated daily BPI TCA returns. The figure shows volatility clustering captured by the freight rate model

The average resulting skewness is slightly positive for all three indices. This is due to the positive estimate of ρ as the correlation between the Wiener processes in the freight rate model, cf. section 5.2.4. The historical skewness showed greater deviation from symmetry for both BSI TCA and BHSI TCA. However, when examining the cumulative path of the historical skewness, all of the indices tend to slowly revert to zero after being shocked (see fig. A.13). Therefore, in the long-term perspective a simulated skewness of around zero need not be completely wrong.

The simulated average excess kurtosis is lower compared to the historical kurtosis for both BSI TCA and BHSI TCA. The BPI TCA kurtosis is more consistent with historical data. Again, looking at the cumulative historical kurtosis, reveals that it is not static, cf. fig. A.14. It is clear that the Financial Crisis of 2008 had a huge impact on the kurtosis for all indices. As the kurtosis is dynamic and changes over time, trying to perfectly match the historical kurtosis when simulating is not necessary a primary goal. The important aspect is that the model captures some excess kurtosis, which would not have been present if the dynamics followed a pure geometric Brownian motion (due to the constant volatility assumed in the GBM, cf. section 5.2.1).

One reason why the simulated returns exhibit excess kurtosis is the *Student's t* copula used in the model. The number of degrees of freedom (DoF) affects the simulated kurtosis because altering the DoF makes the copula draw from a distribution with a different kurtosis. Higher DoF reduces the kurtosis as the *Student's t* distribution converges to the Gaussian one. The copula used in the simulations was estimated to have 7.65 degrees of freedom.[3]

To conclude, the (average) log returns of the simulated freight rate paths exhibit properties which fits the historical data well. The model is able to capture high excess kurtosis and volatility clustering, which was the goal of including stochastic volatility in the freight rate dynamics. In chapter 8, some of the parameters of the freight rate model are tested to examine what effect they have on the simulated returns.

6.3 Cash Flow Effects

Unfortunately, the cash flows from the simulated shipping activity are dismal news, both for the spot- and time-charter business. In total, there are no years with positive cash flows. This is due to the low initial spot freight rate for each index and the high costs incurred by Bulk Invest.[4] In order to break even, the freight rates need to exceed approximately (in $/day) 16 000 for Panamax, 15 130 for Supramax and 12 530 for Handysize vessels.[5] The problem is that Bulk Invest has fixed its vessels on long-time charter contracts at a very high rate compared to the prevailing market. This is also the reason why the company was forced to declare bankruptcy. The cash flow effects would definitely have been different for other companies with another cost structure (e.g. owning the vessels instead of chartering them in), but with the freight market as battered as it is today, any dry bulk shipping company would feel its cash reserves desiccate.

Estimating the Cash Flow-at-Risk when the simulated cash flows only yield negative values may seem wasteful as the company loses money

[3]Fitted by the Copula package in R.

[4]The initial spot freight rates are (in $/day) 2 260, 3 024 and 3 138 for BPI TCA, BSI TCA and BHSI TCA, respectively.

[5]Because the bunker price is in contango (i.e. increasing over the horizon), the break even spot rate is calculated by the average bunker price over the horizon, which is 193.7 $/mt.

anyhow. But to answer the research questions, the takeaways from this thesis should focus on the potential cash flow losses relative to the means and the effects of hedging by time-charter contracts, instead of the actual losses incurred by Bulk Invest.

Total cash flows over the simulated horizon are calculated according to the simple cash flow models presented in section 5.3.2. Figure 6.3 shows the density of the total cash flows for each contract type in million US Dollars (mUSD). It should be noted that although there is evidently some excess kurtosis in the densities, the longest tails are on the right side. This is because the model incorporates that freight rates can never be less than zero (due to simulations of the *logarithmic* freight rates). This implies that the freight rates and hence cash flows, are log-normally distributed.[6]

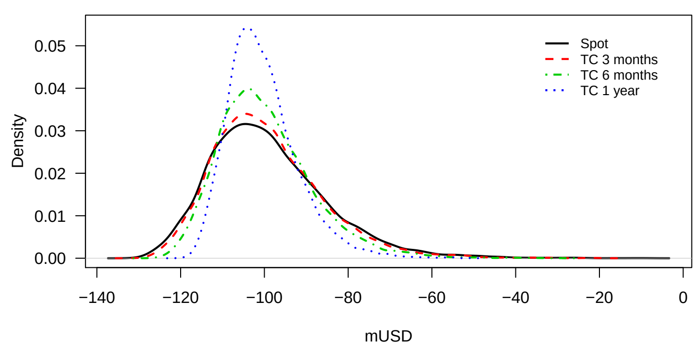

Figure 6.3: Density plot of the total cash flows throughout the horizon of the simulations

6.3.1 Quantifying the Risk

Having identified the risk, the next steps in the Risk Management Process are to measure it and if needed, mitigate it. In this thesis, these two steps collapse somewhat into each other as the quantified risk is compared between the hedged activities (TC) and the unhedged activity (spot) in order to assess which method is the best.

[6]The costs of Bulk Invest are linearly dependent on the revenue (except the bunker cost).

The quantitative risk measures used in comparing the spot- and time-charter business are $CFaR$ and ES as presented in section 3.2. Table 6.4 shows the average potential cash flow losses in million US Dollars (mUSD) over the simulated horizon (Feb 2016 – Dec 2017), as well as the numbers in percentages (relative to the mean cash flow of the horizon). The confidence level is set to 95% for both $CFaR$ and ES (i.e. we are looking at the 5% quantile; α = 0.05).

Table 6.4: Summary of $CFaR$ and ES

α = 0.05	$CFaR$	ES	$CFaR$	ES	$\frac{ES}{CFaR}$	Risk reduction	
	in mUSD		in percentage			$CFaR$	ES
P spot	9.1	11.5	6.8%	8.6%	27%		
P 3	8.5	10.8	6.3%	8.1%	27%	6.3%	6.2%
P 6	7.3	9.4	5.4%	7.0%	28%	19.5%	18.8%
P 1Y	5.6	7.0	4.1%	5.2%	25%	38.6%	39.5%
S spot	6.4	8.2	8.4%	10.8%	29%		
S 3	6.0	7.7	7.8%	10.1%	29%	6.5%	6.1%
S 6	5.1	6.7	6.7%	8.7%	31%	20.6%	19.2%
S 1Y	4.1	5.2	5.3%	6.8%	29%	36.6%	36.8%
H spot	1.8	2.4	8.9%	12.2%	38%		
H 3	1.6	2.3	8.3%	11.5%	39%	6.9%	6.1%
H 6	1.4	2.9	7.2%	9.9%	38%	18.6%	18.5%
H 1Y	1.1	1.6	5.7%	7.9%	39%	36.1%	35.6%
PF spot	14.4	18.6	7.4%	9.5%	29%		
PF 3	13.5	17.5	6.9%	8.9%	29%	6.4%	6.2%
PF 6	11.6	15.1	5.9%	7.7%	30%	19.8%	18.9%
PF 1Y	9.1	11.6	4.6%	5.9%	28%	37.7%	38.2%

The risk measures are calculated with 95% confidence level. To save space, the indices are abbreviated to P, S and H for BPI TCA, BSI TCA and BHSI TCA, and 3, 6 and 1Y correspond to three-month-, six-month- and one-year TC contracts. The last rows named *PF* are the aggregated cash flow risks for a portfolio of the three indices.
The $\frac{ES}{CFaR}$ column shows the relative difference between ES and $CFaR$.
Risk reduction is calculated as the reduction in $CFaR$ and ES from the spot values of each index. On average, operating with one-year time-charter contracts reduces the potential cash flow risk by 37%. For three-month- and six-month TC contracts the average risk reduction is 6% and 19%, respectively (average reductions are approximately equal for both $CFaR$ and ES).
Note that it is standard procedure to report *at-Risk* measures as positive numbers even though they actually are losses (i.e. they are in fact negative numbers).

As mentioned earlier; the simulated cash flows are gloomy, especially for the Panamax vessel class. There is a negative relation between cash flows and vessel size, which is due to the difference in TC expenses for each vessel class (cf. table 4.3). However, if the starting values of the simulations had been higher (e.g. above the break-even level), the cash flows would not have been so dismal. Thus, it is important to look at the percentage risk as well, which provide a relative and more comparable number.

It can be seen that the spot business incurs the highest potential cash flow loss for all vessel classes, as proven by both the $CFaR$ and the ES measure. Furthermore, Expected Shortfall exceeds Cash Flow-at-Risk by 25% – 39% (the $\frac{ES}{CFaR}$ column), with a negative relationship between vessel size and the exceeding ratio (e.g. ES of Panamax exceeds $CFaR$ less than the smaller vessel classes Supramax and Handysize). This implies heavier tails of e.g. Handysize than of Panamax rates.

However, what is more interesting is the risk reduction induced by operating in the time-charter business. For all vessel classes, longer TC contracts reduce both the Cash Flow-at-Risk and the Expected Shortfall. This is naturally also true for the aggregated portfolio of the three indices. The potential cash flow loss is reduced by 37% on average when shifting from the spot rate market to operating by one-year time-charter contracts. Hedging by six-month- and three-month TC contracts induce on average 19% and 6% reduction in potential cash flow losses, respectively.[7]

The risk reduction from TC contracts stems from the truncated freight rate paths. These paths exhibit less volatility than the spot freight rate path, and the reduction in volatility is greater the longer the TC contract. Figure 6.4 illustrates the difference in fluctuations between the different contracts.

[7] The risk reductions are approximately equal in terms of both $CFaR$ and ES.

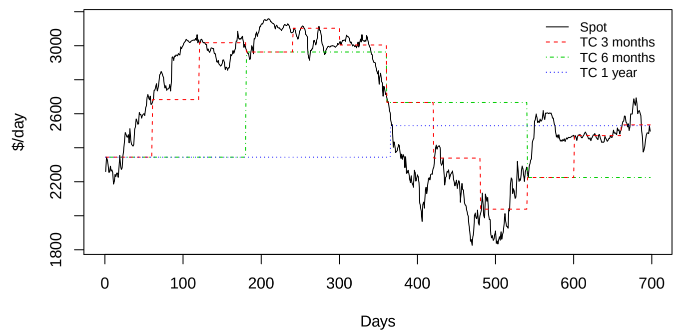

Figure 6.4: Comparison of spot- and time-charter (TC) contracts. Longer TC contracts incur less fluctuations in the freight rate paths

The conclusion of this chapter is that operating purely in the spot market induces a greater potential cash flow loss than operating in the time-charter market. The effect of mitigating risk by time-chartering increases with the length of the TC contracts. From a pure (and somewhat naïve) risk management perspective, the spot market should then be avoided and the vessels be fixed at contracts with longest possible duration. However – building on the metaphor in section 3.1 – it should be stressed that risk management is only the navigating systems of the company ship; other parts and processes are needed to sail the ship economically and profitable. For instance, there is no point in sailing if the hull is too large for the waters the ship is sailing in, and there is similarly no point in running a company if the costs are too high for profitable operations.

Sailing without navigation system may work well if you know your waters, but there is always a risk that a heavy fog sets in and the visibility reduces to zero. In other words, a company may operate in well-known markets, but there is still a chance something unexpected happens. Sailing by a navigation system then reduces the chance of running aground or the ship sailing around perplexed until the fuel tank (the cash reserves of the company) drains out.

The question of whether to operate in the spot- or time-charter market is not a mutually exclusive one. If one expects the market to grow, fixing the vessels at a presumably lower freight rate is then counterintuitive. A profitable shipping company operating by *asset-play* needs to combine operations in both the spot- and time-charter market.[8] But then again, speculations (about the market) and risk management are two different things and from a risk reducing perspective; time-charter contracts provide more stable revenues.

[8]When operating by *asset-play* the earnings of the company are solely based on providing high tonnage when freight rates are high and restricting the tonnage when they are low.

7 Backtesting

When modelling *at-Risk* measures which are meant to be backtested, the models are estimated based on an *in-sample* part (typically the first two thirds) of the data and the *out-of-sample* part (the last third) is then used for backtesting. This is done in order to weed out any bias in the training of the model.[1] The data at hand for this thesis is highly affected by the Financial Crisis of 2008, with two of the indices (BSI TCA and BHSI TCA) starting just before the freight rates rocketed, cf. fig. 5.1. Choosing the *in-sample* data is thus somewhat tricky, since the model will reflect the Financial Crisis if this part is included and then potentially overestimate the risk in a more normal market state. Overestimating the risk is bad because it can induce hedging activities which are not needed and consequently, reduce a potential return.[2]

The *in-sample* periods of the data sets are chosen as the data up to 2012, making the *out-of-sample* the data between 2012 and 2015. The *out-of-sample* periods are set according to the available historical cost data for Bulk Invest. For the BSI TCA and BHSI TCA most of the *in-sample* data consists of freight rates during the Financial Crisis, but there is also a sufficient amount of observations following this turbulent market period.

The following subsection is devoted to backtesting of the Cash Flow-at-Risk estimates from section 6.3. The Expected Shortfall is omitted from backtesting due to the reasons stated in section 3.3.2. Before proceeding, one remark about the simulations and the backtesting is needed. The hori-

[1] Fitting a model to a data set and then testing if the model works on that same set of data is doomed to succeed. Evidently, this backtest is not a test of how the model fits to real virginal data, but more a test of how good the data is fitted to the training data.

[2] This is equivalent to the effect of reduced volatility on the value of an option. Reducing the volatility reduces the possibility of ending in a more profitable market state, which is mirrored by a reduced option value.

zon of the simulated freight rates is 2016 – 2017, but the model is trained over the *in-sample* period up to 2012. This means that some important properties of the freight rates stemming from 2012 – 2015 may be omitted in the model. To check for this, a similar model was fitted over the whole data set (up to 2016) for each index and the risk was measured in a similar way to the original model. Fortunately, the resulting $CFaR$ and ES only differed in the range 0% – 0.7% between those models. Consequently, the *in-sample* model does not omit too much sensitive and important data.

7.1 Cash Flow-at-Risk

Backtesting is performed by calculating two-year rolling cash flows (corresponding to the horizon of the simulations) over the period 2012 – 2015. The rolling window method results in 500 observed two-year cash flow estimates for each index.

One issue which needs to be considered when backtesting $CFaR$ is whether to use the costs implemented in the cash flow model or to use the historical costs corresponding to the backtesting period. In other words: should we keep costs fixed at 2016 level when backtesting between 2012 and 2015 or should the costs be altered to the actual costs in that period? The latter part creates a more realistic picture of the cash flows of that period, but using these costs may also undermine the $CFaR$ level from the simulated model because the inputs are changed. This makes comparison between the simulated- and the historical $CFaR$ somewhat difficult. Keeping costs fixed is also not optimal because the cash flows become less realistic.

Because of the issue regarding fixed or floating costs, backtesting is performed with both alternatives and the results are presented in table 7.1. The numbers of violations are listed in the No columns, with the respective percentages of total observations in parentheses. The p-values of the unconditional coverage tests are presented in the LR_{uc} columns. In accordance with the recommendations of Christoffersen (2012), the significance level is set to 10%.[3] Bold p-values indicate that the null hypothesis is not

[3]Setting a higher significance level (e.g. one or five percent) increases the risk of accepting an incorrect model (i.e. Type II error). In risk management, Type II errors can be very costly (Christoffersen, 2012, ch. 8).

rejected (i.e. the $CFaR$ model passes the statistical criterion) at the 10% significance level.

Table 7.1: Summary of unconditional coverage tests

	Fixed costs		**Floating costs**	
	No	LR_{uc}	No	LR_{uc}
P spot	55 (11%)	0.00%	33 (7%)	**11.68%**
P 3	46 (9%)	0.01%	93 (19%)	0.00%
P 6	14 (3%)	1.38%	199 (40%)	0.00%
P 1Y	111 (22%)	0.00%	264 (53%)	0.00%
S spot	46 (9%)	0.01%	35 (7%)	5.37%
S 3	43 (9%)	0.08%	28 (6%)	**55.25%**
S 6	123 (25%)	0.00%	141 (28%)	0.00%
S 1Y	143 (29%)	0.00%	230 (46%)	0.00%
H spot	72 (14%)	0.00%	20 (4%)	**28.40%**
H 3	92 (18%)	0.00%	4 (1%)	0.00%
H 6	88 (18%)	0.00%	60 (12%)	0.00%
H 1Y	150 (30%)	0.00%	195 (39%)	0.00%
PF spot	57 (11%)	0.00%	19 (4%)	**19.60%**
PF 3	51 (10%)	0.00%	8 (2%)	0.01%
PF 6	67 (13%)	0.00%	121 (24%)	0.00%
PF 1Y	145 (29%)	0.00%	236 (47%)	0.00%

The null hypothesis tested with LR_{uc} is that the average number of $CFaR$ violations is correct.
Bold p-values (percentages in the LR_{uc} columns) indicate that the null hypothesis is not rejected (i.e. the $CFaR$ model passes the statistical criterion) at the 10% significance level.
The number of violations (No) refers to the times the estimated $CFaR$ is exceeded. Figures within parentheses report the frequency of violations in the sample size.

None of the $CFaR$ models pass the unconditional coverage test at 10% significance level when the costs are kept fixed at 2016 level. This is as expected because cash flows are naturally highly dependent on the costs, and the cost level of a company which charters in its vessels depends on the market state (i.e. higher freight rate levels equal to higher TC expenses and vice versa). However, there seem to be a pattern of significant models for shorter time periods when historical costs for the respective years are used in the cash flow calculations (i.e. floating costs). This is interesting and needs further explanation.

First of all, recall that the cash flows for backtesting were calculated as two-year rolling windows. Furthermore, recall that the time-charter contracts were calculated as truncated spot freight rates due to lack of TC data. When backtesting, the cash flows are converted to percentages relative to the mean cash flow of the backtesting period. The longer the time-charter contract, the less ups and downs occur in the cash flows, as can be seen in fig. 7.1. The x-axes show the 500 two-year rolling cash flows from backtesting for the Panamax vessel class. Each cash flow is plotted as a percentage relative to the respective mean cash flow over the backtesting period (y-axis). The shaded areas mark cash flows which violate the $CFaR$ limits (dotted lines), cf. table 6.4. There is an overall positive tendency in the relative cash flows over the backtesting period. This is because the scale between average costs and revenues decreased over the period (i.e. less costs relative to revenues).

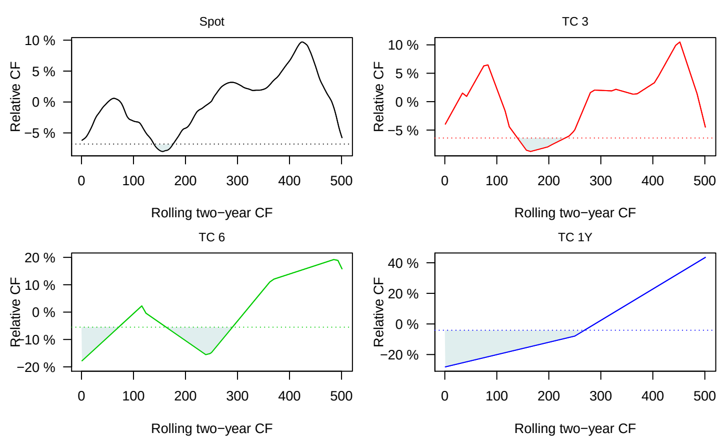

Figure 7.1: Rolling two-year cash flows (CF) over the backtesting period for the Panamax vessel class

Secondly, the cash flows from spot contracts change every day (due to a daily change in revenues), but the cash flows from a one-year TC contract change only 4 times during the backtesting period. Each of the four shifts in cash flow relate to the time when a new TC contract is fixed. If the spot rate at that time differs greatly from the spot rate at the time of the previous TC fixing, the cash flow will incur an equally large shift. Since there are only four observations of cash flows for the one-year TC contract (as opposed to the spot contracts where there are 1000 observations), each observation has greater impact on the average cash flow.[4] Even though, rolling the two-year cash flows forward throughout the backtesting horizon creates more observations, the effect is still present and each observation deviates more from mean in the one-year TC contract than in the spot contract. The magnitudes of these deviations are dependent on the changes in cash flows over the period.

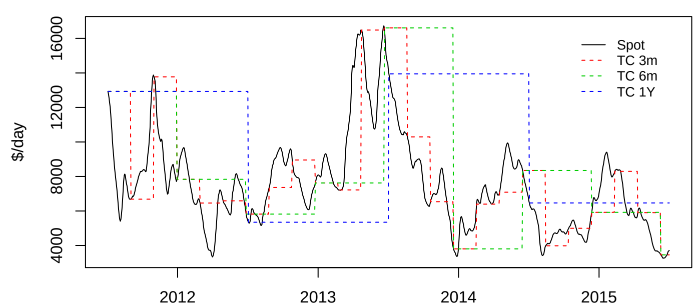

Figure 7.2: BPI TCA spot rates during the backtesting period along with truncated spot rates corresponding to different lengths of time-charter (TC) contracts

To illustrate the changes in cash flows, the spot freight rate along with truncated time-charter rates over the backtesting period (2012 – 2015) are plotted in fig. 7.2. The figure shows that revenues from the one-year contract shift four times during the period and each shift is large in magnitude. Because of this, the relative cash flows from the one-year TC contract

[4]Let Δ be the change in the average of a vector containing some random variables. When Δ goes to infinity, the effect of Δ would be 250 times greater in a vector containing four variables than in a vector containing 1000 variables.

deviates greater from mean than for the shorter contracts, and this result in more violations of the $CFaR$ limit in the one-year TC contract (illustrated by a larger shaded area in fig. 7.1 than for the shorter contract types).

The reasoning of why one-year TC contracts exhibit more violations is also true for the six-month TC contract, although the effect is not as prominent as in the one-year contract (the effect diminishes with shorter time-charter periods). The question is how to deal with the low p-values of the unconditional coverage test for six-month and one-year time-charter contracts? Do the low p-values mean the $CFaR$ measures of longer TC contracts are useless? The answer is: not necessarily. Backtesting is a somewhat ad hoc process and the way the tests are set up affects the potential results (Holton, 2003). In my case, the backtest resulted in few changes of the cash flow for longer time-charter contracts and this affected the number of violations heavily. A backtest with more observations (i.e. a longer period) could give a more reliable and correct result for the longer TC contracts. However, take for instance the one-year TC contract; the cash flow changes once a year which means the backtesting period must be very long in order to get a significant *out-of-sample* period. Such a dense data set is not available for this thesis. The low p-values of the longer TC contracts are therefore best neglected because there is not a valid backtest procedure for them at this time.

8 Sensitivity Analysis

The way Bulk Invest's costs are determined in section 4.4.2 implies a linear relation between costs and cash flows. This means there will be no exciting effect of altering any cost variable; the cash flow would only be altered equal to the change in costs. Because of this, I will not conduct any sensitivity analysis on the costs, but instead focus on variables with a non-linear relation to its determinant. These variables are found in the stochastic freight rate model. The sensitivity analysis is performed by altering one variable at a time and setting consistent seeds in the Monte Carlo simulations.[1]

8.1 Freight Model Variables

Selected variables of the freight rate model are changed successively by either +10% or –10%, and the resulting changes in $CFaR$ and ES are listed in table 8.1. Due to limited space, only changes in a portfolio of Panamax, Supramax and Handysize vessels are presented. These changes are the most relevant ones in a multivariate setting, but it should be noted that none of the individual changes differ much from the aggregated results. Furthermore, only effects in the spot- and one-year time-charter market are presented in the table, since the two shorter time-charter contracts induce similar effects as the one-year TC contract, only at a smaller scale.[2]

[1]Essentially, all random numbers generated from programs like Excel and R are *pseudo-random* due to the algorithms used in the creation of the numbers. Setting a seed means that the random number generation starts at the same point each time, creating the same "random" numbers every time the simulations are run.

[2]The effects are in between the effects of the extremities: spot- and one-year TC contract.

Table 8.1: Results from sensitivity analysis on selected freight model variables

Portfolio	Spot				One-year time-charter			
	$CFaR$		ES		$CFaR$		ES	
Change	+ 10%	– 10%	+ 10%	– 10%	+ 10%	– 10%	+ 10%	– 10%
V_1	2.1%	–1.6%	1.4%	–1.1%	1.1%	–1.6%	1.2%	–0.8%
$\bar{V}$	1.2%	–0.7%	0.8%	–0.6%	0.9%	–1.2%	0.9%	–0.4%
κ	–0.9%	0.5%	–0.5%	0.6%	–1.1%	0.0%	–0.6%	0.7%
ξ	4.1%	–3.0%	3.7%	–3.1%	3.6%	–3.7%	3.9%	–3.2%
ρ	–0.6%	0.3%	–1.3%	1.0%	–1.2%	1.6%	–1.6%	1.0%

As a recap: the freight model variables are V_1 – the starting value of the variance process, $\bar{V}$ – the long-term variance, κ – speed of reversion to the long-term variance, ξ – volatility of the variance process, ρ – correlation between Wiener processes in the variance- and freight path processes.

Each variable from the freight rate model are successively either increased or decreased by 10% from the initial values listed in table 6.1.

Spot contracts are compared to one-year time-charter contracts for a portfolio of Panamax, Supramax and Handysize vessels; cf. the fleet of Bulk Invest in table 4.2.

In the following, some general comments as well as comments on the results of the sensitivity analysis are provided for each variable.

- V_1 – All else being equal it seems that altering the starting variance will not alter the shape of the variance process throughout the simulations. Higher starting values increase the intercept of the variance process, and due to a lag effect (i.e. the change in variance is dependent on the prevailing value of the variance, cf. eq. (5.16)), the standard deviation of the simulated log returns will increase. However, higher standard deviation does not mean higher kurtosis per se, it just alters the scaling of the density. This means that the relative $CFaR$ and ES should not be changed.[3] From table 8.1, it can be seen that there are indeed some changes in the potential cash flow risk. This however, is because the means also incur a slight (non-proportional) change when the starting value is increased. When decreasing V_1, the opposite effect is true.

- $\bar{V}$ – Changing the long-term variance which the variance process reverts to changes only the slope of the variance process. This has similar effect to the standard deviation of the simulated returns as altering the starting variance V_1. The reasoning regarding changes in $CFaR$ and ES is therefore the same, however it seems that the effects are smaller in scale.

- κ – The speed of mean reversion controls how fast the variance reverts to the long-term variance $\bar{V}$ when shocked. Higher values of κ imply that the variance wanders less from the long-term value, i.e. less fluctuations in the variance process. This in turn yields less fluctuations in the corresponding freight rate paths (i.e. less kurtosis), and consequently a reduction in the potential cash flow risk. Similarly, lower values yield higher risk due to increased fluctuations in the variance process. However, a 10% change in the speed of mean reversion in the variance process does not have a huge impact on the potential risk, as can be seen from the results of the sensitivity analysis.

[3] *Relative* $CFaR$ and ES are the risk measures relative to the mean cash flow of the simulations.

- ξ – The volatility of the variance determines how much impact each shock to the variance has and consequently controls the kurtosis of the simulated returns. When ξ is high, the variance process is highly dispersed creating a higher kurtosis and fatter tails than when ξ is small (Rouah, 2013). The value of ξ has a very high effect on the returns and corresponding freight rate paths. For instance, test simulations showed that too high values will cause some freight rate paths to exceed levels observed just before the Financial Crisis of 2008 and this is of course not very likely.[4]

- ρ – Changing the correlation between the Wiener processes of the variance process and the log freight rate process alters the skewness of the simulated log returns and cash flows. For instance, reducing ρ yields higher $CFaR$ and ES due to a fatter left tail (i.e. more observations in the left tail because the distribution is more negatively skewed). This effect is shown in table 8.1, where a 10% decrease in ρ increases both $CFaR$ and ES. The opposite is true for an increase in the correlation between the Wiener processes.

To conclude the sensitivity analysis, it seems that the effects of altering the variables are mostly equal for both spot- and time-charter contracts. The variable with greatest effect on the potential cash flow risk is the volatility of the variance process ξ. This parameter scales the shock of the variance process and consequently the freight rate paths. However, changing the variables by 10% produced no major effects on the potential cash flow risk. This means that the estimated $CFaR$ and ES from simulations (cf. table 6.4) are fairly inelastic and not very affected if the parameters are wrongfully estimated (at least not within a 10% interval).

[4]The test simulations mentioned here were not a part of the final sensitivity analysis results.

9 Conclusion

The shipping industry is a risky business with low expected return compared to the risk taken. Still, some companies occasionally make fortunes and there are plenty of market agents participating in this world scale poker game. The goal of this thesis was to examine how a dry bulk shipping company could manage its freight rate risk in order to limit the potential downside risk. As this is a wide and comprehensive problem, certain additional research questions were listed for the purpose of narrowing the scope of the thesis.

In order to quantify its freight rate risk, a shipping company can model and simulate the freight rates over a suitable horizon. A good starting point for modelling is to examine the distributional properties of the freight rates and their returns. Consistent with previous literature, I found the logarithmic freight rates to be first-difference stationary. Furthermore, an econometric analysis on the log returns (the first differences) showed that they deviate from normality by having excess kurtosis and some skewness. Closer examination revealed that the volatility clustered and varied with time.

Much of previous literatures involving freight rate modelling have assumed mean reversion in the freight rates. This assumption is backed by sound theoretical arguments but lacks backing from empirical research (i.e. the freight rates are non-stationary). Furthermore, the theoretical arguments of mean reversion are mostly valid in the long-term because adjustments in the shipping market may take several years. Consequently, mean reversion models (e.g. the Ornstein-Uhlenbeck process) are not used in this thesis. Instead, log freight rates are proposed to follow a geometric Brownian motion process, which means that they are not bounded to any long-term mean. Moreover, a stochastic volatility is included in the error term in order to model the excess kurtosis and volatility clustering found in the historical returns. The stochastic volatility is modelled as mean reverting according to the Heston model.

When simultaneously modelling several indices instead of one, we move from a univariate- to a multivariate setting. In this setting it is important to preserve the interdependency between the indices. Due to shipment of homogeneous products freight rates for the different dry bulk vessel classes are unlikely to move in complete opposite directions (over time). This correlation is modelled by using a *Student's t* copula on the Wiener processes driving the log freight rate.

By including stochastic volatility and a copula, the proposed freight rate model is able to capture the distributional properties of the historical returns as well as the interdependency between the indices. Sensitivity analysis on the freight model variables showed that the model is not very sensitive to (10%) changes in the variable inputs. It also showed that the parameter controlling the volatility of the variance process ξ is the one having greatest effect on the modelled freight rate paths (and consequently the cash flows).

The modelled freight rate risk is best quantified by the so called Cash Flow-at-Risk and Expected Shortfall measures. $CFaR$ is analogous to the well-known Value-at-Risk but it is more adequate to non-financial firms like shipping companies, because the commercial cash flows are the targeted variable. In addition, Expected Shortfall focuses on the potential loss if one exceeds the $CFaR$ level, answering the question of how bad the worst case scenario is (on average).

The cash flow effect associated with operational activities was examined through two different markets. First, the unhedged spot market, where the company is fully exposed to the freight rate risk and second, the time-charter market where the company hedges by fixing its freight rates in lengths according to the time-charter contracts. These two markets were compared and analyzed in terms of cash flow risk. The main findings were that time-chartering indeed reduces the cash flow risk for all three indices. The longer the TC contract, the less Cash Flow-at-Risk faced by the company. On average, the cash flow risk was reduced by 6%, 19% and 37% for three-month-, six-month-, and one-year TC contracts. The $CFaR$ measures passed backtesting for spot contracts and one three-month TC contract (Supramax), when using the correct historical costs for the *out-of-sample* period. The backtesting was concluded to not be adequately set up for longer TC contracts due to lack of suitable data.

To answer the problem statement of this thesis with one short sentence; a dry bulk shipping company can manage its freight rate risk by operating in the time-charter market. As a last note on risk management, it is not always desirable for a shipping company to hedge away all risk (by for instance time-chartering) because this affects (i.e. reduces) potential revenues. Using Cash Flow-at-Risk and Expected Shortfall is then a good starting point for the top level management to determine whether they should operate in the exposed spot market or hedge by time-charter contracts. At the same time, users should be aware that the risk quantified is by no means an absolute maximum loss for the company in all scenarios. The risk measures should be used the way they are meant to; as frequency measures, and not a max loss function. As George E. P. Box (1987, p. 424) phrased it: "all models are essentially wrong, but some are useful". Correct use of the risk measures along with continuous updating of inputs in the freight rate model may make these models useful as well.

References

Acerbi, Carlo and Balazs Szekely (2014). "Backtesting Expecting Shortfall". In: *MSCI*.

Acerbi, Carlo and Dirk Tasche (2002a). "Expected shortfall: a natural coherent alternative to value at risk". In: *Economic notes* 31.2, pp. 379–388.

— (2002b). "On the coherence of expected shortfall". In: *Journal of Banking & Finance* 26.7, pp. 1487–1503.

Ådland, Roar Os (2000). "A non-parametric model of the timecharter equivalent spot freight rate in the very large crude oil carrier market". In:

Alesii, Giuseppe (2005). "VaR in real options analysis". In: *Review of Financial Economics* 14.3, pp. 189–208.

Alizadeh, Amir and Nikos Nomikos (2009). *Shipping derivatives and risk management*. Palgrave Macmillan.

Alizadeh, Amir H. and Wayne K. Talley (2011). "Microeconomic determinants of dry bulk shipping freight rates and contract times". In: *Transportation* 38.3, pp. 561–579.

Andrén, Niclas, Håkan Jankensgård, and Lars Oxelheim (2005). "Exposure-Based Cash-Flow-at-Risk: An Alternative to VaR for Industrial Companies". In: *Journal of Applied Corporate Finance* 17.3, pp. 76–86.

Angelidis, Timotheos and George Skiadopoulos (2008). "Measuring the market risk of freight rates: a value-at-risk approach". In: *International Journal of Theoretical and Applied Finance* 11.05, pp. 447–469.

Artzner, Philippe, Freddy Delbaen, Jean-Marc Eber, and David Heath (1999). "Coherent measures of risk". In: *Mathematical finance* 9.3, pp. 203–228.

Barndorff-Nielsen, Ole E and Neil Shephard (2001). "Non-Gaussian O–U-based models and some of their uses in financial economics". In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 63.2, pp. 167–241.

Benth, Fred Espen and Steen Koekebakker (2015). "Stochastic modeling of Supramax spot and forward freight rates". In: *Maritime Economics & Logistics*.

Benth, Fred Espen, Steen Koekebakker, and Che Mohd Imran Che Taib (2015). "Stochastic dynamical modelling of spot freight rates". In: *IMA Journal of Management Mathematics* 26.3, pp. 273–297.

Berkowitz, Jeremy (1999). "A coherent framework for stress-testing". In: *FEDS working paper*.

Bjerksund, Petter and Steinar Ekern (1995). "Contingent claims evaluation of mean-reverting cash flows in shipping". In: *Real options in capital investment: Models, strategies, and applications*, pp. 207–219.

Box, George EP and Norman Richard Draper (1987). *Empirical model-building and response surfaces*. Vol. 424. Wiley New York.

Brooks, Chris (2008). *Introductory econometrics for finance*. Cambridge university press.

Castañeda, Liliana Blanco, Viswanathan Arunachalam, and Selvamuthu Dharmaraja (2014). *Introduction to probability and stochastic processes with applications*. John Wiley & Sons.

Cerrato, Mario (2012). *The mathematics of derivatives securities with applications in MATLAB*. John Wiley & Sons.

Cherubini, Umberto, Elisa Luciano, and Walter Vecchiato (2004). *Copula methods in finance*. John Wiley & Sons.

Cherubini, Umberto, Sabrina Mulinacci, Fabio Gobbi, and Silvia Romagnoli (2011). *Dynamic Copula methods in finance*. Vol. 625. John Wiley & Sons.

Christoffersen, Peter F. (1998). "Evaluating interval forecasts". In: *International economic review*, pp. 841–862.

Christoffersen, Peter F (2012). *Elements of financial risk management*. Academic Press.

Colletaz, Gilbert, Christophe Hurlin, and Christophe Pérignon (2013). "The Risk Map: A new tool for validating risk models". In: *Journal of Banking & Finance* 37.10, pp. 3843–3854.

Danish Shipowner's Association (2015). *Danish Shipping Statistics November 2015*. Half-yearly report. Danish Shipowner's Association.

Dickey, David A. and Wayne A. Fuller (1979). "Distribution of the estimators for autoregressive time series with a unit root". In: *Journal of the American statistical association* 74.366a, pp. 427–431.

Emmer, Susanne, Marie Kratz, and Dirk Tasche (2013). "What is the best risk measure in practice? A comparison of standard measures". In:
Fama, Eugene F. and Burton G. Malkiel (1970). "Efficient capital markets: A review of theory and empirical work". In: *The journal of Finance* 25.2, pp. 383–417.
Gatti, Stefano (2013). *Project finance in theory and practice: designing, structuring, and financing private and public projects*. Academic Press.
Geman, Hélyette (2009). *Risk management in commodity markets: from shipping to agriculturals and energy*. Vol. 445. John Wiley & Sons.
Gneiting, Tilmann (2011). "Making and evaluating point forecasts". In: *Journal of the American Statistical Association* 106.494, pp. 746–762.
Heston, Steven L (1993). "A closed-form solution for options with stochastic volatility with applications to bond and currency options". In: *Review of financial studies* 6.2, pp. 327–343.
Holton, Glyn A (2003). *Value-at-risk: theory and practice*. Academic Press.
Ibe, Oliver C (2013). *Elements of Random Walk and Diffusion Processes*. John Wiley & Sons.
Jarque, Carlos M. and Anil K. Bera (1980). "Efficient tests for normality, homoscedasticity and serial independence of regression residuals". In: *Economics letters* 6.3, pp. 255–259.
Jørgensen, Peter Løchte and Domenico De Giovanni (2010). "Time charters with purchase options in shipping: Valuation and risk management". In: *Applied Mathematical Finance* 17.5, pp. 399–430.
Jorion, Philippe (2007). *Value at risk: the new benchmark for managing financial risk*. Vol. 3. McGraw-Hill New York.
Kavussanos, Manolis and Amir Alizadeh (2001). "Seasonality patterns in dry bulk shipping spot and time charter freight rates". In: *Transportation Research Part E: Logistics and Transportation Review* 37.6, pp. 443–467.
— (2002). "The expectations hypothesis of the term structure and risk premiums in dry bulk shipping freight markets". In: *Journal of Transport Economics and Policy*, pp. 267–304.
Kavussanos, Manolis G (1996). "Comparisons of volatility in the dry-cargo ship sector: Spot versus time charters, and smaller versus larger vessels". In: *Journal of Transport economics and Policy*, pp. 67–82.
Kavussanos, Manolis G and Ilias D Visvikis (2004). "Market interactions in returns and volatilities between spot and forward shipping freight markets". In: *Journal of Banking & Finance* 28.8, pp. 2015–2049.

Kerkhof, Jeroen and Bertrand Melenberg (2004). "Backtesting for risk based regulatory capital". In: *Journal of Banking & Finance* 28.8, pp. 1845–1865.
Koekebakker, Steen, Roar Adland, and Sigbjørn Sødal (2006). "Are spot freight rates stationary?" In: *Journal of Transport Economics and Policy*, pp. 449–472.
— (2007). "Pricing freight rate options". In: *Transportation Research Part E: Logistics and Transportation Review* 43.5, pp. 535–548.
Koopmans, Tjalling Charles (1939). *Tanker Freight Rates and Tankship Building: An Analysis of Cyclical Fluctuations, by Dr. T. Koopmans*. 27. De erven F. Bohn nv.
Kupiec, Paul H. (1995). "Techniques for verifying the accuracy of risk measurement models". In: *The J. of Derivatives* 3.2.
Markowitz, Harry (1952). "Portfolio selection". In: *The journal of finance* 7.1, pp. 77–91.
McNeil, Alexander J, Rüdiger Frey, and Paul Embrechts (2015). *Quantitative risk management: Concepts, techniques and tools*. Princeton university press.
Merikas, Andreas, Anna Merika, and Henry Penikas (2013). "Dry bulk time charter rates joint return distribution modeling: Copula-approach". In: *Procedia Computer Science* 17, pp. 1125–1133.
Nomikos, Nikos K, Ioannis Kyriakou, Nikos C Papapostolou, and Panos K Pouliasis (2013). "Freight options: Price modelling and empirical analysis". In: *Transportation Research Part E: Logistics and Transportation Review* 51, pp. 82–94.
Righi, Marcelo Brutti and Paulo Sergio Ceretta (2013). "Individual and flexible expected shortfall backtesting". In: *Journal of Risk Model Validation* 7.3, pp. 3–20.
Rouah, Fabrice D (2013). *The Heston model and its extensions in MATLAB and C#*. John Wiley & Sons.
Said, Said E and David A Dickey (1984). "Testing for unit roots in autoregressive moving average models of unknown order". In: *Biometrika* 71.3, pp. 599–607.
Sklar, M (1959). *Fonctions de répartition à n dimensions et leurs marges*. Université Paris 8.
Sødal, Sigbjørn, Steen Koekebakker, and Roar Aadland (2008). "Market switching in shipping – A real option model applied to the valuation of combination carriers". In: *Review of Financial Economics* 17.3, pp. 183–203.

Sødal, Sigbjørn, Steen Koekebakker, and Roar Adland (2009). "Value based trading of real assets in shipping under stochastic freight rates". In: *Applied Economics* 41.22, pp. 2793–2807.
Stopford, Martin (2009). *Maritime Economics 3e*. Routledge.
Stulz, René M (2003). *Risk management and derivatives*. South-Western Pub.
Tvedt, Jostein (1995). "Market structure, freight rates and assets in bulk shipping". In: *Norges Handelshøyskole*.
— (1997). "Valuation of VLCCs under income uncertainty". In: *Maritime Policy and Management* 24.2, pp. 159–174.
— (2003). "A new perspective on price dynamics of the dry bulk market". In: *Maritime Policy & Management* 30.3, pp. 221–230.
Zumbach, Gilles O (2007). "The riskmetrics 2006 methodology". In: *Available at SSRN 1420185*.

A Appendix

A.1 World Merchant Fleet

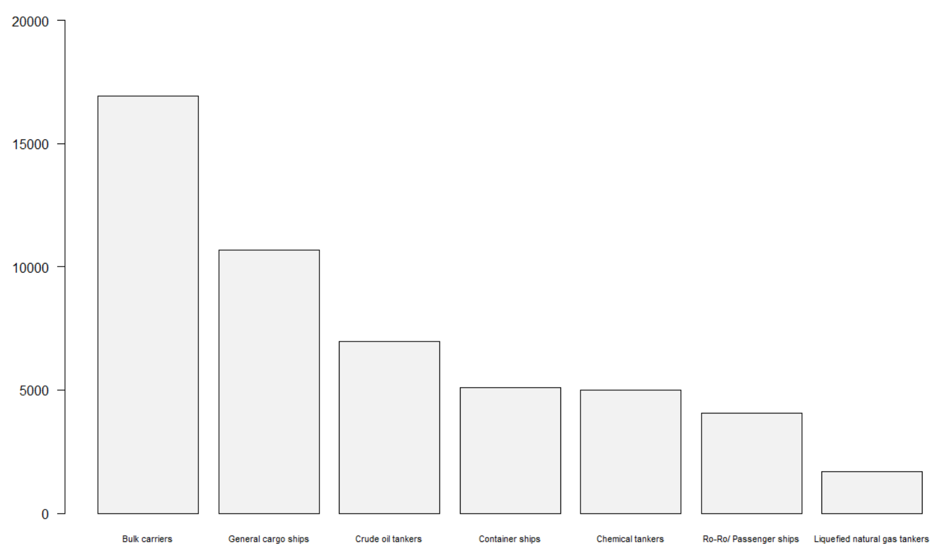

Figure A.1: Number of ships in the world merchant fleet as of January 1, 2015, by type. *Source:* ISL; Marine Flotten Kommando © Statista 2016

A.2 Bunker Forward Curve

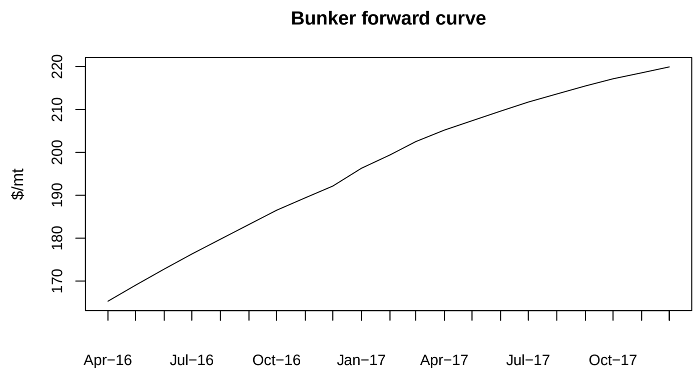

Figure A.2: The bunker forward curve showing the market in contango

A.3 ACF/PACF Plots

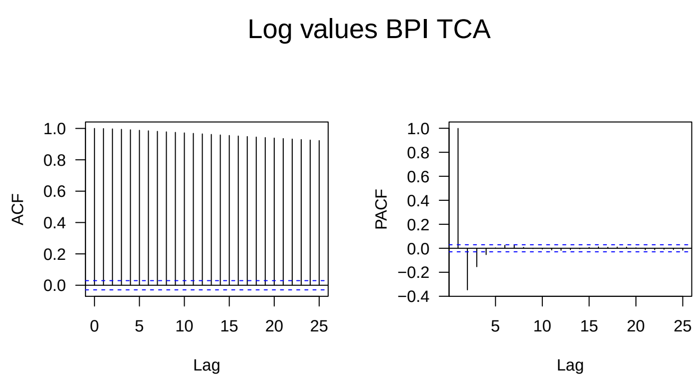

Figure A.3: ACF and PACF for log values of BPI TCA

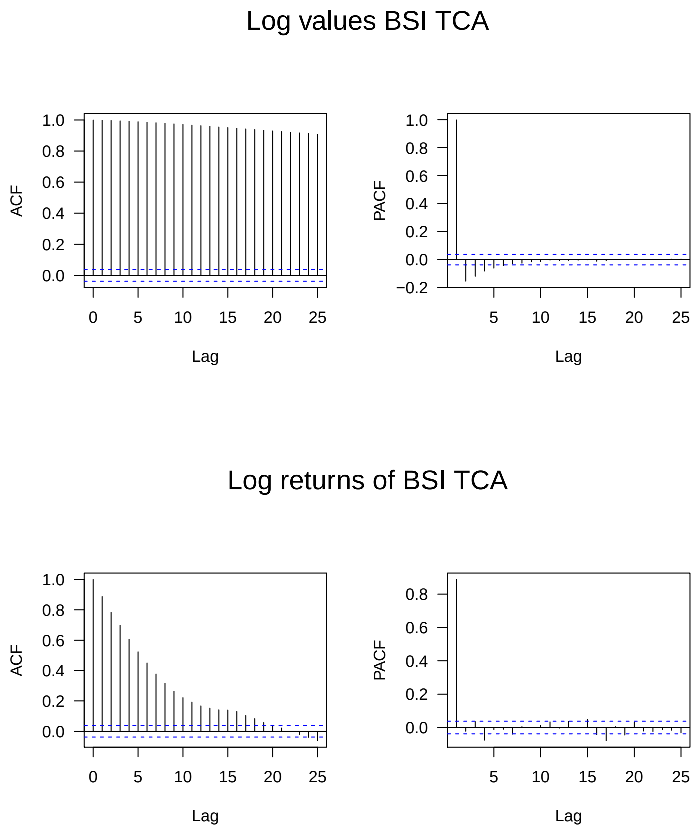

Figure A.4: The ACF and PACF plot of log BSI TCA (above) and first difference of log BSI TCA (below)

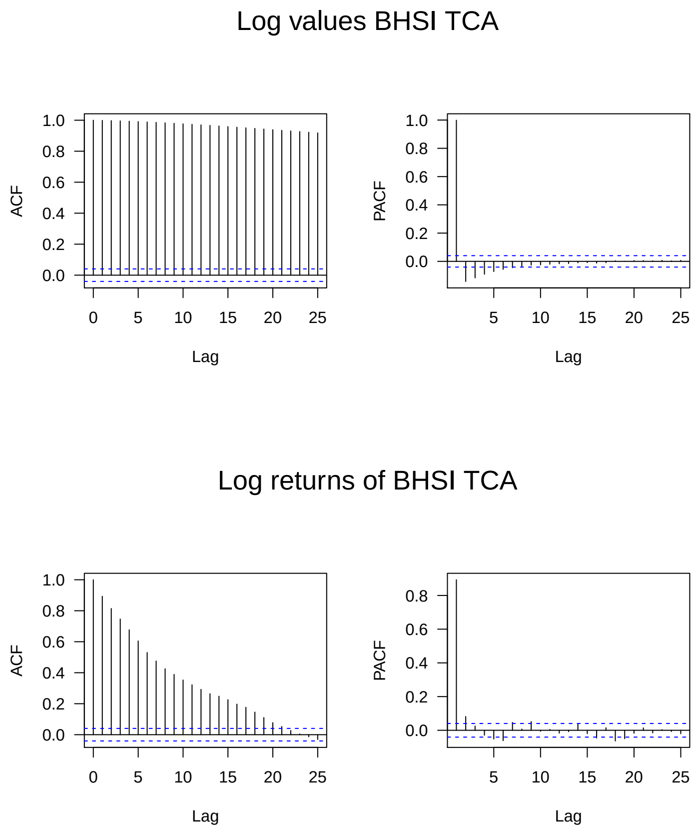

Figure A.5: The ACF and PACF plot of log BHSI TCA (above) and first difference of log BHSI TCA (below)

A.4 Log Returns

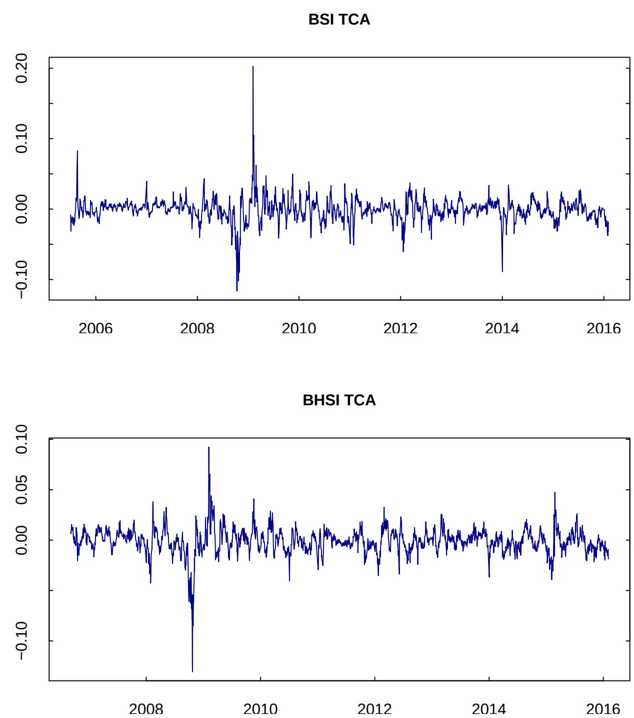

Figure A.6: These figures show the log returns for BSI TCA and BHSI TCA

A.5 Normal Returns

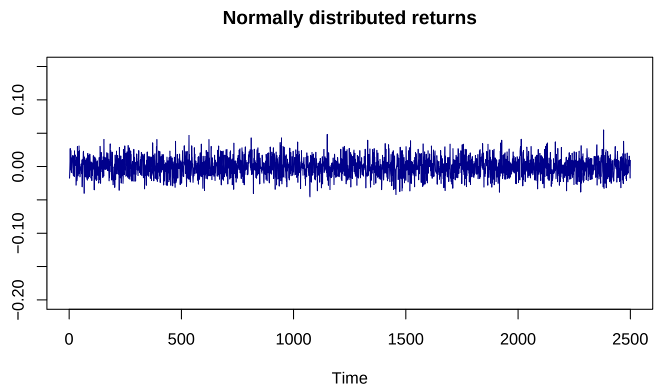

Figure A.7: This figure show returns simulated from a normal distribution. This is added as a comparison with the returns of the freight indices. Note that the parameters of the normal distribution used in the simulations are determined by fitting the model to the returns of the BPI TCA index

A.6 Residual vs. Fitted Values

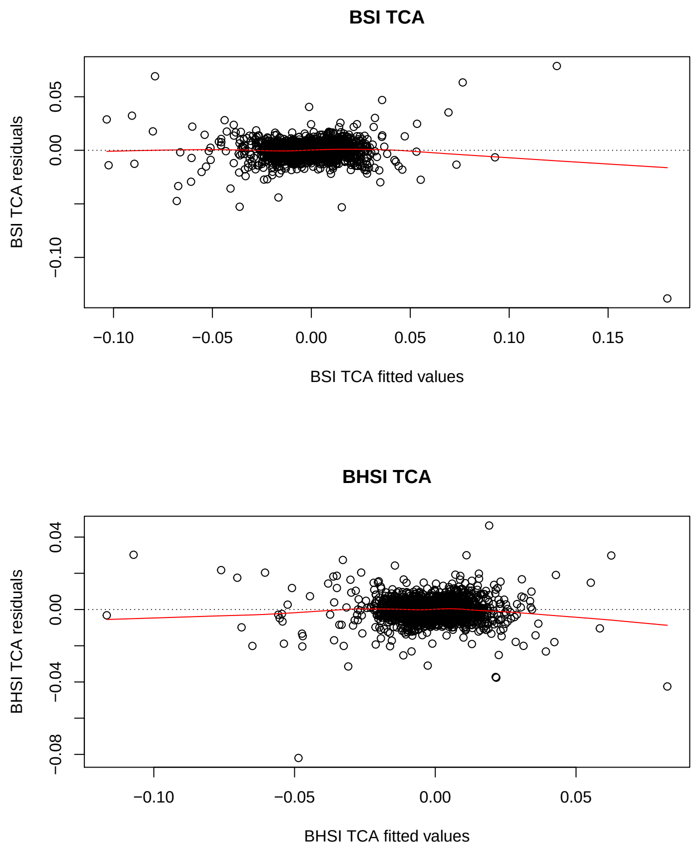

Figure A.8: Resiudals vs. fitted values for BSI TCA and BHSI TCA

A.7 Q-Q Plots

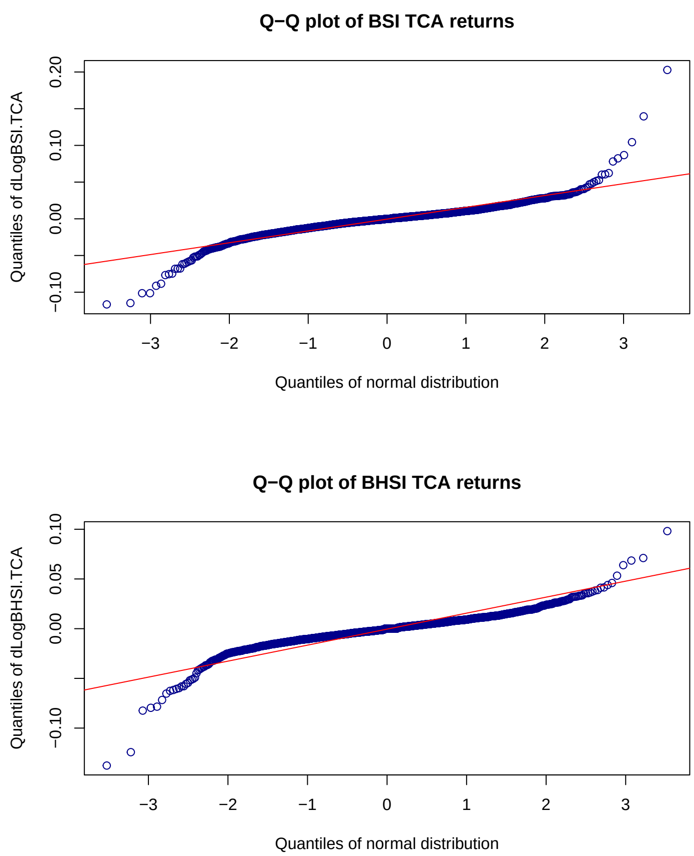

Figure A.9: A quantile-by-quantile plot for the BSI TCA and BHSI TCA indices showing the fat tails of each series

A.8 ACF of Squared Log Returns

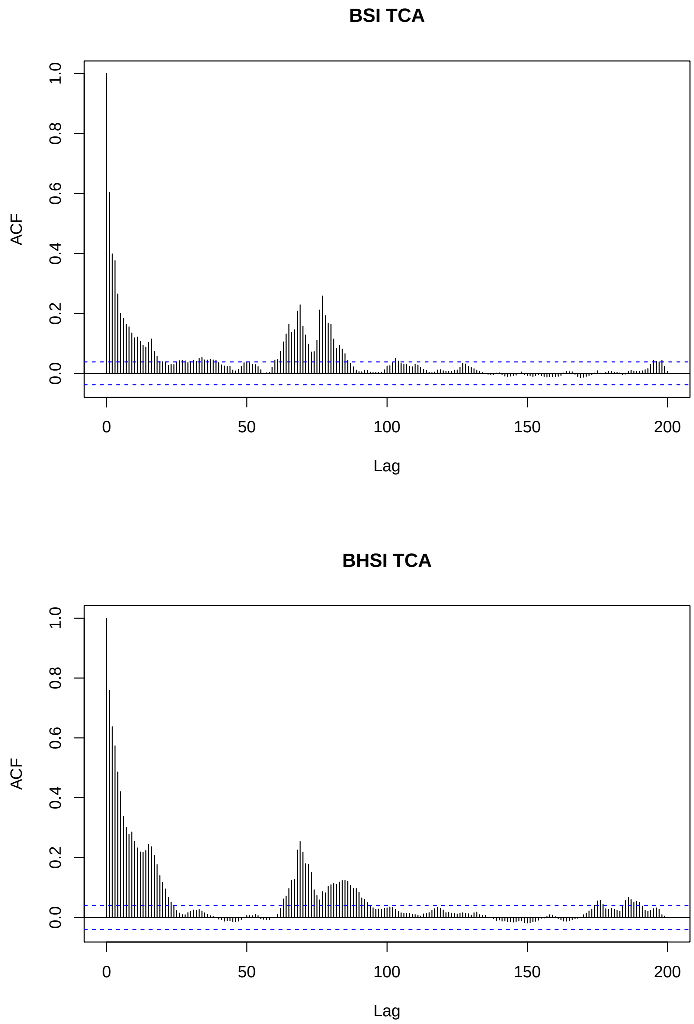

Figure A.10: ACF of Squared Log Returns for BSI and BHSI
ACF of squared log returns for BSI TCA and BHSI TCA

A.9 Density Plots

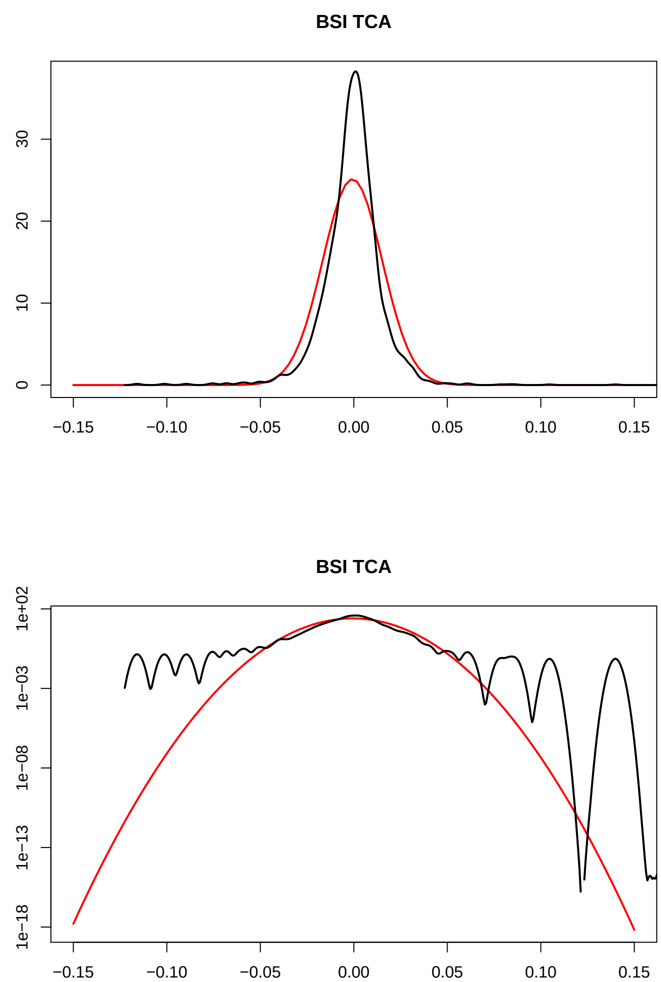

Figure A.11: Empirical density of the BSI TCA log returns (black) along with the fitted normal distribution (red). Un-transformed frequency axis in the top panel and logarithmic frequency axis in the bottom panel

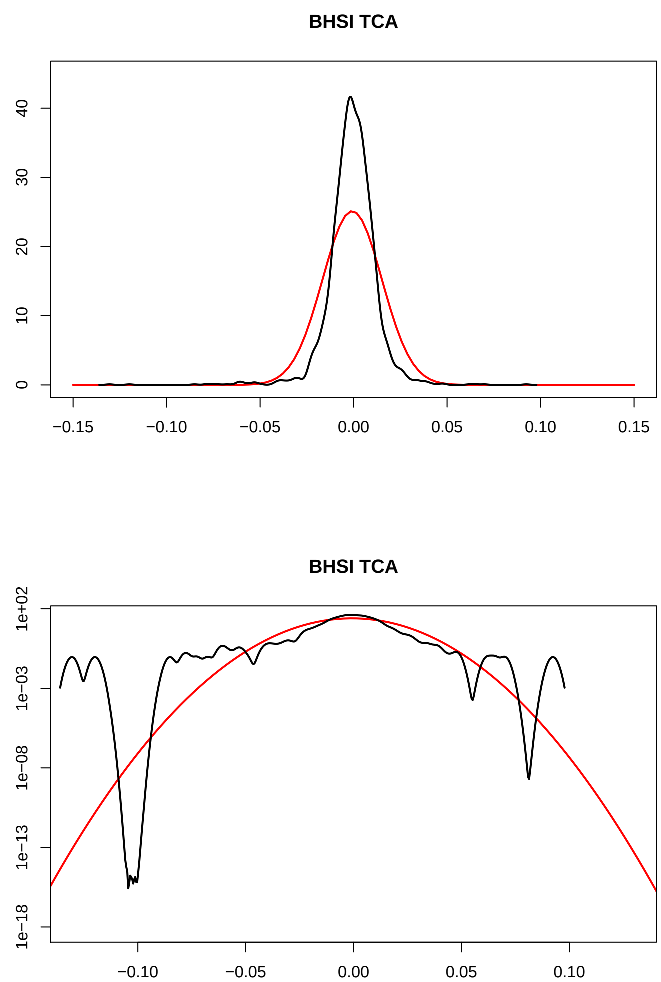

Figure A.12: Empirical density of the BHSI TCA log returns (black) along with the fitted normal distribution (red). Un-transformed frequency axis in the top panel and logarithmic frequency axis in the bottom panel.

A.10 Historical Skewness

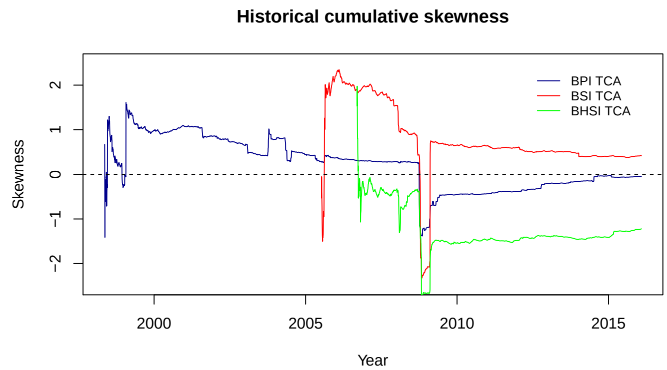

Figure A.13: Historical cumulative skewness. The skewness of each index seems to revert to zero after being shocked. Note that the initial shocks of each index must be ignored because they are due to the small sample size at that time

A.11 Historical Kurtosis

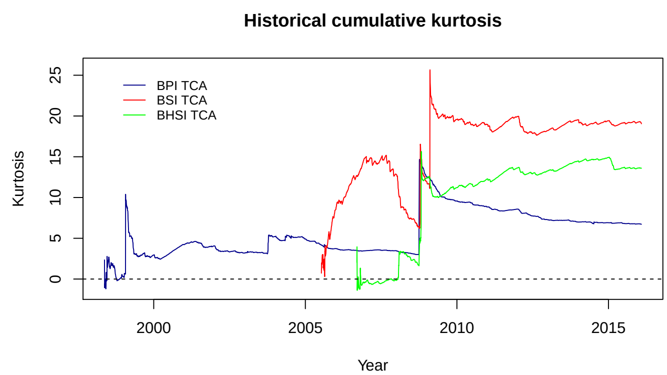

Figure A.14: Historical cumulative kurtosis. There are significant jumps to the kurtosis of each index due to the Financial Crisis of 2008. Note that the initial shocks of each index must be ignored because they are due to the small sample size at that time