University of Stavanger

Faculty of Science and Technology

MASTER'S THESIS

Study program/ Specialization: Industrial Economics/ Project Management and Risk Management	Spring semester, 2013 Open access
Writer: Sunniva Landmark Bjørnstad	………………………………………… (Writer's signature)
Faculty supervisor: Roy Endré Dahl External supervisor: Johan Magne Sollie, Statoil ASA	
Title of thesis: Pricing and Risk Management of Spread Options on Brent and West Texas Intermediate Oil Futures Markets	
Credits (ECTS): 30	
Key words: Oil Futures Markets, Spread Options, Monte Carlo Simulations, Option Pricing, Delta Hedging	Pages: 108 + enclosure: 8 Stavanger, June 13th 2013

Pricing and Risk Management of Spread Options on Brent and West Texas Intermediate Oil Futures Markets

Sunniva Landmark Bjørnstad

June 13, 2013

Abstract

This thesis investigates the price spread between futures on Brent oil from the Intercontinental Exchange and West Texas Intermediate oil from the New York Mercantile Exchange. Historical futures data is calibrated to a multi-factor forward curve model based on Clewlow and Strickland (2000), and the model is fitted, based on Sollie (2013)'s approach, to allow for non-constant volatility. An asymmetric generalized autoregressive heteroskedastic model based on Nelson (1991), and principal component analysis is performed to find key common factor explaining the forward curve dynamics. The model is used to draw realisations of the forward curves for Brent and West Texas Intermediate (WTI) crude oils, and three selected realisations are further analysed. Sensitivity analysis is performed on the expected prices at Day 1, and options are priced on the Brent/WTI futures spread with Monte Carlo Simulations. Each realisation is risk managed with delta hedging, attempting to offset movements in the option prices during their lifetime. The delta hedge is rebalanced one time per day for each contract, which this analysis will find is not sufficient in all cases to capture the extreme volatility in the movement of the underlying assets.

Acknowledgements

This thesis is written as the finalisation of my Master of Science Degree in Industrial Economics at the University of Stavanger, specializing in Risk Management and Project Management.

The thesis has been written in cooperation with Statoil ASA, in the Risk Management department for Crude Oil, Liquids and Products (CLP). I would like to thank Lars Dymbe for welcoming me to the department, and paying interest in my work throughout the semester. A special thank you to Johan Magne Sollie, my supervisor at Statoil, for his patience and support during the work of the thesis. He has truly been a motivator and great teacher, and I wish to extend my deepest gratitude.

I would also like to thank Roy Endre Dahl, my supervisor at the University of Stavanger, for his presence and availability during the thesis. He has been a great supervisor, both by accommodating any request I have had during the semester, as well as by being encouraging and enthusiastic about my work.

Stavanger, June 13, 2013

Sunniva Landmark Bjørnstad

Contents

List of Figures

1 Introduction

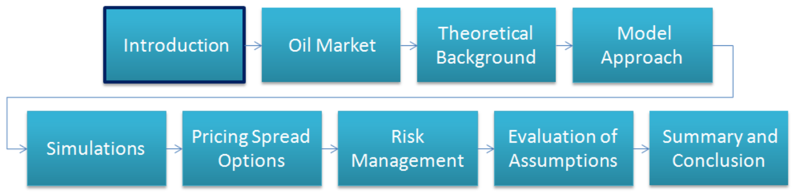

This chapter presents an introduction to the thesis, including the motivation for the choice of topic, the scope of the thesis, and how the thesis is structured.

1.1 Motivation for the Topic

Within Statoil ASA's department of Crude Oil, Liquids and Products (CLP), there is a Risk Management team assessing different types of risk to which Statoil has exposure. Conversations with the quantitative analysis part of the CLP group started in the fall of 2012, and a theme for the master thesis was selected in December 2012. The quantitative analysis group in Statoil presented me with 5 suggestions to subjects they would like further investigated, and after careful consideration I chose to look at valuing flexibility in the oil futures markets, by pricing and risk managing spread option on Brent and West Texas Intermediate oil futures contracts. This topic is of particular interest to me, due to its current presence in both academic circles and the media, and its opportunity to look at both the macroeconomic aspect as well as the detailed characteristics of how the oil market works. My previous experience working with derivatives has been both interesting and challenging, and I saw this thesis as an excellent way of learning more about these financial instruments. Working with spread options also includes a large degree of programming in the modelling and simulations of prices, a skill I see great value of having learned by working with this thesis.

1.2 Scope of the Thesis

This thesis describes how spread options on the difference between futures prices for Brent oil from the North Sea and West Texas Intermediate oil from the U.S. are priced, and how the payoff profile for these options can look for three different realisations of future prices. Historical prices are fitted to the multi-factor forward curve model by Clewlow and Strickland (2000), solved for volatility and variance using Nelson (1991)'s exponential general autoregressive heteroskedasticity model as proposed by Sollie (2013). The fitted model is the basis from which three realisations of the futures prices are drawn. The realisations of the futures prices are risk managed using delta hedging techniques, where the replicated hedge is rebalanced daily.

The thesis addresses a present topic, with the price discrepancies between Brent and West Texas Intermediate crude oil growing since 2010. At present time, there does not exist any standardized product that enables trading options on this price discrepancy. In this thesis, the value of flexibility and optionality is addressed; since the characteristics of the two oil markets are similar, options on the price spread can create value. What is this value, and how can it be kept throughout the lifetime of the options?

When the option values are found, delta hedging is performed as an attempt to risk manage the exposure of holding the spread options. Delta hedging is becoming increasingly popular amongst financial traders, where the frequency of the monitoring of the hedge impacts the level of risk exposure secured by the hedge. This thesis investigates how delta hedging works for the highly volatile oil prices. Can exposure to option values derived from multidimensional assets with high volatility be risk managed with delta hedging?

Choices made when working on the thesis to keep its scope narrow will be presented where appropriate in the following chapters.

1.3 Outline of the Thesis

Chapter 1 introduces the motivation, scope and structure of the thesis.

Chapter 2 contains a presentation of the oil market and the background to trading in oil derivatives. The price of oil is further investigated, and key fundamentals to how the price is determined and quoted in the market are explained. Characteristics of the two most important benchmark prices, Brent oil and West Texas Intermediate (WTI) oil, are introduced, and their current price spread is displayed.

Chapter 3 introduces derivatives theory and pricing theory as the theoretical background essential to the thesis. Although this chapter contains theory essential to understand the remaining chapters, additional theoretical terms are presented throughout the remaining chapters in the thesis, as this fits best with how the data is presented and the results of the analysis are illustrated.

Chapter 4 explains the model approach to the thesis. Historical data on futures prices for Brent and WTI crude oils is presented. Different pricing models for derivatives are explained, and a multi-factor forward curve model is selected as the choice of model going forward with the thesis. The variance is analysed with an EGARCH-model, and principal component analysis is performed to find key factors explaining the evolution of the forward curves. EGARCH-models are thoroughly introduced, as well as the characteristics of principal component analysis.

Chapter 5 uses the model described in chapter 4 to simulate possible future paths for the Brent and WTI futures prices. This chapter shows how the model from the previous chapter is used to draw random paths, and three selected realisations, or samples, illustrate the variety in how the paths may look. The three samples are illustrated with future prices, the corresponding volatility time structure and the simulated price spread between Brent and WTI.

Chapter 6 simulates spread option prices on the three samples drawn in chapter 5. Monte Carlo analysis is performed to find estimates of the option prices, and sensitivity analysis and payoff profiles are illustrated. In this chapter, Monte Carlo simulations as a method is explained.

Chapter 7 presents the results of the risk management, in which the goal is to reduce risk exposure. Delta hedging is selected as a method for re-balancing the risk management one time per day, and how this hedging technique has worked for the three samples drawn in chapter 5 is presented. Delta hedging as a technique is presented, and general theory with regards to calculating delta values and creating a delta hedge is explained. A reality-check is done at the end of the chapter, discussing if risk managing options on such volatile underlying assets really is possible.

Chapter 8 evaluates the validity of the results from the thesis by considering assumptions and scope limitations, and points out opportunities for further research.

Chapter 9 sums up the findings of the thesis.

As a further explanation to the thought process during the thesis, the figure below attempts to illustrate how the working process has been to find answers to the questions asked in part 1.2 (scope of the thesis). The model formulation and calibration to market data is performed based on historical input, and is therefore only necessary to perform one time, as discussed in chapter 4. The part from simulations - pricing - risk management is co-dependent, and each process will depend on its corresponding realisation. The steps from drawing a sample, on which Monte Carlo simulations are performed to find price estimates, to risk managing the prices, are repeated three times, based on three realisations of the forward curves. Optimally, this procedure could be repeated an infinite number of times, but the scope of the thesis have restrained the number of realisations to three due to the time consume of simulations and analysis required for each set. To avoid repetitions, Sample 1 is devoted more attention than the second and third realisation, as this sample is used to illustrate features that are similar for all realisations. Sample 2 and 3 are illustrated in chapter 5-7 in the manners where they differ from Sample 1.

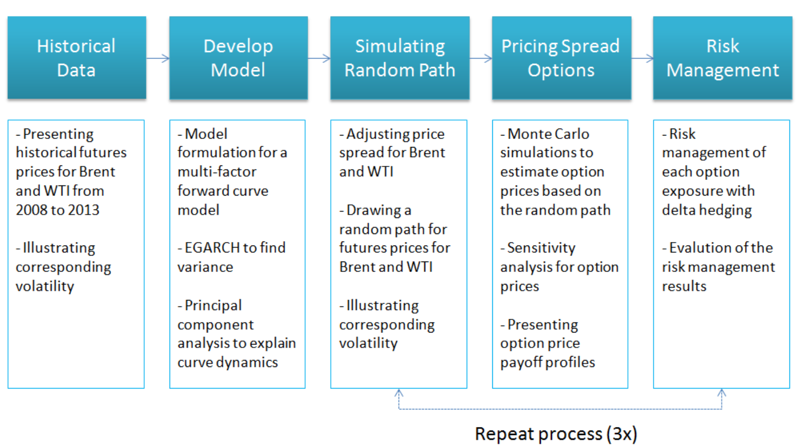

2 The Oil Market

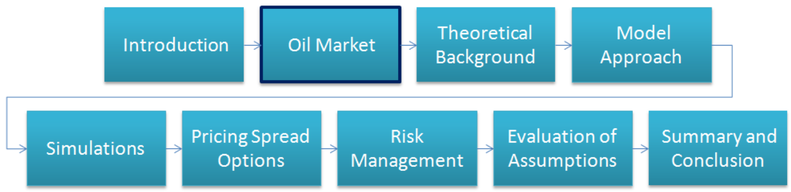

This chapter presents an introduction to the oil market, and explains how risk management in this market have been met with the development of derivatives over the past 30 years. The two largest and most liquid oil markets are introduced, Brent oil from the North Sea and West Texas Intermediate (WTI) oil from North America, as the rest of the thesis studies these two markets. A brief review of how oil prices occur and how they are quoted in the marketplace is also included in this chapter.

2.1 Introduction to Risk in Energy Markets

Risk management in energy companies relates to a set of risk types that can be grouped into five categories. First, operational risk concerns internal risk and includes measures to reduce the risk and consequences of e.g. equipment failure and errors, as well as other risk inside the organization e.g. fraud amongst employees. Second, market risk involves understanding and estimating changes in interest rates, stock prices or commodity prices, and is typically related to external events. Third, the credit or default risk is related to the counterparts ability to withhold the agreement, and includes both suppliers and customers. Fourth, political risk concerns changes in regulation and expropriation. Finally, the last type of risk is extreme risk, or the risk of unpredictable events like the financial crisis, war, depression or technological breakthroughs.

The market risk is of highest importance to this thesis, since the market risk, or price risk, is of highest concern for businesses trading in the petroleum industry. Trading in oil derivatives is mainly based on the price of the underlying asset, where the other risks will affect the price, which again will impact the derivatives. The price risk in the oil market is a consequence of the extreme volatility in oil prices and was introduced after the market deregulation in the mid 1980's.

The price risk introduced in the oil market in the 1980's has been met with the growth of derivatives. A derivative is a financial contract which depends on a certain underlying variable. Trading in derivatives is done both at exchanges with specified contracts, and in "Over-The-Counter" (OTC) markets where the participants negotiate the content of the trade. The next chapter explains thoroughly the three most important derivatives; the futures, the forward and the option contract, as well as the characteristics of the derivatives markets.

2.1.1 The Oil Market

Oil is the world's most important source of energy, meeting almost 35% of global energy needs in 2009 (Herrmann et al. (2010)). Both measured in volume and in value, oil is the world's largest traded commodity. Deutsche Bank estimates the physical crude oil alone to be worth USD 2.2 trillion per year based on a 5 year average historical price.[1]

This thesis focuses on the two most liquid and common benchmark prices, Brent and WTI. Other benchmark crudes exist, but since the Brent and WTI share many of the same characteristics and both have highly liquid markets, these two benchmark crudes are suited for comparison. In both markets, trading occurs in both the physical assets as well as in financial assets (derivatives) based on the prices of the physical assets. When purchasing or selling physical oil, the physical oil is actually bought or sold and delivered. When trading in the financial oil market, also called the paper market, the physical oil is not delivered, and the trades are made based on risk management or speculation. The derivatives traded are settled according to the price of physical oil, providing a close link between the financial and physical aspect of oil trading. As discussed in section 2.2.1 about the oil price determination, the financial aspect of the oil trading have great impact on the price of physical oil, making the link between these two parts of the oil market very intricate.

The oil market has changed dramatically after deregulation the last 30 years, leading to more competition, increased volatility in prices, and an increased volume of participants exposed to potentially higher risks (Hull (2012)). The deregulation of the energy market has lead to higher awareness from both the producer and consumer of a commodity concerning the increased need for risk management. Producers and consumers are naturally exposed to risk in the prices of the commodities they depend on, and they stand on different sides of any trade, where the producers aim to sell the commodity at a high price, and the consumers aim to buy the commodity at a low price. The use of derivatives has become a common means of helping these two participant groups to manage the risk that arise from the high volatility in energy prices, by securing some of their future income/spending. However, the financial participants, such as investment banks, are also a huge part of the energy derivatives market. This is evident in the oil derivatives market, where the derivatives traded on crude oil exceed the physical trading of oil by approximately 14 times (Bruce (2009)). The entry of the financial participants in the oil market has lead to a more volatile oil price, since increased activity and trading on an asset increases its volatility. This results in the following cycle: High volatility in oil prices introduces the need for risk management. Derivatives are used to assess the risk, and the increased financial trading makes the volatility even higher; again increasing the need of risk management.

The Brent oil price from 1985 to 2013 is illustrated in Fig. 1. Notice how the price fluctuations has increased after year 2000, as the popularity of derivatives

[1]Deutsche Bank uses the historical average price of USD 71.5/bbl and the global demand from 2009 of c.85mb/d in their calculations.

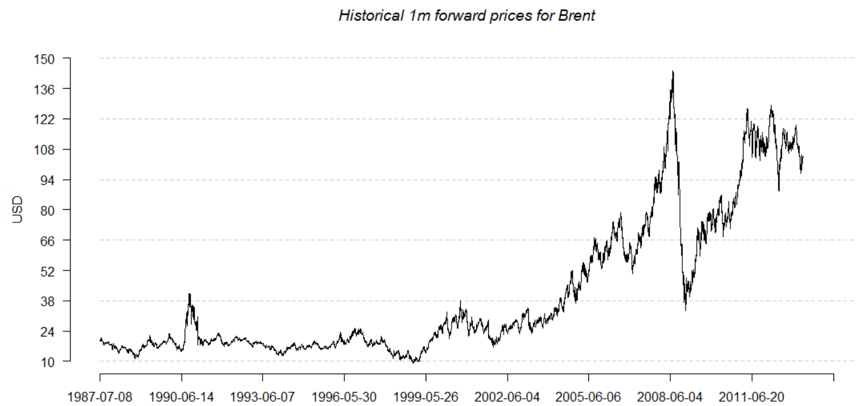

Figure 1: Historical 1m forward prices for Brent illustrates the high volatility in prices of Brent Crude, especially rising after year 2000

increased.

2.1.2 Background to the Oil Derivatives Market

Today's oil traders might take the advanced level of liquidity and complexity in the oil market for granted, although the foundations of the market were laid as late as in the mid 1980's and early 1990's (Bruce (2009)). Before the deregulation, the large oil majors, often called the "Seven Sisters" [2], along with the Organization of Petroleum Exporting Countries (OPEC) set the oil prices by fixed contracts and posted prices. OPEC was, in fact, founded in 1960 by Iran, Iraq, Kuwait, Saudi Arabia and Venezuela, with the principal objective of taking a collective stand against the Seven Sisters.

The first change to this regime of dominance shared between the Seven Sisters and OPEC was the oil shocks of 1973 and 1979. During the Arab-Israeli War of October 1973, the Arab members of OPEC announced an embargo against the United States in response to the U.S. decision to re-supply the Israeli military during the war. OPEC members also extended the embargo to other countries that supported Israel. The embargo both banned petroleum exports to the targeted nations and introduced cuts in the oil production. The second crisis occurred in 1979, after the fall of the Shah of Iran in the wake of the Iranian Revolution. The fall of the Shah lead to a disruption of the Iranian oil sector, causing lower exports and hence higher prices. These two oil crises forced the oil majors to turn away from the fixed contracts, as well as to look elsewhere for exploration and production. This turned the oil majors towards

[2] The "Seven Sisters" consisted of: The Anglo-Persian Oil Company (now BP), Gulf Oil, Standard Oil of California, Texaco (now Chevron), Royal Dutch Shell, Standard Oil of New Jersey and Standard Oil Company of New York (now ExxonMobil).

the North Sea, where discoveries in the early 1970's such as the Brent, Forties and Ekofisk oil fields caught their attention. The discoveries in the North Sea lead to a switch of focus in the global oil market, and since the spot prices had recently been introduced, the oil market looked to the North Sea in need of a benchmark price.

The second large change to the oil market came in 1981, when Ronald Reagan removed all the remaining domestic price controls on crude oil in the U.S. This resulted in a new era of transparency, and dissolved the power of the huge American Oil Companies such as Koch, Exxon and Amoco. Up till that time, two of the largest crude oils in the U.S., the West Texas Intermediate (WTI) and the Louisiana Light Sweet crude (LLS), had been traded at posted prices set by the large American oil companies.

These two changes lead to the rise of the spot price, where, for the first time, the oil price was set transparently by the market forces of supply and demand (Bruce (2009)). After the deregulation of the oil market, derivatives was introduced as a way for consumers and producers to manage their risk exposure to the growing volatility in prices. Derivatives had already been established in the interest rate market and stock market, but this was the first for the energy market. In 1983, two futures contracts were initiated: The WTI futures at the New York Mercantile Exchange with delivery in Cushing Oklahoma, and the LLS at the Chicago Board of Trade with delivery in St. James Louisiana. The Chicago contract collapsed after a month due to delivery problems, whereas the WTI futures contract has become the most liquid crude contract in the world. Around the same time as the crude futures contracts birth in the U.S., the Brent 15-day market traded "Over-The-Counter" (OTC) at the International Petroleum Exchange (IPE) was established. In this OTC forward market, the seller gave the buyer a minimum of 15 days notice (now expanded to 25 days notice) of the intended loading dates for 600,000 barrels of crude oil. However, the IPE wanted to standardize futures contracts for Brent, but the complexity of the physical Brent market made settling a delivered contract difficult. The delivery for the contract was potentially at Sullom Voe or Rotterdam, but the Brent futures was finally established with cash settlement in 1988. However, after the instalment of this contract, all the mechanisms were in place to support an advanced exchange traded derivatives market.

Until the 1990's, the derivatives - and hedging as a risk management technique in general, were viewed with some suspicion by the conservative oil companies. This changed in the first Gulf War crisis in 1990, where the prices rose from USD 21 to USD 46 in two months. At this time, the companies who had hedged their exposure could well manage the increase in the oil price, whereas the price increase had painful consequences for the un-hedged consumers. After this realisation, derivatives became more popular, and financial institutions entered the scene without any physical trading presence. The banks became indispensable providers of liquidity and risk transfer to the oil markets, and the foundations for the derivatives market as seen today were created.

2.2 The Price of Oil

The price of oil has fluctuated significantly throughout the years, from the lows of USD 2.5/bbl seen in 1940-1970 to the highest levels in 2008 with almost USD 150/bbl. Supply and demand[3] are the most important factors affecting the oil prices, but several other factors have great impact on the oil prices as well. The determination of the crude oil price is a complex matter impacted by several different factors. Although it is outside the scope of this thesis to try to solve the puzzle behind the oil price, the current discussion on oil price determinants as well as some established factors are mentioned to get an overview of the fundamental drivers of the oil price.

2.2.1 Oil Price Determination

Disagreement persists regarding the relative importance of oil supply and demand factors in determining oil prices. For instance, Hamilton (2009) emphasizes oil supply disruptions in explaining major run-ups in oil prices, while Kilian (2009) argues that shocks to oil demand have driven oil prices historically. The central message in Kilian (2009) is that oil price increases may have very different effects on the real price of oil, depending on the underlying cause of the price increase. Kilian (2009) states that an increase in precautionary demand for crude oil will cause an immediate, persistent and large increase in the real price of oil; an increase in aggregate demand for all industrial commodities will cause a delayed, but sustained, increase in the oil price; and a production disruption in crude oil will cause a small and transitory increase in the real price of oil within the first year. Hence, when demand drives the prices, the changes are more substantial than if the supply-side drives the price fluctuations. The conclusion in Hamilton (2009), on the other hand, is that the low price-elasticity[4] of short-run demand and supply, the vulnerability of supplies to disruptions, and the peak in U.S. oil production account for the broad behaviour of oil prices over 1970-1997. Hence, the supply will drive the prices according to this approach.

Although difficult to determine whether the supply-side or the demand-side has the highest impact on the oil price, the driving factors behind the market fundamentals of supply and demand are considered. The oil products market, price elasticity, OPEC's spare capacity, inventory levels, geopolitical and political issues, financial trading, available resources and global GDP indicating the general conditions of the global economy, will all have impact (to a varying degree) on the oil price. The factors may affect the oil price at short-term or long-term levels.

The demand for oil is driven by consumers; both individuals and companies who depend on crude oil. Individuals can impact the oil demand for example by the car fleet, which drives the need for gasoline. According to Herrmann et al. (2010), the transportation fuels will account for the majority of the growth in

[3]Read about demand, supply and market equilibrium for example in the textbook for macroeconomics by McConnell and Brue (2008).

[4]Read about price elasticity in Hamilton (2009).

world oil demand in the years to come. Refineries are another driver of oil demand, and the oil demand normally declines as the refineries upgrade their plants or take other production breaks. The price of crude can in some cases drive the prices of crude products such as gasoline, but the situation is often the reverse: If, for example, the refining capacity is tight, an increase in product price can lead to an increase in crude price since the market expects and assumes that demand for crude will increase as companies seek to take advantage of high product prices. Likewise when significant spare refining capacity is evident, or when inventories of oil products are high, this can lead to an incline in the prices of crude oil. The United States is by far the largest single importer of crude oil, but since much of the U.S. imports come from Mexico and Canada, it is in fact the Asia Pacific who holds the position as the largest regional importer (Herrmann et al. (2010)).

The supply of oil is driven by producers; OPEC controls 77% of the total global oil reserves, and was accountable for 41% of the total oil production in 2009 (Herrmann et al. (2010)). The world's largest exporters of oil are Saudi Arabia, Russia and Iran, hence these countries impact the global supply. OPEC's power on the oil price is explained through the OPEC capacity, hence the theoretical volume which OPEC can produce of oil. OPEC has historically tended to restrict supply in order to prop up the price of oil. However, Kilian (2009) states that OPEC's efforts to coordinate production do not influence changes in the real price of oil to a large degree. Kilian (2009) estimates the dynamic effects of supply shocks to the real price of oil during the 1975-2007 period, and finds that there is "little evidence that cartel activities mattered for the sample period in question".

Regarding the global supply of oil, the trend seen today is shifting towards unconventional oil, such as production in deepwater ocean and oil sands.

Inventory levels and financial trading impact the short-term oil price without being either on the supply-side or demand side (Herrmann et al. (2010)). The world's largest storage capacity is in the U.S., which first started storing oil as a response to the 1973 oil embargo in an attempt to mitigate future oil disruptions. Japan and China are other countries with large inventory capacity, where emergency supply can be held. Weekly data is published regarding the U.S. inventory levels, since the inventory in the U.S. is an indicator of current capacity/tightness in the market. The U.S. currency can also impact the prices, since all oil prices are quoted in U.S. dollars, making the strengthening/weakening of the dollar an input to the oil price. The increased trading in oil derivatives has made the price of oil more volatile, and commodity price speculation can change prices both short-term and long-term. If large positions are dumped into the market by for example commodity, this will also create a shock to the prices. In markets for storable commodities, such as crude oil, inventories play a crucial role for the price formation, as changes in inventories will affect market expectations and consequently change prices.

Prices and inventory levels fluctuate considerably every week, in part due to predictable reasons, and in part due to unpredictable reasons. Predictable reasons for changes in prices can be season changes, such as for example maintenance in refineries that normally occur in March/April, which will lower the demand for crude oil. Normally, the expected future changes are considered in the price today. The unpredictable changes are, on the other hand, not considered in the price today and can hence cause large price fluctuations. These unpredictable factors can for example be geopolitical, such as lower than expected growth in China, extreme conditions such as natural disasters or imposed political sanctions affecting trading. Geopolitical events may impact both supply and demand, where changes in the world's economic climate will cause changes in all global prices, not just the price of oil. Examples of geopolitical impacts are the financial crisis of 2008, the Oil Embargo in 1973 and the Gulf War in 1990.

2.2.2 Oil Price Quotation

All commodity prices are set in the market by Pricing Reporting Agencies, agencies such as Platt and Angus that provide information about energy and metal commodities, and quote prices on a daily basis. According to Platts' homepage, their principle is that price is a function of time, and that the most useful price for oil and refined products markets is the value at the close of the market. Platts' process "Market On Close" (MOC) is the assessment of prices for crude oil, petroleum products and related swaps. The MOC process is highly transparent: Bids, offers and transactions are submitted by market participants to Platts' editors and published in real-time throughout the day until the market close. The market participants are buyers and sellers of crude oil, petroleum products and financial instruments that are tied to the value of the physical oil. These are for example major national and international oil companies, financial institutions and trading houses, and end-users such as airlines and utilities. Platts' MOC process was launched in Asia in 1992, in Europe in 2002, and in North America in 2006.

2.3 Brent and West Texas Intermediate

The main international exchanges for trading of oil and oil products (both physical and financial) are the Intercontinental Exchange (ICE) in London and the New York Mercantile Exchange (Nymex).[5] Given the large number of crude oils and the difficulty of following them all, two benchmark crudes are widely used; WTI traded on Nymex and Brent traded on ICE.[6][7] In 2012, ICE Brent became the world's largest crude oil futures contract in terms of volume and the ICE Brent market share has almost doubled since 2008. The Brent and WTI crudes are used as indicative oil prices, and most other crude oil prices will trade at either a discount or a premium to these two benchmark prices depending on

[5] The Intercontinental Exchange is the successor of the International Petroleum Exchange.

[6] Other important benchmark oil prices are Dubai crude, Oman crude and OPEC reference basket.

[7] The majority of Brent is traded at ICE and the majority of WTI is traded on Nymex, but both benchmark contracts are traded elsewhere as well.

their quality. Both ICE and Nymex trade spot contracts for immediate delivery as well as futures contracts for delivery at a later date, providing possibilities for hedging, speculating and price discoveries.

2.3.1 Brent Crude

Brent was originally produced from the Brent oilfield discovered in the late 1960's.[8] Brent crude is a light crude, with an API gravity[9] of 38.06 and a specific gravity of 0.835. It is considered a sweet crude (sulphur level below 0.5% is considered sweet) with a sulphur level of 0.37%. Brent crude is a major trading classification comprising four key crude streams: Brent, Forties, Oseberg and Ekofisk (the BFOE quotation), all sourced from the North Sea. Brent crude oil is the largest underlying physical market of any comparable, traded and transparent benchmark. Besides the existence of a spot market for immediate delivery for specific physical cargoes, there are two widely variable markets: 25-day forward BFOE and Brent futures.

The market for spot prices (prices today) based on the Brent crude oil price is called the Dated Brent market. Dated Brent itself is not an actual spot market, but rather a short-term forward market affected by Contract-for-Difference (CDF)'s[10] derived from the forward curve of Brent futures and short-dated cash market options. The Dated Brent is a price listed in the market once a day and is the basis of 65% of the world's trade in crude oil, including deals done for immediate delivery. The Dated Brent is therefore not an asset able to trade, but rather a benchmark price on which the market relies upon for information. The term "Dated Brent" refers to physical cargoes of crude oil in the North Sea that have been assigned specific delivery dates, with delivery within the next 10-21 days (23 on a Friday). Cargoes that have been assigned loading dates are referred to as dated cargoes, wet cargoes or wet barrels. Cargoes without loading dates are known as paper barrels and are traded for speculative or hedging purposes. The value of the Dated Brent is set every day by ICE at 16:30 GMT, and is assessed by Platts as the value of the cheapest crude in the BFOE group on that day.

The forward market, the 25-day BFOE, is unregulated and consists of private agreements between large oil companies, oil traders, investment banks and others. The private agreements result in different prices within any day for the same contract size, resulting in a need for an index; ICE calculates the Brent Index every day, an index quoting the average forward price with 10-25 days until delivery. The original Brent forward market was assessed on a 7 to 15 day range, i.e. cargoes loading 7 to 15 days forward. As the range of North Sea grades was broadened, the assessment period was also extended to a 10 to 21 day basis in 2002 and finally to a 10 to 25 day basis in January 2012.

[8] The name has its origin after the bird "Brent Goose", since Shell and Exxon named all their fields after birds. Brent is also an acronym for the formation layers of the oil field: Broom, rannoch, etieve, ness and tarbat.

[9] The American Petroleum Institute (API) gravity measures how light/heavy an oil is compared to water.

[10] Contract for difference: Cash swap market.

The Brent futures are tied to the physical forward market, and the futures contracts are settled financially (there is no physical delivery upon delivery) against the Brent Index. Each futures contract on Brent has a size of 1,000 barrels, with the symbol "B" on ICE. Each tick lost or gained equals USD 10. Since the Brent futures are settled daily in cash, the investor/owner of the futures has a margin account where losses/gains are written on a daily basis. The Exchange Futures for Physical mechanism allows for cash-settled futures contracts to be exchanged for physical delivery.[11]

2.3.2 West Texas Intermediate

West Texas Intermediate (WTI) is the crude oil extracted in the U.S., and delivered at Cushing, Oklahoma. WTI is the North American benchmark for crude oil, also referred to as the "light sweet crude". WTI has an API gravity of 39.6, a specific gravity of 0.827 and a sulphur level of 0.24%, and is hence both lighter and sweeter than the Brent crude.

The history of the petroleum industry in the United States goes back to the early 19th century, where it became an important industry following the oil discovery at Oil Creek Pennsylvania in 1859. The WTI has been the largest traded commodity for years, until the traded volume of Brent exceeded that of WTI in 2012. Each contract of WTI has a size of 1,000 barrels, with the symbol "CL" on Nymex, where each tick lost or gained equals USD 10. WTI futures is one of the most liquid crude contracts in the world, and it is settled physically in Cushing, Oklahoma. The futures trading stops on the third trading day prior to the 25th day of the month prior to the deliver month. This is done to inform which producers must make arrangements to have their oil delivered through the pipelines before the end of the month. Although WTI futures are settled physically, less than 1% of the Nymex contracts ever get to delivery; the investors sell or quit their positions before ever getting to the physical delivery (Herrmann et al. (2010)).

The crude oil extracted in the U.S. cannot be exported out of the country due to political sanctions. Hence, the WTI crude oil is only consumed by companies and refineries in the U.S. mid-continent, leading to only a very small part of the world's physical oils being priced against the American domestic oil. However, the WTI futures price remain an important contract due to its high level of liquidity and transparency. [12]

2.3.3 Price Spread between Brent and WTI Crude Oil

European Brent oil and American WTI are the most important crude oils worldwide. Although extracted in geographically distant locations, the chemical composition of Brent and WTI is quite similar, since both of them are considered

[11] Read more about specifications for Brent trading on the Intercontinental Exchange's website, www.theice.com.

[12] Read more about specifications for WTI trading on the New York Mercantile Exchange's website, www.cmegroup.com.

to be sweet crudes. Chemical composition set aside, the Brent and WTI crudes operate under different circumstances; the Brent is exported in the whole of Europe and worldwide, while the WTI does not leave the U.S. The spread between the Brent Crude and the WTI benchmark oil price has changed substantially over the last few years. Historically, ever since the introduction of the spot price, the Brent and WTI prices have stayed relatively close to each other, with the WTI traditionally trading on a 1-2 dollar discount to the Brent, since the WTI is lighter and therefore easier to refine. However, since 2010, the Brent/WTI spread has widened, as illustrated in Fig. 2. Fig. 2 illustrates how the U.S. benchmark oil has been extensively cheaper than the Brent crude oil over the last three years.

The underlying reasons that initiated the price discrepancies between Brent and WTI are challenging to define, and a simple answer explaining the price spread does not exist. Although this thesis considers the value of the flexibility encountered in the price spread rather than how the price spread itself has evolved, commonly used factors to partially explain the price spread are introduced.

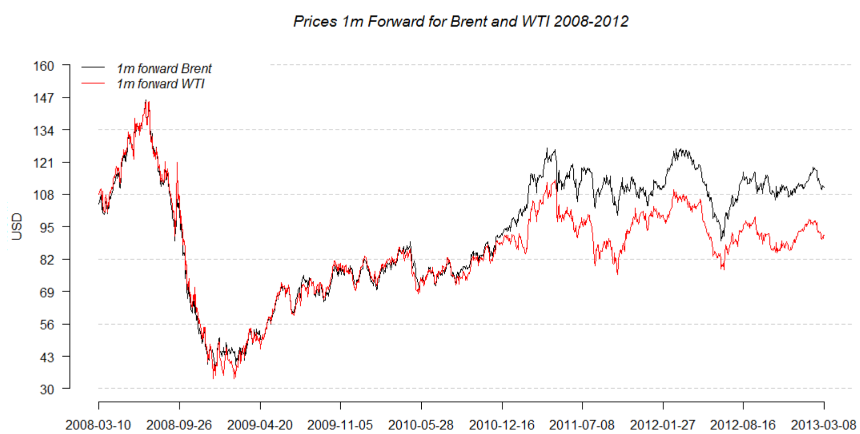

Figure 2: Development of 1m forward price spread between Brent and West Texas Intermediate

The supply side for Brent is easier to control than that of WTI, since Brent oil is transported straight from the oil fields to its destination by ships, as opposed to the WTI which is transported mostly through pipelines first to Cushing, then again through pipelines to its end-destination. It is generally accepted that the large WTI discount to the Brent came about the same time as an oversupply of new crude production from Canada and U.S. domestic shale oil fields, such as the Bakken field in North Dakota, into the Midwest Market. The new production from the U.S. and Canada backed up supplies at the Cushing hub where WTI is settled and traded. An increase in production at the same time as a decrease in U.S. oil demand, lead to an oversupply of the WTI oil, causing WTI prices to fall relatively to international crudes linked to the Brent.

The new oil production in Canada and in several U.S. domestic fields is mostly shale oil, rapidly emerging as a significant and relatively low cost new unconventional resource in the US. The quality of the shale oil is considered lower than the conventional American oil, since the shale oil is more costly to refine. This may also influence the WTI price without impacting the Brent price, making the two benchmarks drift further apart.

The dependency on pipelines in the U.S. can create a "bottleneck" effect, since the pipeline system have a maximum limit of how much oil they can transport every day. This effect in the physical transfer of oil causes the supply side of the market to disentangle from normal supply/demand laws. Changing the infrastructure of oil pipelines is a time consuming process, and can not be adjusted as easily as the ships in and out of the North Sea. Fig. 3 shows the pipeline infrastructure in North America[13], where the pipelines mainly go from Cushing to the Mexican Gulf, the Midwest U.S. and the Western Canada.

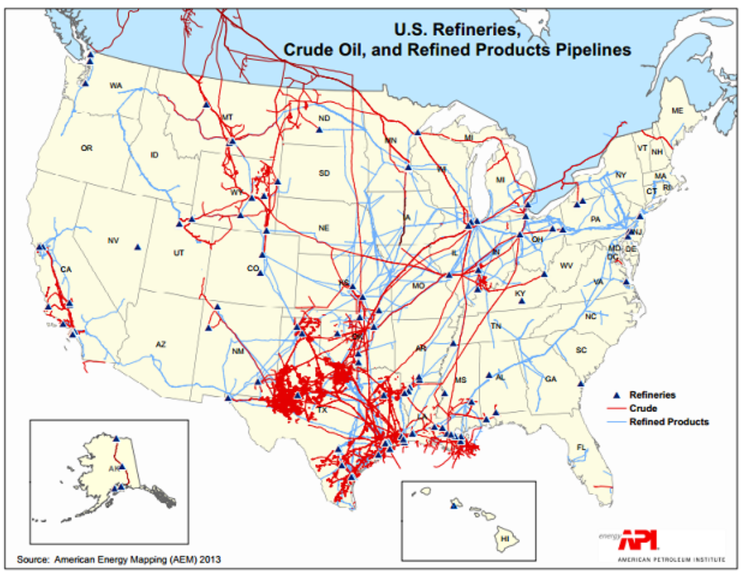

Figure 3: Pipeline infrastructure in North America

Since the price dislocation occurred around 2010, some have assumed that new pipeline infrastructure would remove the bottleneck-effect in Cushing, leading to the WTI prices moving back toward the prices of Brent. A new large pipeline infrastructure, the "Seaway Pipeline" going from Cushing to the Mexican Gulf opened in 2012 without leading to the anticipated effects on the WTI price.

[13]Illustration of North American pipelines are retrieved from American Petroleum Institute (API)'s website www.api.org.

The Arab Spring in 2011 increased the oil demand for the Arab countries involved, since the countries increased public spending in an attempt to appease its citizens. This increased the purchase of Brent oil relatively to the WTI oil, since only Brent can be sold globally. This means that the Arab Spring influenced the prices of the Brent crude oil, while the prices of WTI oil was not influenced by these events.

Later in the thesis, in chapter 5, the present price spread between Brent and WTI will be considered as a constant transportation cost, and adjusted by imagining a buyer located in the U.S. East Coast with the possibility of either buying Brent oil or WTI oil. This approach assumes a well-functioning market, which in theory will even out arbitrage opportunities. Without arbitrage opportunities, the price spread is equal to the difference in transportation costs, meaning that the buyer will be indifferent to whether he/she buys North Sea oil and has it delivered by ships, or buys WTI oil and has it delivered by U.S. pipelines.

The assumed well-functioning market would even out any arbitrage opportunities; consider for example that the prices of Brent are lower than that of WTI (including transportation costs). Then, the buyer could profit from buying the Brent oil and have it delivered to somewhere in the U.S. East Coast as opposed to buying the domestic WTI oil with pipeline delivery. Then, according to the well-functioning market, the storage in Cushing would increase since all investors would buy Brent oil instead, leading to a decrease the WTI oil price, creating a rebalance of the price spread.

This introduction provides background to understanding the Brent/WTI price spread. The above-mentioned factors may influence the price spread to different degrees, and other factors not mentioned may also exist. Section 3.2.2 discusses more theoretically how a price spread can occur between two similar markets.

3 Theoretical Background

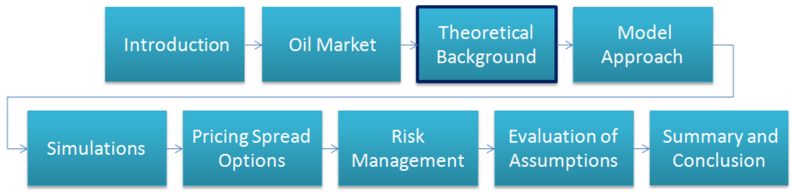

The theoretical background for this thesis is derivatives theory, price theory and an introduction to analysing the price process of a financial asset. Since this thesis is based on the prices of derivatives, the theoretical background will include a description of different types of derivatives, and their purpose and characteristics. The price theory is important to understand the relationship between the spot and futures price in the oil market, and factors affecting the spread between Brent and WTI are included. The chapter also explains how financial time series are analysed with stochastic processes to provide a basis for estimating the future prices and volatility.

3.1 Derivatives

During the last 30 years, trading in the derivatives market has become an increasingly important part of finance. But what exactly is a derivative? According to Hull (2012), "A derivative can be defined as a financial instrument whose value depends on (or derives from) the values of other, more basic, underlying variables". The variable on which the derivatives contract depend, can be almost any possible variable asset. Derivatives are most commonly traded on assets in the stock, commodity and interest rate market. Trading of derivatives is done at both exchanges, where the contracts are standardized and specified by the exchange, and in the OTC markets, where the specification of the trades are negotiated by the participants themselves. While trading in the exchange market is easier due to the high degree of standardization, the trading of derivatives in the OTC market is of higher volume. Hull (2012) presents measurements from 2009, where the OTC market is valued to be 614.7 trillion U.S. dollars, and the exchange-traded market is valued to be 73.1 trillion U.S. dollars. These measures estimate the total principal amounts underlying the transactions in the OTC markets, and the total value of the assets underlying exchange-traded contracts outstanding in December 2009.[14]

Trading in derivatives is very popular and has attracted many different types of traders due to the high liquidity in the derivatives market. When an investor wants to take one side of a contract, there is usually no problem in finding

[14]Statistics collected from The Bank of International Settlements are not exactly comparable for the two markets, since they compare the total principal for OTC markets to the total value for exchange-traded contracts. Read more about these statistics at the The Bank of International Settlements' website, www.bis.org.

someone who is prepared to take the other side. Three broad categories of traders can be identified:

Hedgers use derivatives to manage risk. A company, or an investor, that is exposed to risk in some market variable or asset (say, fluctuations in the oil price), may choose to hedge its exposure to eliminate or reduce the risk. Forward and futures markets are often used by risk managers to hedge risk, and the liquid futures prices help the price discovery mechanisms to determine the fair value for the future delivery. Although the purpose of hedging is to reduce risk, there is no guarantee that the outcome with a hedged position is better than the outcome in an unhedged position. Assume an oil producing company, profiting from increased future oil prices and suffering from decreased future oil prices. This company can hedge its commodity risk by locking in future sales at a future price. Locking in prices can be either profitable or non-profitable for the business; depending on whether the spot price is higher of lower than the agreed price at the time of the sale. Whether or not this hedge is profitable, the company will have a more certain future by selecting some sort of hedging strategy, as opposed to remain in an unhedged, or "naked", position.

Speculators use derivatives to bet on the future direction of a market variable or asset. Whereas hedgers want to avoid exposure to adverse movements in the price of an asset, speculators seek risk, and wish to take a position in the market in an attempt to earn a profit.

Arbitrageurs take offsetting positions in two or more instruments to lock in profit. This is done by simultaneously entering into transactions in two or more markets, where the futures price of an asset gets out of the line with its spot price at maturity. When this happens, it normally does not last long, as markets move continuously to close arbitrage opportunities.

The three pillars in derivatives are futures, forward and option contracts. Other variants of derivatives are structured based on these three. In all derivatives, the parties enter into either a long or a short position. Usually, the contracts are made up of two parties, whereas one holds the long position, and the other holds the short position. In a contract, the party that assumes the long position agrees to buy (or has the choice to buy) an asset at a given time and price, and the party that assumes the short position agrees to sell (or has the choice to sell) an asset at the same time and price.

3.1.1 Futures and Forward Contracts

A futures contract and a forward contract share many of the same features, as they are both contracts in which two parties agree to either buy or sell an asset at a certain future point in time for a certain future price. The value of a futures/forward price $F(t,T)$ can be found by compounding the present value $S(t)$ at time t to maturity T by the rate of the risk free return r:

$$F(t,T) = S(t) \times (1+r)^{(T-t)} \tag{1}$$

Compounding is the ability of an asset to generate earnings, which are then re-invested in order to generate their own earnings. This happens for example when placing money in the bank, where interest is earned on the initial amount. When the earned interest is compounded, new interest will be based on the previous amount + the compounded amount earned on the interest. In other words, compounding refers to generating earnings from previous earnings. Continuous compounding is an extreme case of compounding, which can be thought of as making the compounding period infinitesimally small. The value of the futures/forward price with continuous compounding is

$$F(t,T) = S(t)e^{r(T-t)} \tag{2}$$

This relationship only holds in a perfect market, and does not consider storage costs, dividends or convenience yields. In addition, market imperfections such as transaction costs, difference in interest rates between borrowing and lending, and trading restrictions will impact the price. In the following sections, the simple relationship is considered.

The **futures contract** is traded through organized exchanges, where the exchange standardize the contracts, and specifies the features of the derivative. This standardization helps to create liquidity in the marketplace, enabling participants to close out positions before expiration of a contract. Futures contracts are reported to the futures exchange, as well as a clearing house and at least one regulatory agency on a daily basis; which provides the futures contract with practically zero credit risk. The clearing house guarantees for the default risk by taking both sides of the trade and "marking to market" their positions. Marking to market is a process where daily gains and losses in futures contracts are converted into actual cash gains/losses each day, set when the exchanges close. Where one party has suffered a loss on the contract, its counterpart has gained, and the clearing house moves the payment through the process of marking to market. The futures contracts are regulated at federal government level to ensure that manipulation of prices does not occur.

The **forward contract** is not traded on organized exchanges, and hence have no regulatory agency or clearing house to insure the honoring of a contract. Since the exchange is not present to guarantee for the honoring of the contract, there is credit and default risk involved in the forward contracts. The forward contracts are considered as private transactions, and are usually between two financial institutions, or between a financial institution and one of its clients (Hull (2012). Forward contracts are often of larger size than futures, and they require modelling and customization to meet the user's special needs.

3.1.2 The Futures Market

The futures market is very liquid and transparent, and its prices are recorded and available from pricing services. Due to this high degree of liquidity and transparency, some key features and characteristics are the same for all futures contracts. Historical data on futures contracts for both Brent and WTI are presented in the Model Approach chapter. These historical futures contracts and prices lay the basis for fitting a model for simulating possible realisations of the future. Since the futures prices are the input to the model in chapter 4, the key features and characteristics for the oil futures markets are presented.

- Price

  The prices of futures are normally decided by the market, determined by the volume traded, as well as supply and demand. If, at a particular time, more traders wish to buy rather than to sell a commodity, the price will go up. Then, new sellers will enter the market so that a balance between the buyers and sellers are maintained. If more traders wish to sell rather than to buy the commodity, the price will go down. Then, new buyers will enter the market to maintain the balance between the buyers and sellers. Recall from the previous chapter, where section 2.2.1 discussed the oil price determination, that supply and demand are not the only factors determining the price of oil.

- Contract size

  The contract specifies the amount of the underlying asset that has to be delivered in one contract. In the oil market, one futures contract is normally based on 1,000 barrels of the underlying crude oil.

- Delivery

  Every contract has a delivery month, which is specified in the contract. The place of delivery for the futures that are settled in physical delivery, is also specified by the contract. Recall that the WTI futures are physically delivered in Cushing, Oklahoma, while the Brent futures are settled in cash with no physical delivery (the cash settled Brent futures can be exchanged for physical delivery).

- Asset features

  When the asset is a commodity, there may be a difference in the qualities available in the marketplace. As opposed to for example trading in Japanese yen, where there is no need to specify the features of the asset, the futures traded on commodities must go through a certain form of quality control.

- Daily settlement

  Futures contracts are settled daily for every investor on a margin account. Here, the investor has to enter a deposit fund called the initial margin at the start of the contract period. Every day, the account is adjusted to account for the daily losses/gains.

- Clearing house

  The exchanges work as clearing houses, and guarantee the honoring of each contract.

- Convergence of the spot price to the futures price

  One of most important mechanisms of the futures markets is the convergence of the spot price to the futures price: As the delivery period for the contract is approaching, the futures price and the spot price of the underlying asset will converge. With physical delivery, there will always be a convergence between the futures and spot price at maturity, assuming no arbitrage.

3.1.3 Understanding the Forward Curve

A forward curve is a curve illustrating the futures prices observed in the market at an exact date. The forward curve observed on a Monday is different from that observed the following Tuesday, etc, although forward curves are highly correlated. There is only one observed forward curve at a time, and this curve provides information on all futures contracts existing in the market for each asset. In the context of the oil futures market, an observed forward curve consists of 50 futures contracts at any given point in time. The 50 existing contracts have different expiration dates, where the first contract expires in 1 months, the second in 2 months, etc, up till contract number 50 which expires in 50 months. When the first contract expires, the second contract shifts position and becomes the first contract, and the previous contract number 50 shifts position and becomes contract number 49, giving room for another contract in the far end of the curve. This rolling system makes it possible to observe prices for 50 contracts going forward at any point in time, presenting a forward curve which gives information about the expectations in the market at that time. One month is considered as 21 trading days on the exchanges, and 50 contracts going forward provide information about expectations for the next 4.167 years.

The futures prices are especially important in the oil market, since oil purchased today is in general delivered at some point in the future (cannot be delivered right away). This makes the futures prices more essential than the spot prices in the oil markets, and hence making the forward curve the best way to illustrate the crude oil prices. Since the spot price always will converge to the futures price, one can also look at the observed spot prices (such as the Dated Brent), even though these will not include the cost of transportation, cost of storage, etc. The prices in the forward curve, however, will reflect all the relevant information publicly known in the market at the time. The forward curves can take on different shapes dependent on the level of transaction cost, the concept of convenience yield and changes in supply and demand (Clewlow and Strickland (2000)). The forward curve can be either in backwardation, in contango or in a mix between the two. If the forward curve is in backwardation, the futures prices are lower than the spot price today, as seen in Fig. 4. If the forward curve is in contango, the futures prices are higher than the spot price, as seen in Fig. 5. Normally, when the market is in backwardation, the oil producers will sell their oil right away, whereas in a contango market, it can be profitable to store the oil now, and sell it later at a higher price. Contango and backwardation will be more discussed in the price theory part of this chapter.

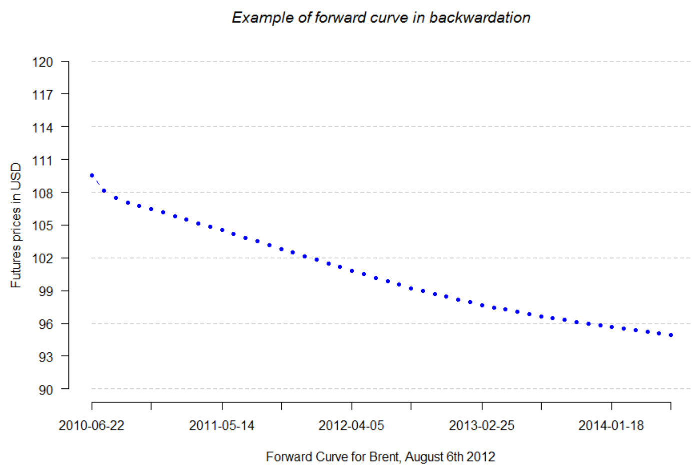

Figure 4: Example of the Brent forward curve in backwardation formation, from August 6th 2012

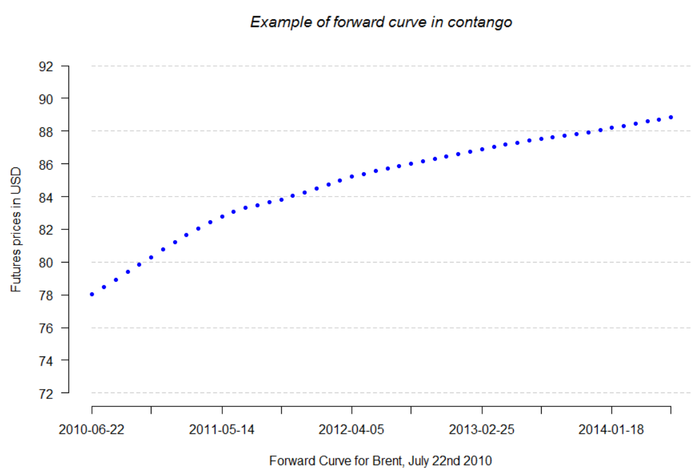

Figure 5: Example of the Brent forward curve in contango formation, from July 22nd 2010

3.1.4 Options and Option Payoff Profiles

An option is a right to buy or sell an underlying asset for a certain price by, or at, a certain date. The price agreed upon in the option contract is called the strike price, and is a fixed price for which the two counterparts agrees to buy/sell the underlying asset. The date in the contract is known as the expiration date or maturity. Options are fundamentally different from futures and forwards with two characteristics:

1. An option gives the buyer the right to either buy or sell an asset, but the buyer is not obligated to exercise this right. By contrast, the two parties in a futures or forward contract have committed themselves to some specified action.

2. In a futures or forward contract, there are no costs of entering into the arrangement (except for margin requirements). The purchase of an option however, requires an up-front payment, which is received by the counterpart selling the option.

Options are traded both at exchanges and in the OTC market. There are two basic types of options; a call option gives the holder the right to buy an asset by, or at, a certain date for a certain price, and a put option gives the holder the right to sell and asset by, or at, a certain date for a certain price. As with the forward and futures contracts, the price and time are agreed and settled as the contract is entered. The two basic types can either be bought (entering a long position) or sold (entering a short position), presenting four different options positions: buying a call option (long call), selling a call option (short call), buying a put option (long put), and selling a put option (short put).

"Moneyness" is an important term when considering options. The moneyness of an option depends on the relationship between the strike price of the derivative and the current price (or the price at expiry) of the underlying variable (Hull (2012)). There are three classifications of moneyness; in-the-money, at-the-money and out-of-the-money. If the option is in-the-money, the option value is positive. If the option is at-the-money, the current price and strike price are equal, and the investor will be indifferent to exercising the option or not. If the option is out-of-the-money, the current price and strike price is in such a relation that the investor would not profit from exercising the option, leading to a profit value of zero.

Options can be either European or American; a European option can only be exercised at the expiration date, whereas an American option can be exercised at any time up to the expiration date. This thesis will only analyse European-style options, and the option prices presented in chapter 6 are long put options. The profit functions for the four main option positions are all explained (ignoring discounting) to present an overview of the option profiles.

Long European Call
Fig. 6 illustrates the profit function for an investor buying a European call option on one share of an underlying asset. The payoff for a long position in a European call is, without the initial cost

$$\max(S_T - K, 0) \tag{3}$$

where S_T is the final price of the underlying asset, and K is the strike price. For the investor who has bought the call option, the option will only be exercised if $S_T > K$, since the investor will only buy if there is a discount to the market. Let's say the price of the asset at maturity is USD 90. The profit function max $(S_T - K, 0)$ gives max $(90 - 80, 0)$ = profit USD 10. See in Fig. 6 that the initial cost for the investor was USD 5, so the net profit will be USD 5. Here, the higher the price of the underlying asset, the higher the profit for the investor. Now consider the price of the asset at maturity to be 70. The profit function gives max $(70 - 80, 0)$ = USD 0. An investor would never purchase the asset for a higher price than what he/she can get in the market, therefore the option expires worthless. The net payoff for the investor is then the initial cost, USD -5.

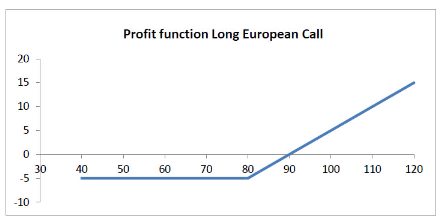

Figure 6: Profit from buying a European call option, Option price = USD 5, Strike price = USD 80

Short European Call
Fig. 7 illustrates the profit function for an investor selling/writing a European call option on one share of an underlying asset. The payoff for a short position in a European call is, without the initial profit

$$\min(K - S_T, 0) \tag{4}$$

where S_T is the final price of the underlying asset, and K is the strike price. For the investor who has sold the call option, he/she will only profit if the option is not exercised. Let's say the price of the asset at maturity is USD 90. The profit function min $(K - S_T, 0)$ gives min $(80 - 90, 0)$ = USD -10. Fig. 7 shows that the initial profit from selling the option is 5, so the net profit will be USD -5. Here, the lower the price of the underlying asset, the higher the losses are for the investor. Now consider the price of the asset at maturity to be 70. The profit function gives min $(80 - 70, 0)$ = USD 0. The net payoff for the investor is then the initial profit, USD 5. Here, the investor who has written the option will never profit more than the initial profit. Also, notice that the profit/loss for the investor who bought the option is exactly the opposite of the investor who sold the option.

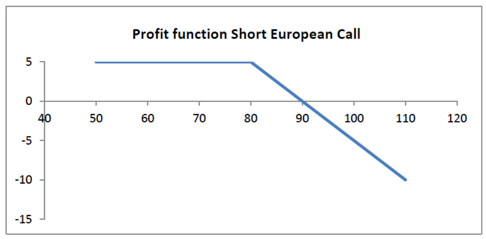

Figure 7: Profit from selling/writing a European call option, Option price = USD 5, Strike price = USD 80

Long European Put

Fig. 8 illustrates the profit function for an investor buying a European put option on one share of an underlying asset. The payoff for a long position in a European put is, without the initial cost

$$\max(K - S_T, 0) \tag{5}$$

where S_T is the final price of the underlying asset, and K is the strike price. For the investor who has sold the put option, the option will only be exercised if $K > S_T$, since the investor will only buy the option if the strike price is at a discount to the market price at maturity. Let's say the price of the asset at maturity is USD 90. The profit function max $(K - S_T, 0)$ gives max $(80 - 90, 0)$ = profit USD 0. An investor would never purchase the asset for a higher price than what he/she can get in the market, therefore the option expires worthless. Fig. 8 shows that the initial cost for the investor was USD 5, so the net profit will be USD -5. Now consider the price of the asset at maturity to be 70. The profit function gives max $(80 - 70, 0)$ = USD 10. The net payoff for the investor is then USD 5, since the initial cost was USD 5. Here, the lower the price of the underlying asset, the higher the profit is for the investor.

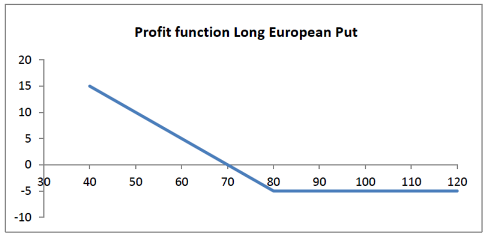

Figure 8: Profit from buying a European put option, Option price = USD 5, Strike price = USD 80

Short European Put

Fig. 9 illustrates the profit function for an investor selling/writing a European put option on one share of an underlying asset. The payoff for a short position in a European put is, without the initial profit

$$\min(S_T - K, 0) \tag{6}$$

where S_T is the final price of the underlying asset, and K is the strike price. For the investor who has sold the put option, he/she will only profit if the option is not exercised. Let's say the price of the asset at maturity is USD 90. The profit function min $(S_T - K, 0)$ gives min $(90 - 80, 0)$ = USD 0. Fig. 9 shows that the initial profit from selling the option is 5, so the net profit will be USD 5. Now consider the price of the asset at maturity to be 70. The profit function gives min $(70 - 80, 0)$ = USD -10. The net payoff for the investor is then USD -5. The lower the price of the underlying asset at maturity, the higher the losses will be for the investor. Here, the investor who has written the option will never profit more than the initial profit. Again, the profit/loss for the investor who bought the option is exactly the opposite of the investor who sold the option.

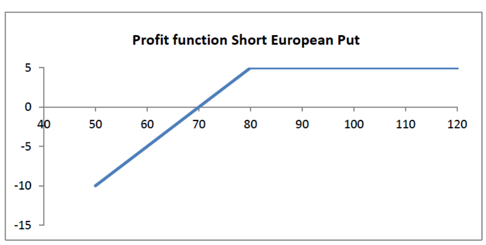

Figure 9: Profit from selling/writing a European put option, Option price = USD 5, Strike price = USD 80

To sum up, the above-mentioned examples illustrate that the long positions have a guaranteed minimum loss, and a limitless potential upside. On the contrary, the short positions have a limited upside, and a limitless potential downside.

3.1.5 Bid-Ask Spread and other Transaction Costs

Most option exchanges use market makers to facilitate trading (Hull (2012)). A market maker is an individual who quotes a bid and an offer (ask) price for a given option or another contract, upon request. The bid is the price at which the market maker is prepared to buy, and the offer is the price at which the market maker is prepared to sell. The offer is always higher than the bid, so that the market maker profits from facilitating the trade for the parties involved. The existence of the market makers ensures that buy- and sell orders can be executed without time delays, and they therefore add liquidity to the market. The amount at which the offer exceeds the bid price is called the bid-ask, or bid-offer, spread. The size of the bid-ask spread is a measure of the liquidity in the market as well as the size of the transaction costs, where a narrow bid-ask spread indicates great liquidity. If the spread is 0, it is a frictionless asset or market (only possible in theory). If the bid-ask spread is large, it implies low liquidity, hence large transaction costs. The size of the spread and the price of the asset are determined by supply and demand.

Although the bid-ask spread is a measure of the transaction costs for buying/selling an asset or option contract, other costs such as front-office, back-office costs and taxation will increase the costs of trading options. An investor trading options would expect the profit of his/her position to exceed the total costs of trading.

3.1.6 Futures Options

Until now, this thesis has introduced options based directly on the value of an underlying asset. It is also possible, and for many assets more common, with options priced on futures contracts, known as futures options. In these contracts, the exercise of the option gives the holder a position in a futures contract. The nature of futures options are similar to that of spot price options; a futures option is the right, but not the obligation to enter into a futures contract at a certain price by, or at, a certain date.

Why trade options on futures when it is possible to trade options directly on the underlying asset? Futures contracts are in many circumstances more liquid than the underlying asset on which they depend, as for example in oil, where the futures contracts are extremely liquid as opposed to trading the physical cargo's of crude oil. Futures prices have the advantage that they are known immediately from trading on the futures exchanges, whereas the spot price can be more difficult to find. In addition, futures on commodities are easier to trade than to trade the commodity itself, since futures options do not usually lead to delivery of the underlying asset. Recall that Brent futures are settled in cash, and that even though the WTI futures are settled physically, more than 99% of the positions in the WTI futures contracts are closed out prior to delivery.

3.1.7 Spread Options

A spread option is an option written on the difference between the prices of two underlying variables[15]; either spot prices or futures prices. Since this thesis will price spread option on futures prices of Brent and WTI, the payoff of a European style spread option is defined such that $F_{1,T}$ is the first underlying futures contract, and $F_{2,T}$ is the second underlying futures contract at time T for the four common option positions:

- Long European Call

$$\max(F_{1,T} - F_{2,T} - K, 0) \tag{7}$$

- Short European Call

$$\min(K - (F_{1,T} - F_{2,T}), 0) \tag{8}$$

- Long European Put

$$\max(K - (F_{1,T} - F_{2,T}), 0) \tag{9}$$

- Short European Put

$$\min(F_{1,T} - F_{2,T} - K, 0) \tag{10}$$

where K is the strike price for the spread options. Since the spread option values are derived from the difference between the prices of the two underlying contracts, the strike price can in this context also be called the strike price spread. An investor holding for example a long put position in spread options for Brent and WTI futures, bets that the spread will increase (or widen) compared to the initial spread. The larger the increase in the spread, the larger the profit from the option. The strike price spread can take any value, and the strike price spread used to price options in this analysis is the initial price spread between Brent and WTI observed from the forward curves at Day 1, valuing the spread options at-the-money.

[15] There exists several types of options which go under the name "spread options", such as calender spreads, production spreads and quality spreads. This thesis uses the definition of the spread option as an option written on the difference between the prices of two underlying variables.

3.2 Price Theory

The difference between the spot price and the futures price is worth explaining, and especially in the context of oil markets is this relationship important to examine. Key terms introduced in this section are the theory of storage, convenience yield and backwardation versus contango formation in the forward curves for oil markets. After this assessment, factors affecting the spread between Brent futures prices and WTI futures prices are introduced.

3.2.1 Relationship between Futures and Spot Prices for Commodities

Debreu (1959) introduces the dual concepts of a commodity and its price, and a mathematical approach to how prices function to balance supply and demand. The two concepts are simplified to the following: A commodity is characterized by its physical properties, the date at which it will be available and the location at which it will be available. The price of a commodity is the amount which has to be paid now for the (future) availability of one unit of that commodity.

The price of a commodity may be positive (scarce commodity), null (free commodity), or negative (noxious commodity). Each commodity has one of these prices, depending on the technology, the resources and the tastes of the economy. An economy is defined by m consumers (characterized by their consumption sets and their preferences), n producers (characterized by their production sets) and the total resources available. Debreu (1959) assumes that there is only a finite number of distinguishable commodities, and that the quantity of any one commodity is a real number. Debreu (1959) introduces a general theory on commodities, with special cases to focus attention on change of time, or change of locations. Meaning, the commodities are supposed distinguished by either the location where they are delivered, or by the date when they are delivered. By focusing attention on changes in location, a theory of location, transportation, international trade and exchange is suitable, and by focusing attention on changes of dates, a theory of savings, investment, capital and interest is more suitable. Read more about Debreu's Theory of Value in Debreu (1959).

Hotelling (1931) separates assets into two classes: Exhaustible assets and non-exhaustible assets. The exhaustible assets are also referred to as irreplaceable assets, and will have other characteristics than replaceable assets when it comes to the pricing process. An exhaustible asset can be any form of a scarce resource, for example minerals such as oil and gas, where great wastes can arise from the suddenness and unexpectedness of mineral discoveries. This suddenness and unexpectedness can lead to wild rushes to secure valuable property, where each owner wishes to collect the property as soon as possible, so that no one else achieves it before him/her. This rush makes storage of great volumes impossible, and may not be as economically profitable as a slower exploitation of the asset.

Hotelling (1931) assumes that the owner of an exhaustible supply wishes to make his/her present value of all future profits a maximum. The force of interest is denoted by γ , so that $e^{-\gamma t}$ is the present value of a unit of profit obtained after time t, interest rates being assumed unchanged in the period. It is a matter of indifference for a resource owner whether he/she receives for a unit of a product at price p_0 now, or a price $p_0 e^{\gamma t}$ after time t. This is the basis of "Hotelling's rule", as quoted in Devarajan and Fisher (1981), which states that the price of an exhaustible resource grow at a rate equal to the rate of interest, both along an efficient extraction path and in a competitive resource industry equilibrium. This means that, in a perfectly competitive market, price should exceed costs in the cases of exhaustible resources.

Hotelling's rule primarily addresses one basic question of the owner or agent involved in the exploitation of the non-renewable resource: How much of the asset should I consume now and how much should I store for the future? In other words, the agent has to choose between the current value of the asset if extracted and sold now, and the future increased value of the asset if left unexploited.

The "Theory of Storage" was first introduced by Working (1933), by describing in detail the futures markets in wheat and calculating the price spread between nearby and distant futures. In Working (1934), empirical research on wheat inventor is applied to the theory developed in Working (1933). Further work on the theory of storage was developed by Kaldor (1939), who states that the spread between spot prices and futures prices are determined by fundamental supply and demand conditions. Kaldor (1939) introduces the term "convenience yield", which he defines as "the possibility of making use of the commodity the moment when it is wanted". Kaldor (1939) separates between the yield for stocks on raw materials with storage possibilities, and all other assets. For storable commodities, the yield of a stock consists of the aforementioned convenience. Kaldor (1939) generalizes a theory for forward markets where the markets have yields consisting of convenience as follows:

$$FP - CP = i + c', \quad \text{hence} \quad FP = EP - r + q \tag{11}$$

where i is the interest cost, r is the marginal risk premium, c' is the carrying cost and q is the convenience yield. CP is the current (spot) price, EP is the expected future (spot) price and FP is the forward price.

Since the 1930's, convenience yield has grown to become an important term when analysing consumption commodities. For investment assets, that does not have storage possibilities, the convenience yield must be zero; otherwise arbitrage will be possible. According to Hull (2012), the convenience yield reflects the market's expectations concerning the future availability of the consumption commodity. It is the benefit from holding the physical asset compared to the financial asset. This benefit counts for those holding the commodity, to be able to meet unexpected demand, and for those dependent on the commodity as input in a production process. If the storage cost per unit is a constant proportion, u, of the spot price S_0, then the convenience yield y is defined so that

$$F_0 = S_0 e^{(r+u-y)T} \tag{12}$$

where r is the interest rate and T is the time at maturity. The greater the possibility that shortages will occur, the greater the convenience yield. If users of the commodity have high inventories, there is little chance of shortages in the near future, and the convenience yield tends to be low. If inventories are low, shortages are more likely to occur, leading to a higher convenience yield.

The relationship between futures prices and spot prices can be summarized in terms of the cost of carry (Hull (2012)). This measures the storage cost plus the interest that is paid to finance the asset less the income earned on the asset. For a commodity that provides income at rate q and requires storage costs at rate u, the cost of carry is $r - q + u$, where r is the interest rate. If the cost of carry is defined as c, the futures price for a consumption asset such as crude oil is

$$F_0 = S_0 e^{(c-y)T} \tag{13}$$

where y is the convenience yield.

Pindyck (2001) discusses the short run dynamics of commodity prices, production and inventories, as well as the effects of market volatility. Pindyck (2001) considers two interconnected markets; a cash market for spot purchases and sales of commodities, and a market for storage. The price of storage cannot be directly observed, but it can be determined from the spread between futures prices and spot prices. The price of storage is equal to the marginal value of storage, which in fact is the benefit for the inventory holder to hold one marginal unit of inventory; also referred to as the marginal convenience yield. Pindyck (2001) calculates the convenience yield from futures and spot prices as follows

$$\xi_{t,T} = (1 - r_T) P_t F_{t,T} + k_T \tag{14}$$

where P_t is the spot price at time t, $F_{t,T}$ is the futures price for delivery at time $t - T$, r_T is the risk free t-period interest rate, and k_T is the per-unit cost of physical storage. The formula assumes no arbitrage, and lets $\xi_{t,T}$ denote the capitalized flow of marginal convenience yield over the period t to $t + T$.

Eq. 14 shows that the futures price could be greater or less than the spot price, depending on the net storage costs. If marginal convenience yield is large, the spot price will exceed the futures price, $P_t > F_{t,T}$, and the futures market will be in strong backwardation. A separation can be made between strong backwardation and weak backwardation, where strong backwardation occurs when the spot price exceeds both the futures price as well as the discounted futures price. Weak backwardation occurs when the marginal convenience yield is positive, but not large. Then the spot price exceeds the futures price but not the discounted futures price, such that $F_{t,T} > P_t > F_{t,T}/(1 + r_T)$. When the marginal yield is zero, the spot price equals the discounted futures price, $P_t = F_{t,T}$, leading to zero backwardation. Pindyck (2001) defines contango as the state when the futures price exceed the spot price, $P_t < F_{t,T}$. Thus contango includes both weak backwardation and zero backwardation.

Crude oil is an extractive resource, and its futures market is therefore expected to be in either weak or strong backwardation most of the time. This is also the case, since holding inventory is equivalent to owning a call option with an exercise price equal to the extraction cost, with a payoff equal to the spot price of the commodity. Without the backwardation formation, producers would have no incentive to exercise the option, and there would be no production. When the forward curve is in contango, it is profitable to store oil today and sell it later. The higher the spot price volatility, the more valuable it is to owning in ground reserves, making producers require strong backwardation in the futures market in order to produce the oil. Although the "normal" conditions for crude oil is backwardation, the forward curve goes into contango on occasion, as illustrated by the forward curve in Fig. 5.

3.2.2 How Does a Spread occur between Similar Markets ?

This thesis defines the spread between Brent and WTI futures at time t with maturity at time T as

$$\text{Spread}_{t,T} = FB_{t,T} - FW_{t,T} \tag{15}$$

where $FB_{t,T}$ is the price of a Brent futures contract at time t with maturity at T, and $FW_{t,T}$ is the price of the WTI futures contract at the same time and with the same maturity as the Brent contract. Since the price spread is the difference of two prices, changes in the price spread will result from non-parallel movements in either FB or FW or both. Below is discussed some possible variables affecting the spread, presented by Milonas and Henker (2002).

When it is assumed more valuable to own a spot commodity contract rather than a distant futures contract, a yield will appear. **Convenience yield** is the incremental value of spot price over futures prices after accounting for carrying costs. If convenience yields are part of both the Brent and WTI futures price, their price spread will be due to the relative changes in the two convenience yields.

The cost of shipping oil from the delivery point to alternative refineries enter into the pricing structure of the oils under consideration. As long as the transportation cost structure does not change, it is not expected that the **transportation costs** itself will influence the price spread. However, over long periods of time, substantial changes in the transportation cost structure will explain part of the volatility in the price spread.

The rate of change in the demand of and supply for oil is not kept constant throughout the year. Disruption in production occurs more often in the winter months compared to the summer months, especially for Brent oil that is more vulnerable to extreme weather conditions. Demand for oil is affected by other conditions, such as "the driving seasonal in the U.S." in the summer months, and by heating oil needs in the winter months. **Seasonal factors** may not affect the prices of Brent and WTI equally and may therefore be an input to the price spread of the two crude oils.

Temporary divergence in demand/supply may explain some variation in the price differential between the two oil markets. If U.S. demand for oil is not moving in the same direction of as the oil inventory build-up, this may affect the prices of WTI and not impact the Brent prices. Similarly, temporary imbalances in the European market are not likely to induce changes in the WTI prices. Therefore, **inventory levels** are likely to explain a portion of the price spread.

Volatility in the price of the underlying cash commodity determines some of the variability in convenience yields.The greater the volatility in cash commodities, the greater chance that the cash price will exceed the futures price, leading to a greater convenience yield. By affecting the size of the convenience yield, the **cash volatility** is expected to affect the price spread directly.

The five above-mentioned variables which are expected to affect the price spread between Brent and WTI oil futures are not necessarily independent from each other. Convenience yield is the premier variable, which includes all the other four variables. All else equal, the convenience yield is expected to decrease by the building of production and inventories, the magnitude by which the supply is greater than demand, the increase in the cost of inventories and the decrease in the volatility of the underlying commodity. The opposite relationship is valid when the convenience yield is increasing.

3.3 Pricing Process for Crude Oil Markets

Analysing financial time series involves analysing historical data and estimating the volatility present in the dataset. The analysis of the historical data requires a thorough approach, since this analysis lay the basis for estimating future prices. This section explains the term volatility as utilised during the rest of the thesis, and how stochastic processes are used to explain the price process for Brent and WTI futures.

3.3.1 Characterising Volatility

Before introducing volatility in financial assets, four important statistical terms in probability theory will be introduced: Expectation, mean, variance and standard deviation. Before introducing these terms, the probability density function will be briefly presented. Adams (2003) defines the probability of an event occurring as the real number between 0 and 1 that measures the proportion of times the event can be expected to occur in a large number of trials. Denote X as the random variable having a distribution of some possible outcomes within some range. X can be either a discrete- or a continuous-time variable; if X is discrete it can take any value of some pre-determined values within a certain range. For example, when throwing a die, the range for X is [1,6] and the possible outcomes are 1,2,3,4,5 or 6, with no values in between possible. If X is a continuous variable, it can take any real value within a certain range. For example, suppose a needle is dropped at random on a flat table with a straight line drawn on it. For each drop, X is the number of degrees in the angle that the needle makes with the drawn line. Then, X can take any real value in the interval [0,90] and is a continuous random variable since there are infinitely many real numbers in the interval. When considering a continuous random variable, Adams (2003) defines the probability density function as a function defined on an interval $[a, b]$ for a continuous random variable X distributed on

$[a, b]$ if, whenever x_1 and x_2 satisfy $a \le x_1 \le x_2 \le b$, giving

$$Pr(x_1 \le X \le x_2) = \int_{x_1}^{x_2} f(x)dx \tag{16}$$

The following definitions by Adams (2003) for expectation, mean, variance and standard deviation all assume X as a continuous random variable.

If X is a continuous random variable on $[a, b]$, with probability density function $f(x)$, the **mean** (denoted μ), or **expectation** of X (denoted $E(X)$), is

$$\mu = E(X) = \int_a^b xf(x)dx \tag{17}$$

The **variance** of a random variable X with density $f(x)$ on $[a, b]$ is the expectation of the square of the distance from X from its mean μ. The variance is denoted σ^2 or $\mathrm{Var}(X)$.

$$\sigma^2 = \mathrm{Var}(X) = E((X-\mu)^2) = \int_a^b (x-\mu)^2 d(x)dx \tag{18}$$

The **standard deviation** of the random variable X is σ, the square root of the variance. Thus, it is the square root of the mean square deviation of X from its mean.

$$\sigma = \left(\int_a^b (x-\mu)^2 f(x)dx\right)^{1/2} = \sqrt{E(X^2) - \mu^2} \tag{19}$$

The standard deviation gives a measure of how spread out the probability distribution of X is. The smaller the standard deviation, the more concentrated is the area under the density curve around the mean, and so the smaller is the probability that a value of X will be far away from the mean.

The term volatility is closely related to the aforementioned statistical terms, and is applied especially in finance as a measure of uncertainty of the investment rate of return (Tsay (2005)). Volatility is a tendency of a security or asset to rise or fall within a period of time, and empirically, this often means the calculated standard deviation over a given time interval. When volatility is calculated as the standard deviation of a set of historical returns[16], the calculated volatility will be a constant number, for example "annual volatility over the last 5 years have been 30 %". Tsay (2005) points out some mutual characteristics for volatilities commonly seen in asset returns. First, there exist volatility clusters, stating that volatility may be high for certain period of time and low for other periods. Second, volatility moves over time in a continuous manner, meaning

[16]The procedure for finding volatility in this thesis is through historical data, based on five years information about the futures prices for Brent and WTI. For option prices, another popular method for analysing volatility is implied volatility - the volatility observed in the market. Since this thesis estimates volatility from futures contracts, implied volatility is not further researched.

that volatility jumps are rare. Third, volatility varies within some fixed range, and does not diverge to infinity. Fourth, volatility in asset returns react differently to positive and negative price drops, where a large price decrease will create a larger volatility than a price increase of the same amount. This is called the leverage effect and will be further explained in chapter 4.

These common characteristics for a financial asset, such as an oil futures contract, introduce us to non-constant volatility. It is evident that simplifying volatility to a constant number is not sufficient to describe the reality of the volatility in financial markets. In chapter 4, the multi-factor forward curve model allows for non-constant volatility, making the model a realistic tool to accept the real dynamics of the oil futures prices.

3.3.2 Analysing Financial Time Series

Hull (2012) defines a stochastic process as "any variable that changes over time in an uncertain way", and Tsay (2005) defines a stochastic process as "a statical term used to describe the evolution of a random variable over time". Stochastic processes are used to explain the price process of many financial assets and to price options. A stochastic process can be categorized as either a or continuous time process. In a stochastic discrete time process, the value of the variable can change only at certain fixed points in time, and the values can only take certain values. In a stochastic continuous time process the value of the variable can change at any time, and the values can be of any value within a certain range. In real life, the continuous time stochastic process does not exist, as the exchanges have opening hours, and every traded asset has some restriction to their values (such as multiples of cents for futures contracts for crude oil). However, looking at price processes that follow a continuous time process presents information on how to price derivatives, and techniques to analyse financial time series.

This thesis considers variables that follow Markovian stochastic processes. A Markov process is a particular type of stochastic process, where the prediction of the future only depends on the current value of a variable. In a price process, only the price today matters for prediction the future, not the price yesterday, or any other days in the past. The stock market is assumed to follow a Markov price process, and it implies that the present price of any stock is a result of all the information available in the market at that time. Meaning, the price today incorporates all information known about the past and present market, making explicit historic prices irrelevant since they are already accounted for in today's price. The same assumptions are valid for the oil futures market (Tsay (2005)).

A Wiener process is a particular type of the Markov stochastic process, where the expected mean change is zero, and the variance rate is 1.0 per year. To be classified as a Wiener process, the following two properties must be met:

1. The change Δz during a small period of time Δt is

$$\Delta z = \epsilon\sqrt{\Delta t} \tag{20}$$

where ϵ has a standardized normal distribution ϕ (0,1).

2. The values of Δz for any two different short intervals of time, Δt, are independent.

These two properties show that Δz itself has a normal distribution, and that z follows a Markov process. As the small changes become closer to zero, it is possible to proceed from small changes to the limit, thus $dx = adt$ indicates that $\Delta x = a\Delta t$ is in the limit as $\Delta \to 0$ (Hull (2012)). The Wiener process can be referred to as dz when assuming that the process has the properties for Δz given above in the limit as $\Delta \to 0$.

The process introduced so far is called the basic Wiener process, dz, and has a drift rate of zero and a variance rate of 1.0. The drift rate is the mean change per unit time for a stochastic process, and the variance rate is the variance per unit time. For a Wiener process, the drift rate of zero means that the expected value of z at any future time is equal to its current value. The variance rate of 1.0 means that the variance of any change in z in a time interval of length T equals T. The basic Wiener process can be defined as a General Wiener process, as illustrated in Fig. 10, for a variable x in terms of dz as

$$dx = adt + bdz \tag{21}$$

where a and b are constants. The adt term implies that x has an expected drift rate of a per unit of time dt. The bdz term adds variability to the path followed by x. The amount of variability added to the path is the constant b times a Wiener process dz, as shown in Fig. 10. The bdz term has a variance rate per unit time of b^2. The first term is therefore the linear direction of the Generalized Wiener process, and the second term is the uncertainty in the path of the process.

So far, this section has only discussed how the stochastic process for a single variable can be represented. Since this thesis models the spread between Brent and WTI futures prices, two correlated price processes are present, which both follow stochastic processes. First, briefly; what is correlation?

Correlation is derived from the term covariance. Walpole et al. (2012) defines covariance as "The measure of the nature of the association between two random variables, X and Y". If X and Y are continuous, then

$$\sigma_{XY} = E[(X-\mu_X)(Y-\mu_Y)] = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}(x-\mu_X)(y-\mu_y)f(x,y)dxdy \tag{22}$$

If large values of X often result in large values of Y or small values of X result in small values of Y, positive $X - \mu_X$ will often result in positive $Y - \mu_Y$ and negative $X - \mu_X$ will often result in negative $Y - \mu_Y$. Thus, the product of $(X - \mu_X)(Y - \mu_Y)$ will tend to be positive, and the opposite if large/small values of X result in small/large values of Y. The *sign* of the covariance indicates whether the relationship between two dependent variables is positive or negative. If X and Y are statistically independent, the covariance is zero. The covariance between two variables provide information regarding the nature of the relationship, but the magnitude of σ_{XY} does not specify the

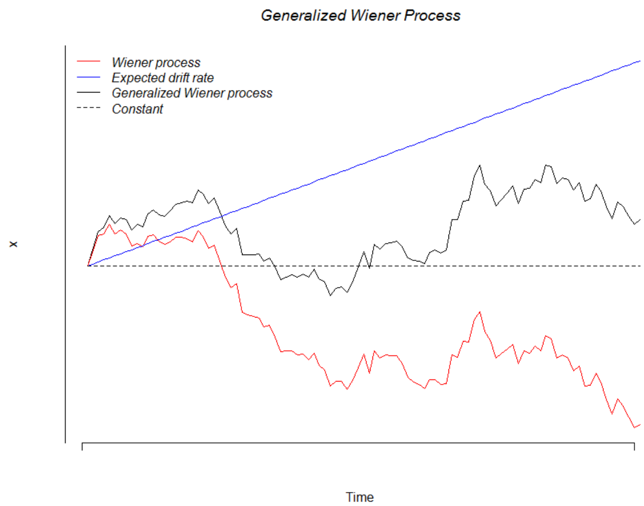

Figure 10: Generalized Wiener process with a = 0.01 and b = 0.05. The red line is a Wiener process dz, the blue line is the expected drift rate adt, the dotted line is the constant b, and the black line is the Generalized Wiener process $adt + bdz$

strength of the relationship, since σ_{XY} is not scale-free. Its magnitude will therefore depend on the units used to measure both X and Y. The correlation coefficient is the scale-free version of the covariance, widely used in statistics. Let X and Y be random variables with covariance σ_{XY} and standard deviations μ_X and μ_Y, respectively. The correlation coefficient of X and Y is

$$p_{XY} = \frac{\sigma_{XY}}{\sigma_X \sigma_Y} \tag{23}$$

Then, p_{XY} is scale-free of the units for X and Y.

Now, consider X_1 and X_2 to be two correlated price processes. The process followed by the two variables are

$$dx_1 = a_1 dt + b_1 dz_1 \quad \text{and} \quad dx_2 = a_2 dt + b_2 dz_2 \tag{24}$$

where a and b are constants, as in Eq. 21, and dz_1 and dz_2 are Wiener processes. Assume ϵ_1 and ϵ_2 as two sample standard variables from a standard normal distribution $\phi(0, 1)$ for x_1 and x_2, respectively, with correlation p. Obtaining samples for correlated standard normal variables for the two prices can be done as follows

$$\epsilon_1 = u \quad \text{and} \quad \epsilon_2 = \sqrt{1 - p^2} v \tag{25}$$

where u and v are uncorrelated variables with standard normal distributions. The model approach for capturing the correlation between Brent and WTI is illustrated in the next chapter.

4 Model Approach

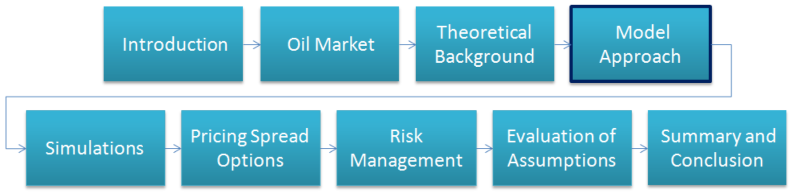

There are different types of existing models used for modelling energy derivatives. The most common method on a historical basis is the spot price model, where spot prices observed in the market lay the basis for modelling future prices. The two-factor model of Schwartz and Smith (2000) is briefly explained as an example of a spot price model. However, the spot price models suffer from some setbacks, which have lead this thesis to use Clewlow and Strickland's (2000) multi-factor forward curve model, calibrated based on the approach of Sollie (2013). This section explains background to the multi-factor forward curve model, and shows how the model is calibrated to market data on historical forward curves for Brent and WTI futures prices. The model allows for non-constant volatility in the forward curves by analysing the volatility time structure of the market with an EGARCH model, and Principal Component Analysis is performed to capture the driving factors of the innovations of the time series.

4.1 Introduction to Derivatives Pricing Models

The two main methods for pricing energy derivatives are spot price models and forward curve models. Spot price models use observed spot prices as input, while in the forward curve models, the observed forward curves are the input to the model. Historically, the majority of work on modelling energy and commodity prices has been focused on spot price models (Hull (2012)). In this case, the forward curve is the output from the model. The output of the model is highly dependent on certain key state variables that might not be observable in the market. Therefore, the output from the model may not correspond to the observable market parameters. Many industry participants now require the forward curve, which is observable in the market, to be an input to the model in order for the model to better replicate the market.

4.1.1 Two-Factor Spot Price Model

Probably the best known assumption to illustrate the dynamics of a commodity spot price is given by this well-known stochastic differential equation:

$$dS = (r - \delta)Sdt + \sigma Sdz \tag{26}$$

This equation represents a risk neutral process, where S is the commodity spot price, r and δ are constants describing the short term interest rate (risk free rate) and the convenience yield, respectively, and σ is the volatility of proportional change in the spot commodity price (Hull (2012)).

An established spot price model is the Two-Factor model introduced by Schwartz and Smith (2000). This model is often referred to as the short-term/long-term model, since the model is based on two variables, one long term and one short term, explaining the dynamics in commodity prices. The model allows for mean reversion in short term prices and for uncertainty in the equilibrium level to which the prices revert. These two factors are not directly observable, but they can be estimated from observed futures prices and spot prices. Movements in prices for long maturities provide information about the equilibrium level, whereas the difference between prices for short and long term contracts provide information about the short term variations in the prices.

Assume S_t is the commodity price at time t. The spot price is decomposed into two stochastic factors as ln(St) $= \chi_t + \xi_t$, where ξ_t is the equilibrium price level, and χ_t is the short term deviation in the price. In Schwartz and Smith (2000), χ_t is assumed to follow an Ornstein-Uhlenbeck process[17] and ξ_t is assumed to follow a Wiener process with drift;

$$d\chi_t = -\kappa\chi_t dt + \sigma_\chi dz_\chi \tag{27}$$

$$d\xi_t = \mu_\xi dt + \sigma_\xi dz_\xi \tag{28}$$

where dz_χ and dz_ξ are correlated increments of the standard Wiener process with $dz_\chi dz_\xi = p\chi\xi dt$, where p is the correlation, as introduced in section 3.3.2. The model can be calibrated in various ways, but the most transparent way of calibration is via state space form and the Kalman filter using maximum likelihood techniques. See Schwartz and Smith (2000) for further details on Kalman filtering and model calibration.

Schwartz and Smith's long-term/short-term model assumes the equilibrium price (the long term variable) to evolve according to Geometric Brownian Motion with drift, reflecting expectations about the existing supply, improving technology, inflation, regulatory and political effects. The short term deviations (the short term variable) are expected to revert to zero (since they describe the difference between the spot and equilibrium price). The short term deviations can, for example, explain short-term changes in demand resulting from variations in the weather conditions, and are tempered by the markets ability to adjust inventory in response to changing markets (Schwartz and Smith (2000)). The short term deviations are therefore not expected to persist, whilst the changes in the equilibrium represent fundamental changes that are expected to persist.

Spot price models have the strength that they are easy to interpret, and that they allow for an explicitly expressed spot price process. The shortcomings of

[17] The Ornstein-Uhlenbeck process is usually considered for the velocities (speed rate) of Brownian particles, as stated in Bibbone et al. (2008).

spot price models are that they may be complex to estimate, and that the model output in general do not match the observed market assets correctly. Since spot prices are formulated explicitly, the dynamics of the calibrated model will rely on stable parameters over time. It is difficult to calibrate the model with fixed inputs, and therefore the spot price model is subject to errors caused from the model parameters. The spot price models will generally not provide a perfect fit to the observed forward curves, leading us to look more closely at forward curve models.

4.1.2 Single Factor Forward Curve Model

Forward curve models explicitly model all the forward prices simultaneously. A single factor model of the forward curve is represented by this stochastic differential equation:

$$\frac{dF(t,T)}{F(t,T)} = \sigma e^{-\alpha(T-t)} dz(t) \tag{29}$$

where $F(t,T)$ denotes the futures/forward price observed at time t for maturity date T, and where $\sigma e^{-\alpha(T-t)}$ is a single factor associated with the source of risk $dz(t)$. The equation does not have a drift term, since both futures and forward contracts can be entered into without any initial investment. The model assumes a risk free world, where the expected return is zero. The single factor $\sigma e^{-\alpha(T-t)}$ can also be explained as a volatility function, and it illustrates that short dated futures/forward returns are more volatile than long dated futures/forward returns. This can simply be explained by the fact that information occurring in the market at present time will have greater impact on the short term rather than on the long term futures/forward prices (Hull (2012)).

The stochastic differential equation for the single factor model can be generalized to:

$$\frac{dF(t,T)}{F(t,T)} = \sigma(t,T) dz(t) \tag{30}$$

where $F(t,T)$ denotes the futures/forward price observed at time t for maturity date T, and $\sigma(t,T)$ is the volatility at time t. The form in which $\sigma(t,T)$ will take can be determined from market data. The single factor model has one major shortcoming: Even though it is established that the volatility will decrease from short maturities to long maturities, the single factor model forces the volatility for the longer maturities to go to zero too quickly. This leads the model to underestimate the volatility of longer maturity futures/forward prices, and this will lead to a downward bias to, for example, option prices that depend on the back-end part of the curve (Hull (2012)). Although futures and forward prices of different maturities are highly correlated, they are not perfectly correlated; the curves generally move up and down together, but they also change shape in several complex ways. Therefore, it is necessary with more than one factor to explain the drivers of all the forward curves simultaneously, introducing the multi-factor forward curve.

4.2 Multi-Factor Forward Curve Model

The multi-factor model in Clewlow and Strickland (2000) allows for two or more common factors as input to the model in order to explain the dynamics of all forward curves. The multi-factor model can be shown as

$$\frac{dF(t,T)}{F(t,T)} = \sum_{i=1}^{n} \sigma_i(t,T)dz_i(t) \tag{31}$$

where $F(t,T)$ denotes the futures/forward price observed at time t for maturity date T, with n independent sources of uncertainty $dz_i(t)$, and where $\sigma_i(t,T)$ are the associated volatility functions. The volatility functions multiplied with the independent sources of risk will drive the evolution of the forward curves. The number of volatility functions can in theory be large, but since the forward curves are highly correlated, not that many common factors are normally necessary in order to describe the price process. The method used for reducing the common factors to a reasonable level in this thesis is Principal Component Analysis, which will be explained in section 4.4.2.

Sollie (2013) shows that the changing rate of mean reversion of spot prices to futures prices can be incorporated into the forward curve model by allowing for non-constant volatility. While the forward process in a forward curve model is simple, the implied spot price dynamics is generally not. In Sollie (2013), a method for simulating spot prices via a simple approximation is introduced. As stated, spot price models allow for straight forward formulation of spot price dynamics, while this convenience is not generally present in forward curve models. Recall that the futures price is the expected future spot price, as explained in section 3.2, adjusted for risk premium and convenience yield. Regardless of how the spot price and futures price evolve over time, the spot price must converge to the futures price at maturity of the futures contract, so that

$$S(t) = F(t,t) \tag{32}$$

meaning that the spot price in the forward curve model only is observed at maturity of the futures contract. In the oil futures market, where one contract expires every 21th day, the spot price will only be observable at one time per month when utilizing a forward curve model. In Sollie (2013), the spot prices are included in the principal component analysis and the volatility time structure. This can be done since the volatility will affect the rate at which the spot price converges to the futures price. The rate of mean reversion of spot to futures prices is not-constant, and since the multi-factor forward curve model allows for non-constant volatility, it also allows for non-constant mean reversion rate of spot to futures prices. The volatility time structure is interpolated between the observed data point, resulting in an approximation giving observed spot prices as one futures contract matures every day. This secures the convergence of spot to futures prices, and providing spot prices with approximately correct values and dynamics. The small errors that may happen due to this approximation are considered negligible for most models and markets in Sollie (2013).

4.3 Introduction of Data: Brent and WTI Futures Contracts

The data used in this analysis are daily observed prices[18] for two futures markets; Brent and WTI crude oil. The observations collected span from 02-01-2008 to 18-01-2013, and are used as the input for the model construction. The collection contain daily observations of the forward curves for Brent and WTI, which at each point in time consists of futures contracts going forward in time with different maturities. Fig. 11 illustrates the historical forward curves for Brent and WTI futures data used as input to the model, with Brent observed in the front of the figure, and WTI observed in the back of the figure.

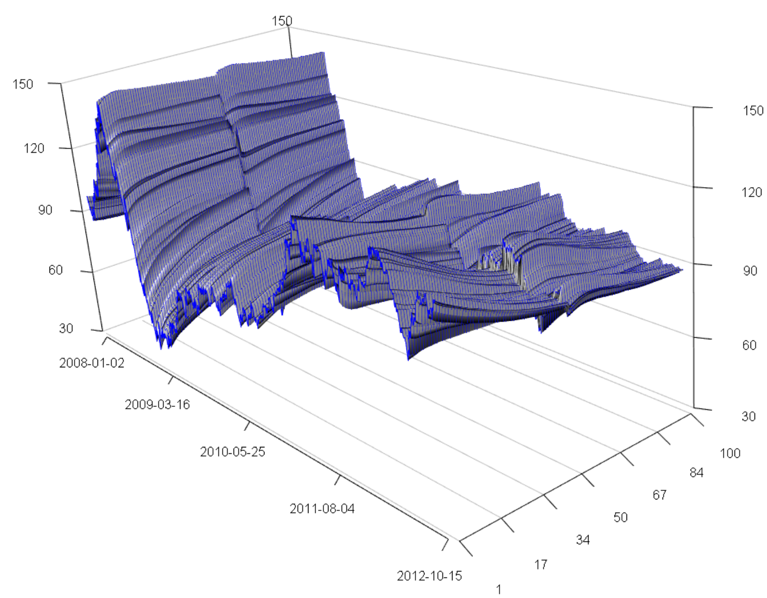

Figure 11: Historical forward curve movement with Brent in the front and WTI in the back

Fig. 11 illustrates how the selected period of time has had great variation in both prices and curvature structure. The elapse of the financial crisis in late autumn 2008 is clearly visible in the forward curves, as the prices dropped from levels of about USD 135 to USD 35 in a matter of months. During the downfall period around late 2008 - early 2009, most of the forward curve observations for both Brent and WTI were in contango formation. This indicates that in that

[18]Input data are retrieved from the Statoil database.

period, investors believed that the future oil price would rise from the observed spot prices at that time. Around Christmas 2010, the Brent and WTI forward curves disentangled from their, up until that time, very similar characteristics. The front-end of the Brent forward curves rose, placing the Brent forward curves in backwardation curvature, a position in which the curves have remained since. In the same period the WTI forward curves have been in contango, or contango-like formations, where at least the front-to-mid-end of the curves have been in contango (recall the discussion in chapter 2 regarding the price spread between Brent and WTI).

Calibrating the multi-factor forward curve model for Brent and WTI allows us to consider historical data and make calculations regarding variance, volatility and correlation in the data set. Fig. 12 shows the annualized volatility for the two time series in the period from 2008 to 2013. Notice how the volatility reached extreme levels of almost 90% during the elapse of the financial crisis.

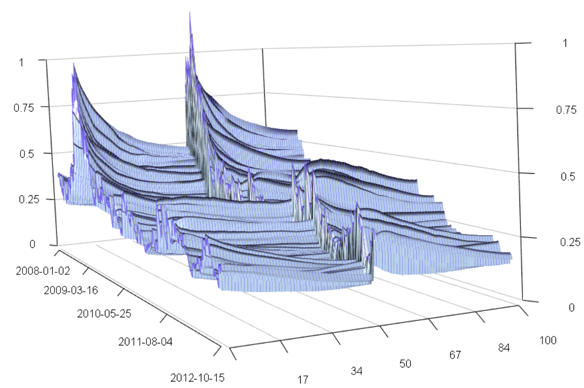

Figure 12: Historical volatility time structure for Brent and WTI

Fig. 12 shows how the volatility in the front-end of the forward curves is normally higher than in the back-end of the curves. An exception is visible in the WTI forward curves during the last two years, where the volatility in some curves are in contango-like shapes. Notice also that the WTI curves have a "spike" at their very front, thus being more volatile at this part of the curve than their Brent counterpart. Regardless, it is obvious from the figure that the volatility for the crude oils can not be considered constant.

4.4 Volatility Time Structure of the Market

The forward curve models rely on the observed forward curves and a volatility time structure, which does not need to be constant. Variance equations can update the local volatility and thus let the volatility time structure evolve. This results in a volatility time structure which evolves according to the variance equations, but is driven by the innovations of the underlying factors found by the principal component analysis. The number of factors to include is determined by carefully examining the covariance matrix of returns for the futures contracts for Brent and WTI. This section contains two parts, where the first part explains how the variance equations are estimated using maximum likelihood techniques to solve for an EGARCH model, and the second part explains how principal component analysis is performed to find a set of common factors to explain the volatility and correlations of the asset returns.

4.4.1 Variance Estimation using EGARCH

Autoregressive conditional heteroskedasticity (ARCH) models are used in econometrics to model and characterize observed time series. The AR means that these models are autoregressive models in squared returns. Heteroscedasticity means non-constant, or changing, volatility (i.e., variance). The conditional comes from the fact that, in these models, next period's volatility is conditional on information in this period. The conditional volatility quantifies the uncertainty about the future observation, given all information available at present time. The ARCH model, first introduced by Engle (1982), captures the tendency in financial data where the volatility move in clusters; large returns are followed by more large returns, and small returns are followed by more small returns, a characteristic of financial data also stated by Pindyck (2001). This suggests that returns are serially correlated, and the ARCH models are used whenever there is reason to believe that, at any point in a series, the terms will have a characteristic size, or variance. ARCH models assume the variance of a current error term, or innovation, to be a function of the actual sizes of the previous time period's error term. Often the variance is related to the squares of the previous innovations.

GARCH models are Generalized ARCH models, first developed by Bollerslev (1986). In an ARCH model, the next period's variance only depend on the last period's squared residuals/errors, so that a crisis caused by a large residual would not have the sort of persistence that has been observed after actual crises, as stated by Bollerslev (1986). The GARCH model is an extension of the ARCH model, where the next period's forecast of variance depend on both the last period's squared residuals, as well as the last period's forecast. Here, the unconditional variance is constant, while the conditional variance of the time series evolves over time (Tsay (2005)). The conditional variance is the weighted sum of past squared residuals, and the weight decrease as you go further back in time. See Bollerslev et al. (1992) and Bollerslev et al. (1994) for more details on ARCH and GARCH models.

The ARCH and GARCH models have one main weakness when considering the datasets for Brent and WTI futures prices; they respond equally to positive and negative shocks (Tsay (2005)). To overcome this weaknesses of the GARCH model in handling financial time series, Nelson (1991) proposes the exponential GARCH (EGARCH) model. In the EGARCH model, the variance equation can be formulated to allow for this asymmetric response of volatility to returns. The EGARCH variant models the logarithm of the conditional variance process, and uses this logged conditional variance to relax the positiveness constraint of the model coefficients. This alteration and adjustment from the ARCH and GARCH models allows for asymmetric effects between positive and negative asset returns. This fits with financial time series on commodities such as forward curves, since there often exists a leverage effect in this type of volatility structure, as noted in section 3.3.1. The leverage effect explains how negative shocks impact the volatility to a larger extent than positive shocks of the same size. This can be explained by a simple example: An investor is in general more upset by loosing, than he/she will be happy by winning the same amount. This effect is also called the endowment effect in behavioural finance, introduced by Thaler (1980).

Based on the findings of Sollie (2013), an asymmetric exponential generalized autoregressive heteroskedastic (EGARCH) is chosen to model the variance. The EGARCH model can be stated as

$$\ln(h_{i,t}) = \alpha_{i,0} + \alpha_{i,1}\frac{|r_{i,t-1}| + \gamma_i r_{i,t-1}}{\sqrt{h_{i,t-1}}} + \beta_i \ln(h_{i,t-1}) \tag{33}$$

for contract i. The equation illustrates how 4 parameters will drive the variance. Notice how the model uses $\ln(h_{i,t})$ instead of just $h_{i,t}$ to avoid the possibility of negative variance. $\alpha_{i,0}$ can be interpreted as the constant of the natural logarithm of the local variance, where the local variance is a mean reverting process. $\alpha_{i,1}$ measures the effect of a change in the previous innovation, standardized by local volatility. The equation illustrates how a positive $r_{i,t-1}$ contributes $\alpha_{i,1}(1+\gamma_i)|r_{i,t-1}|/\sqrt{h_{i,t-1}}$ to the log volatility $\ln(h_{i,t})$, whereas a negative $r_{i,t-1}$ contributes with $\alpha_{i,1}(1-\gamma_i)|r_{i,t-1}|/\sqrt{h_{i,t-1}}$, thus γ_i signifies the leverage effect, which is expected to be negative in real applications due to the endowment effect amongst investors. β_i is the autoregressive parameter of the natural logarithm of local log variance of contract i.

The four parameters in the EGARCH model are unknown since they are not directly observable. However, the parameters can be estimated using maximum likelihood techniques for the observed data points for Brent and WTI, as shown in Fig. 13 and 14. This technique optimizes the value of each parameter to match the variance in the observed data. A value is chosen for each parameter, such that the chance (or likelihood) of the data occurring is maximized (Hull (2012)).

Below is illustrated a general approach for maximum likelihood estimation. Assume observations $x_1, , x_n$ of a phenomenon which is described by a probability function $f(x;\theta)$. The value of θ is unknown. If the observations are independent and identically distributed (i.i.d.), they can be used to estimate

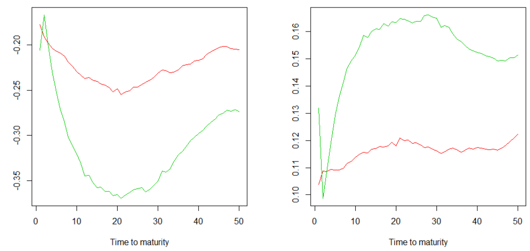

Figure 13: To the left, $\alpha_{i,0}$ is a constant, measuring local variance. To the right, $\alpha_{i,1}$ measures the change in previous innovations, standardized by local variance. Red lines show for Brent dataset, green lines show for WTI dataset

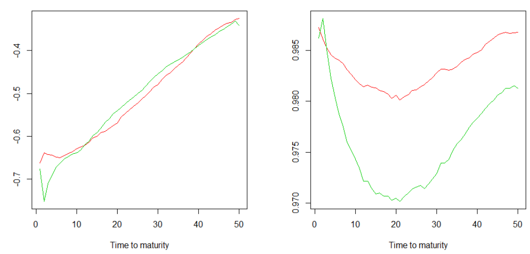

Figure 14: To the left, γ_i is negative, and hence indicates a leverage effect in the datasets. Notice the leverage effect being larger for front month maturities. To the right, β_i illustrates the autoregressive parameter of the natural logarithm of local log variance of contract i. Red lines show for Brent dataset, green lines show for WTI dataset

the unknown parameter θ. If the observations are i.i.d. they are considered as a random sample, or a randomly selected subset, of all possible observations that could have been made. All possible observations are denoted as the population, and how these are distributed is described by the probability distribution $f(x;\theta)$. The parameter θ is estimated to describe the whole population, where the goal is to find an estimated θ as close to the "real" θ as possible. To calculate the maximum likelihood, it is more convenient to use the logarithm of the likelihood function, called the log-likelihood, as shown in the following equations. This is possible since maximizing an expression is equal to maximizing the logarithm of an expression (Hull (2012)). The maximum likelihood estimators (MLE) are found by following these five steps:

1. Define the likelihood function

$$L(\theta) = f(x_1, ...x_n; \theta) \underbrace{=}_{\text{indep.}} f(x_1; \theta)...f(x_n; \theta) = \prod f(x_i; \theta) \tag{34}$$

2. Take ln()

$$l(\theta) = ln[L(\theta)] = \sum_{i=1}^{n} ln[f(x_i;\theta)] \tag{35}$$

3. Take the derivative w.r.t. θ

$$\frac{\partial}{\partial\theta} l(\theta) = \sum_{i=1}^{n} \frac{\partial}{\partial\theta} ln[f(x_i;\theta)] \tag{36}$$

4. Set equal to zero and solve w.r.t. θ

$$\frac{\partial}{\partial\theta} l(\theta) = 0 \quad \text{gives} \quad \hat{\theta} \tag{37}$$

5. Check that

$$\frac{\partial^2}{\partial\theta^2} l(\theta) < 0 \quad \text{for} \quad \theta = \hat{\theta} \tag{38}$$

Keep in mind that the five steps above illustrate the general method for MLE. By performing the MLE procedure on the EGARCH equation, adjustments are made since the returns are lognormal, and not independent. However, the MLE - steps achieve the same goal, and the results of the procedure are the optimized values for the four parameters describing the variance for the datasets. The four parameters are illustrated for 50 contracts for Brent and WTI in table 3 to 5 in Appendix A.

4.4.2 Finding Volatility and Correlation with Principal Component Analysis

Estimates of current levels of volatilities and correlations are necessary to predict the future, and hence be able to value derivatives (Hull (2012)). Clewlow and Strickland (2000) describe the main advantage for the user of the forward curve models as the flexibility in choosing both the number and form of the volatility functions. The most distinct feature of the model is exactly that it allows for non-constant volatility as well as for non-constant correlation.

Principal Component Analysis (PCA) is a method for analysing volatilities and correlations to find a set of factors describing the dynamics of the price evolution. In Aleksander (2008), the PCA is explained as based on the spectral decomposition (or eigenvalue decomposition) of a covariance matrix or a correlation matrix. PCA use the relationship

$$A = W\Lambda W' \tag{39}$$

where A is either a covariance matrix or the corresponding correlation matrix. W is the orthogonal matrix of eigenvector of V (V will be introduced shortly), and Λ is the diagonal matrix of eigenvalues of V. If PCA is performed on a correlation matrix then the results will only be influenced by the correlations of returns, but if the input is a covariance matrix, then the results of the PCA will be influenced by the volatility of returns as well as the correlations of returns. The observations in our case is a covariance matrix filled with all the returns from the historical futures prices. The PCA transforms a set of possibly correlated variables into a set of variables that are linearly uncorrelated called the principal components.

To find these principal components, denote by X the $T \times n$ matrix of time series to be analysed by PCA. T is the number of data points used in the analysis, and n is the number of variables. The variables of interest are financial variables such as asset, portfolio or risk returns; in this case asset returns on futures prices for Brent and WTI. The columns of X are denoted $x_i, x_2, ...x_n$, which are time series of data on a correlated set of returns. For each x_j (for j=1,...,n) is a $T \times 1$ vector, where T is the number of observations. Assuming that each column of X has a zero mean, the sample covariances of the data can be summarized by a matrix:

$$V = T^{-1}X'X \tag{40}$$

where V is the correlation matrix of the returns if the data is normalized so that each time series of observations x_j has variance 1 and mean 0. This normalisation is done by dividing the returns by the sample standard deviation after subtracting the sample mean from each observation.

The output of the PCA is all the eigenvalues of V. The number of components will equal the number of original variables. Each principal component is defined as a linear combination of the columns of X, where the weight of the components are chosen in such a manner that:

1. The principal components are uncorrelated with each other.
2. The first principal component has the largest possible variance. This component accounts for as much of the variability in the dataset as possible. The next component accounts for as much of the variability left in the data set as possible, and so on for all the remaining components.

Recall Λ, which is the diagonal matrix of the eigenvalues of V, and W, which is the orthogonal matrix of the eigenvectors of V. The eigenvalues and corresponding eigenvectors are ordered from largest to smallest, where the largest are the most important. Normally, this leads to only a few factors necessary to explain the dynamics of the curves. PCA can be applied to any set of stationary time series, however high or low their correlation, but it works best on highly correlated systems, such as a set of commodity futures returns of different maturities. The higher the correlation, the fewer components are necessary to explain the variation in the system. One of PCA's most established financial applications include multi-factor option pricing models, which makes the technique suitable for this thesis.

A typical pattern obtained from PCA are three types of risk factors which act to "shift", "tilt" and "bend" the curve (Tsay (2005)). The most important factor is always factor number 1, normally a shift-factor. This includes the majority of the movements to the curve, and is positive for all maturities. Since it is positive for all maturities, a positive (or negative) shock to the system will cause all prices to go up (or down). The second most important factor is normally a tilt-factor, which causes the short and long end of the curve to move in opposite directions. The third factor, normally a bend-factor, bends the curve, causing both the short and long maturities to move in opposite directions from the middle of the curve.

Table 1 shows the standard deviation, proportion and cumulative proportion of explained variance for the first 5 components found by the principal component analysis, where Fig. 15 illustrates factor 1 and factor 2 and Fig. 16 illustrates factor 3 and 4. Note that the principal components for returns are found by standardizing with the local estimated volatility, while the unconditional standard deviation was used for the volatility. The reason for this is that it is not necessary to estimate a time varying variance for the variance process itself.

Table 1: Standard deviation, proportion and cumulative proportion of explained variance for returns

	PC1	PC2	PC3	PC4	PC5
Standard deviation	9.7350	1.6573	1.1676	0.5588	0.4853
Proportion of Variance	0.9511	0.0276	0.0137	0.0031	0.0240
Cumulative Proportion	0.9511	0.9787	0.9924	0.9955	0.9979

Table 1 clearly shows that the first factor is by far the most important factor, explaining 95.11% of all volatility and correlation. This type of factor is a shift factor, as can clearly be seen from the shape of the factor to the left in Fig. 15. The shift factor explains how a shock to the system, for example by some abrupt information entering the market, will cause either a positive or negative shift in both Brent and WTI forward curves, depending on the information. Even though this factor is very similar for the two assets, notice how the Brent and WTI are slightly different at the very front-end of the curve. The second factor, illustrated to the right in Fig. 15, explains 2.76% of the variation in the dataset, and together with factor 1 it explains almost 98% of the Brent and WTI forward curves. The second factor has a different shape, indicating that this factor is a tilt factor. This factor will increase the difference between the front-end of the curve and the back-end of the curve, as it tilts the curve, dragging the front- and back-end further apart. The factor is quite similar for Brent and WTI curves.

The third and fourth factor does not play a crucial role to account for the variation, as seen in Table 1. However, they are included to achieve as realistic simulations going forward as possible. The third factor, illustrated to the left in Fig. 16, accounts for 1.37% of the variation in the dataset, and is another shift factor. This factor can affect the Brent and WTI data differently; as a shock to system might have a positive effect on the Brent data, it can lead to a negative

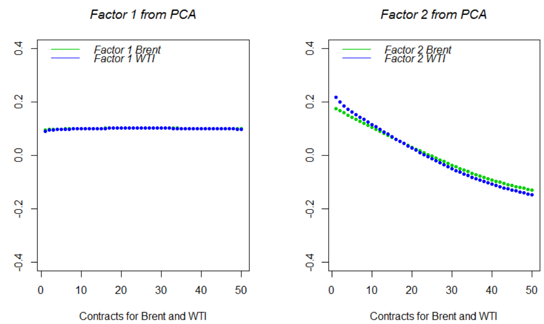

Figure 15: Component 1 and 2 found from principal component analysis. Brent illustrated by green dotted lines, and WTI by blue dotted lines

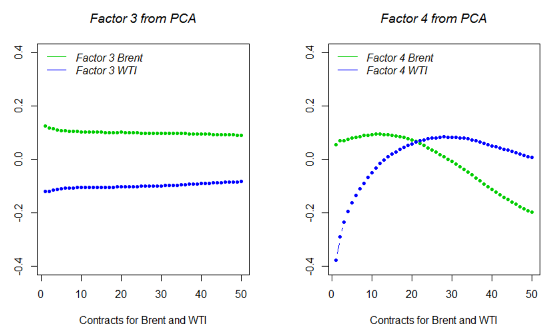

Figure 16: Component 3 and 4 found from principal component analysis. Brent illustrated by green dotted lines, and WTI by blue dotted lines

effect for the WTI data, and vice versa. Or, if the shock is positive (negative) for the Brent, it can be less positive (more negative) for the WTI data. Finally, factor 4 is a bending factor, accounting for only 0.31% of the variation. This factor is almost like a tilt factor, but instead of pulling the start and the end of the curve in separate directions, the bending factor pulls the centre of the curve away from the start and end. This factor is different for the Brent data and the WTI data, where the bending factor will have greater impact on the back-end of the curve for the Brent data, and more impact on the front-end of the curve for the WTI data. Since this factor explains such a small part of the variations, it will be difficult to see how it will affect the simulated forward curves.

The factors exceeding the fourth factor only contribute with noise to the system, and are therefore discarded as not relevant for further simulations.

5 Simulating Forward Paths for Brent and WTI futures

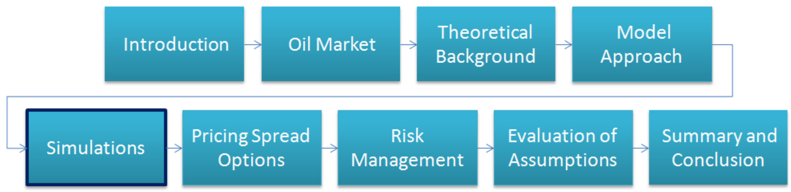

The model introduced in the previous chapter can now be used to sample realisations of future outcome for the prices of Brent and WTI. Remember from section 3.3.2 that the Wiener process is a Markovian process, meaning that each observation only depends on the observation directly prior. The simulation process therefore begins with the last known observation, which is from January 18th 2013. The observation of the two prices at this date, together with the estimated EGARCH parameters and the PCA components, and multiplied with a random number, drive the observation for the next day, and so on going forward.

Due to the high volatility for the datasets, the realisations of the simulated price paths can result in very different appearances. Recall that the prices of both Brent and WTI declined with almost USD 100 per barrel in the downfall of the financial crisis in 2008-2009, leading to a local volatility of almost 90%, as illustrated in Fig. 11 and 12 in the previous chapter. Any random realisations can take extreme levels, with both high or low future prices, and the recent past confirms that these extreme fluctuations in fact are possible. In this thesis, three samples are selected to show the variety of the possible outcomes. Each sample is illustrated with its corresponding volatility and its simulated spread for the two benchmark prices.

5.1 Price Spread Adjusted for Brent and WTI futures

Since this thesis will price spread options on Brent and WTI futures prices, the multi-factor forward curve model has to make sure this is possible. Therefore, the model is adjusted to account for the large initial spread between Brent and WTI. For simplicity reasons, consider the present spread between Brent and WTI as transportation costs. In Europe, the long coastlines simplifies the transportation of oil, due to the possibility of transportation by ships. In North America, on the other hand, the coastline is short compared to the total landmass, and the continent depends heavily on a developed pipeline infrastructure, as illustrated in Fig. 3 in chapter 2. Therefore, by bringing the WTI crude oil out of Cushing (where it has delivery), and transported to somewhere on the East Coast, say Boston, the two crudes are easier to compare. So, imagine a buyer located in Boston in a perfect market, where he/she could by either Brent for a given price and have it delivered by the coast, or he/she could buy WTI somewhat cheaper but pay for transportation of the oil from Cushing to Boston.

A challenge with this approximation is the difference in the curvature. Fig. 17 illustrates both the Brent and WTI forward curves observed on January 18th and the adjusted price spread, where the graph to the left illustrates the initial forward curves, and the graph to the right illustrates the adjusted forward curves used as the starting point for simulations. The figure shows that the Brent forward curve is in backwardation, whilst the WTI forward curve has a small contango at the front of the curve, before it goes into backwardation as well. When both curves are in backwardation, the Brent curve has a faster speed than the WTI. With this information in hand, it can be discussed where it is most logical to align the graphs, whether it makes most sense to align the curves at the front, the back, or somewhere in the middle of the curve. For the analysis in this thesis, the curves are chosen to be aligned in the front of the curves to get a mutual starting point. This is done by "lifting" up the WTI forward curves to match the front month of the Brent forward curves, and this adjustment is evaluated when looking at the results in the next chapter.

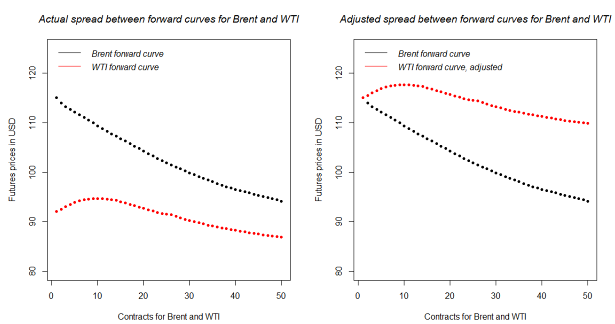

Figure 17: The black line illustrates the Brent forward curve, and the red line illustrates the WTI forward curve. The presentation to the left is the original spread, and the one to the right is the adjusted spread

5.2 Realisations of the Forward Curves

The model introduced in chapter 4 explains how historical volatility and correlation drive the model. For each forward simulation, the outcome of the prices can look very different, since the volatility is so high. It is important to update the volatility when time passes, to make sure that the volatility is adjusted every time a futures contract goes to expiry. For example, when the first contract goes to expiry, and you move from month 1 to month 2, the second month is now the front month, and the volatility structure needs to be adjusted accordingly. Three samples illustrate how different the simulated price paths may look. Optimally, an analysis could look at an infinitive number of price path realisations, but due to the time required to both simulate and analyse the outcomes, three samples are chosen to suffice to illustrate the dynamics of the two markets.

5.2.1 Sample 1

Fig. 18 shows the development of the price path for the first realisations of the forward curves for Brent and WTI. The prices for the two benchmark crudes lie between approximately USD 75 and USD 175, which is a deviation not too far from the levels today and in the recent past. In the curvature seen in the graph, one key characteristic separating the Brent and the WTI futures prices is visible; from 2013 to 2014 the WTI curves are sometimes in contango formation while the Brent curves are in backwardation formation. This can be compared both to the historical forward curves for Brent and WTI illustrated in Fig. 11 on page 11, and in the forward curves observed on January 18th illustrated in Fig. 17. The WTI curves often tend to go into contango-like formations, while the Brent curves stay in backwardation. Another visible characteristic is that, in 2015, when a "shock" to the system increases the prices, the WTI prices increase more relatively to the Brent prices.

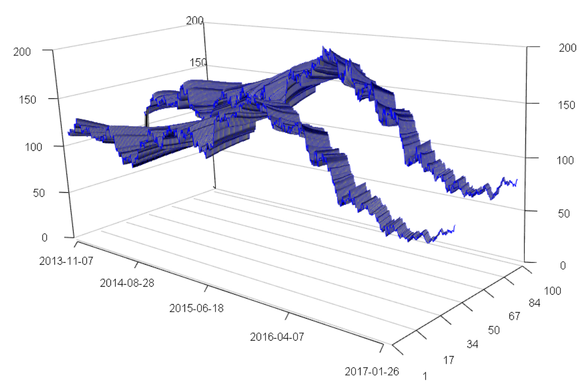

Figure 18: Sample 1: Forward curves for Brent and WTI

Fig. 19 shows the corresponding volatility time structures for both Brent and WTI. As the figure illustrates, the volatility time structures are very similar, but the volatility in the front months of the curves is higher for WTI than it is in the front months of the Brent curves. This is the case historically as well, as seen in Fig. 12. This can also be compared to the factors retrieved from the PCA in the previous chapter. The first, second and fourth factor all have larger movements for the WTI curves than the Brent curves in the front-months.

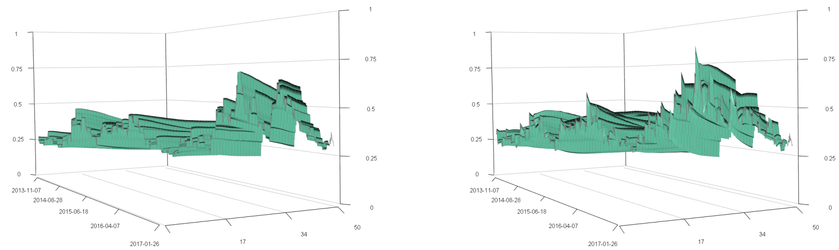

Figure 19: Sample 1: Simulated annualized conditional volatility time structure corresponding to the price realisation in Fig. 18. Brent volatility is displayed to the left and WTI volatility to the right

Fig. 20 illustrates the simulated spread between the Brent and WTI forward curves, which is simply the Brent prices minus the WTI prices, adjusted for the initial price spread. This presents the intrinsic value of the spread, which illustrates how the spread evolves over time, imagining that the spread is 0 for all contracts at Day 1. The initial price spread at Day 1 is the strike price, so it is the development from the initial spread to the simulated spread outwards in time which will decide the value of the spread options. Fig. 20 illustrates that for Sample 1, the spread is sometimes higher and sometimes lower than the initial spread. A simulated spread below zero means that the price spread has widened compared to the initial spread, which will give option values in-the-money. From 2013 to 2016 the spread is below zero, before it turns positive in mid-2016 and ends around zero in 2017. The simulated spread is what will drive the option prices calculated in chapter 6.

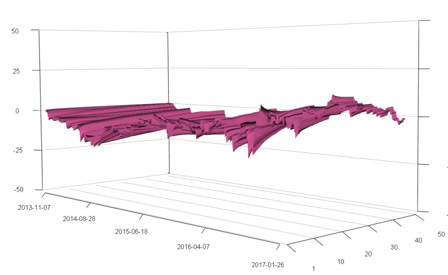

Figure 20: Sample 1: Simulated spread Brent/WTI

5.2.2 Sample 2

The second realisation of the prices for forward curves for Brent and WTI is illustrated in Fig. 21. This price realisation lies between USD 60 and USD 140, and illustrates a more "calm" sample. There are indeed movements in the prices, but the dynamics are more stable, as well as the curves are more alike for both prices in this sample. This illustrates that the curves can take on different shapes from one sample to the next, and when the difference in curvature was evident in the last sample, in this sample the two prices behave more similarly. By looking closely, some contango-like shapes can be spotted for WTI where the prices are highest in 2016 as well as for some days in 2013, but in overall this sample illustrates a reality where the two prices are practically similar throughout the period.

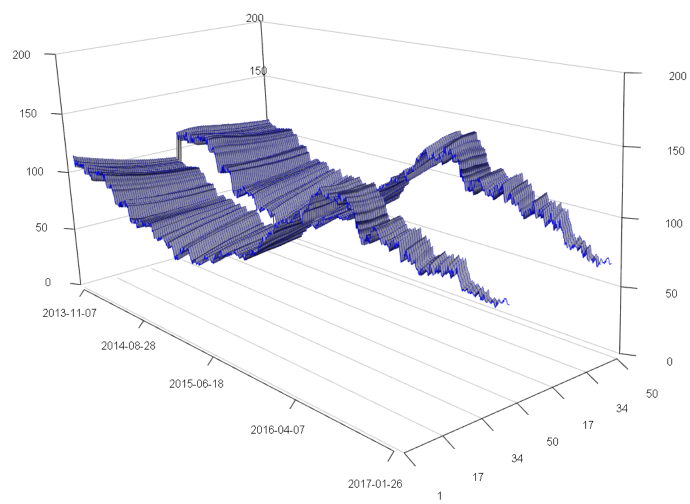

Figure 21: Sample 2: Forward curves for Brent and WTI

The volatility time structure corresponding to Sample 2 for Brent and WTI is illustrated in Fig. 22. Similarly to the previous realisation, the volatility time structures for the two prices move in the same manner, but show signs of higher volatility for the WTI in some areas, as well as overall in the front months for WTI.

The simulated spread for the Brent and WTI forward curves in Realisation 2, adjusted for the initial price spread, is shown in Fig. 23. In this sample, the spread develops in a slow and stable way, staying around zero without any

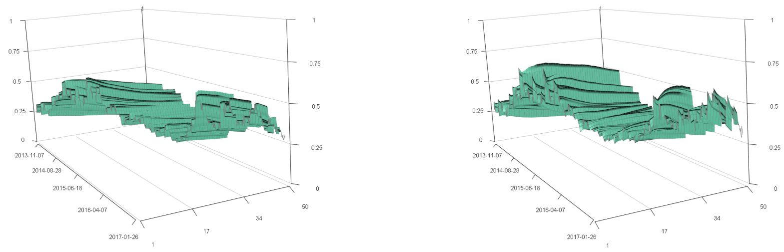

Figure 22: Sample 2: Simulated annualized conditional volatility time structure corresponding to the price realisation in Fig. 21. Brent volatility is displayed to the left and WTI volatility to the right

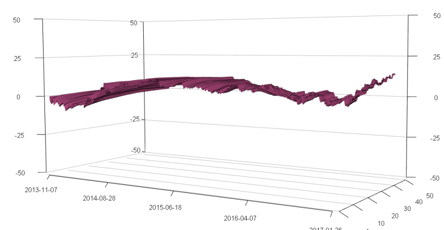

Figure 23: Sample 2: Simulated spread Brent/WTI

drastic movements. Two areas of the time period have a simulated spread below zero, as seen around 2013 and late 2016. In the majority of the time span, the spread lies directly above zero, meaning that the spread is narrower between Brent and WTI than it was initially. Large movements in the simulated spread will lead to option values either deep in- or out-of-the-money. For this sample, when the simulated spread evolves close around zero, the option values stay around at-the-money, with either low option values or options which expire worthless. In the last part of the dataset, the simulated spread peaks above zero, meaning that the option values will expire worthless at this point.

5.2.3 Sample 3

The third realisation of the forward curves for Brent and WTI is shown in Fig. 24. This price realisation is a more "extreme" variant, with future oil price spiking to over USD 300 for WTI in 2016. In the first part of the time period, the curves show stable development, however the curves experience more drastic changes early 2016, where the dynamics between the Brent and WTI curves start to differ. Both the curvature and the values for the two prices are different in the last part of the period, where the WTI prices spike and go into steep backwardation, while the Brent prices increase at a relatively lower pace, and with backwardation shapes of a more relaxed character. Although this outcome is not considered as likely as Sample 1 and 2, the outcome seen in Sample 3 could perhaps occur in a market in some kind of crisis.

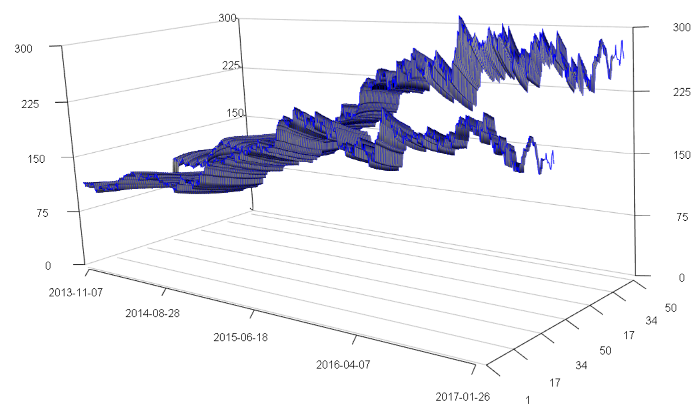

Figure 24: Sample 3: Forward curves for Brent and WTI

The volatility time structure corresponding to Sample 3 for Brent and WTI is illustrated in Fig. 25. The volatility time structures have the same characteristics as in the previous two samples. Notice however that even though the prices peak from the 2015 and forward, the volatility time structure behave rather similarly for the entire period. This finding will impact the risk management of the option prices in this sample, as will be discussed in chapter 7.

The simulated spread for Brent and WTI for Realisation 3 is illustrated in

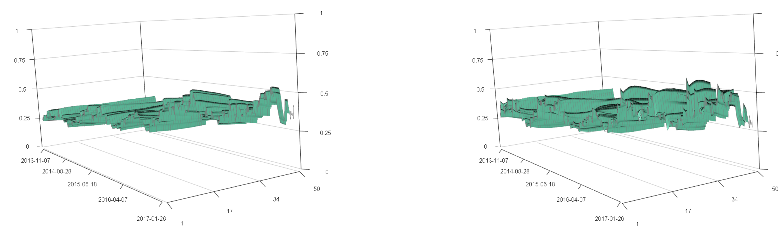

Figure 25: Sample 3: Simulated annualized conditional volatility time structure corresponding to the price realisation in Fig. 24. Brent volatility is displayed to the left and WTI volatility to the right

Fig. 26. Here, the spread starts nice and easy for the first two years, before it drops early 2016. The simulated price spread continues to drop during the rest of the period, staying around values of USD -50 before ending at around USD -75. This large deviation in the simulated spread from the initial spread will create large option values in this realisation.

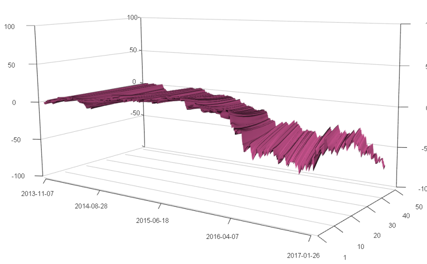

Figure 26: Sample 3: Simulated spread Brent/WTI

6 Pricing Spread Options on Brent/WTI

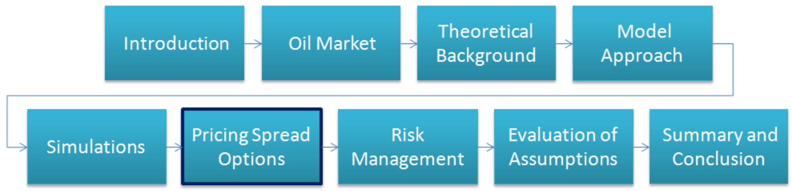

In this chapter the results of calculated option prices on the Brent/WTI futures spread are presented, as well as background to the Monte Carlo method for simulating prices. The first day in any random sample is the day with the largest number of days until maturity, and is therefore the most volatile day and hence subject to extensive analysis. Recall from the previous chapter that the start of any simulation of forward paths is based on the last available observation. The expected option prices when situated at Day 1 in the simulation process are therefore equal for all samples. This chapter illustrates the results of sensitivity analysis for the expected option prices, as well as presenting the payoff profiles for the three drawn samples from the previous chapter.

6.1 Monte Carlo Simulations for Option Prices

Tsay (2005) states that, in general, any quantity that can be written as the expected value of a random variable X defined on a probability space, can be estimated by the Monte Carlo method. The Monte Carlo method is based on the analogy between probability and volume, and samples randomly (or pseudo-randomly) from a universe of possible outcomes. In the Monte Carlo method, the law of large numbers is essential. This law ensures that the estimate from the method converges to the correct value as the number of draws increase. Consider a derivative dependent on a single market variable S (or F if a futures price is considered) that provides payoff at time T. Assume interest rates as constant, and that the derivative is valued in a risk neutral world. Hull (2012) presents the following five steps for valuing a derivative with the Monte Carlo method:

1. Simulate a path of the underlying variable over time according to the given model in a risk neutral world.

2. Calculate the payoff from the derivative given the path of the underlying.

3. Repeat step 1 and 2 to get many sample values of the payoff from the derivative in a risk neutral world.

4. Calculate the average of all the sample payoffs to get an estimate of the expected payoff in the risk neutral world.

5. Discount the expected payoff achieved in step 4 at the risk free rate to get an estimate of the present value of the derivative.

Valuing a derivative security with Monte Carlo typically involves simulating a number of paths from stochastic processes, which are used to describe the evolution of underlying price assets, interest rates and other model parameters. It is important to remember that the Monte Carlo principle cannot be applied to the prediction of the evolution of the underlying asset, but to the simulation of many possibilities/paths once a certain diffusion process is assumed.

There exists several methods for options pricing. The Black-Scholes-Merton (or Black-Scholes) formula for pricing European-style options is perhaps the most common, and this formula assumes that there is a closed-form solution available. When this is not possible, other methods, such as the Monte Carlo method, is used to price options. Monte Carlo is not competitive to other methods when it comes to computing one-dimensional integrals, in this case a measure such as Black Scholes is better suited. However, Monte Carlo is suitable in evaluating integrals in high dimensions and with complex stochastic processes. According to Hull (2012), one of the most distinct advantages of the Monte Carlo method is that it allows for the payoff of the derivative to depend both on the *path* of the underlying variable S as well as on the $final$ value of S. Meaning, Monte Carlo allows for a payoff to be calculated at several times during the lifetime of the derivative, not only at its maturity, which is the case with Black-Scholes (Hull (2012)). Since Monte Carlo values the payoff during the lifetime of an option, any stochastic process for S can be accommodated. The method can also be extended to consider several underlying market variables, where the payoff from the derivative depends on more than one underlying variable. In the case of this thesis, F_1 is considered as the underlying price for Brent futures, and F_2 is considered as the underlying price for WTI futures.

6.2 Sensitivity Analysis of Expected Option Prices

When valuing an option, it is necessary to decide what kind of option to value. Recall from section 3.1.4, about option profiles, that four basic options are available: Long call, short call, long put and short put. In this thesis, a long put option is considered. A long put strategy in options trading is a basic strategy where the investor buys the right to sell some underlying asset with the belief that the price of the underlying security will go significantly below the strike price before expiration date. The investor bets that the price will decrease, so that he/she can sell the underlying assets at a higher price than what will be the market price at expiry. In section 3.1.4, it was explained how long positions in options have a limited downside and a limitless upside. For short positions, the upside is limited, while the downside is limitless. Therefore, compared to short selling, buying put options does not require the investor to borrow stock to short. The risk is capped to a minimum as seen in Fig. 8 (page 31), as opposed to the unlimited risk present in short selling seen in Fig. 7 (page 30). The maximum profit for the investor in a long put position is unlimited, which is appealing for the investor.

In this thesis, when looking at the value of optionality between the Brent and WTI futures prices, it seemes most reasonable to consider a long option profile due to the limitless upside. Preferably, more than one option profile could be evaluated, but due to the scope limitations of the thesis, one option profile is considered as sufficient.

Recall from Eq. 15 that the simulated price spread between Brent and WTI futures can be stated as

$$\text{Spread}_{t,T} = FB_{t,T} - FW_{t,T} \tag{41}$$

To find expected option prices for 50 contract going forward at Day 1, 20,000 price paths for each underlying variable are simulated with Monte Carlo. The payoff for each contract from each price simulation is calculated using equation

$$\max(K - \text{Spread}_{t,T}, 0) \tag{42}$$

where each payoff relies on the long put option shown in Eq. 5 on page 30, and where K is the strike price. Zero initial cost from buying the option is assumed. The option is valued at-the-money, using the initial price spreads of Brent and WTI futures contracts at Day 1 as the strike price for each of the 50 contracts. This initial price spread is simply the Brent forward curve minus the WTI forward curve at the starting point of the simulations. The starting point of the simulations is the last day of the historical forward curves; January 18th 2013. Recall from Fig. 17 on page 57 that the WTI curve was lifted up to the Brent curve at the starting point of simulations, leading to the strike price for contract 1 being zero, and to negative values for the strike prices for the other 49 contracts. The spread could just as easily have been defined in the opposite turn, but this does not really make a difference since the initial spread is taken into account when valuing each of the options; the value of the option does not rely on the strike price itself, just the change in the value of the simulated spread compared to the strike price spread. The holder of a long put option on this spread bets that the simulated spread will widen compared to the initial spread. An increase in the spread will drive up the option prices, and a narrower spread will drive the option prices out-of-the-money.

After calculating the payoff for each option, the result is 20,000 option prices for each contract, where each one is either zero or has a positive value. The positive values are the options that have gone in-the-money, and the worthless options have either ended out-of-the-money or at-the-money. The averages for each contract are found. Each outcome from the 20,000 prices per contract has impact on that specific contract's average value, and the possible outcomes for each option value lie in the range between zero and some value above the contract's average value. Since many of the option values are zero (recall that an option which expires out-of-the-money will not be exercised and hence will expire worthless), many of the outcomes must lie high above the average value. The option values of zero will pull the averages down, when in fact, as seen in one of the samples in the next section, some option values can take on extremely high levels. From the five steps in the Monte Carlo method explained by Hull (2012), the fifth step is to discount the average prices to present value. For a spread option on Brent/WTI futures, it is difficult to decide which level of interest rate to use. Therefore, the analysis in this thesis assumes r=0.

The procedure above results in one average price for each of the 50 contracts priced at Day 1. Looking at average prices is useful, but it does not tell us anything about the uncertainty in the estimates. Information about this uncertainty can be found by looking at the distribution of the option prices, which shows the variation in the outcomes of the simulations. To find a distribution of the price averages, the procedure above is repeated 3,000 times. Fig. 27 illustrates the distribution of the expected option prices. The black line in the middle is the mean value from the 3,000 average prices. The green lines illustrate the 95% and 5% quantiles of the distribution, and the blue lines are the minimum and maximum observation of prices from the 3,000 averages.

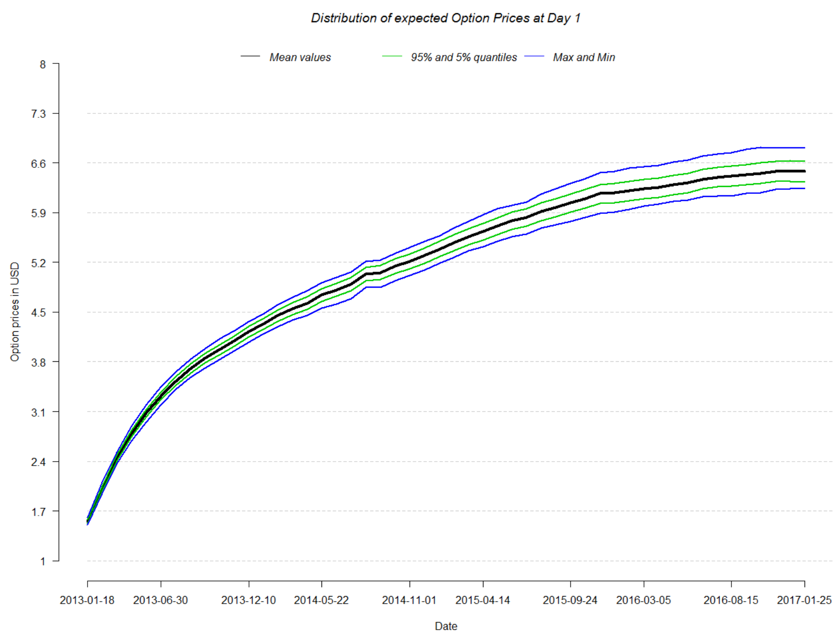

Figure 27: Expected option prices for the 50 contracts, priced at Day 1

The distribution of the prices shows that the variation between the sets of 20,000 prices is not too large, and in the most volatile contract, number 50, the minimum and maximum average is 6.1 and 6.9, respectively. This tells us that using 20,000 draws for each contract per day can be adequate, since the distribution is sufficiently narrow. The accuracy of the results given by the Monte Carlo simulations depends on the number of trials. Given the simulation trials, the standard deviation can be calculated. A good measure of confidence in a distribution is comparing the mean value μ, with the standard deviation σ, by looking at the confidence interval

$$2 \times \sigma \quad < \quad \mu \quad < \quad 2 \times \sigma \tag{43}$$

This confidence interval indicates the reliability of the estimate for expected option prices, and is illustrated for each of the 50 contracts in Table 6 and 7 in Appendix B.

Fig. 27 illustrates how the mean option prices are increasing with every time interval (from each contract to the next), and that the increases are not linear. In the first 10 contracts, the increase in prices is relatively larger than in the next 40 contracts. This reflects that the volatility is higher for the contracts with short durations compared to the volatility in the contracts with longer maturities. This finding is comparable to the historical data illustrated in Fig. 12, which illustrated the historical volatility for Brent and WTI. The option prices start close to USD 1 and grow steadily up to about USD 6, which is mostly caused by the time value of money. The time value of an option is the premium a rational investor would pay over its current exercise value, based on the probability that its value will increase before expiry. The time value can be considered as a risk premium the option seller charges the buyer, since the higher the expected risk, the higher the premium. The expected risk is the volatility times the time to maturity.

6.3 Payoff Profiles for the 50 Contracts

So far, the expected option prices for the 50 contracts priced at Day 1 have been investigated. This part, however, illustrates how the payoff looks for the three realisations of the forward curves during the lifetime of all 50 options. 50 contracts going 50 months forward in time are priced on a daily basis until expiry, using 20,000 simulations [19] for each contract per day to get good estimates. Keep in mind that on Day 1 it is possible to price 50 contracts, but since one contract goes to expiry each month, there are fewer contracts to price as time passes. On Day 22, there are 49 contracts left to price, on Day 43 there are 48 contracts left to price, and so on, since 21 days is categorized as one month by the exchange standards.

To find the payoff both during and at the end of each option's lifetime, Monte Carlo simulations are performed each day until expiry of the 50th contract for the three realisations. In detail, the analysis starts with Day 1 and simulate 20,000 price paths for each of the 50 contracts. Then, each option price is found, using the same equation as before

$$\max(K - \text{Spread}_{t,T}, 0) \tag{44}$$

where K is the strike price, and $\text{Spread}_{t,T}$ is the simulated futures price spread. Then, 20,000 price paths are drawn for each of the 50 contracts when situated at

[19] The exact number of simulations per day per contract is 20,001, because the Monte Carlo analysis is performed simultaneously in three commands, and thus needs the total number of simulations to be dividable by three.

Day 2, giving 20,000 option prices for each remaining contract on which average prices are calculated. Since the underlying price processes for Brent and WTI futures data are assumed to follow a Markov process, the simulations on any given day depend only on the day directly prior. Therefore, after simulating Day 2, it is possible to simulate Day 3, and so on for all days in the three samples. Below are the payoff profiles explained and illustrated, both during the lifetime of, and at the end of the options for the three realisations.

6.3.1 Sample 1

Fig. 27 illustrates the expected option prices at Day 1. These values provide information about what each contract was expected to be worth at that exact day. The goal for an investor is to keep these values throughout the period until maturity. It is therefore interesting to examine how all the 50 contracts ended at maturity, without any form of risk management, to show what really happened to the option values for each contract in the three samples. After all the contracts have gone to expiry, it will be clear whether or not they have been profitable. Remember that contract 1 expires after 21 days, contract 2 after 42 days, contract 3 after 63 days, and so on, up to contract 50, which expires after 1050 days. Fig. 28 displays which value each of the 50 contracts in Sample 1 have at maturity, illustrated together with the expected values priced at Day 1. From Fig. 28, it is evident that the moneyness of the contracts is very varying. The contracts 1-36 all end in-the-money at maturity, but to very different values. The highest profit is USD 22.775 for contract 30 and the lowest profit is USD 1.101 for contract 32. Contract 37-47 all end out-of-the-money, with option values equal to zero. Contract 48 and 49 are back in-the-money, while the last contract, number 50, ends up approximately at-the-money, with an option value of 0.000023.

These findings looks reasonable when comparing Fig. 28 to the simulated spread illustrated in Fig. 20 on page 59. The simulated spread in Fig. 20 is negative compared to the strike price spread/initial price spread until 2016 (which is approximately the duration of 36 contracts), leading to positive option values since the options valued are put options. The simulated spread from 2016 to 2017 (contract 37-47) is above the strike price, leading to option values of zero, since these contracts will not be exercised. Note that the last contract ends up with a simulated spread around zero, meaning around the strike price, which makes this contract mature approximately at-the-money.

For an investor exposed to this realisation, he/she would make more money (roughly) on contract 1-36 than expected, and less money on contract 37-50 than expected. In overall, the payoff for the investor in this set is quite good, especially when considering the first 36 contracts. Risk management of this set would, when viewing in retrospect, diminish some of the value which the investor could obtain if not risk managing at all.

Fig. 28 only illustrates the option values at the maturity of each contract. It is also interesting to look at the option values during their lifetime, since the prices fluctuate a great deal up until they expire. One contract that expires

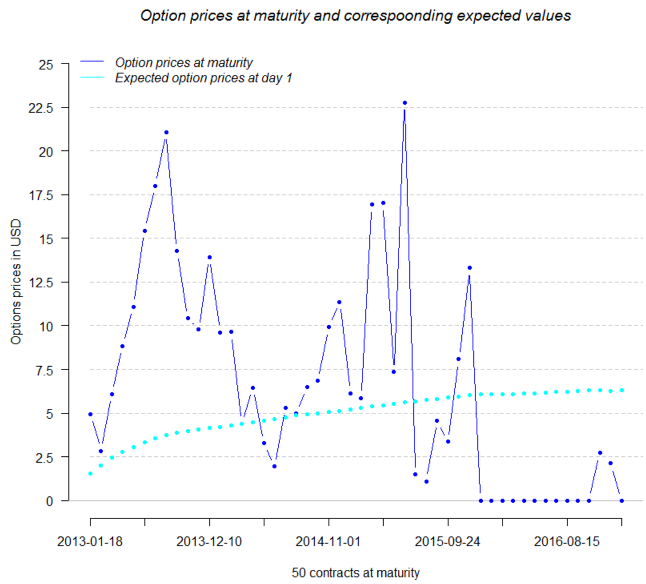

Figure 28: Sample 1: Moneyness of 50 options at expiry. The dark blue points are the moneyness of the 50 options at expiry, the light blue points are the expected option values priced at Day 1

in-the-money and one contract that expires worthless, contract 24 and 50, are illustrated in Fig. 29 and 30 as examples of how the lifetime of an option can look. The expected values of these contracts are about USD 5 for contract 24, and USD 6.5 for contract 50, as seen in Fig. 28. Contract 24 has a sharp increase around July 2013, and in overall the option values stay around USD 10, before the option expires with a value of USD 11 in 2015. Contract 50 is in-the-money throughout most of its lifetime, before it ends approximately at-the-money at expiry. By looking at these two contracts' lifetimes, it is evident that an option's profile during its lifetimes does not necessarily decide the outcome of its moneyness at maturity. Contract 24 is selected for illustration as is it a quite "stable" contract; aside from the sharp increase in 2013 the option value fluctuates around levels of USD 10 during the entire lifetime, before it expires at USD 11. Contract 50 is, on the other hand, selected for illustration as its payoff during, and at maturity, diverges. In this contract, the option value is around USD 10 for the first half, before the prices drop. Notice the small spike at the last part of the contract, where the option value increases to USD 4, before the prices go to zero and the option expires worthless. The end values of both contracts can be cross-checked with Fig. 28, confirming end values of around USD 11 for contract 24 and zero for contract 50.

Although contract 24 and 50 are illustrative examples of the lifetime of two option values, looking at all option prices during their lifetime illustrates more

Figure 29: Sample 1: Payoff from option contract 24. Notice the sharp increase in option values around July 2013. The option expires in-the-money

about the entire sample. Fig. 31 illustrates simultaneously the option prices for all 50 contracts during the entire simulated period. The figure shows each option price at it goes to maturity, starting with 50 contracts in 2013, and ending with only one contract that matures in March 2017. During the 4.2 year period, only one contract will be in front month position at a time. The front of the curves illustrates where each contract goes to maturity, and it is evident that the volatility in the option prices right before maturity is high. In comparison, in the further back-end of the curves, where the contracts have longer time left until maturity, the movements are still large, but they are more "calm" and lack the spikes experienced in the front-end of the curves.

Fig. 31 also illustrates how correlated the option prices are. The high correlation in the dataset can be seen in the option prices, as they move together throughout the time period. The option in the front month position at a given time will be more volatile than the remaining contracts, and the front month contract will drive the prices of the other contracts due to the high correlation. The extreme volatility in the front-end month can create very different option price outcomes going from one contract to the next. Notice especially the price drop in mid-2016, between contract 30 and 31, illustrated in detail in Fig. 32. Contract 30 expires with the highest option price in the set, with a payoff of USD 22.775. Contract 30 is shown with the black line in Fig. 32, and the high

Figure 30: Sample 1: Payoff from option contract 50. The prices are stable until the middle of the period, from where the prices drop. Notice the small jump in prices towards the end, before the option expires worthless

volatility in the months prior to expiry can be seen clearly, as the option price jumps from around USD 7 to USD 22.775 for the last period of the contracts duration. When contract 30 rises at the end of its lifetime, this affects the other remaining contracts, due to the high correlation. Notice how the price for contract 31, illustrated in the light blue line, follows the price for contract 30 in Fig. 32. When contract 30 has expired, contract 31 is the new front month contract, impacting the prices of all remaining contracts. After the 30th contract expires, the price of contract 31 drops from around USD 20 to USD 1, where it expires. This drop pulls down the rest of the contracts still current in the dataset, as can be seen in Fig. 31.

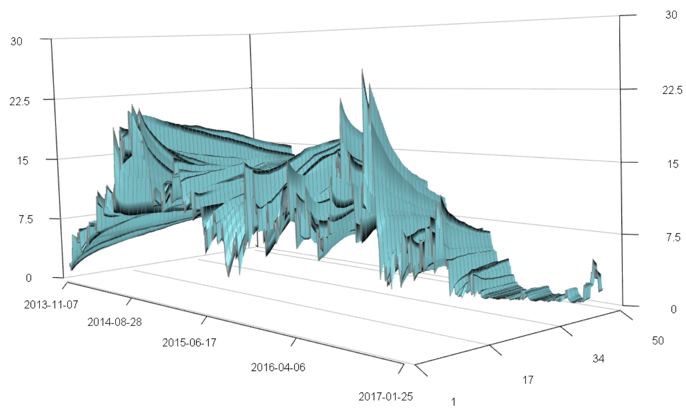

Figure 31: Sample 1: Option prices for all 50 contracts, illustrating correlation between the contracts and high volatility in the front-end months

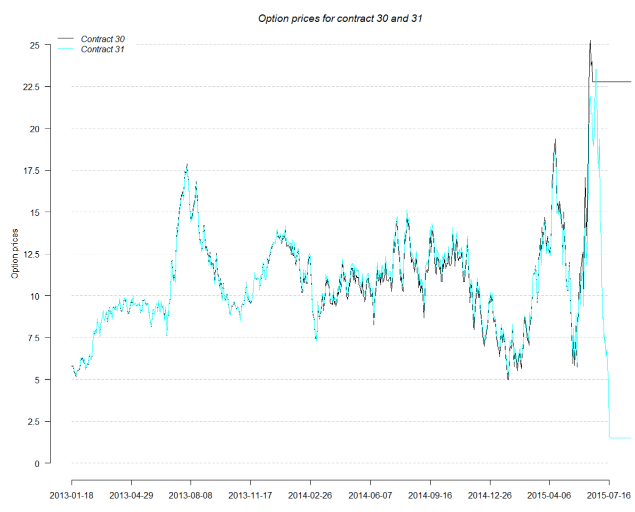

Figure 32: Sample 1: Payoff from option contract 30 and 31. Contract 30, in the black line, pulls up contract 31, in the light blue line, before it expires. Contract 31 drops in its last month before maturity

6.3.2 Sample 2

Looking at the second sample, Fig. 33 displays how the 50 contracts ended at maturity, without any form of risk management. Sample 2 shows a more stable set than the previously illustrated Sample 1. The dark blue points are the values of each contract at maturity, and the light blue points are the expected option prices at Day 1. Note that the expected option values are identical for all samples (since they all are based on the last observation on January 18th 2013), but that the scale in the graphs where they are illustrated differs. The option values at maturity is zero for many contracts, and the highest option value is seen in contract number 3, with an option value of USD 7.80. This sample illustrates a set where most of the contracts have expired out-of-the-money, as seen in contract 8-34, 38 and 44-50 which all expire worthless.

These findings look reasonable when comparing the options payoffs with the simulated spread illustrated in Fig. 23 on page 61. In the simulated spread, it is evident that the spread is rather stable, without any drastic movements. In the period in Fig. 33 where contract number 8-34 expires worthless, the corresponding simulated spread is above zero. Recall that negative development of the price spread between futures on Brent and WTI drives the option prices. After this period with a positive spread, the spread turns below zero, giving positive option values, before it spikes towards the end and sends the last contracts out-of-the-money.

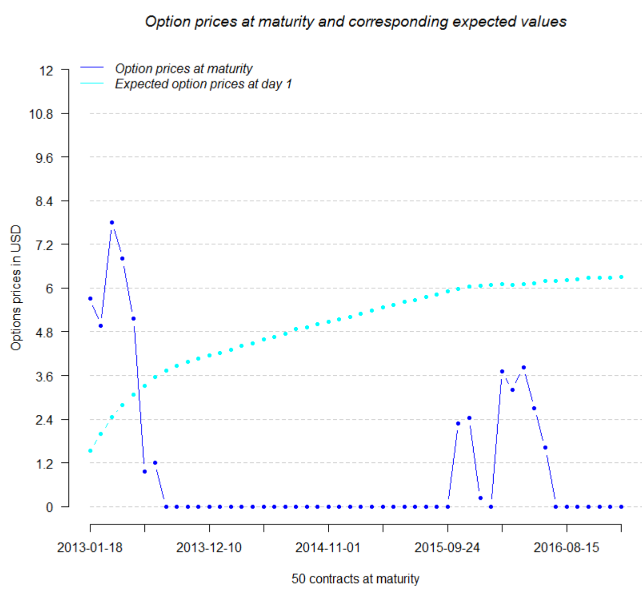

Figure 33: Sample 2: Moneyness of the 50 options at expiry. The dark blue points are the moneyness of the 50 options at expiry, the light blue points are the expected option values priced at Day 1

For an investor exposed to this set, without any form of risk management, the payoff would overall be less than expected. Notice how the expected value is higher than the value at maturity for most contracts in Fig. 33. Few contracts end in-the-money, and the ones who do have relatively low values. Risk management of this set would, when viewing in retrospect, perhaps help keep some of the expected option values throughout their lifetime, and hedge the downside of the fluctuations.

Fig. 34 looks at the option prices both during their life span as well as at their maturity. The maturities are illustrated at the front-end of the figure, and it is evident that this set also has higher volatility in the contracts which are close to maturity. This sample has many contracts expiring out-of-the-money. The contracts going to expiry impact the following contracts due to high correlation in the dataset. Regard how around April 2016, the back-end of the curves are higher than the front-end. Here, the front-end of the contracts correlate with the expired contracts, which pulls the following contracts down to values of zero.

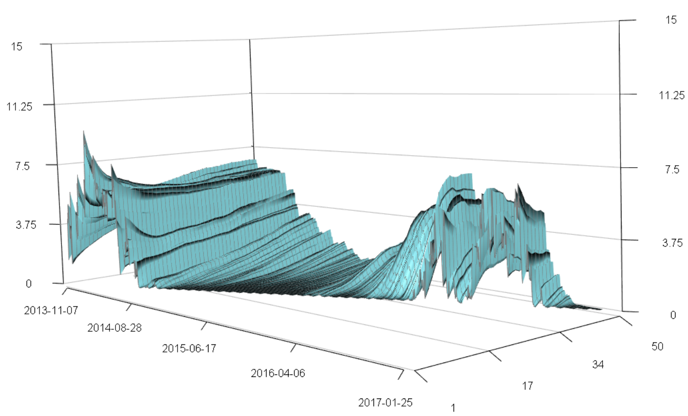

Figure 34: Sample 2: Option prices for all 50 contracts. Regard the correlation between the contracts in the set, especially seen in the contracts which expire out-of-the-money from 2014 to 2016

6.3.3 Sample 3

In Sample 3, the 50 contracts at maturity ends with significantly higher prices than the other two samples. Fig. 35 displays which value each of the 50 contracts have at maturity without any form of risk management. This sample illustrates more of an "extreme" outcome than the previous two, with values sky-rocketing for the second half of the time span. The option values lie around zero for the first half, before they increase in the second half, with contract 50 as the most profitable contract providing a payoff of USD 81.79.

Double checking this result with the simulated spread for this realisation in Fig. 26 on page 63, Fig. 26 shows that the simulated spread drops around 2014-2015. Here, the spread between futures prices for Brent and WTI intensifies, and as the spread widens, the option prices increase drastically. Although this set can be considered as perhaps an extreme outcome, with oil prices reaching USD 300 for WTI (refer to Fig. 24 on page 62), the prices are based on historical data. The model in chapter 4 was calibrated with market data from the 2008-2013 period, hence all outcomes of the model are results of the input to the model. It should be stated, however, that even though all the illustrated samples are possible, they might not be equally likely.

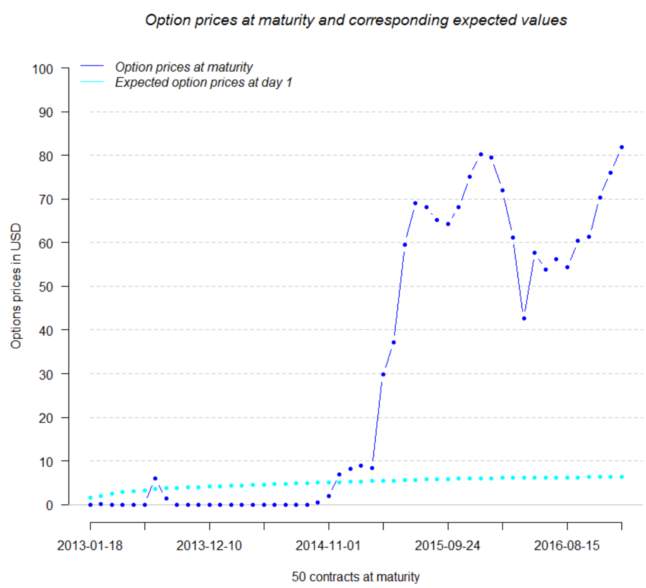

Figure 35: Sample 3: Moneyness of the 50 options at expiry. The dark blue points are the moneyness of the 50 options at expiry, the light blue points are the expected option values priced at Day 1

In Sample 3, the actual payoff at maturity is high compared to the expected payoff. An investor betting on these 50 options without any form of risk management, would be very profitable if this set turned out to be the reality. The investor would get less money than expected in the first half of the set, but would be treated with exceptionally higher prices than expected in the last half. Risk management of this set would, when viewing in retrospect, diminish a great deal of the value possible to achieve in the set.

The option prices for the entire life of the 50 contracts in Sample 3 are illustrated in Fig. 36. This set shows the extreme levels, where Fig. 36 clearly illustrates how the option values peak half-way through the period. As opposed to Sample 1, where the movement in prices goes up and down during the period, Sample 3 shows a unison increase in option prices for all contracts mid-way through the set. The correlation between the contracts is evident in this set as well, as all contracts increase simultaneously around mid-2016.

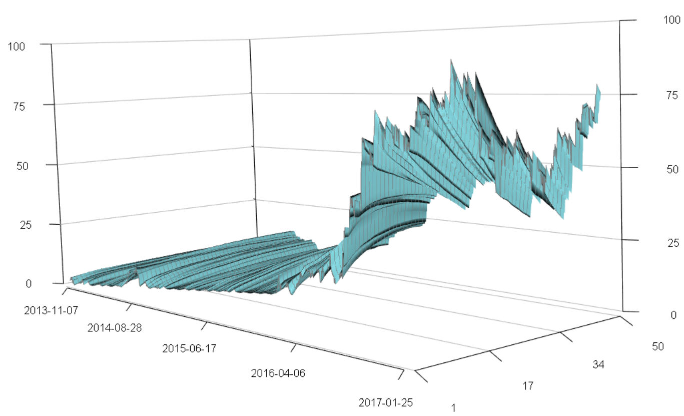

Figure 36: Sample 3: Option prices for all 50 contracts. In the middle of the set, the option values simultaneously increase for all remaining contracts. For the last half of the set, the spread between Brent and WTI futures have widened significantly compared to the spread seen at Day 1

7 Risk Management with Delta Hedging

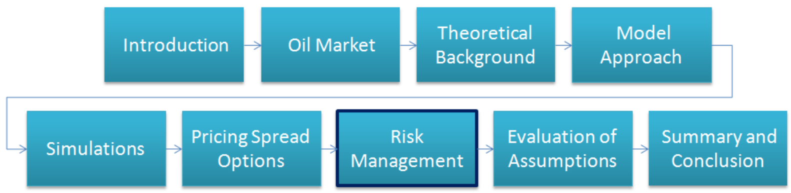

This chapter presents how risk management can be performed on a day-to-day basis to account for price risk in underlying assets. In risk management, portfolios are set up containing both positions in the underlying variable, as well as in the derivative itself. Several different methods are possible to manage risk in options trading, and delta hedging is performed in this thesis as this method is a well established dynamic method. This chapter presents the results of risk managing the three samples with delta hedges that are rebalanced one time per day. The method of delta hedging and the procedure to find delta values with Monte Carlo simulations are explained, as well as a sensitivity analysis of the delta distribution for the 50 contracts at Day 1. The connection between delta values, option values and intrinsic prices is illustrated using selected contracts from Sample 1, as this connection is important going forward with the delta hedge. Finally, the results of the delta hedge is illustrated and analysed for the three samples.

7.1 Why Risk Manage Spread Options?

The previous chapter presented option prices for 50 contracts on three different realisations of the forward paths for Brent and WTI. The results of the samples illustrate that the option values are unpredictable, since the evolutions of the prices of the underlying assets are unpredictable. Due to this high degree of unpredictability and volatility, there is great risk associated with the spread options. From the three samples, it seems like chance if the investor profits or not from holding the options. In Sample 1 and 3 the outcomes were better than expected from an investor's point of view, while in Sample 2 the investor would have suffered a loss compared to the expected payoff. The overall value of flexibility is somewhat up and down in Sample 1, low in Sample 2 and extremely high in Sample 3. This introduces the desire to risk manage the exposure to the options, attempting to lock in the expected value for all 50 contracts.

This thesis performs the delta hedge for the spread options assuming a *frictionless market*, as introduced in chapter 3. This means that transaction costs by purchasing or selling a unit of the underlying asset are not considered, nor is the bid-ask spread (for either the options or the underlying contracts) taken into consideration when technically performing the delta hedge. After the results of the delta hedge are analysed for the three samples, this assumption is discussed in section 7.8.

7.2 Risk Management with Delta Hedging

Institutions that trade energy derivatives are faced with the problem of managing the risk of their positions. For an energy company that sells an option position, the most complete hedge is to buy an exactly offsetting position in the market, as for example long calls bought against short calls. This long position is then held as a static hedge and cannot be changed over time. In most circumstances, this hedge will not be practical or profitable. Hull (2012) refers to static hedging as a "hedge and forget" - method. Dynamic hedges, however, can be changed as often as preferred to encounter new movement in the price of the underlying variable. Delta hedging is an option strategy that aims to reduce (hedge) the risk associated with price movements in the underlying asset by offsetting long and short positions. For example, a long position may be delta hedged by shorting (selling, writing) the underlying asset.

The delta (Δ) of an option is defined as the rate of change, or sensitivity, of the option price with respect to the price of the underlying asset (Hull (2012)). The delta of a call option is positive, and the delta of a put option is negative. If for example the delta of a call option is 0.4, this means that when the price of the underlying asset changes by some small amount, the option price will change by about 40% of that amount. In general

$$\Delta = \frac{\partial V}{\partial S} \quad \text{for calls} \qquad \text{and} \qquad \Delta = \frac{\partial V}{\partial S} - 1 \quad \text{for puts} \tag{45}$$

where Δ is the first derivative of the value of the option, and S is the price of the underlying variable. A position with a delta of zero is referred to as delta neutral. The delta of a position does not remain constant since the value of the underlying asset will change; this means that a traders position remains delta neutral for only a relatively short period of time. Therefore, the hedge needs to be adjusted periodically, also known as rebalancing the delta hedge. Clewlow and Strickland (2000) defines delta hedging an option as dynamically trading a position in the underlying energy such that, over each small interval of time between trades, the change in the option price is offset by an equal and opposite change in the value of the position in the underlying variable. If a portfolio consisting of both an option and a position in the underlying asset is delta hedged, the portfolio should be immune to risk in the underlying variable. Even though the value of the underlying asset as well as the value of the option will change, the delta hedging will make sure that the two changes offset each other, so that the combined value of the portfolio remains the same.

The "Greek letters", or simply the "Greeks"[20], measure different dimensions to risk in options trading. The delta Δ has already been introduced, and is considered the most important Greek letter in risk management of option exposure. The gamma Γ of an option portfolio is the rate of change of the portfolio's delta with respect to the underlying asset price. The gamma is the second derivative of the value of an option, meaning the first derivative of the

[20]The most common "Greeks" are delta, gamma, theta, vega and rho. Read more about the Greeks in Hull (2012).

delta Δ (Hull (2012)). When gamma is small, the delta changes slowly. If the gamma is highly negative or positive, the delta is very sensitive to the price of the underlying asset. Whenever the gamma is high, the delta hedge will need to be rebalanced frequently, as will be further explained when presenting the delta hedge in section 7.7. This thesis performs the risk management based on delta hedging, and does not include gamma as an input to the hedge. This is due to scope limitations of the thesis, and will be evaluated in chapter 8.

7.3 Calculating Delta Values

A derivative is a measure of how a function changes as its input changes. The Greek letter delta in mathematical context is used to describe a change of some kind. In order to delta hedge the portfolio, the delta values need to be calculated. The option price simulations presented earlier returned prices for each of the 50 contracts at any given day from t=0 to t=1050 (when t=1050 there is only one day left to price). The prices of the options can be used to calculate the delta's of the options. Then, additional information is necessary. In the model used for this thesis, the only way to compute the deltas is by Monte Carlo simulations. The equation below is called Newton's quotient

$$\frac{\Delta F(P)}{\Delta P} = \frac{F(P + \Delta P) - F(P)}{\Delta P} \tag{46}$$

where $F(P)$ is the future value of P and ΔP is a small change in the underlying asset for P. In Newton's quotient, in order to find the delta, two values are necessary: The original value of the option based on the underlying asset and another value of the option which is based on the underlying variable added a small increment, as described by ΔP in the equation. Then, the original value is subtracted from the value calculated on the asset with a small incremental change, and divided by that small change, giving us the delta of the option.

Since the prices of the options are based on underlying values of both Brent and WTI, it is necessary to find the delta values with regards to both the Brent data and the WTI data. Therefore, the equation is partial derived, first with regards to Brent data, and second with regards to WTI data. This will return two deltas, as given in the equations below

$$\frac{\partial V(\text{Brent, WTI})}{\partial(\text{Brent})} \approx \frac{V(\text{Brent} + h, \text{WTI}) - V(\text{Brent, WTI})}{h} \tag{47}$$

$$\frac{\partial V(\text{Brent, WTI})}{\partial(\text{WTI})} \approx \frac{V(\text{Brent, WTI} + h) - V(\text{Brent, WTI})}{h} \tag{48}$$

In the first equation there has been added a small change to the value of all Brent data, and in the second equation there has been added a small change to the value of all WTI data. Since calculated as a spread option, where the underlying value is defined as Brent minus WTI, the first change will increase the spread by h, and the second change will decrease the spread with h. In this thesis, the delta values are calculating adding by 1 cent to the underlying Brent and WTI futures contracts.

7.4 Sensitivity Analysis of Delta Values at Day 1

The delta values are calculated one time per day throughout the life of each of the 50 options. A sensitivity analysis is performed to test the accuracy in the delta values on Day 1. Recall that the delta values are calculated based on the prices of the underlying assets, so when calculated on Day 1, the delta values are equal for all possible realisations of the forward paths, similarly as the expected option prices at Day 1 presented in chapter 6. To find a distribution of the delta values at Day 1, the 5 steps of the Monte Carlo process introduced in the previous chapter is followed in the same manner as when the option price distribution was found. For Day 1, 20,000 simulated price paths are found for the tree scenarios required to find the delta values: The original scenario, the scenario with the increase in Brent prices by h, and the scenario with the increase in WTI prices by h. The simulated option values $V(\text{Brent} + h, \text{WTI})$, $V(\text{Brent}, \text{WTI}+h)$, and $V(\text{Brent}, \text{WTI})$ are calculated simultaneously using the same random numbers, making sure that errors in the simulations are taken into consideration. These simulations result in one average delta value for the 50 contracts for each of the three scenarios, as calculated with Eq. 47 and 48. The procedure is repeated 3,000 times to find estimates for the distribution of the delta averages, similarly to the process for finding the distribution of the expected option prices in chapter 6.

The results of the delta value distribution at Day 1 for Delta Brent are presented in Fig. 37. The black line in the middle is the mean value of the 3,000 average delta values per contract, the green lines illustrate the 95 % and 5% confidence levels, and the blue lines are the maximum and minimum averages from the 3,000 delta values.

The delta values lie around -0.5, which is expected for a put option valued at-the-money. The difference between maximum and minimum delta values for the 50 contracts are approximately 0.03 for all 50 contracts, meaning that the number of trials to find delta values seem sufficient due to the narrow distribution outcome. The confidence interval for the mean delta values μ and its standard deviations σ are calculated for $2 \times \sigma < \mu < 2 \times \sigma$. This confidence interval indicates the reliability of the estimated delta values, and is illustrated for each of the 50 contracts in Table 8 and 9 in Appendix C.

7.5 Delta Profiles during the Lifetime of the Options

The delta values are calculated each day for each contract from Day 1 until the last contract goes to maturity. When a contract goes to expiry, its corresponding delta value is always either zero or -1 for Delta Brent, and either zero or +1 for Delta WTI. The delta of an option will vary between 0 and -1 for a put option and between 0 and 1 for a call option during the option's lifetime. If the option matures in-the-money, the delta value will be -1 or +1, in this case -1 for the Delta Brent and +1 for the corresponding Delta WTI. If the option matures at-the-money or out-of-the-money, the delta value will be zero at maturity for both Delta Brent and Delta WTI.

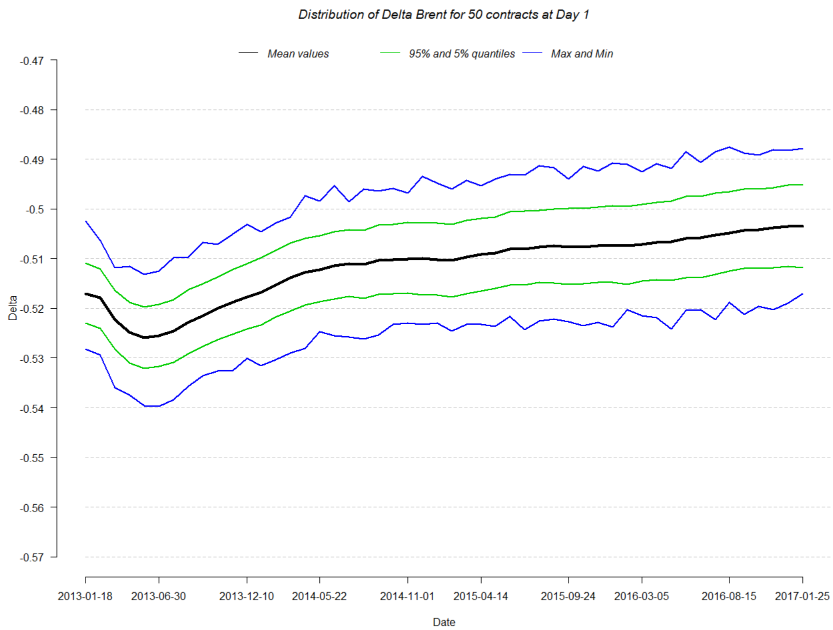

Figure 37: Distribution of Delta Brent for the 50 contracts on Day 1. The spread between maximum and minimum is close to the mean value of the distribution

7.5.1 Sample 1

Two contracts are selected from Sample 1 to illustrate how the delta values evolve during the lifetime of an option. Contract number 24 and 50, which were illustrated in the option pricing part in section 6.3.1, are displayed for both Delta Brent and Delta WTI[21] in Fig. 38 and 39. Contract 24 and 50 both have a net positive sum delta throughout the lifetime of the option, before the sum ends at zero, regardless of whether the option ends in-the-money or not. The sum of the delta values are always positive during the life of the option if the investor holds a long position. Recall that the option value for contract 24 was in-the-money both during the lifetime and at maturity, and that contract 50 had positive option values during the first half of its lifetime, before the prices fell and the option expired worthless. In Fig. 38 and 39, both contracts start with delta values of approximately -0.5 and 0.5 for Delta Brent and Delta WTI, respectively. These start values are the mean values for the two contracts from the distribution illustrated on the previous page.

In contract 24, the delta values have a trend towards moneyness (maturity levels of -1 and +1 for Delta Brent and Delta WTI, respectively) during the

[21]The Delta WTI is the opposite of Delta Brent, or more specifically Delta Brent $\times$ -1. How Delta Brent and Delta WTI correspond is illustrated in Fig. 38 and 39, before only Delta Brent is illustrated for simplicity reasons going forward in the thesis.

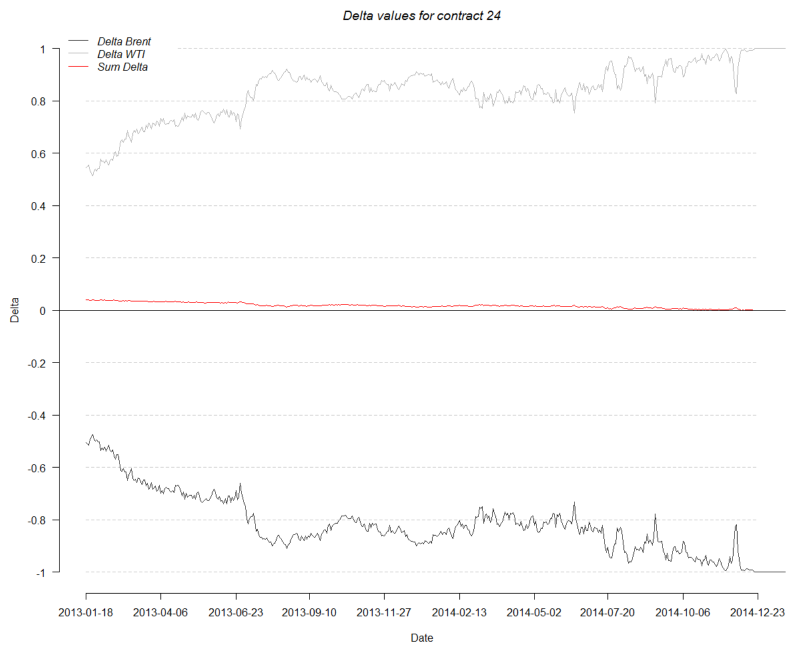

Figure 38: Sample 1: Delta Brent and Delta WTI for contract 24

lifetime of the contract. Notice the change in the delta around July 2013 where the Delta Brent goes from -0.7 to -0.9, and the Delta Brent goes from 0.7 to 0.9. This change in the delta towards moneyness corresponds with the sharp increase in option values illustrated in Fig. 29 for contract 24. Since the delta values go steadily towards moneyness for this contract, the expectation is that the option will expire in-the-money during the entire lifetime of the option.

In contract 50, the delta values fluctuate more during the lifetime of the option. The delta values vary around -0.5 to -0.8 and 0.5 to 0.8 for Brent delta and WTI delta, respectively, until 2016, where the delta values are pulled towards zero. This corresponds with how the option values decrease at this point, as illustrated in Fig. 30. Notice the jump towards -1 and +1 and again back towards zero in the final months of the contract. This jump corresponds with the small spike noticed at the end of contract 50. In the contract, the value goes towards zero, leading to the delta illustrating an expectation of the option expiring out-of-the-money with values towards zero. Then, when the option value suddenly increases, the change is captured by the delta, which moves fast towards -1 and +1. When the option value goes back towards zero, so does the delta value. This large fluctuation in the delta so close to maturity indicates that the delta has a high gamma, as previously introduced.

Contract 24 and 50 have illustrated how the delta values can look during the lifetime of the option. Fig. 40 shows all delta values for Delta Brent in

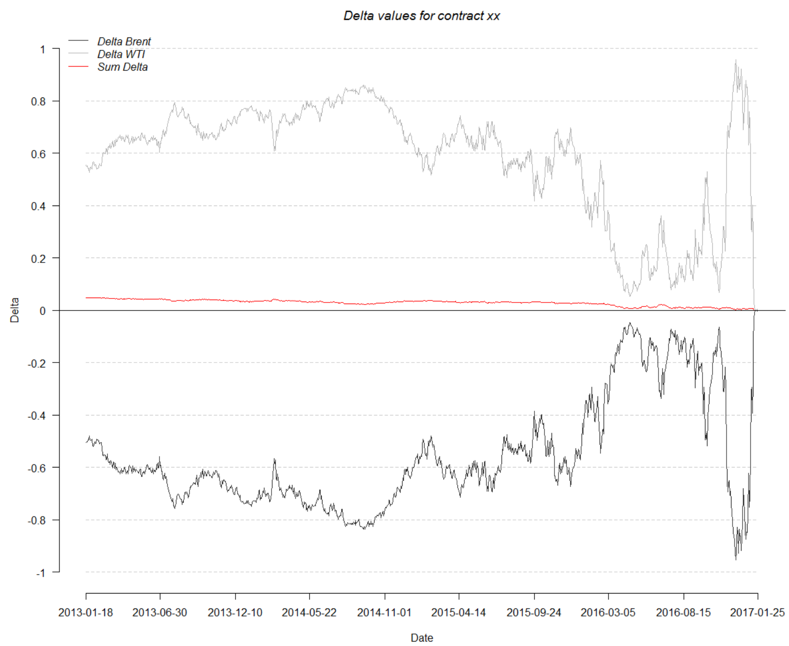

Figure 39: Sample 1: Delta Brent and Delta WTI for contract 50

Realisation 1, both during the lifetime and at maturity of the options. This figure corresponds to the option values illustrated in Fig. 31 on page 73, and, similarly to the options values, the delta values are also highly correlated. The delta values are quite stable from 2013 up until the start of 2016, where the delta values experience larger marginal changes. An investor would appreciate small changes in delta values, as this simplifies the risk management of the options, which will be thoroughly explained in a later section.

The Delta Brent values for the corresponding option values of contract 30 and 31 introduced in section 5.3 are illustrated in Fig. 41. The delta values of the two options are highly correlated, and both delta values face large changes in the last part of the options lifetimes. The delta drops from -0.7 to almost -1 from late 2014 to early 2015. These large changes in delta values within the short period of time will have great impact when attempting to hedge the risk associated with these option prices, as will be discussed when implementing the hedging strategy for this sample.

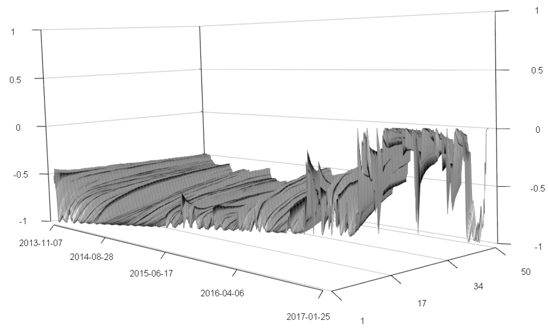

Figure 40: Sample 1: Delta Brent values for all 50 options. The delta values are stable in the first half of the set, before they face larger marginal changes

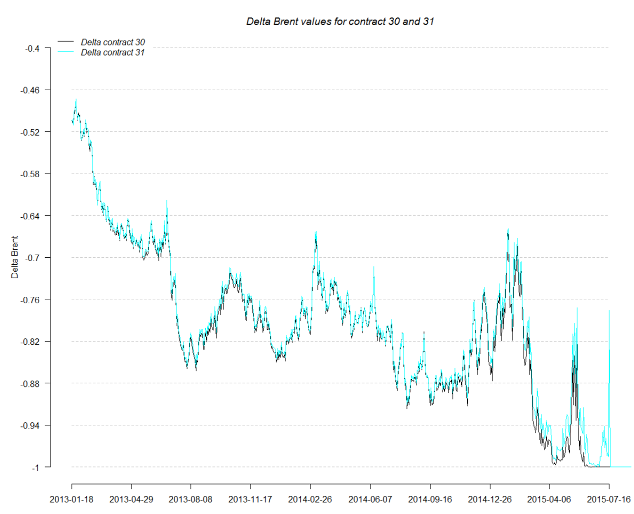

Figure 41: Sample 1: Delta Brent values for option contract 30 and 31. Contract 30, illustrated by the black line, and contract 31, illustrated by the light blue line, show high levels of correlation

7.5.2 Sample 2

Fig. 42 shows the delta values for Delta Brent in Sample 2. This figure corresponds to the option values illustrated in Fig. 34 on page 75. Fig. 42 shows how the delta values end at either zero or -1 at maturity for each contract. The first year of the set has delta values steadily maturing at -1, recall Fig. 33 where the option values at maturity are illustrated. The first options mature in-the-money, succeeded by a long period where the option values expire out-of-the-money. In the last period of the set, the delta values are fluctuating a great deal, illustrating volatile option prices for the last approximately 15 contracts.

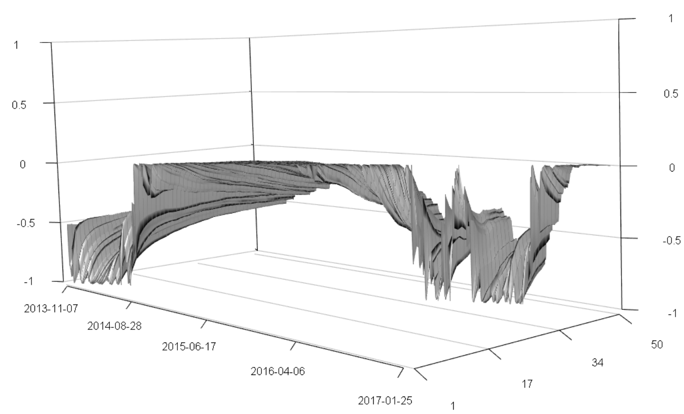

Figure 42: Sample 2: Delta Brent values for all 50 options. The long period where the delta is relatively calm in the middle of the set is when the option values are all out-of-the-money. The delta values vary a great deal towards the end of the set, when the corresponding option values fluctuates

7.5.3 Sample 3

Fig. 43 shows the delta values for Delta Brent in Sample 3. This figure corresponds to the option values illustrated in Fig. 36 on page 77. The delta values fluctuate in the beginning of the set, before they stay around -1 during the last half of the set. Recall from Fig. 36 how high the option values went in this part of the simulation period. The delta values stay close to -1 for all contracts during the last half of the period, meaning that it is a clear expectation that all the contracts will expire in-the-money.

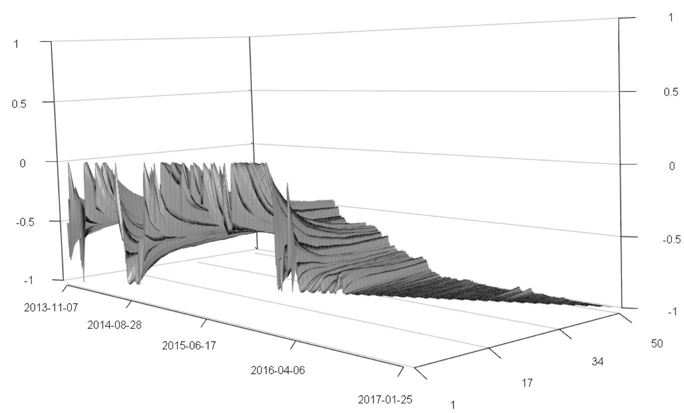

Figure 43: Sample 3: Delta Brent values for all 50 contracts. Notice the correlation for the delta values at the last half of the set, where it is expected that all contracts expire in-the-money, giving delta values of -1

7.6 Connection Intrinsic Value, Option Value and Delta Value

The interconnections between the delta value, the option value and the value of the underlying asset are thoroughly explained before starting to delta hedge the options. Recall from Eq. 47 and 48 that the delta values are derived from the value of the underlying asset. Since the delta value explains the relationship between the underlying asset price and the option price, the connection between the three variables are presented using examples from Sample 1.

Fig. 44 is an illustration of the simplicity in the connection between the delta values and the options moneyness at maturity. When an option expires in-the-money, the Delta Brent expires at -1. When the option values expire at-the-money or out-of-the-money, the Delta Brent expires at zero. If the option expires exactly at-the-money, the investor will be indifferent to whether or not to exercise the option, disregarding the initial cost of buying the option. As seen in Fig. 44, the delta values are insensitive to the degree of moneyness for the option value at expiry; they are either zero or -1.

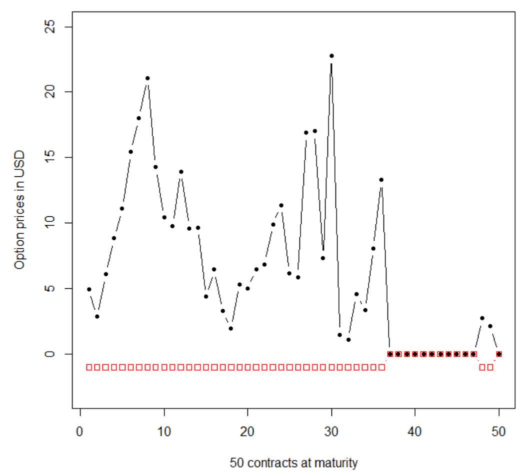

Figure 44: Sample 1: Option- and delta values at maturity for all 50 contracts

More interesting is to examine the evolution of the intrinsic value of the underlying variable together with both the option values and the delta values throughout the lifetime of a contract. Looking again at contract 24 and 50 in Sample 1, Fig. 45 and 46 illustrate how the intrinsic value, the option value and the delta value of each option are connected. The intrinsic values and the option values are linked to the left y-axis, and the delta values are linked to the right y-axis. The intrinsic value shows how the spread has evolved during the lifetime of the option. As opposed to the simulated spread illustrated previously, the intrinsic value in this graph is displayed as a positive value: An intrinsic value above zero indicates that the price spread between Brent and WTI futures has widened compared to how the spread was initially (at Day 1). Positive intrinsic values here will lead to positive option values, as opposed to a negative intrinsic values, which indicate options going out-of-the-money. The red line is the option value, which is always above, or at, the intrinsic value. The difference between the intrinsic value and the option price at t=0 is the expected option price of a particular contract at Day 1, as can be recalled from Fig. 27, showing the expected option prices. In both contract 24 and 50, when the intrinsic value increase, so does the option value at the same time as the delta go closer to -1. This confirms that the price spread in the futures contract between Brent and WTI drive the option prices. Notice how for contract 50, the intrinsic value pulls up the option price in 2017, causing a quick move in the delta value towards -1.

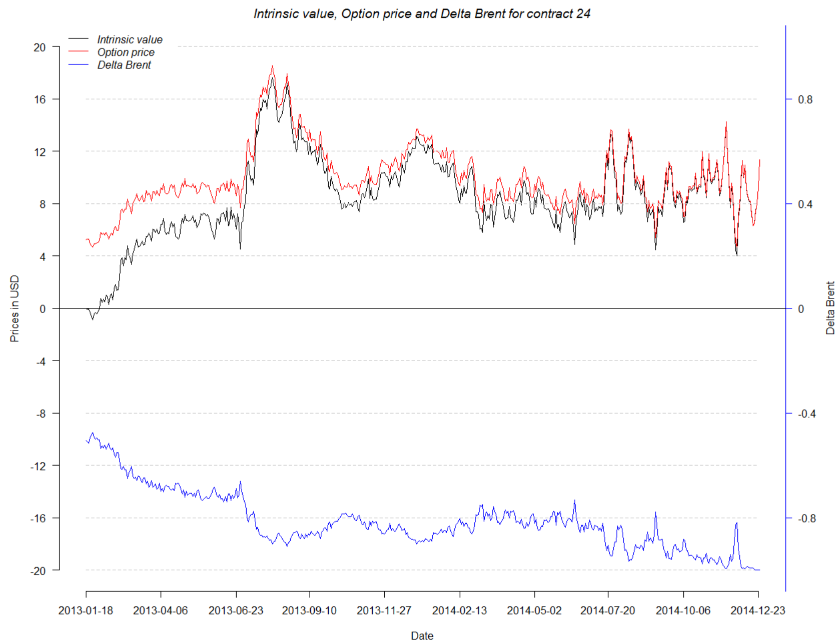

Figure 45: Sample 1: Contract 24: Intrinsic-, option- and delta values

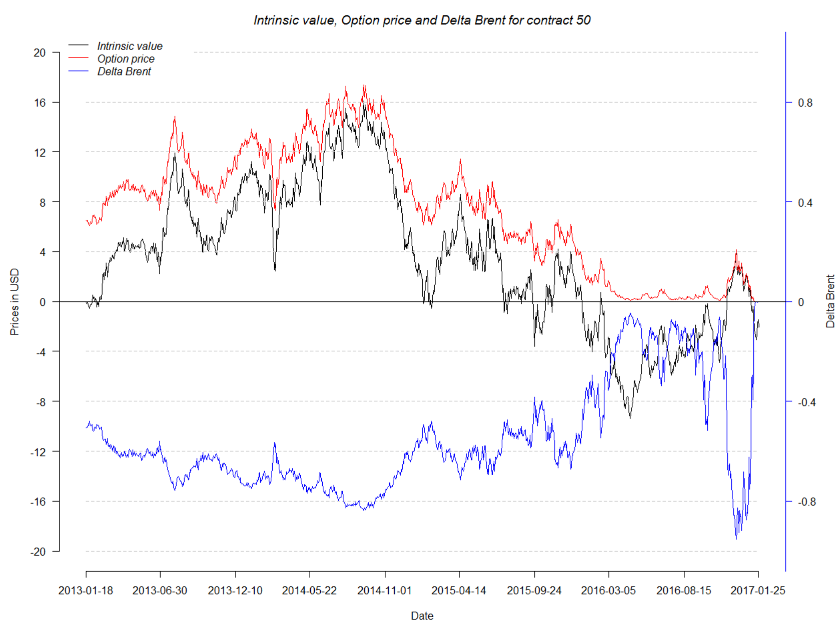

Figure 46: Sample 1: Contract 50: Intrinsic-, option- and delta values

7.7 Performing the Delta Hedge for the Spread Options

The delta hedge is rebalanced every day throughout the life of the options. After calculating the delta values for the 1050 days (recall 50 contracts × 21 days per month = 1050), the delta hedge can be constructed. Day 1 cannot be hedged, so the delta hedge begins for Day 2 in each of the samples. For Day 2, the following steps are performed to find the cash flow from the hedge

$$\text{Delta Hedge Brent} = (F_{2,\text{Brent}} - F_{1,\text{Brent}}) \times \frac{\partial V_2(\text{Brent, WTI})}{\partial(\text{Brent})} \tag{49}$$

$$\text{Delta Hedge WTI} = (F_{2,\text{WTI}} - F_{1,\text{WTI}}) \times \frac{\partial V_2(\text{Brent, WTI})}{\partial(\text{WTI})} \tag{50}$$

$$\text{Cash Flow(CF)} = \text{Delta Hedge Brent} + \text{Delta Hedge WTI} \tag{51}$$

where $F_{2,Brent}$ and $F_{2,WTI}$ are the simulated futures prices at Day 2, and $F_{1,Brent}$ and $F_{1,WTI}$ are the simulated futures prices at Day 1 for Brent and WTI, respectively. For each of the hedge values, the price at Day 1 is subtracted from the price at Day 2 to find the price change, then multiplied with either Delta Brent or Delta WTI for Day 2, as calculated with Eq. 47 and 48. Eq. 51 shows how the replicated portfolio is a result of both the delta values calculated from the underlying asset and the changes in the simulated underlying values (the spread between Brent and WTI futures prices). The cash flow calculates the replicated payoff from the portfolio, so that when loss occurs on the option, the payoff from the cash flow will offset this loss with an equally sized profit. Then, the value of the net position stays equal to before the change in the option price. The cash flow is calculated each day for each of the 50 contracts in the three samples.

For an investor risk managing the exposure to the spread option, the total payoff depends on both the result of the hedge as well as on the option values. The total payoff is calculated as follows:

$$\text{Total Payoff}_t = \text{Exp. Option Price - Acc(CF)}_t + \text{Option Price}_t \tag{52}$$

where t represents the days from 1 to 1050 per sample, and the expected option prices (illustrated in the distribution for the option prices in Fig. 27) are identical for the three samples. The total payoff is then calculated each day for every contract, as will be illustrated shortly. Notice that the payoff at any given day t depends on the accumulated cash flow from the hedge, meaning all losses and gains from the hedge until t. If the payoff is zero, the delta hedge has worked perfectly up until t.

7.7.1 Sample 1

Fig. 47 illustrates the cash flow from the hedge/replicated portfolio during the time span of the 50 contracts in Sample 1. Each day, an investor would experience either a loss or a gain, as visualized by the cash flow evolving around zero. The front of the graph is where the contracts go to maturity. The payoff from the hedge is volatile, and the figure shows large differences in payoff from day to day. Keep in mind that the changes in payoff illustrated here are only from day to day, and not accumulated throughout the lifetime of each option. Although large fluctuations in the set, the cash flow looks similar during the entire sample. The large fluctuations in the cash flow from the hedge indicate large movements in the underlying variable, which can make the risk management process challenging.

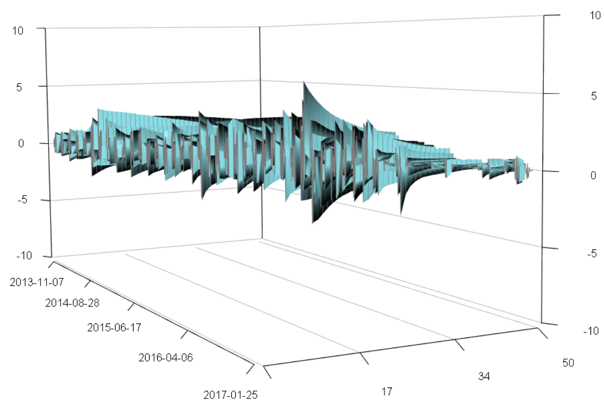

Figure 47: Sample 1: Cash Flow from the replicated hedged portfolio

Fig. 47 only illustrates the cash flow from the hedge, and does not say anything about the total payoff for the risk manager. Consider again contract 24 and 50 in Sample 1 to see the total payoff from the delta hedging. Fig. 48 displays the total payoff for contract 24 and Fig. 49 displays the total payoff for contract 50. In both graphs, the black line is the intrinsic value of the underlying asset; the simulated spread between Brent and WTI adjusted for the strike price as explained earlier. The red line is the option price, which will always be above or equal to the intrinsic value. For any option contract, the difference between the option price and the intrinsic value when t=0 is the expected value of the option at Day 1. The blue line is the accumulated cash flow from the hedge, as calculated and accumulated each day with Eq. 49 to 51. The light blue line is the total payoff from the delta hedged position. The intrinsic value, the option value and the value of the accumulated cash flow are linked to the left y-axis, while the total payoff is linked to the right y-axis to be able to see the evolution of all variables better. The expected values that the investor attempts to hold (keep) by risk managing the exposure can be seen in Fig. 28, and is approximately USD 5 for contract 24 and USD 6.5 for contract 50.

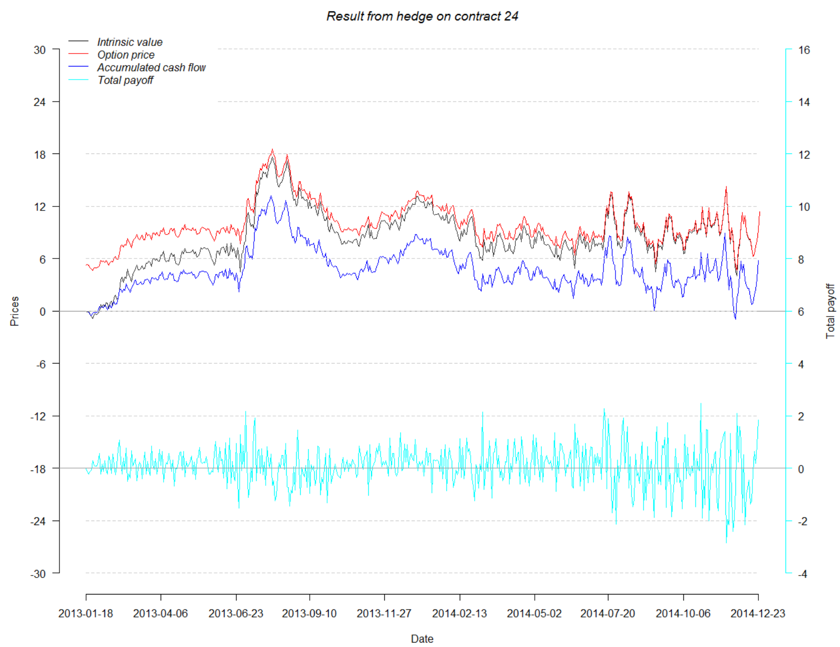

Figure 48: Sample 1: Contract 24: Intrinsic value, option price and accumulated cash flow from the hedge linked to the left y-axis. Total payoff from the risk managed position is linked to the right y-axis

For contract 24 in Fig. 48, the light blue line is steady during the lifetime of the option, and fluctuates between -2 and 2. For a well-functioning delta hedge, it is expected that the light blue line stays around zero. If it is exactly zero, there has been a perfect hedge. The total payoff for contract 24 is more unstable as it reaches maturity. This can be seen in correspondence with the increasing volatility in the months before the expiry of a contract. Contract 24 ends with a payoff of about USD 2 (exactly USD 1.8273), which is coincidental. The payoff could might as well have ended below zero, and the graph shows that the total payoff was below zero only a few days before expiry. By comparing the expected value of the option with the results of the hedge, it can be stated that the variations between USD -2 and USD 2 are large, since the expected value of the option is "only" USD 5.

For contract 50 in Fig. 49, the total payoff is fluctuating around zero and is stable until summer 2015. Then, the payoff falls below the expected average of zero, and stays with a negative sign for the remainder of the lifetime of contract

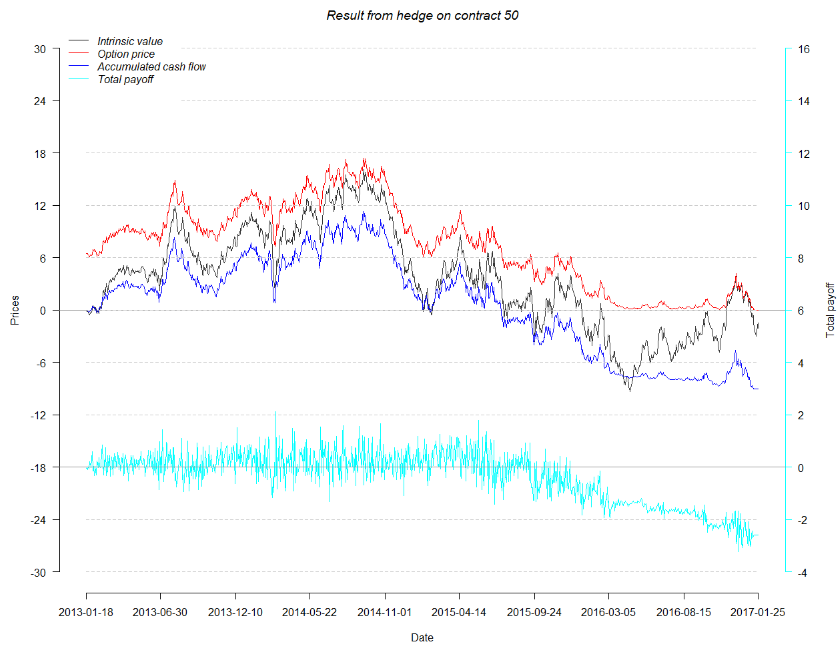

Figure 49: Sample 1: Contract 50: Intrinsic value, option price and accumulated cash flow from hedge linked to the left y-axis. Total payoff from the risk managed position is linked to the right y-axis. Notice how the total payoff endures losses from early 2016 to the maturity of the contract

50. The total payoff ends at about USD -3 (exactly USD -2.5888), which for this contract does not appear as coincidental as for contract 24. Compared to the expected value of 6.5, almost half the value is lost by delta hedging. The average of the total payoff during the last year of the lifetime of the option lies around USD -2. This means that the replicated portfolio/the cash flow from the hedge is too low to offset the change in option value. Fig. 49 illustrates that the total payoff drops to lower values around the start of 2016, implicating a sudden change in the intrinsic value around that time. If the value of the underlying variable changes fast, the delta value may be a bit ”off”, if it does not manage to capture the change. Since the portfolio is changed only one time per day, sudden changes can appear in the value of the underlying asset before the delta hedge is rebalanced. The spread between Brent and WTI futures has such a high volatility that the delta hedge has not been able to respond to the sudden changes in the value of the spread. When the underlying asset changes faster than the delta responds, the delta may be ”off” either by being too low or too high. Whether the delta is too low or too high is coincidental, and it is random that the payoff in this case has dropped instead of risen. However, if the delta is too low at some point in the hedged portfolio, it will affect the entire remaining

part of the contracts due to the high correlation for the dataset. In contract 50, when the total payoff turns below zero in 2016, the payoff stays below zero due to the high correlation in the set. Therefore, this impact on the total value will affect all contracts with maturity some time after 2016.

The same finding can be cross-checked with the delta values illustrated for contract 24 and 50 in Fig. 38 and 39. The delta values for contract 24 was quite stable throughout the duration of the contract, and the total payoff from the delta hedged portfolio for this contract stayed symmetrical around zero. The delta values for contract 50 were calm during the first half, then the delta values had large marginal changes for the rest of the set. Large fluctuations in the delta values mean that the delta value is sensitive to the changes of the option value, meaning that it has a high gamma. Perhaps introducing the gamma variable could increase the functionality of the hedge.

Even though it is random in the case for contract 24 whether or not the total payoff turns below or above zero; the goal is to have the payoff surrounding and averaging zero the entire time and for all contracts, and this has not worked for contract 50. Fig. 50 illustrates where the total payoff from the hedged portfolio of each contract ends. From contract number 1 - 29 the results from the total payoff lie around zero, then all the remaining results are low below zero. The large drop between contract 28 and 29, from a payoff of USD -2 to a payoff of USD -7 corresponds with the time where the hedge in contract 50 falls, around Summer 2015. This point in time affect all the remaining contracts (contract 29-50), due to the high correlation. This means that when contract 24 proved to have a well-functioning hedge, this is because the sudden change in the underlying had not yet occurred, since contract 24 expired early 2015.

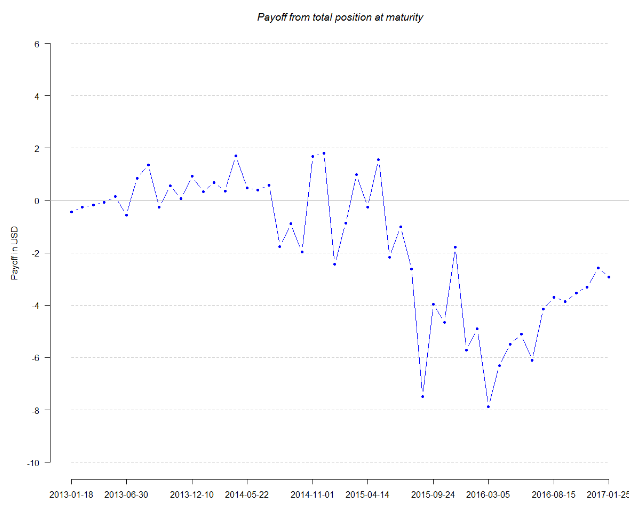

Figure 50: Sample 1: End value for total payoff in the 50 contracts, clearly illustrating that the delta hedge has not had the desired effect for this sample

Fig. 50 only shows the ultimate finishing point, and does not illustrate what happened in each contract in the months before maturity. Fig. 51 illustrates all total payoffs both during, and at the end, of each contract. The graph confirms the fact that the hedge seem to loose its function for the last 20 contracts in the set. For contract 1-29, the payoff is centred around zero, with random outcomes on both sides of zero. After contract 29, the trend for the payoff goes downwards, leading to losses from the total hedged position for all the remaining contracts. This high level of correlation can be partially explained by the first component found by the principal component analysis in chapter 4, and Fig. 50 clearly illustrate that the forward curves are exposed to a "shift" factor highly correlated for all contracts.

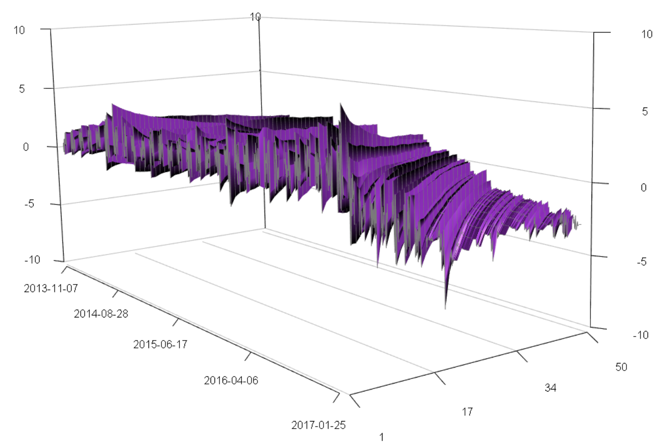

Figure 51: Sample 1: Total payoff from hedged portfolio, illustrating the correlation in the total payoff for the 50 contracts. Notice how the total payoff is symmetrical around zero until 2016, before it falls to an average below zero for the final 20 contracts

7.7.2 Sample 2

Fig. 52 displays the total payoff profile for contract number 50 in Realisation 2. The black line is the intrinsic value, which is the simulated spread between Brent and WTI futures prices adjusted for the strike price. The red line is the option price, which fluctuates between USD 2 and USD 6 before it declines and matures out-of-the-money. Recall that the option price will always be equal to, or higher than, the intrinsic value. The difference between these curves at the starting point is the expected option price. The blue line is the accumulated cash flow from the replicated portfolio. The light blue line is the total payoff from the hedged position.

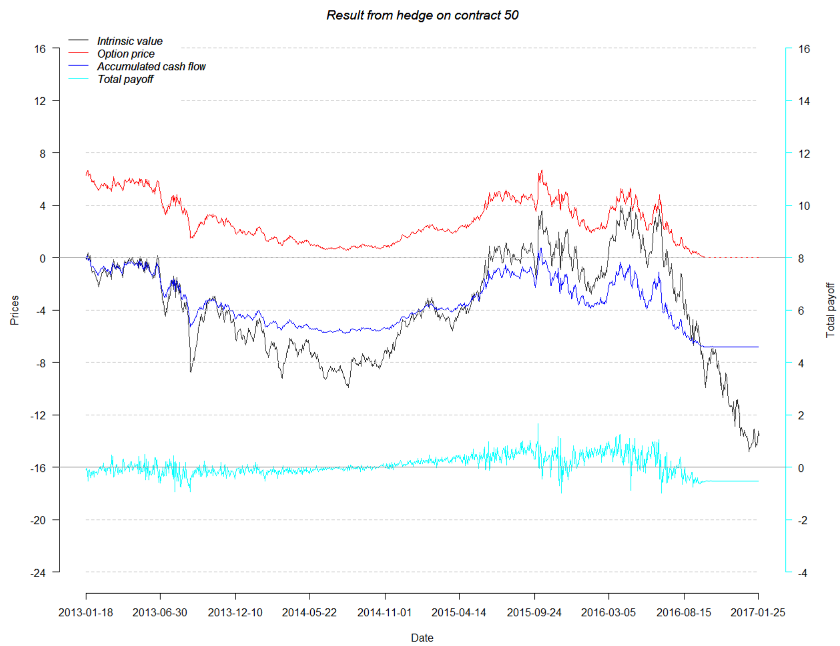

Figure 52: Sample 2: Contract 50: Intrinsic value, option price and accumulated cash flow from the hedge linked to the left y-axis. Total payoff from the risk managed position is linked to the right y-axis. The total payoff from the hedged position stays calm throughout the contract

Due to the high correlation in the dataset, contract 50 in Fig. 52 illustrate how the hedging strategy has worked for the whole set in Sample 3, since this contract has the longest duration. The payoff profile for contract 50 implies how the strategy has performed for the previous contracts in the sample. The figure shows that when the intrinsic value (the simulated spread) decreases, so does the value of the option. The accumulated cash flow from the hedge replicates the option value, and manages to offset the changes in the option value. The hedging strategy has worked quite well for this sample, with payoff values fluctuating close around zero for this contract. Notice the area between 2013 and 2015 when the payoff is very stable; this is the time period when contract number 8-34 expire out-of-the-money. At this time the hedge will be very stable, since the delta values are zero because the options are expected to expire out-of-the-money, as presented in the delta values for Sample 2 in Fig.

42.

The stability in the total payoffs in Sample 2 can also be seen in Fig. 53, showing the total payoff at maturity for the 50 contracts in this sample. In total, the total payoff is always in between the interval USD ± 2, as opposed to Sample 1, where the total payoff fluctuated between USD 2 and USD -8. Although the total payoff at the beginning and end of the set fluctuates to a certain degree, the results of the delta hedge for this set can be considered successful.

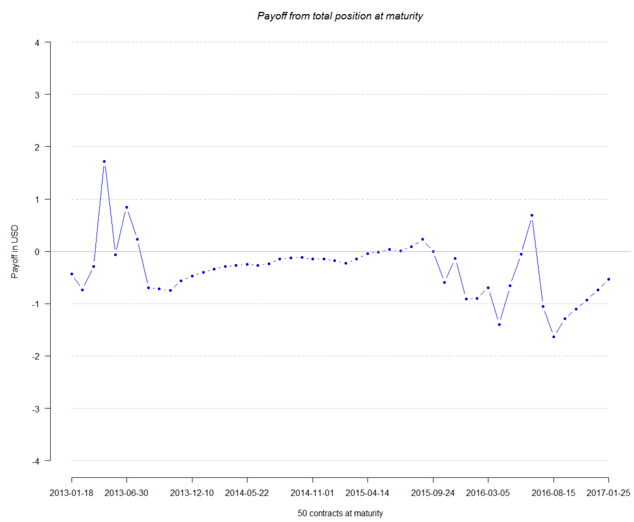

Figure 53: Sample 2: End value for total payoff in the 50 contracts, illustrating results close to zero

Fig. 54 illustrates a combination of the previous two figures: Payoff both during and at the maturity of each contract in Sample 2. Notice how flat the total payoff is when contract 8-34 goes to expiry, with values of zero for both the expiring contracts and for the contracts with long maturities at that time. The contracts with long maturities are correlated with those having short maturities, so that when contracts expire worthless in this period, the upcoming contracts are expected to do the same. For the final months of the sample, some contracts expire in-the-money and some expire out-of-the-money, causing fluctuation seen in the total payoff in 2016-2017.

The calm marginal movements in this sample makes it possible for the delta values to be adjusted in time to offset changes in the value of the options. Especially in the period where the options expire worthless, the total payoff from the hedge lies extremely close to zero, indicating a well-functioning hedge for these contracts. However, although the risk management has worked rather well for this set, the changes in the price spread for Brent and WTI have not been large; meaning that the exposure in this set has been "easier" to hedge than the previous set.

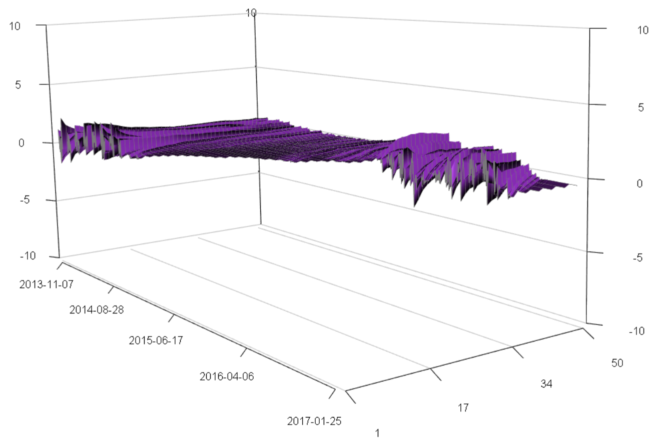

Figure 54: Sample 2: Total payoff from hedged portfolio. Payoff close to zero indicates a well-functioning hedge for this sample

7.7.3 Sample 3

Fig. 55 illustrates the payoff profile for contract number 50 in Realisation 3. The black line is the intrinsic value, the red line is the option price, the blue line is the accumulated cash flow from the hedge and the light blue line illustrates the total payoff for the hedged position. This sample is very different from both Sample 1 and Sample 2, especially notice the major fluctuations in the total payoff in the last half of the set. In the first half, the payoff stays close to zero, moving back and forth between USD ± 1.5. Then, the payoff profile becomes unstable, although still being symmetrical around zero. Recall that the option prices reach extreme levels for this set, and at the same time as the option prices sky-rocket, the total payoff starts fluctuating; the total payoff has its highest level around USD 6 and its lowest level around USD -6, with most levels fluctuating between USD $\pm 3 - 4$.

Although the total payoff fluctuates more than an investor would appreciate, the fact that the payoff remains symmetrical around zero indicates that the delta values have not been offset by quick changes in the value of the underlying, in contrast to what happened in Sample 1. How can the large variation in the last half of the contract be explained? Fig. 56 illustrates variation and variance for the total payoff in contract 50 for Realisation 3. At the top left in the figure is

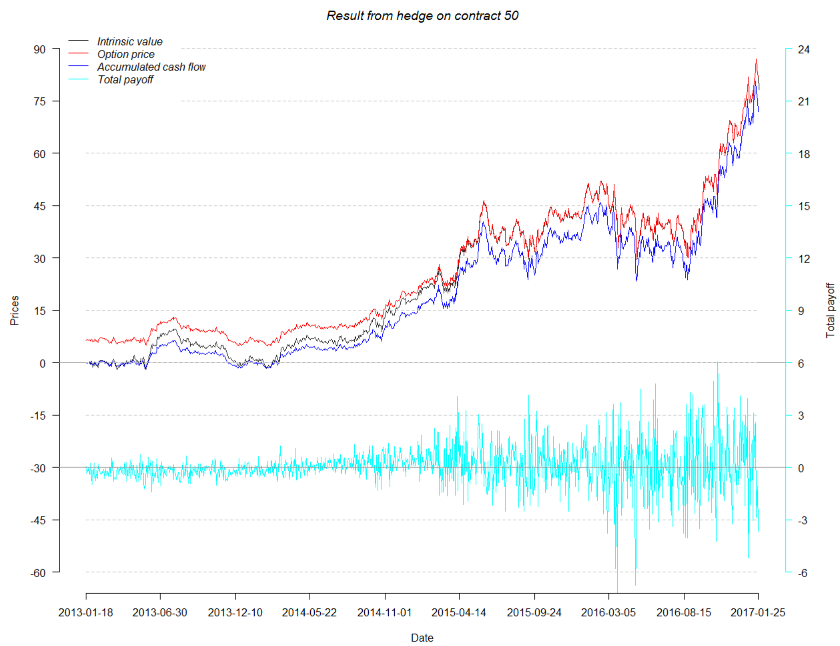

Figure 55: Sample 3: Contract 50: Intrinsic value, option price and accumulated cash flow from the hedge linked to the left y-axis. Total payoff from the risk managed position is linked to the right y-axis. Notice the large variations in the total payoff from mid-2016

the cumulative payoff for the contract, which is exactly the same as the light blue line illustrated in the payoff for contract 50 in Fig. 55. The graph to the top right in the figure is the relative, or non-cumulative, payoff from the contract. This is the cumulative payoff divided by the option values, returning the total payoff relative to the option prices. Local instantaneous variance for the payoff are observed in the two graphs at the bottom of Fig. 56. The graph to the bottom left is the local instantaneous variance in the accumulated payoff and the graph to the bottom right is the local instantaneous variance in the non-accumulated payoff.

Fig. 56 illustrates that even though the local variance is high for the last half in the set for the cumulative payoff, the relative local variance is more "flat" over the duration of the contract. Notice also in the bottom right graph how the variance move in clusters; periods of high variance are followed by periods of high variance, and periods of low variance is followed by periods of low variance. Since the variance is fitted with an EGARCH model, the clustering of variance and returns are captured by the model introduced in chapter 4.

Looking at the total payoff for all 50 contracts in Sample 3 at maturity, Fig. 57 illustrates that the payoff for the last half of the set varies a great

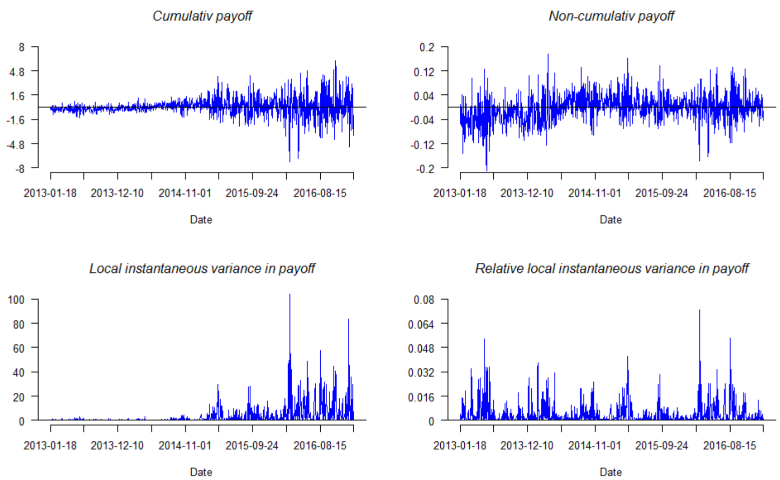

Figure 56: Sample 3: Variation and variance in the total payoff for the hedged position for contract 50. The two graphs to the left show cumulative payoff and variance and the graphs to the right show non-cumulative payoff and relative variance

deal between each contract, fluctuating between almost minus 4 to plus 4 from one contract to the next. Recall that the futures prices simulated for Brent and WTI in Sample 3 reached extreme levels, with WTI prices rising above USD 300 per barrel, and Brent prices almost to USD 250 per barrel. The large price spread that develops between the prices for the two variables create large option values, but this does not necessarily makes it more difficult to hedge the exposure to this set. Fig. 57 illustrates that the hedge did not work as good as in Sample 2, but not as bad as in Sample 1. The variations in the total payoff are large, but the average value of the total payoffs remains around zero for the entire set. The high variations in the last half of the set depend on an effect not previously observed in the first or second sample; the effect of rising prices observed in dollar terms. An easy example can simplify this effect: A 1% change in a price of oil at USD 300 is larger than a 1% change in a price of oil at USD 100. This means that the large variations in the payoff appear when the prices peak, around the middle of the contract. This results in a variation in the payoff in dollar terms, while the relative variation stays the same as before the price increase. The same observation is seen in the variance of the cumulative and non-cumulative payoff. Looking back at the simulated annual volatility for both Brent and WTI in this set, in Fig. 25 at page 63, this figure also confirm that the volatility in the last part of the period is not any higher than it is in the first part of the set, confirming the impact of the dollar effect in this sample.

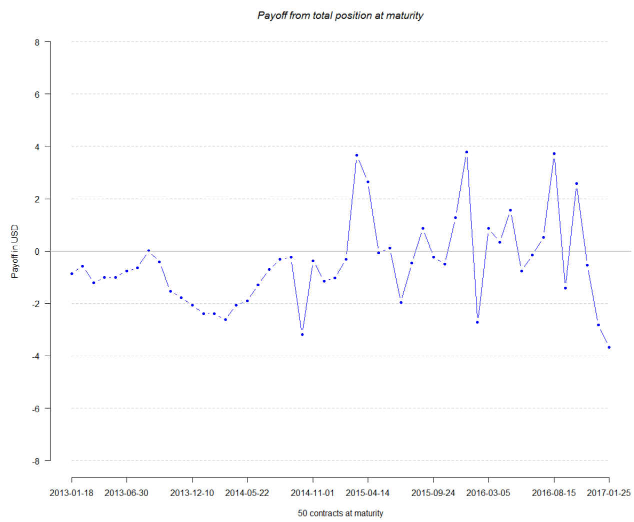

Figure 57: Sample 3: End value for total payoff in the 50 contracts. Note the large variations between each contract in the last half of the set

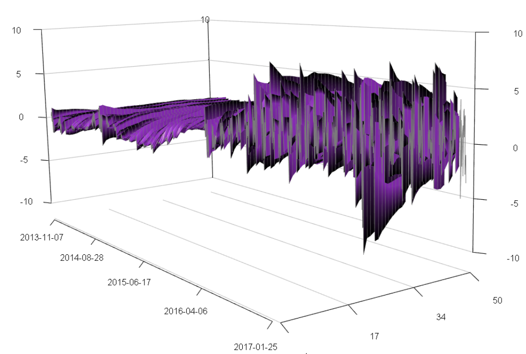

Figure 58: Sample 3: Total payoff from hedged portfolio. The large variations in the last half are too large compared to the expected values for the option contracts

7.8 Considering Bid-Ask Spread and other Transaction Costs

Fig. 58 illustrates how the payoff during each contract has evolved. The same trend observed earlier can also be spotted here, with calm movements in the first half, before the payoff starts fluctuating.

The delta hedge has proved to be both successful and less adequate for the three selected samples of futures prices for Brent and WTI. A characteristic of the market that has been set aside for this analysis is the bid/ask spread existing in the marketplace, as well as transaction costs associated with every purchase or sale. If the investor holding the portfolio has a small number of options on the futures spread, this is liable to be prohibitively expensive due to the transaction costs incurred on trades (Hull (2012)). For a large portfolio of options it is more feasible, since only one trade in the underlying is necessary each day when rebalancing the hedge. The bid/ask spread, as introduced previously, is the difference between the prices quoted by a market maker for immediate purchase (bid) and immediate sale (ask). The size of the spread from one asset to another will differ mainly because of the difference in liquidity of each asset. When the delta hedge is rebalanced once a day, the investor buys/sells shares in the underlying asset one time per day. Then, the investor suffers from the bid-ask spread for every trade performed, proving that the risk management can be expensive. This can make the risk management of the exposure to the options less appealing, especially when considering the fact that the risk management did not work satisfactory for all three samples in the first case.

Another downside with delta hedging is that it requires the investor to buy or sell an exactly offsetting position in the underlying asset whenever the hedge is rebalanced; however, it is not always possible to be completely exact. It is only possible to buy whole units of the underlying asset, not for example 1.5 underlying futures contracts. Logically, this will have greater impact on a portfolio with a small number of options than on a portfolio with a large number of options.

In addition to the aforementioned costs with delta hedging, transaction costs such as front-office costs, back-office costs and taxation have not been included in the delta hedge performed in this thesis. These additional costs will increase the costs associated with the risk management technique in practice.

8 Evaluation of Assumptions and Scope Limitations

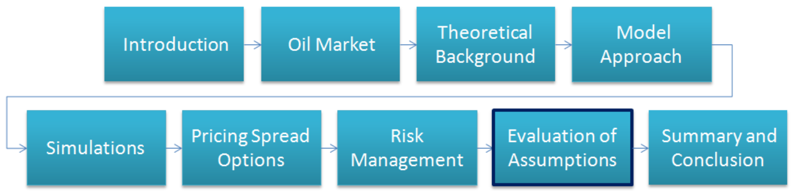

The assumptions and scope limitations made when working with the thesis should be evaluated to test the validity of the results.

The historical data used as input for the model was futures prices observed in the oil market from 2008 to 2013. Perhaps this period is too short to include all the dynamics of the forward curves. In addition, the past five years have been extremely volatile in the oil markets due to the financial crisis. This sets the basis for the simulations going forward, and results in an extremely volatile dataset. This might have overestimated the range of possible extreme outcomes in the model, since the realisations of forward paths are based on the model's input. Thus, the upcoming 5 years are estimated to be as volatile as the period used for input data, which may or may not be the case in reality. However, the financial crisis came surprising to the global energy market, and a new crisis in which prices of USD 300 per barrel occur, as found in Sample 3, could be feasible.

The adjustment made for the Brent and WTI forward curves, as presented in chapter 4, manipulated the WTI data series so that it would be able to compare the two benchmark prices. This was a simple adjustment, done by lifting the WTI forward curve up from its present level to the level where the first month was similar to the first month of the Brent forward curve. The adjustment simplifies the transportation of the crude oils, by considering the shipment of Brent oil from the North Sea, and the WTI from Cushing, to the East Coast of U.S. as a constant cost. Perhaps this simplification understated the importance of transportation costs and made unrealistic outcomes.

The results could have been tested against real market data. However, this was difficult as the price spread of Brent and WTI was adjusted in the beginning of the simulations, making it difficult to find similar situation in the marketplace. In addition, spread options on futures of Brent and WTI are not a standardized product at any exchange, making the existing information in the market inadequate to compare to the results of this analysis. Read Sollie (2013) for his testing of the model with call and put options on Brent futures with different maturities, where the multi-factor forward curve model is proven to capture the observed option prices in the market.

When pricing the options, an interest rate of zero was assumed. This was seen in the expected option prices at Day 1, as the value depended mostly on the time-value of the options. The choice to keep the interest rate at zero throughout the analysis was both for simplicity reasons and also the fact that an appropriate interest rate for discounting future value was difficult to determine.

At the beginning of the thesis, the plan was to consider only one realisation of the forward curves. In the middle of analysing the first sample, the interest and curiosity to look at other sets appeared, and two more sets were simulated and priced. Three sets have obviously been better as an analysis basis to capture the functionality of the delta hedge than only regarding the initial first set, but even more realisations of the forward paths could provide information about the distribution of the outcomes which is difficult to find when only considering three samples.

Delta hedging was chosen as a hedging technique due to its functionality and popularity in the financial markets. Implementing more Greeks, such as gamma, could have helped explaining more of the underlying price movements, and perhaps helped creating a better functioning hedge.

8.1 Further Research

The scope limitations naturally lead to considering opportunities for further research. The spread between Brent and WTI futures itself is interesting to investigate, especially after the discrepancies between the two prices that evolved in 2010. Spread options on the two contracts can clearly be highly profitable, but the high volatility makes such trading to an activity for the risk seeking investor.

For further research, it would be interesting to examine the distribution of outcomes for the forward paths, to provide a more precise explanation of the causes of the inadequate hedging strategy. Further investigations to the spread could also be performed with different assumptions; more focus on transportation costs, different interest rates, more historical data as input, and different structure on the technical risk management. Including the gamma in the risk management technique would be a direct continuing of this thesis.

Especially the transportation costs, which are set constant in this thesis, would be interesting to consider. The uncertainty in international freight rates and the pipeline structure in the U.S., and how this impacts the price spread between Brent and WTI, could be an interesting approach to capture the dynamics of the price spread.

Other types of spread options could also be interesting to examine further, especially for WTI, since this crude oil is experiencing changes from its historical behaviour. For the last observation of the forward curves for Brent and WTI, the Brent forward curve was in normal backwardation, while the WTI had a small contango in the front months of the curve, and backwardation in the last half of the curve. Both spread options on different maturities for WTI, and spread options on crude WTI and WTI products would be interesting for further research, by looking more into convenience yields and the storage structure for WTI.

9 Summary and Conclusion

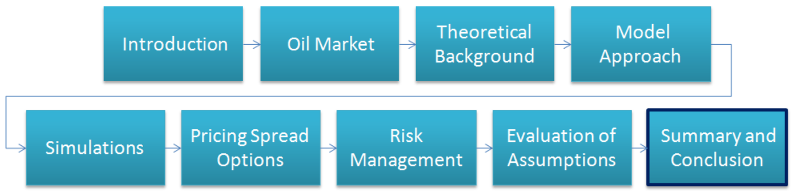

This thesis has investigated the spread between prices of Brent and West Texas Intermediate oil futures. Although the price benchmarks share common features and characteristics, a discrepancy in the prices between the two has developed since 2010. Historical data from futures prices on Brent and WTI from 2008 to 2013 was used as input in a multi-factor forward curve model based on Clewlow and Strickland (2000). The data was calibrated and fitted to the model based on a technique developed by Sollie (2013); using an EGARCH approach as well as principal component analysis to allow for non-constant volatility in the evolution of prices.

The calibrated model was used to draw realisations of possible outcomes for the future prices of Brent and WTI in a period going 50 months forward. Three samples were selected and presented to illustrate future prices on which spread options were simulated and priced. The simulations were performed with the Monte Carlo method, also returning delta values for each day in the three samples.

The option prices for the three samples had large differences; Sample 1 was highly fluctuating, both above and below expected prices, Sample 2 had in general lower prices than expected, and Sample 3 had extremely high option prices during the last half of the set. Interpreting the outcome of the option prices, it is evident that trading and betting on spread options on futures prices of Brent and WTI is "risky business", where it is pure chance whether or not an investor gains or looses from the option exposure.

To risk manage the option exposure, delta hedging was performed with daily rebalancing for the three samples. The total payoff from the hedging gave poor results in Sample 1. During the first half of the set, the replicated portfolio managed to offset any changes to the option values, but halfway through the set the underlying variable changed too fast, making the delta unable to capture the fluctuations. Since all contracts in any sample are highly correlated, the inability of the delta to be adjusted appropriately halfway through the set, influenced the remaining contracts, and lead to a negative total payoff from the hedged position for the last 20 contracts. The risk management worked fine in Sample 2, however since this was a set without much movement, the hedging strategy did not have as many "challenges" as in the other sets. In Sample 3, the results of the hedge stayed symmetrical around zero, but the variations were high. This was explained by the dollar effect, since the relative local variance in the set proved to be equal throughout the 50 contracts.

To sum up the results of the risk management, it is evident that hedging options derived from two such volatile assets as Brent and WTI crude oil futures is difficult. The fact that the hedge worked satisfactory for Sample 2 may be explained by the fact that this set had calmer movements in the underlying variable, thus perhaps being easier to hedge. When discovering that rebalancing the delta hedge once per day is not sufficient to capture the volatility in the underlying asset, it will be necessary to update the hedge several times per day, which will be practically almost impossible and more expensive than the potential benefits from hedging.

To increase the efficiency of the hedging strategy, including more years than five as input to the model would decrease the level of volatility of the outcomes, since the time period from 2008-2013 was extremely volatile. The volatility in the recent past proves that adverse outcomes can occur, and the findings in this thesis show that these outcomes are difficult to hedge. Perhaps introducing gamma to the hedging strategy would increase the efficiency of the hedge in general, by ensuring a hedge more efficient over a wider range of underlying price movements.

When including real market conditions, the cost of hedging occurs by both transaction costs as well as the bid-ask spread working against the holder of the options, making the hedging strategy expensive.

References

Adams, R. (2003). *Calculus - A Complete Course.* Number 5. Department of Mathematics University of British Columbia: Addison Wesley Longman, Pearson Education Canada Inc.

Aleksander, C. (2008). *Market Risk Analysis*, Volume 2. London: Wiley and Sons.

Bibbone, E., G. Pafilo, and P. Tavella (2008). The ornstein uhlenbeck process as a model of a low pass fltered white noise. *Metrologia* (45), 117–127.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroscedasticity. *Journal of Econometrics 31*(3), 307–327.

Bollerslev, T., R. Chou, and K. Kroner (1992). Arch modelling in finance - a review of the theory and empirical evidence. *Journal of Econometrics 52*, 5–59.

Bollerslev, T., R. Engle, and D. Nelson (1994). Arch models. *Handbook of Econometrics*, 2961–2984.

Bruce, R. (2009). Making markets. *energyrisk.com*.

Clewlow, L. and C. Strickland (2000). *Energy Derivatives: Pricing and Risk Management.* London, England: Lacima Publications, London.

Debreu, G. (1959). *Theory of Value; An Axiomatic Analysis of Economic Equilibrium.*

Devarajan, S. and A. Fisher (1981). Hotelling's "economics of exhaustible resources: 50 years later. *Journal of Economic Litterature 19*(1), 65–73.

Engle, R. (1982). Autoregressive conditional heteroscedacticity with estimates of the variance of united kingdom inflations. *Econometrica 50*(4), 987–1007.

Hamilton, J. (2009). Understanding crude oil prices. *The Energy Journal*, 179–206.

Herrmann, L., E. Dunphy, and C. Jonathan (2010). *Oil and Gas for beginners: A Guide to the Oil and Gas Industry.* Deutsche Bank AG/London.

Hotelling, H. (1931, April). The economics of exhaustible resources. *The Journal of Political Economy, The University of Chicago Press 39*(number), 137–175.

Hull, J. (2012). *Options, Futures, and other Derivatives*, Volume volume. Pearson Education Limited 2012.

Kaldor, N. (1939). Speculation and economic stability. *The Review of Economic Studies 7*, 1–27.

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review 99*(3).

McConnell, C. and S. Brue (2008). *Macroeconomics - Principles, Problems and Policies.* Number 7. McGraw-Hill Irwin International Edition.

Milonas, N. and T. Henker (2002). Price spread and convenience yield behaviour in the international oil market. *Applied Financial Economics* (11), 23–36.

Nelson, D. (1991). Conditional heteroscedasticity in asset returns: A new approach. *Econometrica 59*(2), 347–370.

Pindyck, R. (2001). The dynamics of commodity spot and futures markets. *The Energy Journal 22*(3).

Schwartz, E. and J. E. Smith (2000). Short term variations and long term dynamics in commodity prices. *Management Science 46*(7), 893–911.

Sollie, J. (2013). Modelling of commodity prices - dynamics and risk premiums, doctoral thesis, working paper.

Thaler, R. (1980). Towards a positive theory of consumer choice. *Journal of Economic Behavior and Organization* (1), 39–60.

Tsay, R. (2005). *Analysis of Financial Time Series.* Hoboken, New Jersey: John Wiley and Sons.

Walpole, R., R. Myers, S. Myers, and K. Ye (2012). *Probability and Statistics - For Engineers and Scientists.* Pearson Education Inc.

Working, H. (1933). Survey of the wheat situation april to august 1933. *Wheat Studies of the Food Reserch Institute*.

Working, H. (1934). Price relations between may and new-crop wheat futures at chicago since 1885. *Wheat Studies of the Food Reserch Institute 10*(5).

10 Appendix A

Table 2: Factors from EGARCH, first 25 Brent Contracts

Contract	Factor 1	Factor 2	Factor 3	Factor 4
1	-0.18	0.10	-0.66	0.99
2	-0.19	0.11	-0.64	0.99
3	-0.20	0.11	-0.64	0.99
4	-0.20	0.11	-0.64	0.98
5	-0.21	0.11	-0.65	0.98
6	-0.21	0.11	-0.65	0.98
7	-0.21	0.11	-0.64	0.98
8	-0.22	0.11	-0.64	0.98
9	-0.22	0.11	-0.64	0.98
10	-0.23	0.11	-0.63	0.98
11	-0.23	0.11	-0.63	0.98
12	-0.24	0.12	-0.62	0.98
13	-0.24	0.12	-0.61	0.98
14	-0.24	0.12	-0.60	0.98
15	-0.24	0.12	-0.60	0.98
16	-0.24	0.12	-0.59	0.98
17	-0.24	0.12	-0.59	0.98
18	-0.25	0.12	-0.58	0.98
19	-0.25	0.12	-0.57	0.98
20	-0.25	0.12	-0.57	0.98
21	-0.25	0.12	-0.56	0.98
22	-0.25	0.12	-0.55	0.98
23	-0.25	0.12	-0.54	0.98
24	-0.25	0.12	-0.53	0.98
25	-0.25	0.12	-0.52	0.98

Table 3: Factors from EGARCH, last 25 Brent Contracts

Contract	Factor 1	Factor 2	Factor 3	Factor 4
26	-0.24	0.12	-0.51	0.98
27	-0.24	0.12	-0.51	0.98
28	-0.24	0.12	-0.50	0.98
29	-0.24	0.12	-0.49	0.98
30	-0.23	0.12	-0.48	0.98
31	-0.23	0.12	-0.47	0.98
32	-0.23	0.12	-0.46	0.98
33	-0.23	0.12	-0.45	0.98
34	-0.23	0.12	-0.44	0.98
35	-0.23	0.12	-0.43	0.98
36	-0.22	0.12	-0.43	0.98
37	-0.22	0.12	-0.42	0.98
38	-0.22	0.12	-0.41	0.98
39	-0.22	0.12	-0.40	0.98
40	-0.22	0.12	-0.38	0.98
41	-0.22	0.12	-0.38	0.99
42	-0.21	0.12	-0.37	0.99
43	-0.21	0.12	-0.36	0.99
44	-0.21	0.12	-0.35	0.99
45	-0.20	0.12	-0.35	0.99
46	-0.20	0.12	-0.34	0.99
47	-0.20	0.12	-0.34	0.99
48	-0.20	0.12	-0.34	0.99
49	-0.20	0.12	-0.33	0.99
50	-0.21	0.12	-0.32	0.99

Table 4: Factors from EGARCH, first 25 WTI Contracts

Contract	Factor 1	Factor 2	Factor 3	Factor 4
1	-0.21	0.13	-0.68	0.99
2	-0.17	0.10	-0.75	0.99
3	-0.20	0.11	-0.71	0.99
4	-0.23	0.12	-0.69	0.98
5	-0.25	0.13	-0.67	0.98
6	-0.27	0.14	-0.66	0.98
7	-0.29	0.14	-0.65	0.98
8	-0.30	0.15	-0.65	0.98
9	-0.31	0.15	-0.64	0.98
10	-0.32	0.15	-0.64	0.97
11	-0.33	0.15	-0.63	0.97
12	-0.34	0.16	-0.62	0.97
13	-0.34	0.16	-0.61	0.97
14	-0.35	0.16	-0.60	0.97
15	-0.36	0.16	-0.59	0.97
16	-0.36	0.16	-0.58	0.97
17	-0.36	0.16	-0.57	0.97
18	-0.36	0.16	-0.56	0.97
19	-0.37	0.16	-0.55	0.97
20	-0.37	0.16	-0.54	0.97
21	-0.37	0.16	-0.53	0.97
22	-0.37	0.16	-0.52	0.97
23	-0.36	0.16	-0.52	0.97
24	-0.36	0.16	-0.51	0.97
25	-0.36	0.16	-0.50	0.97

Table 5: Factors from EGARCH, last 25 WTI Contracts

Contract	Factor 1	Factor 2	Factor 3	Factor 4
26	-0.36	0.16	-0.49	0.97
27	-0.36	0.17	-0.48	0.97
28	-0.36	0.17	-0.47	0.97
29	-0.35	0.17	-0.46	0.97
30	-0.35	0.16	-0.45	0.97
31	-0.34	0.16	-0.45	0.97
32	-0.34	0.16	-0.44	0.97
33	-0.34	0.16	-0.43	0.97
34	-0.33	0.16	-0.43	0.98
35	-0.32	0.16	-0.42	0.98
36	-0.32	0.16	-0.42	0.98
37	-0.31	0.15	-0.41	0.98
38	-0.31	0.15	-0.40	0.98
39	-0.30	0.15	-0.40	0.98
40	-0.30	0.15	-0.39	0.98
41	-0.29	0.15	-0.38	0.98
42	-0.29	0.15	-0.37	0.98
43	-0.29	0.15	-0.37	0.98
44	-0.28	0.15	-0.36	0.98
45	-0.28	0.15	-0.36	0.98
46	-0.28	0.15	-0.35	0.98
47	-0.27	0.15	-0.34	0.98
48	-0.27	0.15	-0.34	0.98
49	-0.27	0.15	-0.33	0.98
50	-0.27	0.15	-0.34	0.98

11 Appendix B

Table 6: Confidence levels option prices

Contract	Mean - 2 x σ	Mean	Mean + 2 x σ
1	1.520	1.551	1.582
2	1.978	2.020	2.061
3	2.395	2.446	2.496
4	2.735	2.792	2.848
5	3.015	3.078	3.140
6	3.254	3.321	3.388
7	3.455	3.527	3.599
8	3.624	3.701	3.777
9	3.770	3.851	3.931
10	3.888	3.971	4.054
11	4.004	4.091	4.178
12	4.132	4.223	4.313
13	4.233	4.328	4.422
14	4.350	4.448	4.545
15	4.446	4.548	4.650
16	4.521	4.625	4.728
17	4.633	4.740	4.848
18	4.702	4.812	4.922
19	4.779	4.892	5.005
20	4.917	5.035	5.153
21	4.935	5.054	5.173
22	5.024	5.147	5.269
23	5.090	5.215	5.339
24	5.166	5.294	5.241
25	5.254	5.384	5.515

Table 7: Confidence levels option prices

Contract	Mean - 2 x σ	Mean	Mean + 2 x σ
26	5.343	5.477	5.612
27	5.426	5.564	5.702
28	5.492	5.634	5.776
29	5.563	5.708	5.853
30	5.636	5.783	5.931
31	5.681	5.829	5.978
32	5.761	5.913	6.066
33	5.819	5.973	6.126
34	5.881	6.037	6.193
35	5.937	6.096	6.254
36	6.006	6.167	6.327
37	6.013	6.176	6.339
38	6.044	6.208	6.372
39	6.066	6.232	6.399
40	6.087	6.255	6.422
41	6.119	6.288	6.457
42	6.149	6.318	6.487
43	6.200	6.371	6.543
44	6.225	6.398	6.571
45	6.240	6.415	6.590
46	6.256	6.431	6.606
47	6.277	6.453	6.630
48	6.301	6.478	6.655
49	6.302	6.479	6.656
50	6.300	6.476	6.652

12 Appendix C

Table 8: Confidence levels delta values

Contract	Mean - 2 x σ	Mean	Mean + 2 x σ
1	-0.524	-0.517	-0.51
2	-0.525	-0.518	-0.511
3	-0.53	-0.522	-0.515
4	-0.532	-0.525	-0.517
5	-0.534	-0.526	-0.518
6	-0.533	-0.526	-0.518
7	-0.532	-0.525	-0.517
8	-0.531	-0.523	-0.515
9	-0.529	-0.522	-0.514
10	-0.528	-0.52	-0.512
11	-0.527	-0.519	-0.511
12	-0.526	-0.518	-0.51
13	-0.525	-0.517	-0.509
14	-0.524	-0.515	-0.507
15	-0.522	-0.514	-0.506
16	-0.521	-0.513	-0.505
17	-0.52	-0.512	-0.504
18	-0.519	-0.511	-0.503
19	-0.519	-0.511	-0.503
20	-0.519	-0.511	-0.503
21	-0.519	-0.51	-0.502
22	-0.519	-0.51	-0.502
23	-0.519	-0.51	-0.502
24	-0.519	-0.51	-0.501
25	-0.519	-0.51	-0.501

Table 9: Confidence levels delta values

Contract	Mean - 2 x σ	Mean	Mean + 2 x σ
26	-0.519	-0.51	-0.502
27	-0.519	-0.51	-0.501
28	-0.518	-0.509	-0.5
29	-0.518	-0.509	-0.5
30	-0.517	-0.508	-0.499
31	-0.517	-0.508	-0.499
32	-0.517	-0.508	-0.499
33	-0.517	-0.507	-0.498
34	-0.517	-0.508	-0.499
35	-0.517	-0.508	-0.498
36	-0.517	-0.507	-0.498
37	-0.517	-0.507	-0.498
38	-0.517	-0.507	-0.498
39	-0.517	-0.507	-0.498
40	-0.516	-0.507	-0.497
41	-0.516	-0.507	-0.497
42	-0.516	-0.506	-0.496
43	-0.516	-0.506	-0.496
44	-0.515	-0.505	-0.495
45	-0.515	-0.505	-0.495
46	-0.514	-0.504	-0.495
47	-0.514	-0.504	-0.494
48	-0.514	-0.504	-0.494
49	-0.513	-0.503	-0.494
50	-0.514	-0.504	-0.494