University of Stavanger

Faculty of Science and Technology

MASTER'S THESIS

Study program/ Specialization: Industrial economics; contract management and material technology	Spring semester, 2012 Open
Writer: Frithjof Vassbø	……………………………………… (Writer's signature)

Faculty supervisor: Roy Endré Dahl

External supervisor(s): Johan Magne Sollie (Statoil ASA)

Title of thesis:

A Stochastic Model for Correlated Commodity Prices

Credits (ECTS): 30

Key words:

- Stochastic price models
- Simulating correlated price development
- Oil futures trading

Pages: 105

+ enclosure: 0

Stavanger, June 13, 2012

A Stochastic Model for Correlated Commodity Prices

Frithjof Vassbø

University of Stavanger
Department of Industrial Economics, Risk Management and Planning

Abstract

Stochastic models of commodity prices play an integral role in the risk management of companies exposed to commodity price risk. By applying price models, one can obtain expected values for the future prices of the commodity, and also a measure of the uncertainty related to the future price. These figures are crucial for risk management, for example in assessing the need for price hedging.

In this thesis, we propose a model for the price development of two correlated products. The model can be used for forecasting future prices for two correlated products simultaneously, and hence it also allows us to simulate the price spread between the products. The model can be a useful tool for companies seeking to hedge price spread risk, or for investors seeking to speculate on the price spread. Providing a real-life example from the oil market, we will use genuine data from Brent and WTI futures trading.

This thesis utilizes the Schwartz and Smith (2000) model as a basis for developing the model for two correlated products. Also, a three-factor model is proposed in order to describe observed price data more precisely.

Preface

This thesis marks the finalization of my Master of Science program in Industrial Economics at the University of Stavanger (UiS) with Civil Engineering and Contract management as specializations.

My master thesis has been written at Statoil ASA, department of Crude oil, liquids and products (CLP). I wish to thank Lars Dymbe for giving me the opportunity of writing the thesis within his group. It has been interesting for me to observe the daily work at Statoil's CLP Risk Management.

I would also like to thank my supervisor at Statoil, Johan Magne Sollie, for his great support and encouragement during the work with this thesis. His professional expertise within the field of stochastic price models and risk management has been essential for me to learn the theoretical basis for this thesis.

Roy Endré Dahl, my supervisor at UiS, has been a good support and motivator, helping me with approaching the comprehensive work of a master thesis and also showing great interest in my work. I would like to thank him for the time and effort he has spent helping me with this thesis.

Finally, I would like to thank co-student Svein Grude for inspiring me to work diligently through the entire master program. I am glad to have benefited from his enthusiasm, deep professional skill and friendship during the years at UiS.

Stavanger, June 12, 2012.

Frithjof Vassbø

Contents

Chapter 1 Introduction

Before proceeding to the thesis itself, the motivation behind the research performed is presented. In this introductory chapter we will also describe the structure of the thesis, and declare the scope of subsequent chapters.

1.1 Scope of the Thesis

This thesis will describe the price development of Brent and WTI futures using stochastic models rooted in the Schwartz and Smith (2000) model. Stochastic models of commodity prices play an integral role in the risk management of companies exposed to commodity price risk. By applying price models, one can obtain expected values for the future prices of the commodity, and also a measure of the uncertainty related to the future price. These figures are crucial for risk management, for example in assessing the need for price hedging.

Several models for the development of commodity prices exist. Due to its intuitive appeal and simplicity, the model presented by Schwartz and Smith (2000) was chosen as a basis for this thesis. The Schwartz-Smith model can be used for modeling the stochastic development of futures prices for commodities, such as crude oil. However, it can only model the price development of one commodity at a time.

In this thesis, we propose a model for the price development of two correlated products. The model can be used for forecasting future prices for two correlated products simultaneously, and hence it also allows us to simulate the price spread between the products. The model can be a useful tool for companies seeking to hedge price spread risk, or for investors seeking to speculate on the price spread. Providing a real-life example from the oil market, we will use genuine data from Brent and WTI futures trading.

The simplicity of the Schwartz-Smith model makes it unable to capture all the variations in futures price data. In order to describe the observed market development more precise, and thereby get a better starting point for price forecasting, we propose a three-factor model which outperforms the Schwartz-Smith model in accurately reproducing market prices. Another main advantage of the three-factor model is that we are able to explain deviations between Brent and WTI prices by deviations in just one of the factors. The ability to assign deviations between two products to just one factor is an appealing feature of the three-factor model.

Further, the thesis includes an extensive description of the Schwartz-Smith model; its features, underlying assumptions and limitations. We provide detailed descriptions on a spreadsheet procedure for calibrating the model, and how to perform simulations using the model. Microsoft Excel and macros written in Visual Basic for Applications (VBA) provide the platform for all data processing operations referred to in this thesis.

A re-estimation of the three-factor model's parameters is performed and described. Re-estimation is done in order to get an uncertainty measure for the parameter estimates.

1.2 Overview of Thesis

When structuring the thesis, the aim has been to always provide the reader with enough information to understand the next step. This section provides an overview of the report's content.

As this thesis develops and compares stochastic models for oil price, we start with an introduction to what the oil price really is and how it is determined. We will also present the crude oil benchmarks Brent and WTI. This is the scope of chapter two.

In chapter three, we will give an introduction to the concept of futures trading, including also the forward contract in addition to futures contracts. The forward curve and its determinants will be explained, and lastly we will discuss modeling of futures price development.

The Schwartz-Smith model is given a formal presentation in chapter four, where we also explain theory underlying the building blocks of the model. This includes concepts such as Brownian motions and risk-neutral processes.

Chapter five provides a spreadsheet procedure for how to calibrate the Schwartz-Smith model, and discusses problems related to estimation of market risk premiums. At the end of the chapter a overview of the variables and parameters of the model is given, along with interpretations.

Results from calibrating the Schwartz-Smith model is presented and interpreted in chapter six. We compare the present Brent and WTI market configurations, and also provide a historical comparison with data from twenty years ago used by Schwartz and Smith.

Chapter seven explains how to use the Schwartz-Smith model for simulation, and introduces the required method for making correlated draws from the standard normal distribution.

The model for price development of two correlated products is proposed in chapter eight. We give a formal description of the model, and describe how it can be calibrated and used for simulating joint outcomes for future price development of Brent and WTI. We identify a possible application for the model by a contract for future price spreads between Brent and WTI.

In chapter nine we propose the three-factor model. Also for this model, we show how it can be calibrated and used for simulations. The model parameters are also re-estimated in order to investigate the reliability of parameter estimates. We compare the three-factor model to the two-factor model proposed in chapter eight, and discuss the underlying assumptions. As the forecasting horizon becomes very long, the models give unrealistic values for the spread between Brent and WTI. This phenomenon is discussed, and a possible solution is discussed.

Conclusions are drawn in chapter ten. Some thoughts about further work are also noted.

There are no appendices to this thesis. All relevant sources of information are listed in the bibliography. Excel files and VBA source code used for calibrations and simulations can be found on the CD which, together with the printed version of this report, is handed in to the institute administration at UiS. A PDF version of the report itself is also found on the CD.

Chapter 2 The Oil Price

This thesis is focused on modeling the development of oil prices. In order to grasp what we are dealing with, we need to clarify what we mean when referring to the oil price.

When speaking of oil prices we mean the price of one barrel[1] of crude oil. Crude oil is naturally occurring in reservoirs beneath the earth's surface, from where it is extracted. The crude has distinct characteristics, such as density, viscosity and chemical composition, according to what reservoir it has been produced from. These characteristics determine the usefulness of the crude for refining purposes; hence the crudes will be given different prices reflecting their quality. Therefore, a wide variety of oil prices appear side by side. In order for the phrase "oil price" to be precise, we have to specify which crude we are referring to.

In what follows, we will give a brief introduction to how these different oil prices are determined. We will also see that there is both a physical and a financial layer surrounding benchmark crudes. Trading in both the physical and financial layers is used by Price Discovery Agencies (PRAs) in order to assess the market price of oil.

The Brent and WTI futures contracts are examples of derivatives which have emerged in the financial layers. These contracts will play an integral part in the rest of this thesis, thus a presentation of these contracts will be given.

The Brent and WTI futures are founded on the price of Brent and WTI benchmark crudes. In the concluding sections of this chapter, we will briefly describe the characteristics of these crudes. We will also touch on the price relationship between these two benchmarks.

2.1 Who Determines the Oil Price?

The current *"main method for pricing crude oil in international trade"*[2] is the so-called *market-related pricing system*. This has been the prevalent pricing system since the late 1980's. The current system is the successor of past systems where oil prices were determined by oligopolistic price makers. Up until the late 1950s, this role was played by large multinational oil companies called "the Seven Sisters". Then power was shifted to the OPEC[3] countries, through the nationalizing of oil production. Finally, as new producing countries and oil companies have entered the pitch, the traditional price makers lost their market shares and thereby their grip on the oil prices. Since then, oil prices have been determined by "the market".

The crude market has both a physical and a pure financial layer. The physical layer is the one where buyers actually buy crude oil. The financial layer (the paper markets) is where you trade derivatives which settle according to the price of physical oil. When trading in the financial layers you can buy and

---

[1] 1 barrel ≈ 159.0 liters.
[2] See Fattouh (p. 20) which also gives a more thorough description of the events leading to the emergence of the market-related pricing system.
[3] Organization of the Petroleum Exporting Countries. An organization currently consisting of 12 member countries from the Middle East, Africa and Latin-America.

sell contracts to speculate on the oil price development, without actually receiving physical crude oil. There is an intricate web of financial instruments keeping all of this together, and linking the financial layer to the physical crude oil world.

Due to differences in crude oil quality, the crudes have different prices. Some crudes have been chosen as *benchmark* crudes, to which the price of other crudes relate. Several benchmarks exist, with the most famous being Brent (North Sea), WTI (U.S.) and Dubai/Oman (the benchmarks represent different geographical regions). Prices of other crudes are set at a differential to the benchmarks. These differentials are adjusted according to changes in supply and demand for the various crudes.

The benchmark prices are *reported* prices, meaning someone has to determine the value of the benchmark crude. This is done by pricing reporting agencies (PRAs), the two most important being Platts and Argus. The reported prices of the benchmarks play a significant role in the market, and are for instance used *"by oil companies and traders to price cargoes under long-term contracts or in spot market transactions; by futures exchanges for the settlement of their financial contracts; by banks and companies for the settlement of derivative instruments such as swap contracts; and by governments for taxation purposes"*(Fattouh, p. 7). As the benchmark value is very important in determining the revenues of the participants in the oil market, the role of the reporting agencies has to be regarded as crucial. The trustworthiness of the price reporting system heavily relies on the independency and integrity of the PRAs.

The PRAs use sophisticated techniques in order to assess the current market price. An important part of the PRAs' assessments are deals concerning physical delivery of oil, concluded between market participants operating outside the exchanges (over-the-counter). The deals are not revealed to the public, but some market participants allow PRAs to use their deals in assessing the market's state. Of course, a PRA will try to get a sample of deals as large as possible when assessing the market price.

In addition to over-the-counter deals, information from the trade on exchanges or other sources such as "market talk" are used for assessing the market price. Trades in financial derivatives are utilized for assessing the price of physical crude oil through the links between the financial and the physical layers. According to Fattouh (p. 51), *"identifying the oil price relies heavily on information derived from the financial layers"*.

Price differentials between a benchmark and various crudes may also be assessed by a PRA. Some oil-exporting countries choose to set these differentials themselves (some countries also choose not to use benchmarks, but set their own official selling prices).

In determining price differentials, one has to consider differences in quality, supply/demand situation for the relevant crude in the particular market in which it shall be sold, and transportation costs.

2.2 Brent and WTI Futures

In this thesis, we will work with oil prices from trade in ICE Brent Futures and Light Sweet Crude Oil (WTI) Futures. The following presentation will therefore narrow in on these two products. An introduction to the concept of futures is given in Chapter 3.

2.2.1 Brent Futures

Brent futures are traded at the InterContinental Exchange (ICE) in London. In 2010, the daily trade exceeded 400,000 contracts, which equals more than five times the volume of global oil production. The Brent crude future is a cash-settled contract, meaning you don't receive physical oil at maturity; instead you receive the monetary value of the contracts you have bought. Each contract has a size of 1,000 barrels, meaning you can only trade in multiples of 1,000 barrels. At expiry of the contract, the value of the contract is determined according to the ICE Brent Index. It is possible to exchange the cash-settled futures contract for a physical delivery through the EFP (Exchange Futures for Physical) mechanism.

The ICE Brent Index is calculated based on observations from the 21 day BFOE market in the relevant delivery month. The 21 day BFOE[4] market is an over-the-counter forward market, where you buy physical oil for delivery in a specified month. So, in determining the price of the Brent Index, reports from the over-the-counter market of Brent is needed, and from these the Brent Index is derived. Since the 21 day BFOE is a market for physical oil, the Brent Index and therefore also the Brent futures are anchored in the price of physical Brent crude. As the futures contracts approach maturity, prices will have to converge to the Brent Index.

2.2.2 WTI Futures

The WTI futures are traded at the New York Mercantile Exchange (NYMEX). They are even more popular than the Brent futures, with an average daily trade of more than 475,000 WTI futures contracts (2010). In contrast to the Brent futures, WTI futures have physical settlement. Place of delivery is Cushing, Oklahoma. Thus, if not special action is taken before the contract expires (i.e. selling it to someone else), you will have to pick up 1,000 barrels of crude oil at Cushing, Oklahoma. However, only a small fraction of the traded volume is physically settled.

The fact that WTI is physically settled means that the price of the futures contract at expiration has to converge to the spot price of physical WTI crude.

2.3 The Brent/WTI Spread

Both Brent and WTI are light crudes, meaning they have low densities. This makes them easy to refine. Both of them also have low sulphur content, making them *sweet* crudes. This makes them attractive, since sulphur is considered a pollutant and needs being removed during refining. WTI is sweeter than Brent, which is the reason why Brent traditionally has been traded at a $1 to $2 discount to WTI[5].

The similarity in physical characteristics is the reason why the prices normally lie close to each other. However, there are significant differences between the benchmarks regarding logistical aspects. At certain periods of time, these logistical differences result in significant divergency in prices of Brent and WTI.

---

[4] The exact date of delivery isn't decided when the parties enter the contract. The name "21 day BFOE" comes from the fact that the seller has to notice the buyer about the date of loading at least 21 days in advance. BFOE is just an abbreviation for Brent Blend, Fortier, Oseberg and Ekofisk which are the different crudes incorporated in the Brent benchmark.

[5] See for example (Gue, 2011). The WTI-Brent spread is graphed for years 1995-2010 on p. 60 of Fattouh (2011).

While Brent is waterborne crude, meaning it is transported via ships; WTI is transported in and out of Cushing via pipelines. As the pipelines have limited capacity, bottleneck effects may occur. While the problem earlier was to get enough oil into Cushing, yielding low supply of WTI and increasing prices, the problem is now reversed (Fattouh, 2011). The transport capacity into Cushing has increased significantly, while the infrastructure out of Cushing isn’t able to cope with the large supplies. Therefore, crude oil inventories in Cushing are growing large, resulting in lower WTI prices. The logistical matters result in WTI prices being dislocated from the global supply/demand situation, leading to great price differentials between WTI and other benchmarks, such as Brent. This is a major concern for WTI trying to maintain its status as one of the leading international crude oil benchmarks.

Chapter 3 Futures Contracts

As this thesis deals with modeling the development of Brent and WTI futures prices, we need to give a proper introduction to the concept of futures trading. Providing the required knowledge of what a futures contract is and how it can be utilized is the aim of this chapter.

In the first section of this chapter, an introduction to the basic principles of the futures contract is given. Succeeding the futures basics is a presentation of the futures' closest relative; namely the forward contract. The forward contract is primarily explained in order to better understand the features of the futures contract. This is followed by discussing the *forward curve*, which is the cross-section of futures prices prevailing at a certain date. The determinants of the forward curve are discussed, first from a purely theoretical perspective and subsequently by looking at the market operators' beliefs about the future. The impact of various market participants is also discussed.

At the end of this chapter we give an introduction to, and motivation for, modeling the development of futures prices. Specifically, we explain the principal assumptions underlying the Schwartz-Smith model utilized in the remainder of this thesis.

3.1 Basic Principles of the Futures Contract

A futures contract is a derivative, meaning its value is derived from some underlying asset (in our context: crude oil). In a futures contract, terms are determined today for a trade that will take place on a future date. This means that quantity and quality of the asset, time and place of delivery, and also the price to be paid is set today. However, you don't pay the contract price until the time of delivery is reached.

Futures can be used both for risk management and price speculation. For example, an oil producing company wanting to protect itself against price fluctuations, can sell a future contract and obtain a predetermined price. This way, the contract works as a hedge. A speculator, on the other hand, might want to buy the future contract in order to make money on it. If the spot price at contract expiry turns out above the contract price, he can sell the crude oil at a profit. If he is unlucky, the spot price ends up below the contract price, and he will incur a loss. From this perspective, futures trading is a gamble where you bet on what the prices of an underlying asset will be in the future.

Futures contracts are standardized and traded on exchanges. The exchange specifies key aspects regarding the contract, such as the character and quantity of the underlying asset, and place, method and time of delivery. The price is the only field left blank, so to speak, and has to be decided by the buyers and sellers interacting via brokers at the exchange.

3.2 Forward Contracts

The fact that futures contracts are traded on exchanges and heavily standardized, is the main difference between the futures contract and its closest relative; namely the *forward contract*. According to McDonald (p. 142) *"futures contracts are essentially exchange-traded forward contracts".* Both are agreements of a future delivery at predetermined terms.

When concluding a forward contract, the parties meet over-the-counter (OTC) rather than at an exchange. This means that the parties are free to negotiate on all aspects of the contract, not constrained by the standardization of the exchange. The negotiation process provides the participants with flexibility, but in return it is time-consuming and increases the contract's complexity. The tailor-made specifications and potential complexity of the forward agreement makes it difficult to find other buyers for it if you want to get out of the deal.

The futures contracts are much easier to get in and out of. Because they are traded at exchanges, and everybody feels safe about the terms and conditions applying, it is easy to find new counterparts for a futures contract. This makes the futures contract a very *liquid* derivative. Since entering and leaving positions is easy, speculators and arbitrageurs are attracted to the market. This further boosts the volume of trade in the derivative.

3.3 The Forward Curve

As future contracts are traded for the various delivery dates specified by the exchange, we get a strip of observed futures prices. The *forward curve* is a plot of these observed prices against the time axis. The forward curve tells you at what price the contract is traded for specific future dates of delivery. Thus, the forward curve reveals the market's expectation about future spot prices. Indeed, according to Gabillon (1995) the futures price can be regarded as the "*forecast for the spot price prevailing at maturity*".[6] However, the same author states that *"many historical studies have shown that the futures price of oil for a given maturity, taken at a given date, is as bad a predictor for the spot price prevailing at maturity as is the spot price of oil taken at the same initial date"*. Even though the forward curve reflects expectations, the course of history often seems to ignore these expectations. Thus, the forecasting power of the forward curve is low.

3.3.1 The Determinants of the Forward Curve – a Theoretical Approach

Trying to explain some of the determinants of the forward curve, we can imagine a situation where we didn't have prices from observed deals, and just had to construct a curve from explanatory parameters. A procedure for doing this is presented by both Gabillon (1995) and McDonald (2006). McDonald ends up with the forward curve being restricted to

$$S_t e^{(r+C_S-C_Y)\Delta t} \leq F_{t+\Delta t,t} \leq S_t e^{(r+C_S)\Delta t} \qquad \textbf{Eq. 1}$$

Here, S_t represents the spot price at time t, which is the time at which the futures contract is entered. Δt is the time interval from t to the expiry time of the contract. r is the riskless interest rate, C_s is the marginal cost of storage of oil, and C_Y is the convenience yield. These terms will be explained more in the following. In our theoretical discussion we will ignore transaction costs which, according to (Gabillon, 1995), in the trading of crude oil are relatively high.

---

[6] However, as argued by McDonald (p. 172) the presence of a risk premium in the determination of futures prices results in futures prices being a biased estimate of the expected spot price. The difference between the observed futures price and the *true* expected spot price occurs due to the risk premium. Risk averse buyers will cause the futures price to be lower than the true expected spot price.

To start with, forget about the storage cost and the convenience yield. Imagine you were selling a forward contract (which in principle is the same as a futures contract) for a purely financial (non-physical) asset, like a stock. How would you price it? To start with, let's consider if you wanted to be paid today, at the time you entered the contract. How would you determine the price of selling your stock in the future? The way to do it is calculating the expected value of the stock at time t+Δt, and then find the present value of this by discounting it at an appropriate rate of return. Since the stock price at time t+Δt is uncertain (has some risk attached to it), we can't use the risk-free rate. In order to compensate for the risk, we have to use a risk-adjusted rate of return, α, which is higher than r. In other words; due to the riskiness of the stock's future value, the expected future value needs to be higher than what we could have obtained by lending our money at the risk-free rate.

In order to find the expected future value of the stock we need to use α, which also can be interpreted as the *expected* rate of return. This is the rate of return required by a risk-averse investor for investing in it, and can be calculated using for example the CAPM[7] model. Thus, the expected future value of the stock becomes

$$E(S_{t+\Delta t,t}) = S_t e^{\alpha \Delta t} \qquad \textbf{Eq. 2}$$

The use of $e^{\alpha\Delta t}$ implies a *continuously compounded*[8] return rate α, which means that the return is calculated and added continuously ("all the time"), instead of only at the end of the year.

The prepaid (paid today) price of the forward contract becomes the present value of this expected value. When calculating the present value, we discount at the risk-adjusted rate α. This yields

$$F^P_{t+\Delta t,t} = e^{-\alpha\Delta t} E(S_{t+\Delta t,t}) = e^{-\alpha\Delta t} S_t e^{\alpha\Delta t} = S_t$$
$$F^P_{t+\Delta t,t} = S_t \qquad \textbf{Eq. 3}$$

where $F^P_{t+\Delta t,t}$ is the price of the prepaid forward. It turns out that if you want payment for the forward contract today, a fair price would be the current stock price. But what if you change your mind, and rather want to receive the contract price in the future (as in a normal forward contract)? Assuming your buyer won't default, there is no risk associated with postponing the payment. Therefore, you can only require the risk-free rate in determining the future price to be paid. Hence, the fair price to sell your stock in a forward contract becomes

$$F_{t+\Delta t,t} = S_t e^{r\Delta t} \qquad \textbf{Eq. 4}$$

It can also be shown that all other prices of the forward contract would allow for *arbitrage*, which is a situation where you can earn money on trading with no net investment of funds and with no risk. If the forward price is higher than implied by Eq. 4, one could borrow money at the risk-free rate to buy the stock today, then sell it in a forward contract and earn the return rate implied by the forward price. This

---

[7] Capital Asset Pricing Model.

[8] See Appendix B of (McDonald, 2006).

return rate is, given that the forward price is too high, higher than the risk-free rate, and the differential will provide a risk-free positive cash flow with no net investment of funds.

In the opposite case, if the forward price is too low, you can *short*[9] the stock and buy a forward contract. At the time of expiry, you use the stock acquired from the forward contract to close the short position. The forward price which you pay is lower than the future value of the money you earned from shorting the stock, meaning you have earned money without assuming any risk and without making any initial investment.

A market which allows arbitrage is out of balance. As arbitrageurs exploit the arbitrage opportunity, prices will be adjusted and finally market prices will reach equilibrium where arbitrage is impossible. For example, arbitrage will increase the demand for an under-priced contract, thus pushing the price up towards its no-arbitrage equilibrium (which, from the discussion so far, is given by Eq. 4)..

We have now covered the foundation of Eq. 1. Let's further consider the situation where cost of storage and convenience yield plays a part. This happens in the commodity trade, where we are dealing with physical goods.

Oil is possible to store, and therefore sellers are faced with the option of either selling the oil today, or storing it for a future sale. This is equal to a so-called cash-and-carry situation, where you simultaneously buy an asset and sell it forward. The cash-and-carry of oil is only reasonable if the present value of the forward sale is at least as great as the price you could sell the asset at today. Now, if there is storage costs associated with holding the asset, these costs will have to be included in the present value calculation. Suppose the future value of the accumulated storage costs at time of expiry is $C_S(t, t + \Delta t)$. Then, in order to make storage reasonable, we get the following expression:

$$F_{t+\Delta t,t} \geq S_t e^{r\Delta t} + C_S(t, t + \Delta t) \qquad \textbf{Eq. 5}$$

Let's further assume storage costs are being paid continuously, and that they can be measured as a certain fraction of the commodity's value. Then, the expression can be written as

$$F_{t+\Delta t,t} = S_t e^{(r+C_S)\Delta t} \qquad \textbf{Eq. 6}$$

This means that the seller is indifferent between selling today and selling forward with storage as long as the forward price satisfies Eq. 6.

If we analyze Eq. 6, we find that it has some implications that don't match with the reality. As both the risk-free rate and the marginal cost of storage are restricted to positive values, this implies that the forward price always will be higher than today's spot price. The situation where forward prices are higher than the current spot price is called *contango*. However, when reviewing observed forward

---

[9] Meaning you borrow a stock from someone and sell it to someone else, while guaranteeing to replace the borrowed stock later.

curves, we find that the market isn't always in contango. *Backwardation*[10], which is the opposite of contango, commonly appears in the crude oil market.

From our discussion so far, backwardation is highly illogical. As holders of physical oil incur the opportunity cost of the risk-free rate and have to pay storage costs, it is hard to understand why they are willing to sell it at a lower future price than what they would have obtained by selling today. At the other side of the table, market participants buying crude oil at spot price instead of at a lower future price also seem to act irrational.

However, we have to assume that there's some kind of logic underlying the behavior of storing oil during backwardation. In the quest for a rational explanation to this phenomenon, the last factor of Eq. 1, namely the *convenience yield*, emerges.

The behavior of keeping oil inventories through backwardation indicates that there is some kind of benefit from holding the physical oil instead of holding a contract for future delivery. This benefit is known as the "convenience yield" which is defined by Brennan (1989) as *"the flow of services which accrues to the owner of a physical inventory but not to the owner of a contract for future delivery".* An example of convenience yield is the necessity for e.g. a refinery to hold physical oil. If he doesn't have physical access to oil, his activity will entirely stop resulting in great losses. The same principle prevails for all market participants who for business reasons have a critical dependency on holding physical oil.

When implementing it into Eq. 6, by regarding the convenience yield as a dividend being continuously paid to the holder of inventories, we get the lower limit of Eq. 1. But why is there also an upper limit neglecting the convenience yield? This can be explained by looking at the situation from the perspective of an average investor, with no specific business reason to hold physical oil. For him, the convenience yield won't make any impact on the value of holding an inventory of oil. Therefore, reasonable forward prices will be in the interval given by Eq. 1:

$$\mathrm{S_t e^{(r+C_S-C_Y)\Delta t}} \leq F_{t+\Delta t,t} \leq \mathrm{S_t e^{(r+C_S)\Delta t}}$$

However, as long as there are oil-dependent businesses active in the trade, performing operations of buying physical oil (in order to maintain a buffer of inventories) and selling it forward, forward prices will be determined by the lower limit of Eq. 1. As this is the case in the oil market, we conclude that the expression for theoretical forward prices given by Gabillon (1995) applies:

$$F_{t+\Delta t,t} = \mathrm{S_t e^{(r+C_S-C_Y)\Delta t}} \qquad \textbf{Eq. 7}$$

If the convenience yield outdoes the risk-free rate and costs of storage, the market will be in backwardation. If there is no or little convenience yield (in times of low demand and stable supply of crude oil), the market will be in contango.

According to Eq. 7, the risk-free rate (financing costs), costs of storage and the convenience yield specify theoretical limits for contangos and backwardation.

---

[10] Backwardation is the situation where future prices are lower than the current spot price.

The greatest contangos occur in market situations where the convenience yield can be neglected. Thus, contangos are limited by financing and storage costs. If future prices become too high, there will be a possibility for "cash-and-carry" arbitrages. People could lock in a risk-free profit by buying oil at the spot price, financing it and storing it, before delivering it at a forward price which more than covers financing and storage costs. However, if a lot of people do this arbitrage, it will increase the demand for immediate delivery of oil and also boost the supply of oil delivered in the future. This will put upward pressure on spot prices and downward pressure on forward prices, easing the contango situation so that in the end prices will reach equilibrium as given by Eq. 7.

The greatest bacwardation situations occur when the convenience yield is dominant. The convenience yield of crude oil can get pretty large in times of low or insecure supply, as the elasticity of demand for petroleum products is close to zero in the short term (Gabillon, 1995). Consumers depending on supply of petroleum products can't switch to using substitutes on short notice, meaning prices can get pretty high without affecting the demand. The dependency on immediate delivery of petroleum prices puts an upward pressure on the prices of physical oil.

The above discussion has given us an expression for the theoretical forward prices of crude oil (or any other storable commodity). The forward prices are the risk-adjusted expected future spot prices (see footnote no. 6). The theoretical discussion assumes that forward prices can be determined when we know the spot price, the risk-free interest rate and also get a measure of storage costs and the convenience yield (the latter being very hard to observe). In other words, we have claimed that expectations about the future spot prices rely solely on today's spot price and information about financing costs, storage costs and the convenience yield.

3.3.2 Other Factors Affecting the Forward Curve

Gabillon (1995) states that the factors mentioned above are the *"essential determinants of forward curves"* (p. 32). However, there are also other, *"less rational"*, beliefs of market participants affecting the forward curve. The presentation in subsequent sections (including this) relies heavily on Gabillon's article.

The explanatory factors of Eq. 7 above are not the only determinants of the forward curve. Based upon historical prices, market operators might have expectations about future price development. Anticipated future supply/demand configurations, and guesses about moves from influential oil producers such as OPEC, impact the market operators' expectations about future oil prices. All these aspects should be included when interpreting the observed forward curve.

Also, an important feature of the market operators' expectations is the assumption that oil prices are *mean reverting*. Mean reversion imply that oil prices subject to relatively large fluctuations eventually will return to some equilibrium level. The market's assumption of mean reversion results in relatively stable long-term futures prices even if the spot price fluctuates.

3.3.3 The Impact of Various Market Participants

The participants in the market impact the forward curve differently, due to the purposes for which they trade on the curve.

The upstream operators (producing companies or countries) hold reserves, and are therefore exposed to price declines. Therefore, they are “short hedgers”, meaning they want to sell futures contracts in order to secure a decent price on their oil. Their supply of future deliveries puts a downward pressure on future prices.

Refiners are typically “long hedgers” for crude oil, meaning they want to protect themselves from price increases in the main input of their business. They are also short hedgers on the forward curves of refined products, and in effect they are short hedgers of their refining margins (just as oil producers are short hedgers of their margins from crude oil sale). As a matter of fact, refiners are primarily concerned about hedging the spread between crude oil and refined products. The absolute level of the forward curves doesn’t affect them much. Gabillon states that *“the net effect of refiners on forward curves is fairly neutral, since the absolute level of prices is not crucial to their economics”*.

Traders and distribution companies, like refiners, are more concerned about their margins (price differentials) than absolute level of forward curves. They operate with relatively thin margins, inducing them to hedge their price risk. Distribution companies benefit from contango situations, since what they do is essentially selling forward products. Therefore, they will try to sell forward contracts in contango situations in order to ensure a good price for their forward sale. As the distribution companies are eager to lock in high future prices, the net effect of their hedging actions is a downward pressure on the forward curve, mainly concentrated on the short-term part of the curve as their operate with relatively short horizons.

Consumers, who are exposed to upward movements of prices, buy forward contracts in order to protect themselves from price increases. They represent a demand for forward contracts, and therefore put an upward pressure on the forward curve. Like producers, consumers may have long horizons (up to several years) on their hedging operations, in order to lock in their oil price during the whole period of a project.

Investors use the forward curve for speculation, and can, for instance, try to make money on backwardation in the market by buying futures contracts and rolling them forward before expiration. *Rolling forward* a contract is done by first buying a contract, for example the 3rd month futures contract, then as it gets close to expiration you sell it and then buy the new 3rd month contract. If the market has been in backwardation all the time, the price of the contract will increase as maturity approaches, and you can sell it at a profit. As the new contract approaches expiry, you perform the operation over again. This strategy will work as long as the market is in backwardation. However, the speculation puts an upward pressure on the forward curve (in our example the 3rd month contracts) which reduces the backwardation of the market.

Arbitrageurs play an important role in discovering risk-free arbitrages or other obviously profitable operations, and by exploiting the arbitrage they finally bring the prices to equilibrium. Arbitrageurs and speculators are also needed in bringing liquidity to the market. Their influence on the forward curves is complex, and whether they put upward or downward pressure on the forward curve is hard to evaluate.

The net result of all market participants on the forward curve depends on which side is most desperate to hedge their risk.

3.3.4 The Short and the Far End of the Forward Curve

The forward curve can broadly speaking be divided into two parts, the first being made up of maturities up to 18 months, the other covering the subsequent maturities. Up to 18 months, the curve is in connection with the physical market and the short-term expectations prevailing. The price is determined by supply/demand relations, level of inventories and the fear of supply disruptions.

On the far end of the curve, the futures market is more linked to financial markets than the market of physical crude oil.

3.3.5 Modeling the Development of the Forward Curve

Eq. 7 provides us with a tool for explaining the effect of some of the major forward curve determinants. Also, if we found a way to describe the time-development of spot prices, we could use Eq. 7 to simulate the future development of spot prices and draw new forward curves. This way we would be able to simulate how forward curves could look in the future.

However, a model founded on Eq. 7 would be far too simplistic, for instance in making the assumption that both the financial and storage costs, in addition to the convenience yield, are constants. This is very unrealistic. Also, as pointed out by Gabillon (1991), the limit value of the futures prices for an infinite maturity would approach zero for backwardation and infinity for contango. This is a clear shortcoming of such a model. As proven by Gabillon, the model also implies that the volatility of futures prices equal the volatility of spot prices, which doesn't reflect reality. Data from the market shows that volatility of futures prices decreases as time to maturity of the contracts increases. Gabillon compares this to the movements of a cantilever subject to forces on its free end. The deflections on the free end will be large, but as you move further away from the free end the deflections will get smaller and smaller until you reach the fixed end where deflections are zero. Analogous to the behavior of the cantilever, future price fluctuations are greatest at the short end and then decrease towards the longer maturities. This effect is closely related to the assumption of mean reversion, which also is ignored by Eq. 7.

The development of more realistic models for the term structure of prices advanced greatly at the end of the 20th century, and many models have been proposed by various researchers. Among others, models assuming stochastic processes for spot and long-term prices (Gabillon, 1991), convenience yield (Gibson & Schwartz, 1990) and interest rates (Schwartz, 1997) have been proposed.

According to Schwartz and Smith (2000), *"Stochastic models of commodity prices play a central role when evaluating commodity-related securities and projects."* By modeling the development of oil prices, we obtain expected future prices and also a measure of the related uncertainty (variance). Business companies dependent on future oil prices can use this information to make well-founded investment decisions assuming the price risk of oil. A measure of future price uncertainty is also crucial in assessing the need for price hedging. This way, stochastic price models are integral in managing oil price risk.

Speculators, on the other hand, can exploit the information from a model in order to make buy and sell decisions. Comparing prevailing market prices to model implied prices, they can search for apparent under- or overpriced contracts from which they can make profits.

3.3.6 The Schwartz-Smith (2000) Model

As a basis for the remainder of this thesis, we will employ the so-called "Short-term/Long-term" model presented in Schwartz and Smith (2000). This model is shown by the authors to be equivalent to the model of stochastic spot prices and convenience yields presented in Schwartz (1997), although convenience yields aren't explicitly referred to in the Schwartz-Smith model. Thus, following from the equivalence to the model of stochastic convenience yields, the Schwartz-Smith model corresponds to the assumption of variable convenience yields. However, also following from the equivalence to this specific model in Schwartz (1997), the risk-free interest rate is assumed to be constant[11]. This simplification limits the complexity of the model.

The Short-term/Long-term model is a so-called two-factor model, meaning the development of oil prices are explained by two variables; one short-term and one long-term variable. The model allows for mean-reversion of the short-term prices towards the long-term prices (equilibrium level), where both the short-term deviations and the long-term equilibrium level develop via stochastic processes. For contracts of far maturities, the long-term variable will be the most influential in determining the price, but for spot prices and short maturities the sum of the short-term deviations and the equilibrium level will determine prices. Hence, for temporary supply disruptions or increases in demand, higher prices in the front end of the forward curve (backwardation) can be explained by a positive short-term variable. The short-term variable will also cover contango situations, for which the short-term variable takes on negative values.

The authors justify the inclusion of mean reversion in the model by declaring it to be *intuitive*. The reasoning goes like this: in times when the price of a commodity is higher than the equilibrium price level, the supply of the commodity will increase because higher cost producers will enter the market. By the increased supply, prices are pushed downwards. In the opposite situation, when prices are low, some high-cost producers will leave the market thus putting upward pressure on prices. As entering and leaving the market takes some time, prices may be temporarily high or low, but will eventually revert toward the equilibrium level.

However, there might be fundamental changes in the market that will not only change the short-term prices, but rather shift the entire forward curve. On the supply side, such changes might be: exhaustion of existing supply; new oil field discoveries; cheaper production methods; increased recovery from existing fields; inflation; and political/regulatory effects.

Long-term changes in demand also influence on the oil price level. According to IEA (2011) oil consumption is expected to grow during the coming years. The main reason for this is increased energy demand from emerging economies such as China, which is anticipated to consume nearly 70% more energy than the US by 2035. Attempts to substitute petroleum products with other (renewable) alternatives might dampen the expected demand growth, but is not expected to prevent a net growth in oil demand.

[11] For the record, a model including stochastic interest rates is also presented in Schwartz (1997).

The Schwartz-Smith model captures fundamental changes such as the above mentioned by shifts in the long-term variable (the equilibrium price level). Shifting the equilibrium price level will affect the entire forward curve.

Chapter 4 Formal Description of the Schwartz-Smith Model

In Chapter 3, we explained some of the principal assumptions underlying the Schwartz-Smith model. Through this chapter, we will give a more thorough mathematical description of the model.

In the first section, we will present the basic equation of the model, declaring the log spot price as the sum of two *state variables*. These are the short-term and long-term variable introduced at the end of Chapter 3.

After presenting the equation for the log spot price, we will proceed with describing the development of the state variables as stochastic processes. The long-term variable has elements of both constant drift and *random walk*, while the short-term variable is assumed to be mean reverting in addition to exhibiting random walk. In order to describe random walk mathematically, we utilize a process called Brownian motions. From the processes for state variable development, we get the expectation and variance of future spot prices.

Before we can draw the forward curves from expected future spot prices, we need to make an adjustment to consider the risk aversion of market participants. This is done by applying the *risk-neutral measure*, which explained and described before proceeding to the expression for futures prices.

Finally, we discuss the volatility curve of futures prices and present the equation for the volatility curve. A brief explanation of the Samuelson effect is given.

4.1 The Log Spot Price Equation

The Schwartz-Smith is a stochastic model describing the development of commodity futures prices. The spot price is the special case of the futures prices where time to maturity is zero. Spot prices are explained by the model as a function of two stochastic variables; the short-term factor χ_t and the long-term factor ξ_t.

Mathematically, we express the *logarithm* of the spot price (S_t) as:

$$s_t = \ln(S_t) = \chi_t + \xi_t \qquad \textbf{Eq. 8}$$

The short-term factor represents the short-term deviations of the oil price, while the long-term factor represents the equilibrium price level which oil prices are assumed to revert to. Basically, the equilibrium level is what the spot price would have been in the absence of short-term deviations.

The short-term and long-term factors are referred to as *state variables* – they are variables expressing which state the oil price is in today. Both the short-term and the long-term variable change from day to day. If the concept of a changing equilibrium price level seems confusing, it might be helpful to refer to it as this day's *implied* equilibrium price level. It is a measure of the assumed equilibrium price level underlying the settlement of today's futures prices.

So, for each new day, or more generally; for each change in time, both the short-term and the long-term variable changes.

4.2 Continuous Time Development of State Variables

The development of the short-term and long-term factors is described as

$$\mathrm{d}\chi_t = -\kappa\chi_t\mathrm{dt} + \sigma_\chi dz_\chi \qquad \textbf{Eq. 9}$$

$$\mathrm{d}\xi_t = \mu_\xi\mathrm{dt} + \sigma_\xi dz_\xi \qquad \textbf{Eq. 10}$$

The processes for the development of χ_t and ξ_t are correlated through the relation

$$\mathrm{dz}_\chi\mathrm{dz}_\xi = \rho_{\chi\xi}dt \qquad \textbf{Eq. 11}$$

where dz_χ and dz_ξ are the correlated increments of standard Brownian motion processes. The inclusion of a correlation factor to the development of the two variables means that they can't develop independently. Depending on the sign of the correlation factor, the state variables will develop in phase or out of phase. The absolute value of the correlation factor tells us how pronounced this in-phase/out-of-phase relation is.

The differential equation for the short-term factor is a so-called Ornstein-Uhlenbeck process, while the long-term factor follows pure arithmetic Brownian motion.

4.2.1 Brownian Motions

The term Brownian motion needs further explanation[12]. Brownian motion is random walk in continuous time, with continuous movements. The term continuous just means that there is no downtime or pauses – things happen all the time. If we were to graph Brownian motions, we would have to do it without ever lifting our pencil from the paper. And each time we moved the pencil a tiny interval along the time axis, we would have to do a tiny up or down movement. Indeed, the word tiny isn't good enough for describing how small the time interval must be. The correct term is "*infinitesimally* small". It's what you get when you divide one by infinity.

Imagine a standard XY chart with time along the x-axis and where Y is the value of the Brownian motion process. Let's start at (0, 0). For each infinitesimally small time interval, a new random draw is made. The random draw determines the Y direction of the next movement, and the size of each Y movement is infinitesimally small. So your pencil has to move steady towards right, following the time axis, and *all the time* it has to move either up or down one step on the Y axis.

The Brownian process is a *martingale*, meaning the expectation of the Y movement is zero, thus we always expect the variable to stay at the value it had before the draw (its current position). But zero never occurs; we just have ups and downs. So for the first draw we will for certain end up somewhere just above or just below zero, although we expect it to end up at zero. Let's assume it ended up just above zero. For the second draw, we expect the value to end up at its current value (just above zero). But, as the draw is made, the value has to move; either down (to zero) or another step upwards. So, for each new draw, the value can move either back to where it came from, or even further away from where

[12] A brief presentation of Brownian motion and the Ornstein-Uhlenbeck process is given in McDonald (pp. 649-655).

it came from. Note also that the process doesn't care about whether the previous draw was an up or down. In fact, the result of each new draw is independent of *all* preceding draws.

4.3 Discrete Time Development of the State Variables

The discussion about Brownian motion has so far been limited to continuous time. But, as we are not able to perceive continuous time, we need to give discrete time solutions to the phenomenon of Brownian motion. Namely, we are not looking for the development of a Brownian motion process over an infinitesimal time interval, but over a time interval of, say, one day.

Each draw in the Brownian motion process can be looked upon as a random draw from a binomial distribution, where each move is either +1 or -1 with equal probabilities. The distribution then has expectation 0 and variance 1. If we want to evaluate Brownian motions over a finite time interval, then it will be the sum of an infinite number of random, independent binomial draws. Applying the Central Limit Theorem, the distribution of this sum is the normal distribution. Over a time interval of Δt, the sum of Brownian motion increments will have a $\mathrm{N}(0, \Delta t)$ distribution[13], meaning expectation is zero and variance is Δt.

Discrete time solutions to Eq. 9 and Eq. 10 are needed. Solving for discrete time involves some heavy mathematics; therefore we will just give the solutions here[14]:

$$\chi_t = \chi_{t-\Delta t} e^{-\kappa \Delta t} + \sqrt{\frac{1 - e^{-2\kappa \Delta t}}{2\kappa}} \sigma_\chi z_\chi \qquad \textbf{Eq. 12}$$

$$\xi_t = \xi_{t-\Delta t} + \mu_\xi \Delta t + \sqrt{\Delta t} \sigma_\xi z_\xi \qquad \textbf{Eq. 13}$$

Here, z_χ and z_ξ are correlated draws from the $\mathrm{N}(0,1)$ distribution. We see that the random draw relating to the long-term component, ξ_t, is scaled by $\sqrt{\Delta t}$. This is in line with the above conclusion that the sum of Brownian motion increments will have a $\mathrm{N}(0, \Delta t)$ distribution. In fact, $\mathrm{N}(0, \Delta t)$ is equal to $\sqrt{\Delta t} \times \mathrm{N}(0,1)$. For the short-term component, the mean reversion messes the expression up a bit. But the expression $\sqrt{\frac{1-e^{-2\kappa\Delta t}}{2\kappa}}$ still represents a time scaling, with the volatility increasing as the time interval increases.

4.4 Interpreting the Development of the State Variables

In what follows, we will investigate the features of Eq. 9 and Eq. 10.

$$\mathrm{d}\chi_t = -\kappa \chi_t \mathrm{dt} + \sigma_\chi dz_\chi \qquad \textbf{Eq. 9}$$

In Eq. 9, the first term ($-\kappa \chi_t \mathrm{dt}$) describes mean reversion towards zero. To see this, note that the term can be expressed equivalently as $\kappa(0 - \chi_t)\mathrm{dt}$. (If we replace zero with another number then this will be the level which the process will revert to.) So, if the short-term deviation one day is above zero, then the

[13] We use the notation $\mathrm{N}(\mu, \sigma^2)$.
[14] The relevant derivations are given in the Appendix of Schwartz and Smith (2000).

development the next day will have a drift downwards. Likewise, if the short-term deviation one day is below zero, then the development the next day will have a drift upwards. κ is the rate at which χ_t reverts to zero. A high κ value yields fast reversion. κ must be positive in order to represent mean reversion. According to Schwartz and Smith, the "half-life" of the short-term deviations (the time in which the deviations are expected to halve) can be calculated by –ln(0.5)/κ.

The mean reversion of Eq. 9 is expressed in Eq. 12 as $\chi_{t-\Delta t}e^{-\kappa\Delta t}$. $\chi_{t-\Delta t}$ is the value of the short-term variable one time interval ago. No matter the sign or size this "yesterday value", the term $e^{-\kappa\Delta t}$ will result in a factor between 0 and 1 and provide development of today's χ_t towards zero.

So, will the short-term factor always develop towards zero? The answer is no, and here's where the second term of Eq. 9 comes into play. The second term ($\sigma_\chi dz_\chi$) includes a Brownian motion draw, which can push the short-term factor either upwards or downwards, with equal probabilities. The draw is scaled by σ_χ, the standard deviation of increments of χ_t (see Eq. 35 and Eq. 36).

As already mentioned, in the discrete time solution the binomial Brownian draw has mutated into a draw from the $\mathrm{N}(0,1)$ distribution, scaled by time-scaling factor $\sqrt{\frac{1-e^{-2\kappa\Delta t}}{2\kappa}}$ in addition to σ_χ.

To sum up, the development of the short-term factor has one drift component always dragging the short-term factor towards zero, and one random component which can push it in either direction.

$$\mathrm{d}\xi_t = \mu_\xi \mathrm{dt} + \sigma_\xi dz_\xi \qquad \textbf{Eq. 10}$$

The long-term component also has one drift component and one random component. The drift term $\mu_\xi\mathrm{dt}$ gives a mean, linear rise to the long-term factor at a steepness of μ_ξ (see Eq. 40). Here also, the second term $\sigma_\xi dz_\xi$ includes a Brownian motion draw. The scaling factor σ_ξ is the standard deviation of increments of ξ_t (see Eq. 37 and Eq. 38). So for the long-term factor we have one drift component $\mu_\xi\mathrm{dt}$ providing a linear rise, and one random component $\sigma_\xi dz_\xi$ pushing the long-term factor in either direction. And when moving from continuous to discrete time, the random binomial draw of the Brownian motion turns into a draw from the $\mathrm{N}(0,1)$ distribution, scaled by $\sqrt{\Delta t}$.

4.5 Expectation and Variance of the Log Spot Price

The discrete time solution also gives us the covariance between χ_t and ξ_t given as

$$\mathrm{Cov}(\chi_t, \xi_t) = (1 - e^{-\kappa\Delta t})\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} \qquad \textbf{Eq. 14}$$

where $\rho_{\chi\xi}$ is the correlation between the increments of Brownian motion in Eq. 9 and Eq. 10. This is also the correlation between the random draws in Eq. 12 and Eq. 13, and can hence be estimated as the correlation between the scaled increments of χ_t and ξ_t (see Eq. 7).

From what we have done so far, we can set up joint mean vector and covariance matrix for the two state variables:

$$\mathrm{E}[\chi_t, \xi_t] = [\chi_{t-\Delta t} e^{-\kappa \Delta t},\ \xi_{t-\Delta t} + \mu_\xi \Delta t] \qquad \textbf{Eq. 15}$$

$$\mathrm{Cov}[\chi_t, \xi_t] = \begin{bmatrix} \frac{1 - e^{-2\kappa\Delta t}}{2\kappa} {\sigma_\chi}^2 & (1 - e^{-\kappa\Delta t}) \frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} \\ (1 - e^{-\kappa\Delta t}) \frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} & {\sigma_\xi}^2 \Delta t \end{bmatrix} \qquad \textbf{Eq. 16}$$

We can, with reference to Eq. 12, verify that the variance of χ_t is $\frac{1-e^{-2\kappa\Delta t}}{2\kappa}{\sigma_\chi}^2$ by the relation $\sqrt{\frac{1-e^{-2\kappa\Delta t}}{2\kappa}}\sigma_\chi \times \mathrm{N}(0,1) = \mathrm{N}\left(0,\ \frac{1-e^{-2\kappa\Delta t}}{2\kappa}{\sigma_\chi}^2\right)$. The same logic can be applied to Var(ξ_t).

According to $\ln(S_t) = \chi_t + \xi_t$, we then have that the log spot price $\ln(S_t)$ is a combination of two correlated normally distributed variables, and is thus normally distributed itself with

$$\mathrm{E}[\ln(S_t)] = \chi_{t-\Delta t} e^{-\kappa\Delta t} + \xi_{t-\Delta t} + \mu_\xi \Delta t \qquad \textbf{Eq. 17}$$

$$\mathrm{Var}[\ln(S_t)] = \frac{1 - e^{-2\kappa\Delta t}}{2\kappa}{\sigma_\chi}^2 + {\sigma_\xi}^2\Delta t + 2(1 - e^{-\kappa\Delta t})\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} \qquad \textbf{Eq. 18}$$

As the log spot price is normally distributed, it follows that the spot price is lognormally distributed. From the relation between the normal and the lognormal distribution, it then follows that

$$\mathrm{E}[\mathrm{S_t}] = e^{\mathrm{E}[\ln(S_t)] + \frac{1}{2}\mathrm{Var}[\ln(S_t)]} \qquad \textbf{Eq. 19}$$

We can now solve for $\ln(\mathrm{E}[\mathrm{S_t}])$ which is different[15] from $\mathrm{E}[\ln(S_t)]$. In words: the log of the expected spot price is different from the expected log spot price given in Eq. 17. The log of the expected spot price is

$$\begin{aligned} \ln(\mathrm{E}[\mathrm{S_t}]) &= \mathrm{E}[\ln(S_t)] + \frac{1}{2}\mathrm{Var}[\ln(S_t)] \\ &= \chi_{t-\Delta t}e^{-k\Delta t} + \xi_{t-\Delta t} + \mu_\xi\Delta t + \frac{1}{2}(\frac{1 - e^{-2\kappa\Delta t}}{2\kappa}{\sigma_\chi}^2 + {\sigma_\xi}^2\Delta t \\ &\quad + 2\left(1 - e^{-\kappa\Delta t}\right)\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa}) \end{aligned} \qquad \textbf{Eq. 20}$$

We can use this to find the expected spot price in the far future, as the time interval Δt from today to time t gets very large (approaches infinity). All terms including $e^{-\kappa\Delta t}$ and $e^{-2\kappa\Delta t}$ approaches zero and therefore vanishes. Therefore, the expression for the log of the expected spot price in the far future is $\xi_{t-\Delta t} + \frac{{\sigma_\chi}^2}{4\kappa} + \frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} + (\mu_\xi + \frac{1}{2}{\sigma_\xi}^2)\Delta t$. From this expression, we see that for far futures, the log of the expected spot price will rise according to $(\mu_\xi + \frac{1}{2}{\sigma_\xi}^2)\Delta t$. As a comparison, the log of the expected equilibrium price has the same growth rate and in the long run only differs by $\frac{{\sigma_\chi}^2}{4\kappa} + \frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa}$ (see Eq. 36).

[15] This is an example of *Jensen's inequality* (see Appendix C of McDonald (2006)).

$$\ln\left(\mathrm{E}\left[\mathrm{e}^{\xi_t}\right]\right) = \ln\left(e^{E(\xi_t)+\frac{1}{2}\mathrm{Var}(\xi_t)}\right)$$
$$= \ln\left(e^{\,\xi_{t-\Delta t}+\mu_\xi\Delta t+\frac{1}{2}\sigma_\xi{}^2\Delta t}\right)$$
$$= \xi_{t-\Delta t} + (\mu_\xi + \frac{1}{2}\sigma_\xi{}^2)\Delta t \qquad \textbf{Eq. 21}$$

The expectation of the short-term deviations is for the long run equal to zero.

We can find the variance of the spot price by using equation (18.14) in McDonald (p. 595):

$$\mathrm{Var}(S_t) = e^{2E[\ln(S_t)]+Var[\ln(S_t)]}\left(e^{Var[\ln(S_t)]} - 1\right) \qquad \textbf{Eq. 22}$$

Now that we have developed an expression for the expected spot price (Eq. 19), one might think that valuing the futures contracts is straightforward. A reasonable assumption about the buyer and seller of the forward contract is that they will settle with the expected spot price. However, this doesn't conform to the prices we observe in reality. The reason is that we have neglected an important aspect: the reluctance by market participants to assume price risk.

4.6 Risk-Neutral Processes

In a real world setting most trading participants are risk-averse, meaning they will require some risk premium to do a risky trade. A futures contract guarantees a fixed price, and therefore there isn't any uncertainty regarding the payment at maturity. However, we don't know what the spot price at the time of maturity will be. The spot price at maturity doesn't influence on the paid price, but it will decide whether or not the respective parties have lost money on the futures contract. For the seller, a spot price above the settlement price means he has lost the opportunity to sell the contract at a higher price. For the buyer, a spot price below the agreed price will mean he could have got the contract cheaper. So there is risk involved for both parties in a forward deal. By introducing risk premiums in our deal, we allow for some net risk compensation from one party to the other. In what follows, let's assume it is the buyer who needs being compensated for taking on the risk of a futures contract. We will then have to subtract a risk premium from the expected spot price to settle on the futures price. So the futures price will be lower than the expected spot price. If the risk premium turns out to be negative, the result is a futures price that is higher than the expected spot price. In this case, the seller is being compensated for the spot price risk.

The above discussion reflects the question of whether or not the forward price is a predictor of the future price. This matter is discussed in McDonald pp. 140-141 for purely financial assets, and 171-172 for commodities. The conclusion is that *the prediction is biased by the amount of the risk premium*.

To get a more realistic expression for the futures price, we have to adjust our equations to account for risk. In our model, we introduce the risk premium in the appearance of two risk premium parameters; λ_χ

and λ_ξ related to the short-term component and the long-term component, respectively. We introduce the *risk-neutral*[16] processes

$$d\chi_t = (-\kappa\chi_t - \lambda_\chi)dt + \sigma_\chi dz_\chi^* \quad \textbf{Eq. 23}$$

$$d\xi_t = (\mu_\xi - \lambda_\xi)dt + \sigma_\xi dz_\xi^* \quad \textbf{Eq. 24}$$

where still

$$dz_\chi^* dz_\xi^* = \rho_{\chi\xi} dt \quad \textbf{Eq. 25}$$

We have introduced an asterisk (*) above the z's representing the Brownian motion processes. This is because we have changed the probability distribution of the process. We manipulate the random term so that positive outcomes are more likely than negative. This manipulation of the random term needs to be done to counter the risk-aversion of the buyer. If the random term had equal probabilities of positive and negative outcomes, the expected utility of the draw would be negative (as the utility value of a $1 loss is greater than the utility value of a $1 gain). *Manipulation of the random term gives us an expected utility change of zero* for each draw (we add the risk premium to the random draw).

In order to achieve this, we do the following transformation:

$$\begin{aligned} d\chi_t &= -\kappa\chi_t dt + \sigma_\chi dz_\chi \\ &= -\kappa\chi_t dt + \sigma_\chi dz_\chi + (\lambda_\chi - \lambda_\chi)dt \\ &= (-\kappa\chi_t - \lambda_\chi)dt + \sigma_\chi\left(dz_\chi + \frac{\lambda_\chi}{\sigma_\chi}dt\right) \\ &= (-\kappa\chi_t - \lambda_\chi)dt + \sigma_\chi dz_\chi^* \end{aligned} \quad \textbf{Eq. 26}$$

An identical operation is performed on the long-term factor.

Note that we had to reduce the drift terms by the risk premiums, which are defined as constants[17]. If we had not done this reduction, we wouldn't have got correct spot price values due to the increased expected value of the random draws.

The proof that such a transformation can be done, is called Girsanov's theorem, which states that a Brownian motion process can be transformed into a new Brownian motion process that is a martingale under a different probability distribution. In our context, the transformed process is not a martingale with respect to spot price, but it *is* a martingale with respect to the investor's utility.

[16] The *risk-neutral measure* is the most common measure for valuing derivatives, and enables us to always discount future cash flows at the risk-free rate. This makes derivatives valuation a lot less complicated. It is also possible to discount the cash flows at the required rate of return, but this makes calculations a lot more complicated. For more on the risk-neutral measure see McDonald chapters 10, 11 and 20.

[17] This is an approximation. In reality, risk preferences will vary over time.

So the difference between the random element of the "true" process and the risk-neutral process is that the true process has a zero expected change in spot price, while the risk-neutral process has a zero expected change in utility for the risk-averse investor.

To sum up, in order to cope with the problem of investors being risk-averse, we have manipulated the development of state variables so that the expected spot price is reduced via the risk premiums, while the random term is more likely to give positive outcomes.

4.7 Solving for Futures Prices

For the following analysis, we define the reduced drift of the long-term variable as

$$\mu_\xi^* \equiv \mu_\xi - \lambda_\xi \qquad \textbf{Eq. 27}$$

Discrete time solutions to Eq. 23 and Eq. 24 give us an expression for the expected state variables from the risk-neutral perspective:

$$\mathrm{E}^*[\chi_t, \xi_t] = [\chi_{t-\Delta t} e^{-k\Delta t} - (1 - e^{-k\Delta t})\frac{\lambda_\chi}{\kappa},\ \xi_{t-\Delta t} + \mu_\xi^* \Delta t] \qquad \textbf{Eq. 28}$$

We see that the short-term component now will revert towards $-\frac{\lambda_\chi}{\kappa}$, while the long-term component has a reduced drift of $\mu_\xi^* \Delta t = (\mu_\xi - \lambda_\xi)\Delta t$.

From the risk-neutral perspective, the expected utility change of the random term is zero (as the expected value change is for the true process), so there is no change in the covariance matrix:

$$\begin{aligned}\mathrm{Cov}^*[\chi_t, \xi_t] &= \mathrm{Cov}[\chi_t, \xi_t] \\ &= \begin{bmatrix} \frac{1 - e^{-2\kappa\Delta t}}{2\kappa} {\sigma_\chi}^2 & (1 - e^{-\kappa\Delta t}) \frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} \\ (1 - e^{-\kappa\Delta t}) \frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} & {\sigma_\xi}^2 \Delta t \end{bmatrix}\end{aligned} \qquad \textbf{Eq. 29}$$

The fact that the volatility remains the same is a property of the risk-neutral concept.

From Eq. 28, by the relation $\ln(S_t) = \chi_t + \xi_t$, basic addition yields

$$\mathrm{E}^*[\ln(S_t)] = [\chi_{t-\Delta t} e^{-k\Delta t} + \xi_{t-\Delta t} - (1 - e^{-k\Delta t})\frac{\lambda_\chi}{\kappa} + \mu_\xi^* \Delta t] \qquad \textbf{Eq. 30}$$

We also know that

$$\begin{aligned}\mathrm{Var}^*[\ln(S_t)] &= \mathrm{Var}[\ln(S_t)] \\ &= \frac{1 - e^{-2\kappa\Delta t}}{2\kappa} {\sigma_\chi}^2 + {\sigma_\xi}^2 \Delta t + 2(1 - e^{-\kappa\Delta t}) \frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa}\end{aligned} \qquad \textbf{Eq. 31}$$

Now it is easy to solve for the value of a futures contract. The futures value is just the expected spot price under the risk-neutral measure. So, if we are currently at time t,

$$\ln(F_{t+\Delta t,t}) = \ln(E^*[S_{t+\Delta t}]) \qquad \textbf{Eq. 32}$$

$$= E^*[\ln(S_{t+\Delta t})] + \frac{1}{2}Var^*[\ln(S_{t+\Delta t})]$$
$$= \chi_t e^{-\kappa\Delta t} + \xi_t + A(\Delta t)$$

where

$$A(\Delta t) = -\left(1 - e^{-\kappa\Delta t}\right)\frac{\lambda_\chi}{\kappa} + \mu_\xi^*\Delta t$$
$$+\frac{1}{2}\{\frac{1-e^{-2\kappa\Delta t}}{2\kappa}\sigma_\chi{}^2 + \sigma_\xi{}^2\Delta t$$
$$+2(1 - e^{-\kappa\Delta t})\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa}\}$$

This expression for the log future price can be used to determine the future price at various times to maturity, given the parameter configuration of the market and the state variables' values today.

4.8 Instantaneous Volatility

From the expression for Var[ln(S_t)] we can get the *instantaneous* variance of the futures contracts. The instantaneous variance tells us the variance of futures prices at different time to maturities. When time to maturity is short, the variance of futures prices is larger than for longer maturities[18]. This is because the price movements upwards and downwards are much greater for short maturities (e.g. spot price) than for deals with long time to maturity. This coincides with the model's assumption that futures prices are affected by short-term shocks and a long-term equilibrium level.

Var[ln(St)] is a measure of accumulated (integrated) uncertainty from t to t+Δt, and grows when Δt increases. It's like we're standing at time zero looking into an unknown future where the horizon of possible price development outcomes just widens and widens. Instantaneous volatility, on the other hand, is the time derivative of the integrated uncertainty, so that

$$\text{vol} = \sqrt{\frac{dVar^*[\ln(S_T)]}{d\Delta t}} \qquad \textbf{Eq. 33}$$
$$= \sqrt{e^{-2\kappa\Delta t}\sigma_\chi{}^2 + \sigma_\xi{}^2 + 2e^{-\kappa\Delta t}\rho_{\chi\xi}\sigma_\chi\sigma_\xi}$$

The volatility of futures contracts are greatest when time to maturity goes toward zero, in which case the price of the futures contract is greatly affected by the short-term deviations in addition to the equilibrium price level. With longer time to maturity, the futures price is almost solely dependent on the equilibrium level, and hence the volatility just becomes the standard deviation of the equilibrium level.

[18] This is known as the Samuelson effect, as described in Samuelson (1965).

Chapter 5 Calibrating the Schwartz-Smith model

The following procedure for calibrating the Schwartz-Smith model is heavily inspired by Andresen and Sollie (2011). The procedure's general principles are also utilized and quoted by Lucia and Schwartz (2002) and by Cortazar and Schwartz (2003).

The parameters and state variables in Eq. 32 are not given alongside the futures price quotations. In fact, the quoted prices are a result of parties negotiating over the contract's price, so when calibrating the model we just try to interpret the agreed prices through *unobservable* parameters and variables. The fact that our parameters and variables are unobservable means that we have to estimate them. Knowing the value of the parameters will tell us a lot about the market configuration, and is also all we need for using the model to simulate oil prices in the future.

We use historical futures price data to calibrate the model. When calibrating the model, we estimate the model's parameters and, for each day we possess quoted futures prices, state variables.

In this chapter we will also discuss problems in estimating the risk-premiums of the model, and how this affects the usefulness of the model for forecasting purposes. Finally, we will summarize this chapter by giving an overview of the variables and parameters, and their interpretation, in tabular form.

5.1 A Spreadsheet Procedure for Calibrating the Model

Historical futures price data provides a matrix/table where each row represents a date for which we have obtained futures prices, and for each date the prices of the futures contract for each maturity month is given in the columns. For each new day a row of futures prices is added, and the prices in the columns make up the forward curve. As a result we have a new forward curve each day. For each day, we model the forward curve using our parameters (which are static) and the state variables (which change day by day). So our goal is to find the parameters that describe the fundamental conditions of the oil futures market and for each day select the state variables (long-term and short-term components) that are implied by that day's observed forward curve.

Figure 1 Illustration of the dataset. Prices are given in the cells of the table. The model interprets the forward curves as the risk-neutral expected prices based on the prevailing state variables and the model parameters.

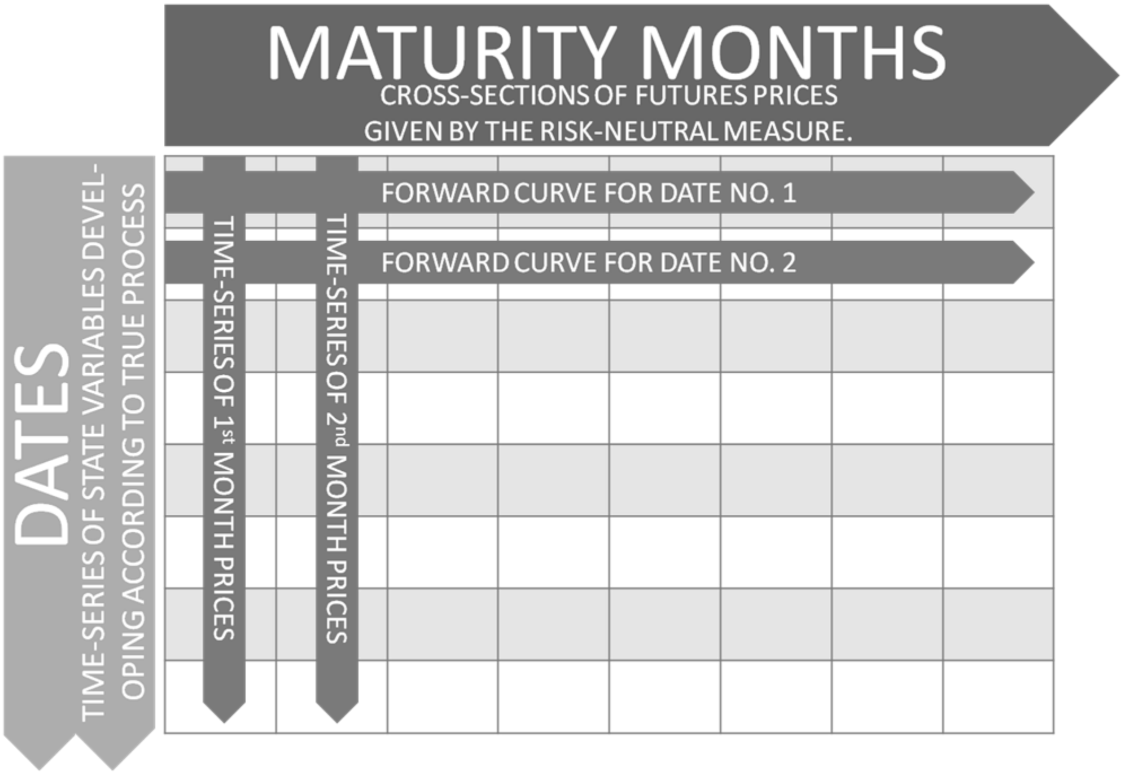

Estimating the parameters and state variables can be done using a tool called the *Kalman filter*. This is a quite advanced method for data analysis. In this thesis, however, we have used an iterative routine that can be easily implemented in MS Excel, using Solver and macros written in Visual Basic for Applications (VBA). Had we applied the Kalman filter, then standard deviations of the parameter estimates would be obtained. Using the iterative routine we don't get any measure of the estimates' precision.

The procedure can be summarized like this:

1. We make a guess about the parameters based on what we believe or parameter estimates from similar markets. Starting parameters close to the true parameters will give faster convergence of the procedure. We leave the state variables for each day equal to zero.
2. For each day of data, based on the parameters we currently have, we fit that day's modeled forward curve to the observed futures prices that day. We do this by selecting the combination of the two state variables that minimizes the squared distance (SSE – sum of squared errors) between the forward curve and the observed prices. For "picking" state variables we use Solver.
3. We have now obtained two vectors, one containing the short-term components for each day, the other containing the long-term components. These two vectors are used to obtain estimates of σ_χ, σ_ξ, $\rho_{\chi\xi}$ and μ_ξ. The vectors can be interpreted as the realized development of the state variables, and can therefore be used to estimate the parameters of Eq. 9 and Eq. 10.

4. Now we have estimates of the state variables for each day, and also three out of six parameters of the expression for the futures price[19]. The last three parameters are solved for by minimizing the squared distances between *all* estimated forward curves and observed prices. Solver "picks" the κ (the rate of mean reversion), μ_ξ* (the equilibrium price's risk-neutral drift term) and λ_χ (the short-term risk premium) that best fits the whole dataset.
5. We have now estimated both state variables, and all parameters. But most likely, for the new configuration of parameters, some new state variables will give an even better fit of the forward curves to the observed data. So we have to repeat the procedure from step 2 to 4 over and over again until the iterative procedure reaches convergence. We define convergence as the situation where the differences in total SSE between two iterations is so small that

$$\ln\left(\frac{SSE_{j-1}}{SSE_j}\right) < \delta \qquad \textbf{Eq. 34}$$

where δ is the *convergence criterion*. We set this equal to 0.0001.

At convergence we have the parameters that best describe the development of futures prices, and the state variables implied by the observed futures prices and the estimated parameters.

Below is an overview of the equations used for estimating the parameters of step 3. The equations are really just the equations for state variable development (Eq. 12 and Eq. 13) solved for the respective parameters.32

$$\boldsymbol{\pi} = \frac{\widehat{\boldsymbol{\chi}}_i - \widehat{\boldsymbol{\chi}}_{i-1} e^{-\hat{\kappa}\Delta t}}{\sqrt{\frac{1 - e^{-2\hat{\kappa}\Delta t}}{2\hat{\kappa}}}} \qquad \textbf{Eq. 35}$$

$$\hat{\sigma}_\chi = \sqrt{Var(\boldsymbol{\pi})} \qquad \textbf{Eq. 36}$$

Where **π** is a vector/column of the scaled increments of $\hat{\chi}_t$; $\widehat{\boldsymbol{\chi}}_i$ and $\widehat{\boldsymbol{\chi}}_{i-1}$ are the estimated state variables at time t = i and t = i – Δt, respectively. Δt is the time interval between two estimates of χ_t. $\hat{\kappa}$ is the estimated rate of mean reversion.

$$\boldsymbol{\psi} = \frac{\hat{\boldsymbol{\xi}}_i - \hat{\boldsymbol{\xi}}_{i-1} - \hat{\mu}_{\xi,j-1}\Delta t}{\sqrt{\Delta t}} \qquad \textbf{Eq. 37}$$

$$\hat{\sigma}_\xi = \sqrt{Var(\boldsymbol{\psi})} \qquad \textbf{Eq. 38}$$

Where **ψ** is a vector/column of the scaled increments of $\hat{\xi}_t$; $\hat{\boldsymbol{\xi}}_i$ and $\hat{\boldsymbol{\xi}}_{i-1}$ are the estimated state variables at time t = i and t = i – Δt, respectively. Δt is the time interval between two estimates of ξ_t. $\hat{\mu}_\xi$ is the estimated drift parameter related to ξ_t.

[19] Note that the parameter μ_ξ isn't used in calculating the futures price. However, it needs to be calculated in order to estimate the long-term risk premium. Problems in estimating this parameter precisely are discussed later.

$$\hat{\rho}_{\chi\xi} = corr[\boldsymbol{\psi}, \boldsymbol{\pi}] \quad \textbf{Eq. 39}$$

$$\hat{\mu}_{\xi} = \frac{1}{N}\sum_{i=1}^{N}\frac{\hat{\xi}_i - \hat{\xi}_{i-1}}{\Delta t} \quad \textbf{Eq. 40}$$

μ is the averaged equilibrium price level change per time change. This equation assumes that there is some underlying, linear drift in the development of equilibrium prices, and estimates it based on the vector of estimated state variables. In taking the averaged change per time interval, we hope to isolate the effect of the drift parameter from the influences of the random draws.

There are some issues regarding the calibration process that need being addressed. First, the iterative routine may fail in giving realistic values for the parameters if the upper and lower boundaries for Solver aren't set properly. The limits must not be too strict, meaning they must give the parameters freedom to take on any reasonable value, but we need some sensible limits to prevent the parameters from taking on totally unreasonable values and causing the calibration procedure to come off the track. Fortunately, as long as the parameters are restricted somewhere around the true values, they will converge to the true values.[20]

5.2 The Risk Premiums Causing Trouble

Another quite important issue related to the model calibration is the fact that we can't trust the risk premiums it recovers. Said more precisely, the routine *does* manage to estimate the risk-neutral process satisfactorily, but the values of the risk premiums are very uncertain. The short-term risk premium always end up very close to the start value of the calibration, while the long-term risk premium simply just exhibits a great deal of uncertainty. As shown by Andresen and Sollie (2011), these are systematic errors of the iterative routine. Neither does the Kalman filter manage to recover "true" values of the risk premiums, the authors conclude.

Therefore, trying to interpret the long-term and short-term risk premiums (λ_ξ and λ_χ, respectively) is worthless. Knowing that the risk-neutral expression for the futures price perform satisfactorily, this means that the problem with the risk premiums are passed on to the true process. And if, as concluded by Andresen and Sollie, all other parameters are satisfactorily recovered, then there must be some adjustment to the state variables to ensure that Eq. 32 yields correct futures prices. So the falseness of the risk premiums contaminates the estimation of the state variables.

Andresen and Sollie don't explicitly attempt to investigate *why* the estimates of the risk premiums turn out wrong, although they do find that the uncertainty of the estimates decrease as the data amount increases, whilst it increases for larger volatility parameters. In the following, an explanation to the apparent impossibility of estimating the risk premiums satisfactorily will be presented[21].

[20] See Andresen and Sollie (2011) for more on the accuracy of the iterative routine.
[21] Actually, this subject is also treated in Schwartz & Smith (pp. 906-907).

The risk premiums are estimated in two very different ways. The short-term risk premium is estimated as part of the SSE-minimizing operation using Solver, while the long-term risk premium is estimated on basis of the long-term state variable development. Let's first look at a possible explanation why the estimate of λ_χ can't be trusted.

When $\Delta t \to \infty$ the expression for the futures price approaches

$$\ln\left(F_{t+\Delta t,t}\right) = \xi_t - \frac{\lambda_\chi}{\kappa} + \mu_\xi^* \Delta t + \frac{{\sigma_\chi}^2}{4\kappa} + \frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} \qquad \textbf{Eq. 41}$$

Let's assume that the calibration manages to recover true values of $\ln\left(F_{t+\Delta t,t}\right)$ and all parameters except λ_χ have true values. Then, the difference between the estimated and the true value of λ_χ will have to be offset by a corresponding falseness in the estimate of ξ_t. Formally, we can define

$$\mathrm{e}_{\lambda_\chi} = \hat{\lambda}_\chi - \lambda_\chi \qquad \textbf{Eq. 42}$$

$$\mathrm{e}_\xi = \hat{\xi}_t - \xi_t \qquad \textbf{Eq. 43}$$

When calibrating the model based on observed $\ln\left(F_{t+\Delta t,t}\right)$ the procedure just offsets the error in the estimate of λ_χ by a corresponding error in the time-series of the long-term state variable, so that

$$\mathrm{e}_\xi = \frac{\mathrm{e}_{\lambda_\chi}}{\kappa} \qquad \textbf{Eq. 44}$$

When $\Delta t \to 0$ the expression for the futures price approaches

$$\ln\left(F_{t+\Delta t,t}\right) = \xi_t + \chi_t \qquad \textbf{Eq. 45}$$

Hence, an estimation error in ξ_t will have to be offset by an equal but opposite error in the short-term variable, so that

$$\mathrm{e}_\xi = -\mathrm{e}_\chi \qquad \textbf{Eq. 46}$$

Now, we need to confirm if these undesired offsetting errors may occur for all Δt's. Let's take the general expression for the futures price and leave out all variables and parameters we expect to have the true value. We are then left with

$$0 = \chi_t e^{-\kappa\Delta t} + \xi_t - \left(1 - e^{-\kappa\Delta t}\right)\frac{\lambda_\chi}{\kappa} \qquad \textbf{Eq. 47}$$

If introduction of estimation errors to the expression on the right side doesn't change the resulting value of the expression, then we have proved that offsetting errors are possible for all Δt's. The zero impact of the errors can be shown this way:

$$\begin{aligned} 0 &= \mathrm{e}_\chi e^{-\kappa\Delta t} + \mathrm{e}_\xi - \left(1 - e^{-\kappa\Delta t}\right)\frac{\mathrm{e}_{\lambda_\chi}}{\kappa} \\ 0 &= -\frac{\mathrm{e}_{\lambda_\chi}}{\kappa}e^{-\kappa\Delta t} + \frac{\mathrm{e}_{\lambda_\chi}}{\kappa} - \frac{\mathrm{e}_{\lambda_\chi}}{\kappa} + \frac{\mathrm{e}_{\lambda_\chi}}{\kappa}e^{-\kappa\Delta t} \end{aligned} \qquad \textbf{Eq. 48}$$

$$0 = 0$$

This result explains why the iterative routine is unable to discover errors in the short-term risk premium estimate. The reason is that there is redundancy of factors in the model. There seems to be no reason at all to include this parameter in the model, unless a perfectly identical short-term risk premium can be estimated by using some other technique.

What are the implications of an error to the estimated λ_χ? The error will be transferred to each individual of the estimated state variables. Therefore, the absolute value of the state variables will be affected by the error, but the size of the increments of the state variables will be unaffected. This makes it possible to use the time series of the state variables in estimating the long-term risk premium, correlations and standard deviations of state variables, as calculations are based on increments of ξ_i's and χ_i's, not absolute values.

The problems in estimating the long-term risk premium, $\lambda_\xi = \mu_\xi - \mu_\xi^*$ is of a different nature. The risk-neutral drift constant μ_ξ^* is concluded by Andresen and Sollie to be estimated with satisfying accuracy. We then get

$$e_{\lambda_\xi} = e_{\mu_\xi} \quad \textbf{Eq. 49}$$

The problems of estimating the long-term risk premium comes from problems with estimating the true drift parameter. The expression used to estimate μ_ξ is given in Eq. 40, and hopes to isolate the drift constant from the increments of ξ_t just by taking the average of the increments. The reason why this in theory could work is that the random term of ξ_t increments has expectation zero (see Eq. 13). In practice it seems very difficult to isolate the drift parameter, at least for realistic data amounts. Three simulations of 2500 dates were made in order to confirm this. Using the simulated (true value) long-term state variables, the drift parameter was calculated by applying Eq. 40. σ_χ, σ_ξ and $\rho_{\xi\chi}$ were also calculated to show that other parameters are estimated with satisfying precision. The results are shown below:

Table 1 Test showing great uncertainty in long-term risk premium estimates.

	True	Sim. 1	Sim. 2	Sim. 3
μ_ξ	0.06	-0.04	0.004	0.07
σ_χ	0.24	0.24	0.23	0.24
σ_ξ	0.24	0.24	0.24	0.24
$\rho_{\xi\chi}$	0.20	0.21	0.20	0.18
Avg. ε_1[22]		-0.02	-0.01	0.00
Avg. ε_2		-0.04	-0.01	-0.01
SD. ε_1		1.00	0.98	0.99
SD. ε_2		1.02	0.97	1.01

The results give a clear indication that even though the random draws has an average close to zero, there is still great uncertainty in estimates of the long-term risk premium. For the time perspectives we

[22] See Eq. 54 below.

are dealing with here, it seems that the random movements of the long-term variable are superior to the drift term.

5.2.1 Implications of Uncertain Risk Premium Estimates on Forecasting

According to Schwartz and Smith (p. 906), the errors in risk premium estimates don't affect the robustness of the model for use in valuation problems, as the risk-neutral model fits observed forward curves. However, the great uncertainty in estimates of the true process drift term reduces the model's robustness for forecasting (simulation) purposes.

Schwartz and Smith report the standard deviation of the drift parameter at 0.0728. This is pretty high, considering that their parameter estimate is -0.0125. The size of the standard deviation indicates that the real drift parameter underlying the formation of oil prices may be quite different than their estimate.

As the true process drift term is used to calculate expected future prices, and plays an integral part in the simulation of price development, the drift parameter μ_ξ should preferably have been estimated with great accuracy in order to perform reliable simulations.

This uncertainty should be kept in mind when utilizing the model for simulation of oil price development, and especially when interpreting simulation results in order to calculate expected future prices and assessing risk related to futures trading.

5.3 Overview of Model Parameters and State Variables

We end this chapter by listing all the parameters and variables of the Schwartz-Smith model, also giving a brief explanation of what they represent.

Table 2 The parameters and state variables of the Schwartz-Smith model.

		Interpretation	*Estimated by*
State variables	ξ	Long-term variable, representing the equilibrium price when short-term impacts have faded.	Minimizing SSE between observed forward curves and the modeled forward curves, subject to the most recent model parameter estimates (see step 2, page 32).
	χ	Short-term variable, representing short-term deviations (shocks) to the oil price.	Minimizing SSE between observed forward curves and the modeled forward curves, subject to the most recent model parameter estimates (see step 2, page 32).
Parameters	κ	The rate of mean reversion, a measure of how fast short-term deviations are assumed to fade. Half-life of short term deviations can be calculated by $-\ln(0.5)/\kappa$.	Minimizing SSE between observed forward curves and the modeled forward curves, subject to the most recent state variable estimates (see step 4, page 33).
	λ_χ	Short-term risk premium, reduces the expected future value of the short-term deviations. See above discussion about the risk-neutral measure.	Minimizing SSE between observed forward curves and the modeled forward curves, subject to the most recent state variable estimates (see step 4, page 33).
	μ_ξ^*	Risk-neutral expected growth rate of the long-term variable. As the market is influenced by risk-averse investors, the drift parameter used for determining future prices deviates from the true growth rate by the long-term risk premium. See above discussion about the risk-neutral measure.	Minimizing SSE between observed forward curves and the modeled forward curves, subject to the most recent state variable estimates (see step 4, page 33).
	μ_ξ	Drift parameter (average growth rate) in the true development of the long-term variable.	See Eq. 40.
	σ_χ	Standard deviation of the random element in the short-term variable's development.	See Eq. 35 and Eq. 36.
	σ_ξ	Standard deviation of the random element in the long-term variable's development.	See Eq. 37 and Eq. 38.
	$\rho_{\chi\xi}$	Correlation between the random elements of short- and long-term variable's development.	See Eq. 39.

Chapter 6 Results from Calibrating the Schwartz-Smith Model

The results from calibrating the model provide us with tangible information about the market configuration. With a small number of parameters, it is easy to compare today's Brent and WTI markets to other commodity markets.

From looking at time-series of 1st month contracts, we perceive a great price gap between Brent and WTI during 2011. The estimated state variables explain this difference principally by deviations between the short-term variables. The long-term variables of Brent and WTI are highly correlated, and follow each other closely through the whole time period.

The Brent market turned from contango to backwardation in 2011. WTI followed about the same development, however during the time of large price gaps WTI sometimes found itself in contango while Brent experienced backwardation.

Parameter estimates suggest great similarities in the Brent and WTI markets. WTI prices are slightly more volatile than Brent, and also exhibit faster mean reversion.

The calibration procedure struggles more in fitting the model to WTI data than Brent.

Negative risk-neutral drift parameters imply that equilibrium prices are expected by buyers to become lower in time. As true process drift is estimated at positive values, it becomes apparent that there is a significant share of risk aversion in the market.

Comparing our WTI estimates with estimates from Schwartz and Smith (2000), we identify some differences between today's market and the market from 1990-1995. Our estimate indicates a lower rate of mean reversion today, meaning price shocks have longer impact on oil prices. Also, the volatility of the equilibrium level has increased remarkably. Risk-premium estimates seem to indicate greater risk aversion today compared to twenty years ago.

In addition to estimates of state variables and parameters, this chapter explains where the prices in the datasets come from. The price development during the observed time period is analyzed and commented on. We describe the market's transition from contango to backwardation, and investigate how good an indicator of contango/backwardation conditions the short-term variable is.

Finally, the estimated volatility curves of Brent and WTI are presented. They fit well with the volatility in the data. The shape of the volatility curves is in accordance with the Samuelson effect, with increasing volatility as maturity approaches.

6.1 Presenting the Datasets Used for Calibration

We calibrate the Schwartz-Smith model on datasets of observed prices for both Brent and WTI futures. The observed prices are the closing prices for each trading day. For each date we have futures prices for 50 consecutive months.

For Brent we have utilized a dataset consisting of prices for 784 dates. The dates are in the time interval from January 2009 to January 2012. For the Brent futures the specifications from the ICE state[23]:

Settlement Price is *"The weighted average price of trades during a two minute settlement period from 19:28:00, London time."* This settlement price is the same as the quoted closing price for the ICE Brent Futures.

"Trading shall cease at the end of the designated settlement period on the Business Day (a trading day which is not a public holiday in England and Wales) immediately preceding:

(i) Either the 15th day before the first day of the contract month, if such 15th day is a Business Day
(ii) If such 15th day is not a Business Day the next preceding Business Day."

When a contract expires, the old 2nd month contract becomes the new 1st month contract. For Brent, this happens on the 16th day before the contract month (the Business day *immediately preceding* the 15th day before the first day of the contract month), provided that all relevant dates are Business Days. For example, the December contract expires on November the 15th. From 19:30 November 15th the January contract becomes the 1st month contract.

The settling price at expiry (delivery/settlement basis) is the ICE Brent Index price for the day following the last trading day of the futures contract. For the December contract, the settlement basis will be the Brent Index as of November 16th (provided that the last trading day is November 15th).

For WTI we have utilized a dataset consisting of 765 dates within the same time interval as for the Brent data.

The daily settlement of the WTI Futures is calculated as the volume-weighted average price ("VWAP") of trades made between 14:28 and 14:30 Eastern Time (ET).[24]

For the WTI futures specifications from NYMEX state[25]:

"Trading in the current delivery month shall cease on the third business day prior to the twenty-fifth calendar day of the month preceding the delivery month. If the twenty-fifth calendar day of the month is a non-business day, trading shall cease on the third business day prior to the last business day preceding the twenty-fifth calendar day."

Settlement of the WTI futures at expiry is made by delivery of physical crude oil. For the December contract, trading ceases on November 22nd (provided that none of the dates from the 25th to the 22nd are non-business days). From 14:30 ET November 22nd the January contract becomes the 1st month contract.

The settlement period of Brent futures is between 19:28 and 19:30 London time. The settlement period for WTI is between 14:28 and 14:30 Eastern Time. As Eastern Standard Time (EST) is five hours behind Greenwich Mean Time (GMT), this means that the quoted prices are recorded at the same time.

---

[23] https://www.theice.com/productguide/ProductSpec.shtml?specId=219#
[24] http://www.cmegroup.com/trading/energy/files/NYMEX_Energy_Futures_Daily_Settlement_Procedure.pdf
[25] http://www.cmegroup.com/trading/energy/crude-oil/light-sweet-crude_contract_specifications.html

As expiry of the 1st month contract occurs at different times for Brent and WTI, the contracts will have different time to maturity. For the case of the December contract, on November 14th time to maturity will be one day for Brent and 8 days for WTI. On November 16th, however, time to maturity for Brent will be almost one month (as the front contract is now the January contract) while WTI expires in six days. This may lead to (temporary) disturbances in the price relation between Brent and WTI prices, as the 1st month contracts of the two products have deliveries in different months. However, when reviewing the obtained data it seems that this *rollover*[26] *effect*, if present, can be neglected. See Figure 2 for data from the first five months of 2010.

When using the obtained data for calibration purposes, we have chosen to neglect this difference in time to maturity for the two products. In our calculations, we have assumed expiry for both products at the last day of the month preceding the contract month. As there doesn't seem to be any significant shifts in prices at rollover, this assumption won't have critical effect on the results from calibration.

Figure 2 First month prices of Brent and WTI futures from Jan 15th to June 15th 2010.

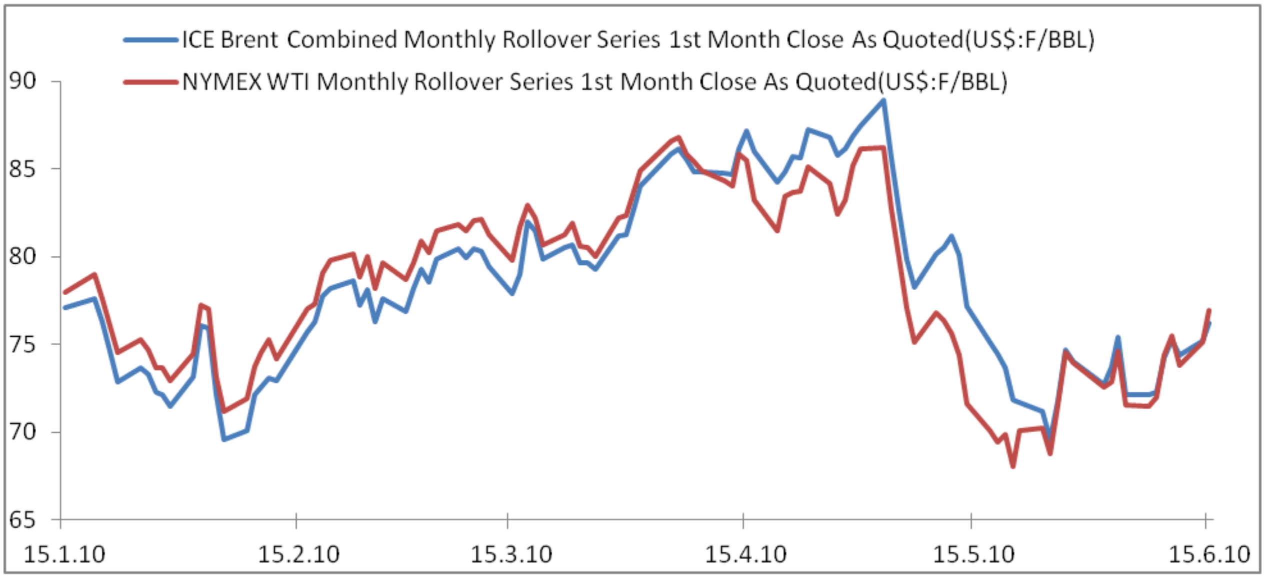

When utilized for model calibration, we have transformed the observed prices into log prices. This is because the Schwartz-Smith model operates with futures prices on log level (or, as formulated by Schwartz (1997): "*the logarithm of the futures price is linear in the underlying factors*")**.**

When taking the logarithm of prices, we get the following plot for the sample above:

[26] *Rollover* is and idiom for the event when the previous 2nd month contract becomes the new 1st month contract (an effect which propagates through the entire forward curve, eventually leading to a new contract being added at the far end of the curve).

Figure 3 First month prices of Brent and WTI futures (log scale).

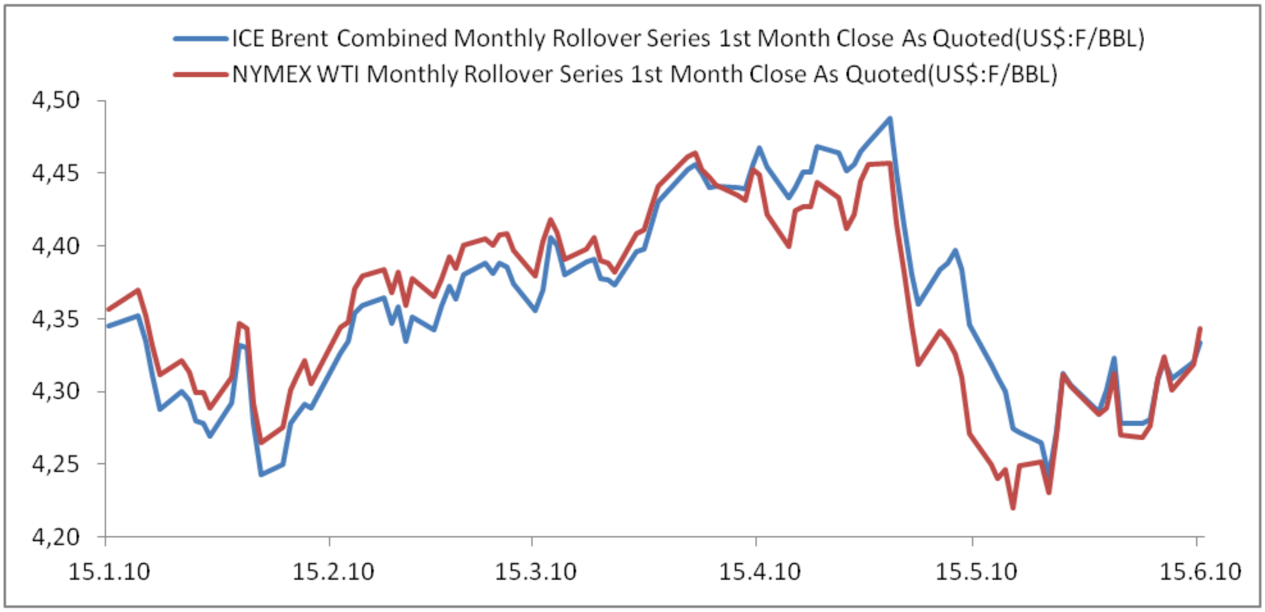

6.2 Plots of Observed and Model Implied Prices

Plotting the observed 1^{st} and 50^{th} month Brent prices yields the following curves:

Figure 4 Plot of Brent 1st and 50th month prices (log scale).

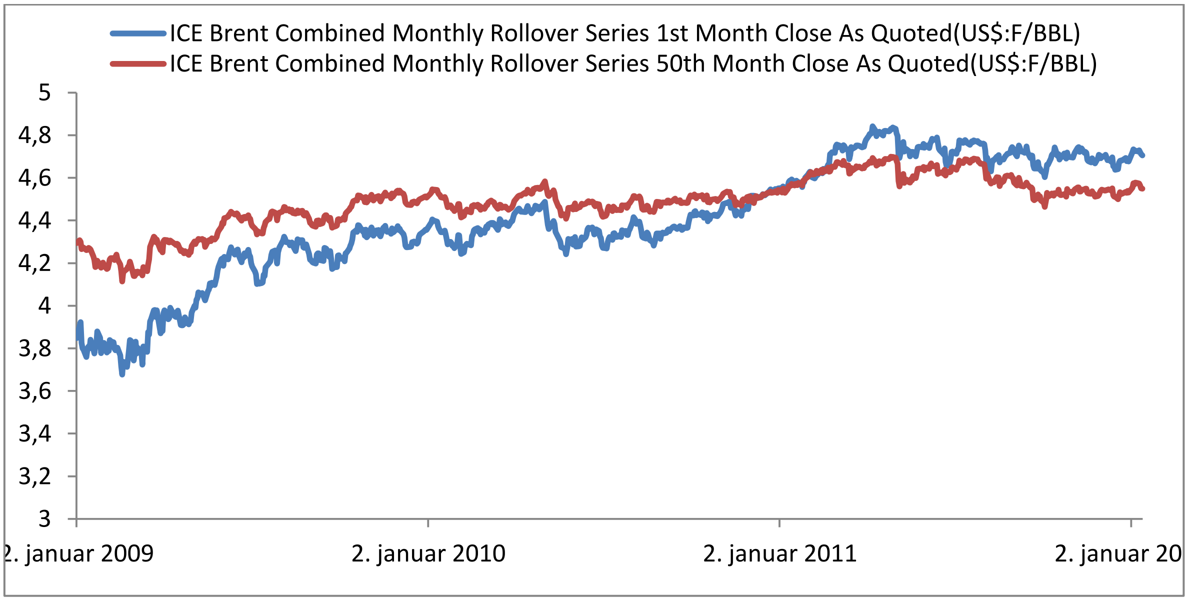

We see that the market goes from contango (price higher at later maturities) to backwardation (price lower at later maturities) around the beginning of 2011. We also see that the deviations of the 1^{st} month price are larger than the deviations of the 50^{th} month. This supports the theory that volatility is largest in the near end of the forward curve, while it decreases at time to maturity gets longer. The plot also indicates the presence of an equilibrium price level which short-term prices revert around, with this equilibrium price exhibiting some movement. The upward movement of the 50th month price indicates some positive drift constant relating to the long-term variable.

Seeing how well the model is able to reconstruct this price development is interesting. Below, we present a plot of the modeled 1st and 50th month prices together with the observed prices. Figure 5 shows the Brent data along with the time series implied by the model calibration. In Figure 6 we present the corresponding data and calibration results for WTI.

Figure 5 Model constructed time series vs. observed Brent prices (log scale).

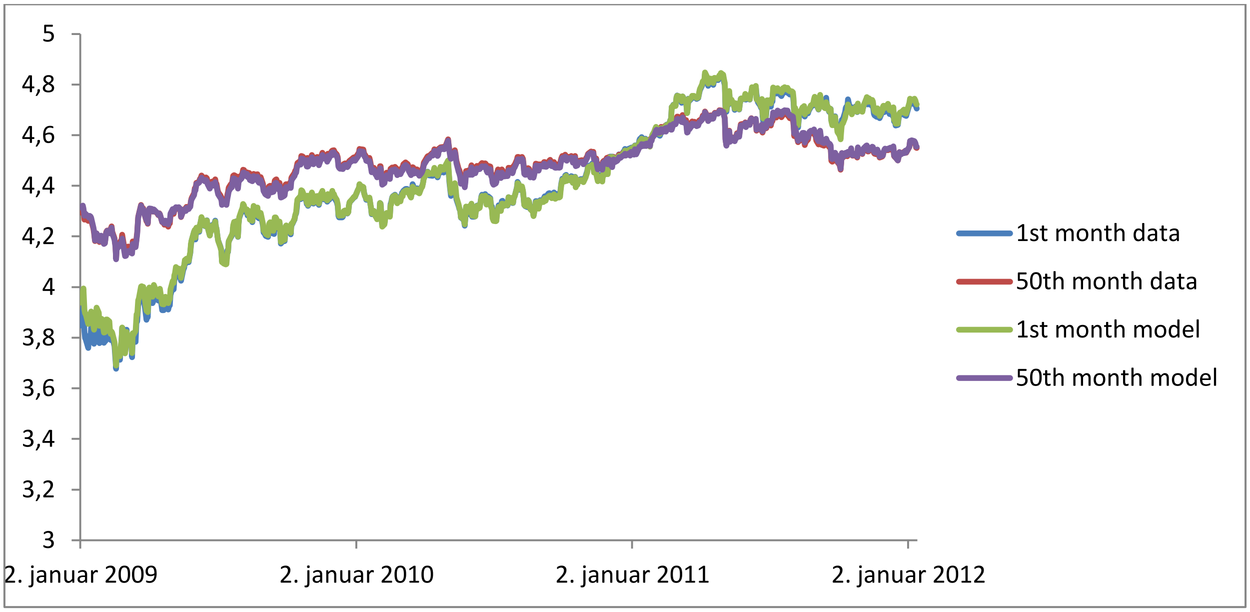

Figure 5 shows that the model reconstructs the historical Brent data well, but it seems to struggle a bit around the beginning of 2009, at least for the 1st month prices.

Figure 6 Model constructed time series vs. observed WTI prices (log scale).

Figure 6 indicates that the fit for WTI isn't as good as for Brent. This is confirmed by the sum of squared errors from the calibration. These figures are presented below. By looking at the above graphic, we note that the WTI model is, as Brent, struggling in the start of 2009, but the WTI calibration also has some visible problems during the more recent periods of our sample.

It is interesting to see how WTI prices develop relative to Brent. The next plot shows time series of Brent and WTI prices for 1st and 50th month:

Figure 7 Plot of Brent vs. WTI prices (log scale).

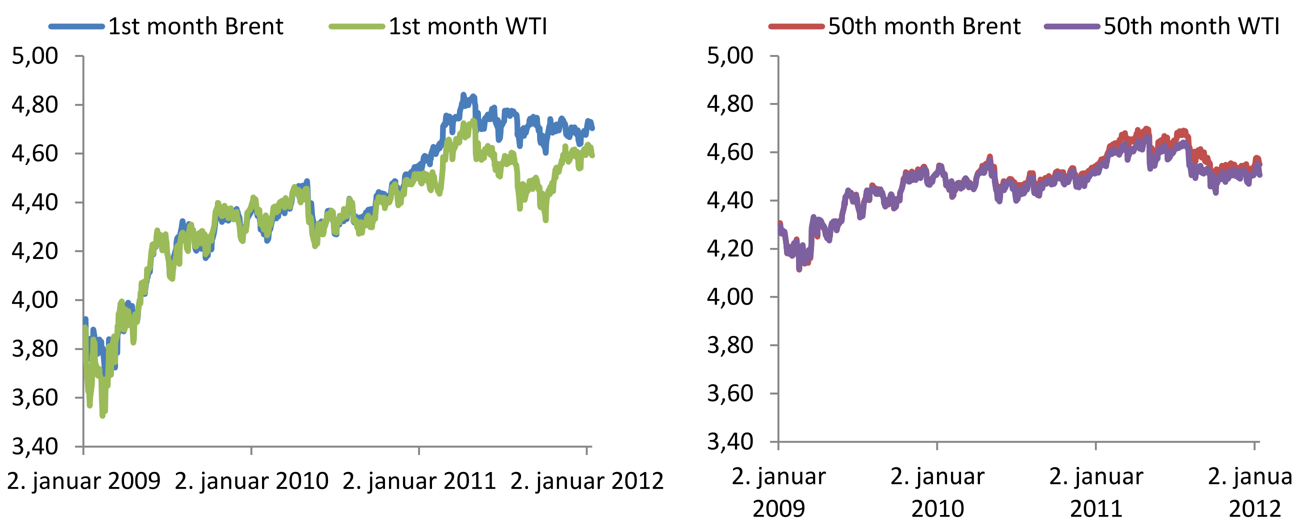

From these graphs we see that Brent and WTI are highly correlated products. It also seems as if the long-term prices are closer to each other than what the case for the short-term prices is. As explained in the earlier discussion about the Brent/WTI spread, great deviations may occur due to bottlenecks at the pipelines transporting crude oil (WTI) out of Cushing, Oklahoma. This is what happened during 2011 and the start of 2012.[27] Actually, during 2011 the markets found themselves in different states; as Brent experienced backwardation during 2011, WTI stayed in contango for great parts of 2011.

6.3 Plots of Estimated State Variables

Plotting the estimated long-term state variables reveals great long-term correlation between Brent and WTI:

[27] See for example (Pickrell, 2012).

Figure 8 Estimates of long-term variables.

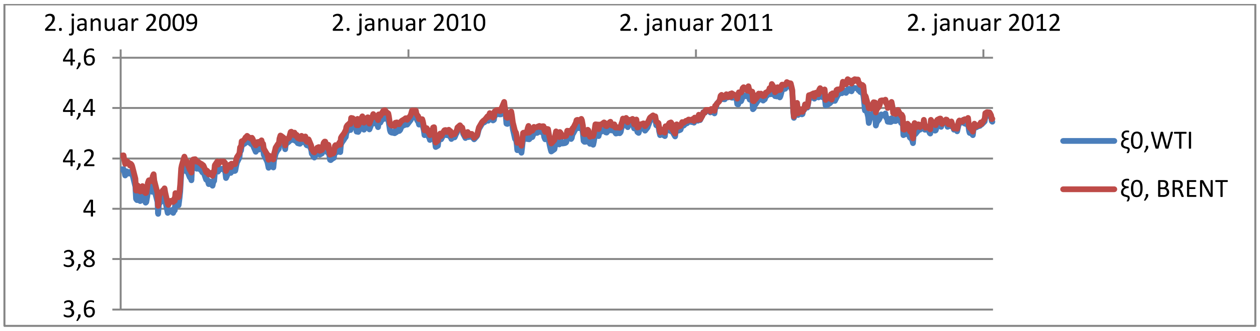

From Figure 8 we also note that the long-term variable of Brent consistently seems to be above WTI; however the difference pretty small. From our discussion about the impact of the errors in short-term risk premium on state variable estimates, we should be careful in over-interpreting the state variables' absolute values.[28] Nevertheless, the relationship between estimates of long-term variables for Brent and WTI doesn't seem to contradict the plots of observed prices, which tell us that the prices are very close to each other with Brent being maybe a bit more expensive on average.

Plotting the estimated short-term variables reveals great correlation between the two futures contracts also for the short-term variable. However, the increasing short-term price gap of 2011 materializes in this plot of short-term variables. This indicates that the model is able to assign short-term price deviations to the short-term variable.

Figure 9 Estimates of short-term variables.

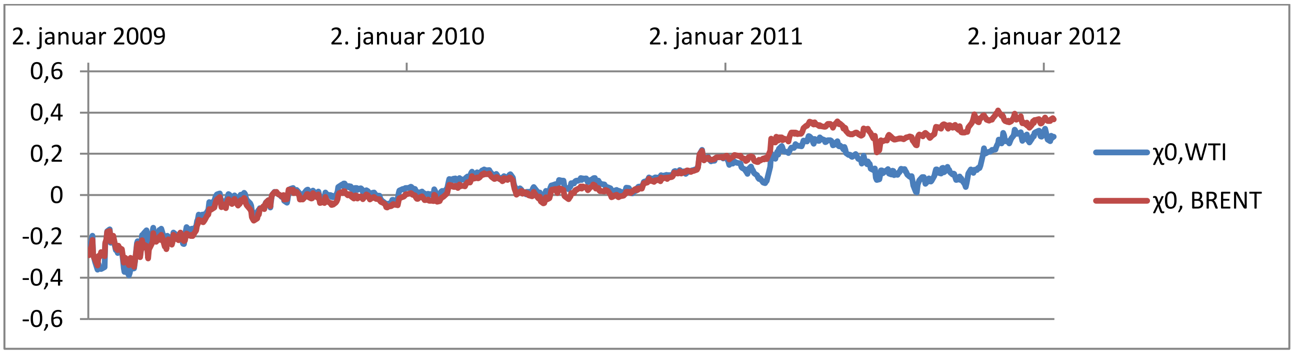

6.4 Is the Short-Term Variable an Indicator of Contango/Backwardation?

One of the strengths of the Short-Term/Long-term model compared to other price models is the simple intuition behind it, which makes it easy to interpret. The introduction of a short-term and a long-term variable gives an intuitive understanding of how expectations in the market changes.

[28] But we can still use the estimated time series of variables to make examinations that rely on the increments to state variables development, as these aren't affected by the error to the short-term risk premium estimate.

The short-term variable is a good pointer as to whether the market is in contango or backwardation. For instance, in the above discussion about the great Brent/WTI spread during 2011, we observed that the Brent market was backwardated while WTI experienced contango. This coincided with Brent exhibiting much higher values than WTI for the short-term variable. In the following, we will examine whether we can determine contango/backwardation configurations solely based on looking at the short-term variable.

The difference between the log prices[29] of contracts maturing at $\Delta t \to 0$ and $\Delta t > 0$ is

$$\ln\left(F_{t,t}\right) - \ln\left(F_{t+\Delta t,t}\right) = \chi_t - \chi_t e^{-\kappa \Delta t} - A(\Delta t) \qquad \textbf{Eq. 50}$$

Where

$$\begin{aligned} A(\Delta t) = -\left(1 - e^{-\kappa\Delta t}\right)\frac{\lambda_\chi}{\kappa} + \mu_\xi^* \Delta t \\ +\frac{1}{2}\{\frac{1-e^{-2\kappa\Delta t}}{2\kappa}\sigma_\chi{}^2 + \sigma_\xi{}^2 \Delta t \\ +2(1-e^{-\kappa\Delta t})\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa}\} \end{aligned}$$

As $\Delta t \to \infty$ we get

$$\ln\left(F_{t,t}\right) - \ln\left(F_{t+\Delta t,t}\right) = \chi_t - A(\Delta t) \qquad \textbf{Eq. 51}$$

Where

$$A(\Delta t) = -\frac{\lambda_\chi}{\kappa} + \frac{\sigma_\chi{}^2}{4\kappa} + \frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} + (\mu_\xi^* + \frac{\sigma_\xi{}^2}{2})\Delta t$$

The market is in backwardation if Eq. 50 or Eq. 51 returns a positive number, meaning that the spot price is higher than the future price. How good an indication of backwardation/contango the short-term variable gives is dependent on the size of the model parameters, but also the value of the short-term state variable. The only factor that doesn't affect backwardation/contango is the long-term variable. A plot of the short-term variable versus the contango/backwardation situation would help us understand the relationship between these variables. We will use our Brent data sample and model for the following presentation of results.

[29] We should keep in mind that $\ln\left(F_{t,t}\right) - \ln\left(F_{t+\Delta t,t}\right)$ just tells us the difference between the log spot prices, it doesn't give us the log of the price difference. In order to get the log of the price difference, we would have to calculate $\ln(e^{\ln(F_{t,t})} - e^{\ln(F_{t+\Delta t,t})})$. So we can't tell the value of the difference using Eq. 50, but we can tell the sign of the price difference. The sign tells us whether we have contango or backwardation.

Figure 10 Short-term variable vs. differentials of log prices on the modeled forward curve.

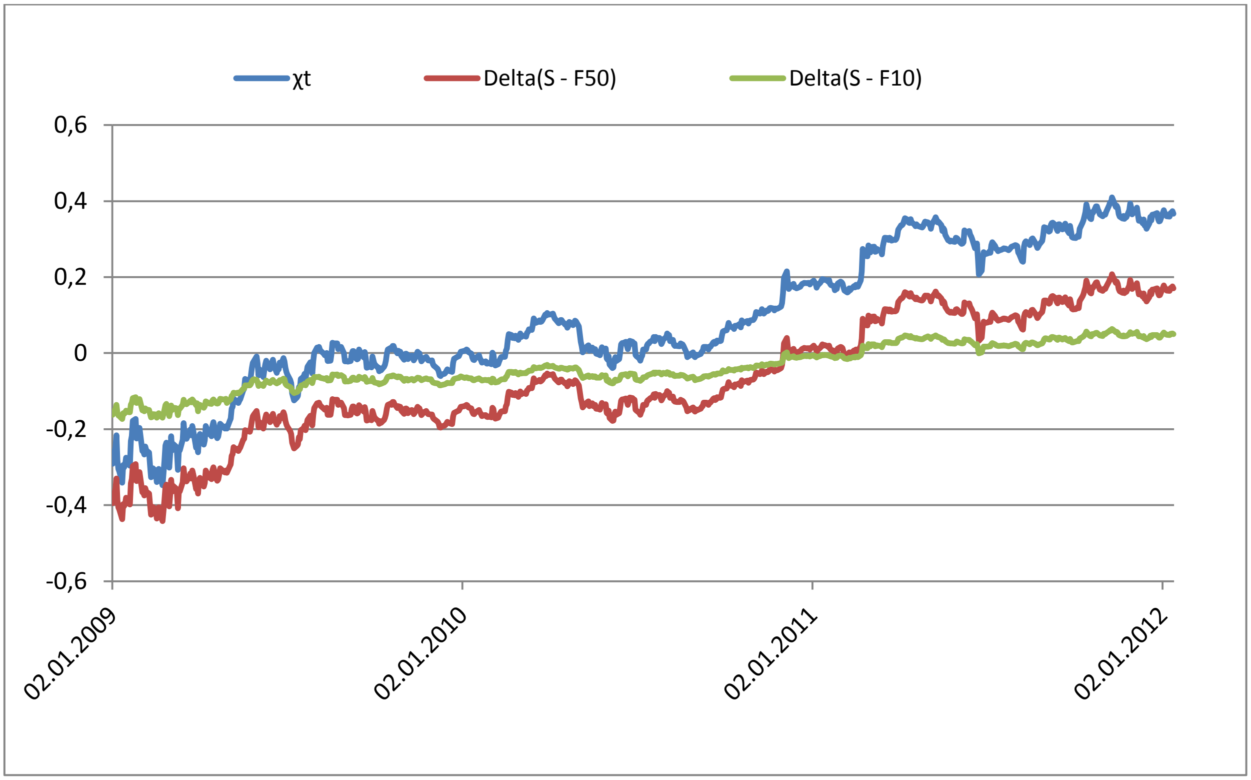

This plot in Figure 10 requires some explanation. The time series of the short-term variable is plotted together with two price differentials along the modeled forward curve. We have plotted the difference between the estimated spot price and ten month price, and between the spot price and fifty month price. It is easy to see the connection between these variables, especially between the short-term variable and the price difference of spot price and fifty month price. The tenth month's price differential looks like a damped version of the fiftieth month's differential.

The market is in contango as long as the price differences along the forward curve are negative and in backwardation when the graphs of price differences have positive values. Also by this view, we get the conclusion that the market shifts from contango to backwardation around the start of 2011. However, what we are investigating here, is the relation between the short-term variable and contango/backwardation situations. Look at the curves for the short-term variable and the price difference between the first and fiftieth month. We see that there is a clear connection between the two variables, and if the distance between the two curves was constant it would be easy to figure out at which value of χ_t the market shifts from contango to backwardation. Unfortunately, the gap between the curves isn't constant. There's some constant difference based on the size of the model parameters, but also since we cannot neglect the effect of the short-term variable on the fifty month price, the gap is affected by the size of the short-term variable (see Figure 11). We need to carry out some analytical calculations to examine exactly at what value of the short-term variable the market shifts from contango to backwardation. This is time-consuming, and therefore we might as well just graph the time series of spot prices vs. fifty month prices to determine whether or not the market is in contango or backwardation.

Figure 11 Gap between curves in Figure 10vs. χ_t values. Gaps are not constant.

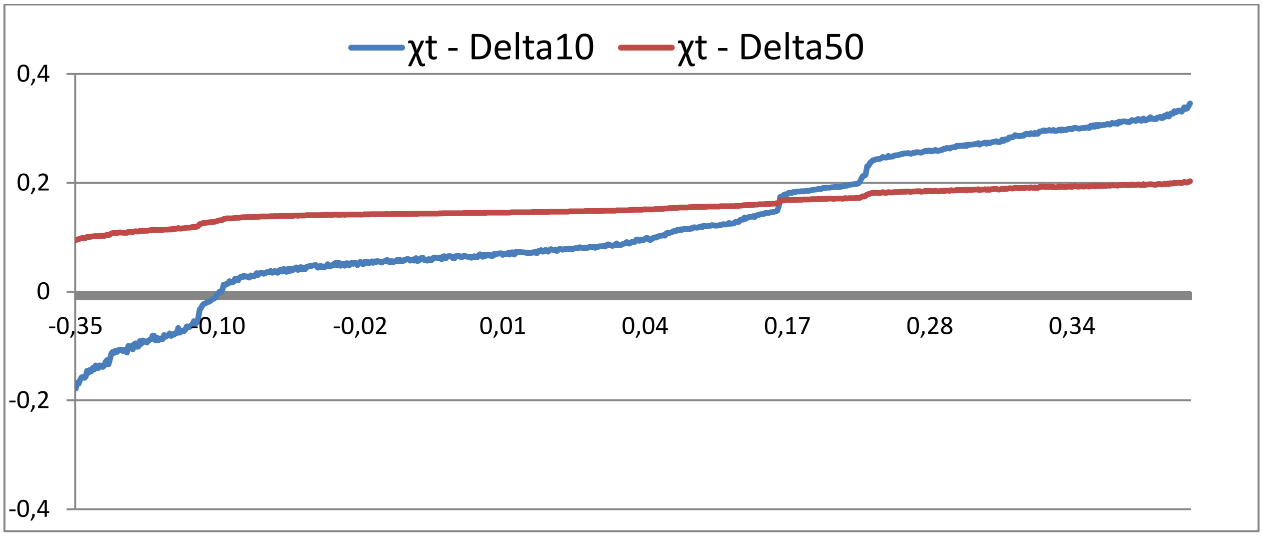

Nevertheless, as a rule of thumb, when the short-term variable has a highly positive value, the market is likely to be in backwardation, and for negative values of the short-term variable the market will be in contango. However, all of this depends on the value of the estimated parameters and the size of the short-term variable itself.

6.4.1 Does the Schwartz-Smith Model Assume Contango for Equilibrium Situations?

From Figure 10 we read that the market is in contango when the short-term variable is zero. We want to investigate whether this is a systematic feature of the Schwartz-Smith model, or if this relies on the market configuration.

In order to answer this question, we have to look at Eq. 50, which, if yielding a negative result, indicates contango. If the short-term variable is zero (market in “equilibrium condition”), then this solely relies on the term

$$A(\Delta t) = -\left(1 - e^{-\kappa\Delta t}\right)\frac{\lambda_\chi}{\kappa} + \frac{1 - e^{-2\kappa\Delta t}}{4\kappa}{\sigma_\chi}^2 + \left(1 - e^{-\kappa\Delta t}\right)\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} + (\mu_\xi^* + \frac{{\sigma_\xi}^2}{2})\Delta t$$

If $A(\Delta t)$ turns out positive, then we have contango. We see that for distant maturities, the result will heavily rely on the growing term $(\mu_\xi^* + \frac{{\sigma_\xi}^2}{2})$. If this sum is a positive number, then the situation way ahead In the future will be contango. Looking at our estimates in Table 3, we conclude that the sum actually is slightly negative, so for very distant maturities (one can maybe call them irrelevant?) we have backwardation. For shorter maturities the other estimated parameters have a say. Since both the correlation factor, the short-term risk premium and the risk-neutral drift term can yield negative values of $A(\Delta t)$, we have to conclude that whether or not the market is systematically in contango or backwardation relies on the market parameters, hence the model can give results in either direction.

6.4.2 Going from Contango to Backwardation

It is interesting to see what happens during the shift from contango to backwardation in 2011. During this period of transformation, as is possible to see in Figure 10 also, the fifty month price is in backwardation while the ten month price is in contango relative to the spot price. Time series plot of the observed data reveals this development.

Figure 12 Model constructed Brent time series.

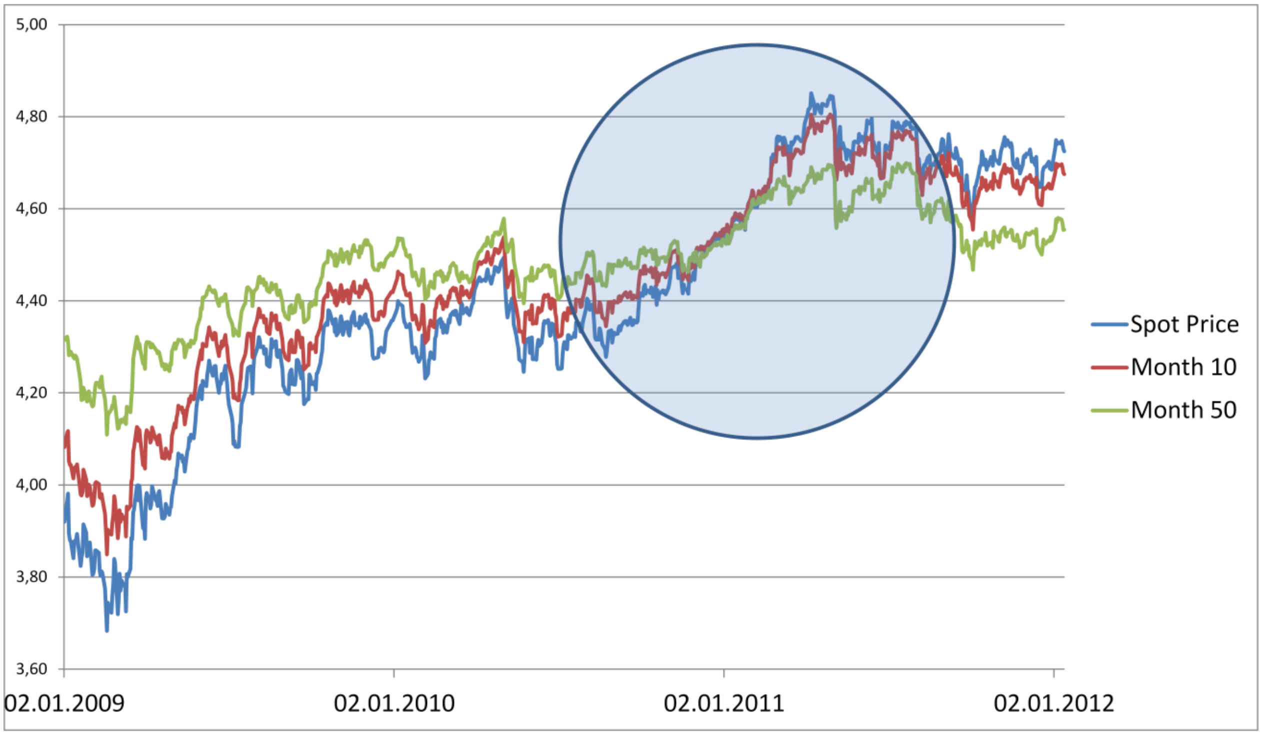

At first (when the whole forward curve is in contango) the fiftieth month price is above the tenth month price which again is above the first month price. During the period of transformation, this changes so that the order is reversed. But, during the transformation period, the ten month price remains in contango while the fifty month price enters backwardation. This event is captured also in the model, but we will use observed data (where the 1st month contract serves as a proxy for the spot price) to show that this happens in real life. As we see from Figure 13 the fifty month price sneaks below the first month price in Dec 2010, while the ten month price remains in contango until the middle of Feb 2011.

Figure 13 Time series showing how prices move relative to each other during transformation from contango to backwardation.

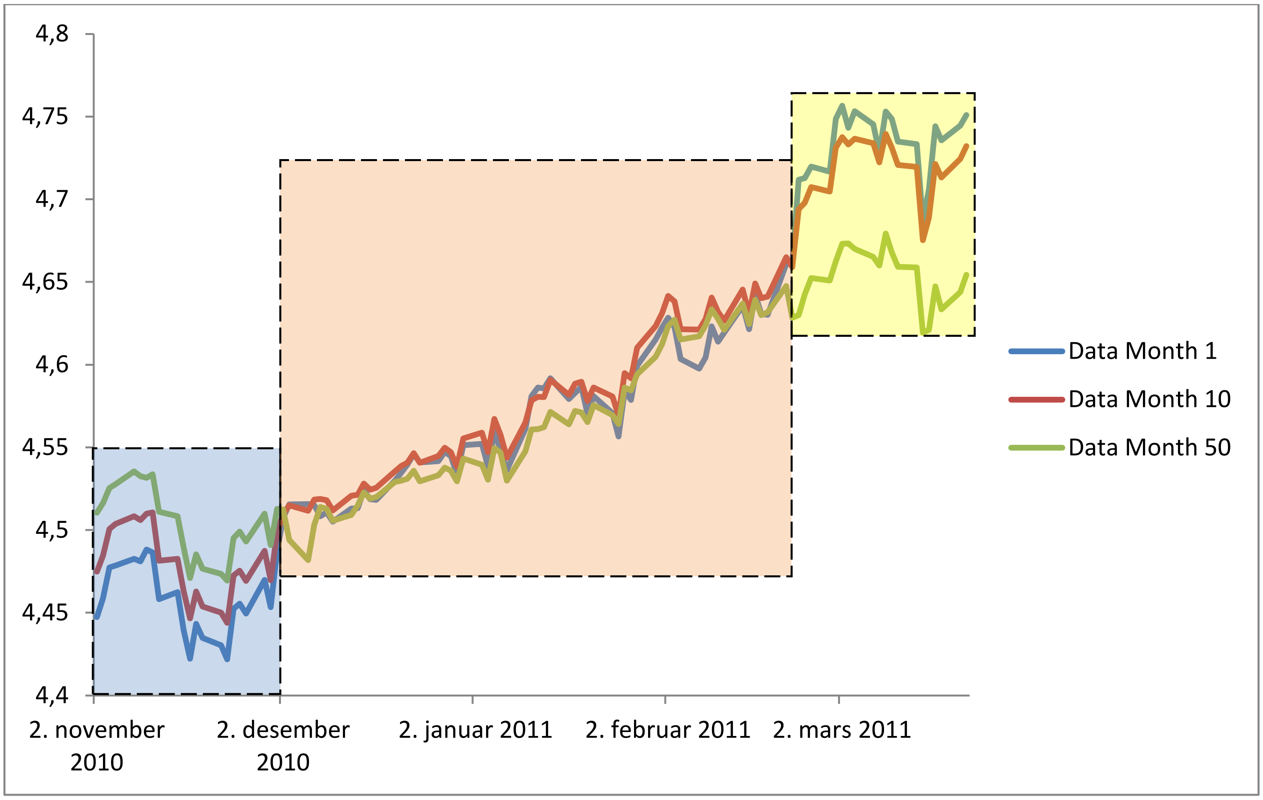

This results in three distinguished shapes of the forward curve, namely pure contango, the mixed state and pure backwardation, as shown below:

Figure 14 Graphs representing the three distinct market states in our sample. Results from the Brent market.

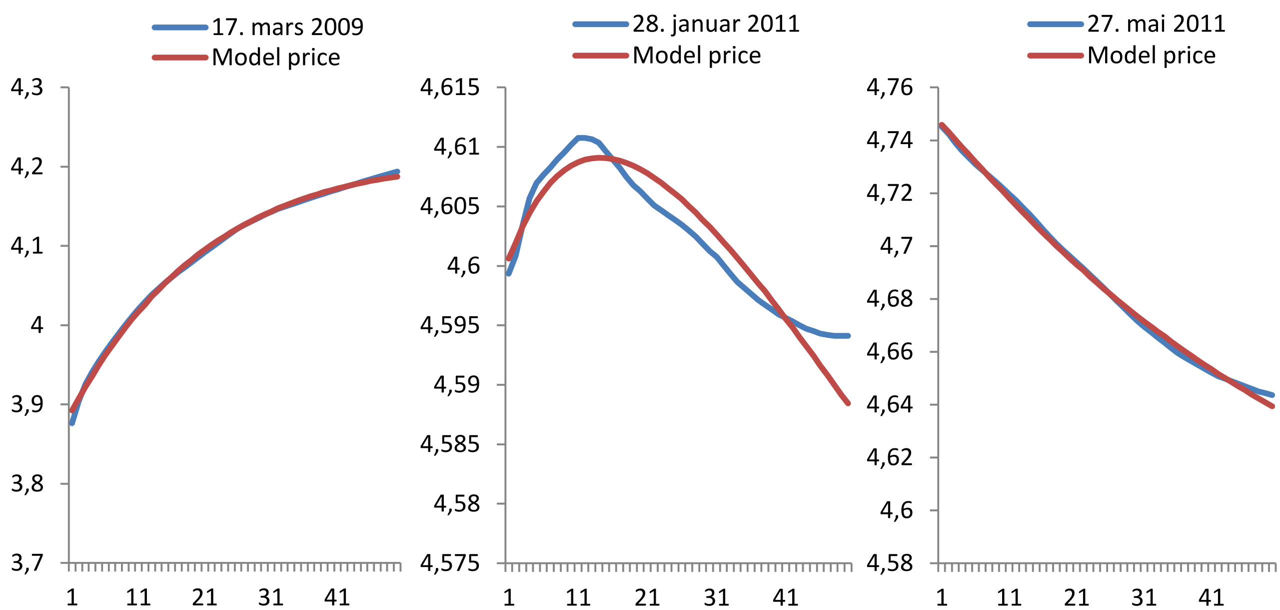

6.5 Obtained Parameter Estimates

The most interesting result from the model calibration is the parameter estimates. They give us information about market fundamentals, such as expected growth rate of equilibrium prices and the rate of mean reversion in the market. The parameters also provide information about the volatility of both short-term deviations and long-term prices, along with the correlation between changes in short-term and long-term variables. For a brief explanation on the meaning of the various parameters and state variables, see Table 2 on page 38.

Below we give the calibration summary for both Brent and WTI:

Table 3 Calibration summary, Schwartz-Smith model.

	Brent	WTI
Data sample	[784x50]	[765x50]
κ	0,4718	0,65584
λ_χ	-0,07652	-0,08126
μ_ξ^*	-0,04002	-0,03401
μ_ξ	0,05629	0,09251
σ_χ	0,23687	0,2773
σ_ξ	0,24222	0,25367
$\rho_{\chi\xi}$	0,19596	0,19918
$\rho\chi_B\chi_W$	0,84943	
$\rho\xi_B\xi_W$	0,96178	
$\rho\chi_B\xi_W$	0,25313	
$\rho\chi_W\xi_B$	0,19322	
SSE	0,95625	2,00027

The following constraints were added to the Solver add-in of Excel:

Table 4 Calibration constraints, Schwartz-Smith model.

	Lower limit	Upper limit
χ	-6	6
ξ	0	6
κ	1x10^-7	4
λ_χ	-4	4
μ_ξ^*	-4	4

Comparing the results in Table 3 with the estimates obtained by Schwartz and Smith (Table 2), we see that our estimates seem to be within a reasonable order of magnitude. Schwartz and Smith calibrated the model using NYMEX crude oil data stemming from 1990-1995; therefore results should be comparable to our NYMEX WTI data.[30]

However, there are differences between our estimates and the estimates for 1990-1995. Our estimate indicates a lower rate of mean reversion today, meaning price shocks have longer impact on oil prices. Also, the volatility of the equilibrium level has increased remarkably, from 14.5% in the early 90's to 25.4% for our dataset.

Today's market participants also seem to have slightly lower expectations about future price growth. Our estimate of risk-neutral drift parameter is -3.4%, while the true process drift parameter indicates a positive growth of 9.3%. The Schwartz-Smith estimates indicate a different attitude towards future growth; they obtained an estimate of risk-neutral drift at 1.15% while the true process drift is estimated at -1.12%. This yields a higher long-term risk premium today than twenty years ago.

From the estimates in Table 3 we read that the Brent and WTI markets are very similar, with WTI prices being a bit more volatile than Brent and also exhibiting faster mean reversion. The short-term risk premiums are close to each other, which can be assumed to be true in reality as well. If then the error in estimated short-term risk premiums is about the same for each product, then it is possible to make sense of comparisons of state variables between the products (ref. the above discussion about risk premium estimates).

The negative (risk-neutral) drift parameters imply that the equilibrium prices are expected by buyers to become lower in time, which shows that the buyers have lower expectations than the true process (which has positive drift terms). Although we should be careful in interpreting μ_ξ, we see from the time series of observed prices and estimated long-term variables that the real μ_ξ probably is positive. The negative risk-neutral drift, implying a positive risk premium, is then an indication of risk-averse buyers.

6.6 The Estimated Volatility Curve

From Eq. 33 we obtained an expression for the instantaneous volatility of the forward curve. As already discussed, the volatility of the forward curve will increase as maturity approaches. We can plot the calculated volatility using Eq. 33 versus the volatility estimated from the dataset.

In order to estimate the volatility using the dataset, we need to make the dataset *stationary*[31]. The prices we have recorded are assumed to be the result of a *non-stationary* process, as the expected value of next day's price oil price changes from day to day (ref. our discussion about Brownian motions). Today's change in oil price will heavily impact tomorrows expected oil price. This change of expected oil price gives us a new probability distribution each day (as the expected value is altered). Calculating the

[30] However, Schwartz and Smith only utilizes 259 observed sets of five futures prices, namely the 1st, 5th, 9th, 13th and 17th month prices. Hence, their estimates may be less reliable than our estimates based on 765 observations of 50 futures prices.

[31] For a brief explanation of stationary and non-stationary processes, see (Iordanova, 2007).

standard deviation of a data vector consisting of data points created using different probability distributions won't make sense.

Instead of finding the standard deviation of the recorded price vector, we use time scaled price increments to find the forward curve's instantaneous volatility. The probability distribution of price increments is the same for each day, with the process' drift term representing the mean, and a constant variance. Therefore, we calculated the dataset's volatility by first creating vectors of time-scaled price increments, one vector for each of the fifty maturity months, and then calculate the standard deviation of each vector. Mathematically, this can be represented as

$$\mathbf{\Phi_i}(j) = \frac{\ln(F_{i,j}) - \ln(F_{i,j-\Delta t})}{\Delta t} \quad \textbf{Eq. 52}$$

$$\widehat{\mathrm{vol}_i} = StDev(\mathbf{\Phi_i}) \quad \textbf{Eq. 53}$$

where $F_{i,j}$ is the future price at time *j* for contract maturing at month no. *i* and $F_{i,j-\Delta t}$ is the future price at time *j*-Δt for contract maturing at month no. *i*. $\mathbf{\Phi_i}$ is the vector of price increments for the contract maturing at month no. *i*.

By performing the operation for all i=1, 2, …, 50 months we can construct a graph of the estimated volatility based on fifty data points. We also graph the results of Eq. 33 based on estimated parameters from model calibration. Comparing the two graphs we can check whether the model calibration yields a volatility curve similar to the actual volatility curve. Below are plots of volatility curves for both Brent and WTI:

Figure 15 Volatility curve for Brent.

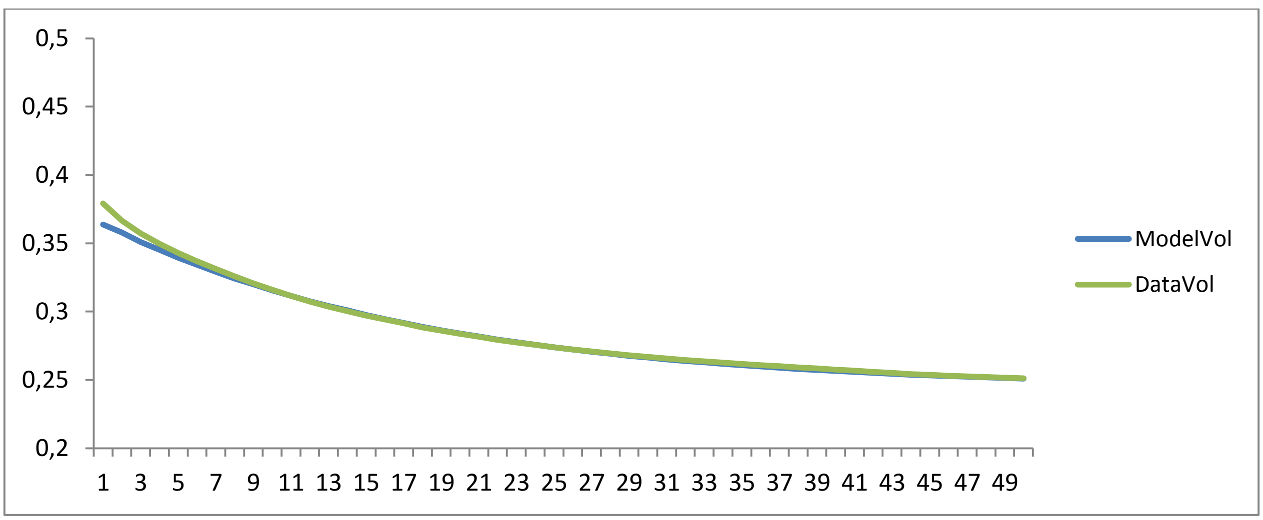

Figure 16 Volatility curve for WTI.

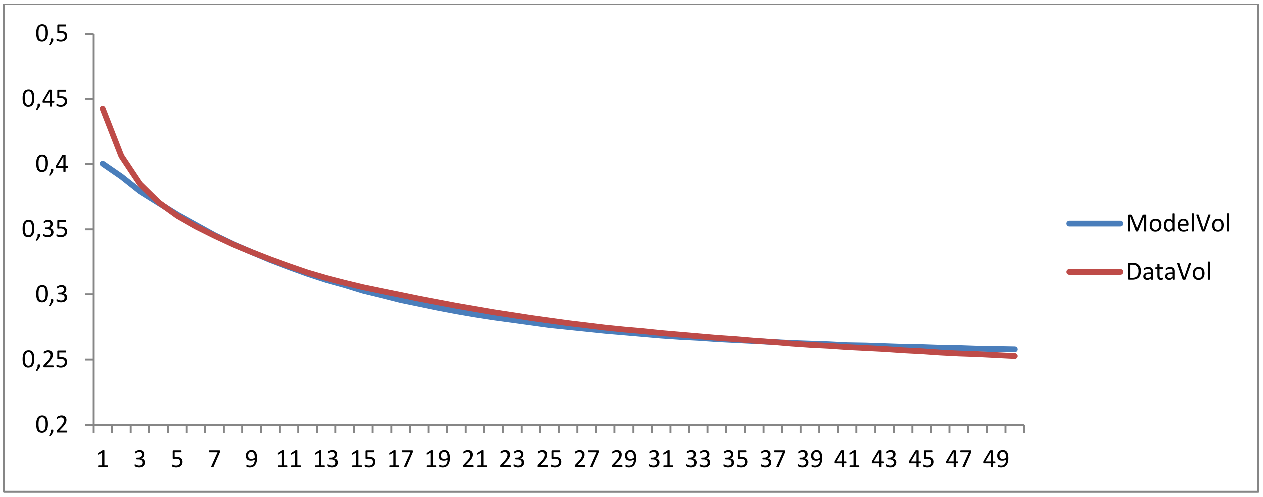

From looking at the volatility curves, we see that the volatility curve constructed by the model fits pretty well to the volatility curve implied by observed data. But for both products we have problems in the front end of the curve – the model gives too low volatility here. Especially for WTI, deviations between model and data implied volatilities are large.

We also note that WTI has larger volatility than Brent in the front end, and, while both tend to stabilize at around 0.25, WTI has slightly higher volatility also for long maturities. This corresponds to the estimated parameters given in Table 3.

Chapter 7 Simulating Using the Schwartz-Smith Model

When we have calibrated the Schwartz-Smith model, and by that have all the parameters estimated, we can use the model for simulating oil price development. We use Eq. 12 and Eq. 13 for simulating the development of the state variables. When state variables are obtained, we retrieve future prices using Eq. 32. Retrieving future prices from the state variables is a straightforward process.

The part of the simulation process that demand the most concern, is simulating the state variables' development. The drift terms of the development are self-explanatory, but the draw of random variables needs some attention. Our goal is to draw *correlated* random values z_χ and z_ξ from the standard normal distribution $\mathrm{N}(0,1)$. The technique used to draw a random value from the normal distribution is quite easy to understand. In the simulations we have just drawn a random number *u* from a uniform distribution between 0 and 1, and interpreted this as the percentile of the cumulative normal distribution. We then use the inverse normal distribution function $\mathrm{N}^{-1}(u)$ to determine which value from the normal distribution this percentile corresponds to. If the value from the normal distribution turns out to be positive (for u>0.5), the random term will drive the state variable upwards, and vice versa for u<0.5. If u = 0.5, it is like drawing 0 from the normal distribution and hence the random term becomes zero.

7.1 Drawing Correlated Random Variables

The greatest complication comes from the requirement that the random draws for the two state variables need being *correlated*. In order to achieve correlation between our variables, we make two independent draws from the standard normal distribution (as explained above), and then make one of the z's a perfect copy of the first independent value while the other z is a weighted combination of the two independent values. This way, the second z is partially dependent on the first z, and that is how correlation between the variables is ensured.

The technique we apply by doing this is called *Cholesky decomposition*, which can be expressed mathematically. For the case of two correlated variables, we make two uncorrelated draws ε_1 and ε_2 from the standard normal distribution. Then,

$$z_\xi = \varepsilon_1$$
$$z_\chi = \rho_{\chi\xi}\epsilon_1 + \varepsilon_2\sqrt{1-\rho_{\chi\xi}}$$

Eq. 54

This way, $\mathrm{Corr}(z_\xi, z_\chi) = \rho_{\chi\xi}$ and both z's are distributed $\mathrm{N}(0,1)$[32].

7.2 The Development of State Variables

In order to simulate time series of the state variables, we just draw the set of two correlated variables for each day we want to simulate, and insert these random draws into Eq. 12 and Eq. 13. By doing this we produce vectors containing the time series of state variables.

[32] Proof for this can be found in McDonald (pp. 643-644).

In order to create the time series of state variables, we have to select start values for the state variables. If you want to simulate from your last observed day and into the future, you can use the state variables estimated for that day as starting values. The development of state variables will then have its origin at the starting values, and deviate from the origin as time goes by. Which direction the state variables moves in is rather random, but in the long run there should be an average drift to the equilibrium price as given by $\mu_{\xi}\Delta t$ and the short-term deviations should revert around zero. Also, even though the development of the state variables is quite random, the amplitudes of the movements are restricted by the model parameters. Therefore, in the short run it is impossible to say anything about which direction the oil price will take, but we can say something about how *large* deviations that are likely to occur.

Note that we use the “true” process for simulating state variables, and the risk-neutral process as premise for valuing the futures contract. So, in “reality” the state variables develop according to the true process, while the valuation of futures contracts is based on risk-neutral expectations.

Chapter 8 A Model for Two Correlated Products

So far we have investigated the features of the Schwartz-Smith model, how to calibrate the model and how to use it for simulations. We have also discussed how to draw two correlated random variables to simulate the development of long-term and short-term state variables. The technique we used for making correlated draws, Cholesky decomposition, is not restricted to two random variables. We can use it for drawing many random variables. The idea of the following part of this thesis is to exploit this to develop a model for two correlated products.

From looking at our datasets, we see that Brent and WTI oil are highly correlated products. The price deviation between the two products is rarely large, but the size of the deviation varies with time. A model that can be used for joint simulations of Brent and WTI prices would be a useful tool in valuing cross-product derivatives. An example of an imaginary cross-product contract, the Spread Contract, is given in the end of this chapter.

The chapter starts with proposing a model describing the price development of two correlated products. We explain how the joint model can be calibrated and used for simulations. Specifically, we describe the required theoretical framework for making correlated random draws.

Calibration results for a joint development of Brent and WTI are presented and compared to previous individual calibrations. We also show the outcome of forecasting oil prices one year ahead, and explain this with reference to the underlying simulation parameters. Finally, we show how forecasting results can be used for speculating in the Spread Contract.

8.1 Model Proposal

We propose a model which, in discrete time, can be expressed as

$$s_t^B = \ln(S_t^B) = \chi_t^B + \xi_t^B \qquad \textbf{Eq. 55}$$

$$\chi_t^B = \chi_{t-\Delta t}^B e^{-\kappa^B \Delta t} + \sqrt{\frac{1 - e^{-2\kappa^B \Delta t}}{2\kappa^B}} \sigma_\chi^B z_\chi^B \qquad \textbf{Eq. 56}$$

$$\xi_t^B = \xi_{t-\Delta t}^B + \mu_\xi \Delta t + \sqrt{\Delta t} \sigma_\xi z_\xi \qquad \textbf{Eq. 57}$$

$$s_t^W = \ln(S_t^W) = \chi_t^W + \xi_t^W \qquad \textbf{Eq. 58}$$

$$\chi_t^W = \chi_{t-\Delta t}^W e^{-\kappa^W \Delta t} + \sqrt{\frac{1 - e^{-2\kappa^W \Delta t}}{2\kappa^W}} \sigma_\chi^W z_\chi^W \qquad \textbf{Eq. 59}$$

$$\xi_t^W = \xi_t^B + c_\xi \qquad \textbf{Eq. 60}$$

where z_ξ, z_χ^B, z_χ^W all are correlated draws from the standard normal distribution, and c_ξ is a constant meant to cover some average difference in the equilibrium price level for Brent and WTI.

When creating this joint model, we have made some simplifications. The most fundamental assumption of the model, stated in Eq. 60, is that the development of long-term state variables is perfectly correlated. Hence, in order to simulate long-term factors for Brent and WTI, we only need drawing one random variable z_ξ. The randomness in the difference between Brent and WTI prices is then expressed only through the development of short-term deviations. We need to make draws z_χ^B and z_χ^W for the respective short-term factors. This gives us a total of three correlated, random draws: z_ξ, z_χ^B and z_χ^W.

8.2 Drawing *n* Correlated Variables

In order to simulate using the joint model, we need drawing three correlated variables. We can expand and generalize the technique used for two variables, and get *n* correlated variables.

Recall from Eq. 54 that we first draw uncorrelated random $\mathrm{N}(0,1)$ variables. Let's call them $\varepsilon_1, \varepsilon_2, \dots, \varepsilon_n$. We need to turn these uncorrelated draws into correlated random variables Z(1), Z(2), …, Z(*n*) where E[Z(*i*)Z(*j*)] = $\rho_{i,j}$. We will turn to Cholesky decomposition to do this.

It turns out[33] that Z(i) can be calculated as

$$\mathrm{Z}(i) = \sum_{j=1}^{i} a_{i,j}\, \varepsilon_j \qquad \textbf{Eq. 61}$$

Z(*i*) is then depending on all the uncorrelated draws in the range ε_1 to ε_i. So for Z(1) only one random draw comes into play, as for Z(3) three random draws make a weighted impact on the value of the draw.

$a_{i,j}$ are the weights given to each uncorrelated draw in determining the value of the correlated value Z(*i*). The trickiest part, namely the operation termed Cholesky composition, is to calculate each of the weights associated with the Z's. The formula for $a_{i,j}$ is

$$\begin{aligned} a_{i,j} &= \frac{1}{a_{j,j}}\left[\rho_{i,j} - \sum_{k=1}^{j-1} a_{j,k} a_{i,k}\right] \qquad for\ i > j \\ a_{i,i} &= \sqrt{1 - \sum_{k=1}^{i-1} a_{i,k}^2} \end{aligned} \qquad \textbf{Eq. 62}$$

When we only have two correlated variables, the result is equal to Eq. 54.

So what we basically need is a matrix/collection of correlation factors[34], and use it as input to get the Cholesky decomposition matrix of $a_{i,j}$'s .

To get the values for the correlation matrix we need to calibrate the joint model on our dataset. We use the procedure described for the Schwartz-Smith model, with some small adjustments.

[33] See McDonald (p. 644).

[34] Generally, for *n* random variables we have $\frac{1}{2}n(n-1)$ pairwise correlation factors that need being accounted for.

8.3 Calibrating the Joint Model

When calibrating the joint model, we make the following modifications to the procedure applied for the standard Schwartz-Smith model:

- When using Solver to minimize the squared errors between the estimated forward curve and the observed data, for each day we minimize the combined SSE's for Brent and WTI. This means for each day *i* we have a value $SSE_{Brent+WTI,i} = SSE_{Brent,i} + SSE_{WTI,i}$ which we minimize by picking a long-term factor (which is "shared") and then a short-term factor for both Brent and WTI. This way the estimated long-term variable is a trade-off between fitting to both Brent and WTI observed data, and the short-term variables (that are estimated independently) are chosen so that the forward curve fits observed data best.
- When using Solver to estimate the parameters κ^B, λ_χ^B, κ^W, λ_χ^W and $\mu_\xi{}^*$ we minimize the SSE for the whole dataset. This SSE value can be represented as $\mathrm{SSE}_{\mathrm{Brent+WTI}} = \sum_i^N \mathrm{SSE}_{\mathrm{Brent+WTI,i}}$ where *N* is the total number of days for which we have observed data.

Calibrating our new model will give us long-term and short-term state variables for both Brent and WTI, and also associated parameters. A restriction is made in the calibration so that $\xi_t^W = \xi_t^B + c_\xi$. We also restrict the parameter $\mu_\xi{}^*$ to be the same for both products. It follows that the long-term variables have identical developments; hence all parameters exclusively associated with the long-term variable (μ_ξ, σ_ξ) will be equal for the two products. Thus we only need one set of parameters relating to the long-term variable. The difference between Brent and WTI prices is captured by the short-term variables and their associated parameters (κ, λ_χ, σ_χ, $\rho_{\chi\xi}$). We need unique sets of short-term parameters for each product.

We find estimates for the state variables and parameters κ, λ_χ and $\mu_\xi{}^*$ via Solver. The remaining parameters are estimated on basis of the already obtained variables and parameters. Expressions for these estimated parameters are given here:

$$\boldsymbol{\pi}^{\mathbf{B}} = \frac{\hat{\boldsymbol{\chi}}_i^B - \hat{\boldsymbol{\chi}}_{i-1}^B e^{-\hat{\kappa}^B \Delta t}}{\sqrt{\frac{1-e^{-2\hat{\kappa}^B t}}{2\hat{\kappa}^B}}} \quad \textbf{Eq. 63}$$

$$\hat{\sigma}_\chi^B = \sqrt{Var(\boldsymbol{\pi}^{\boldsymbol{B}})} \quad \textbf{Eq. 64}$$

$$\boldsymbol{\pi}^{\mathbf{W}} = \frac{\hat{\boldsymbol{\chi}}_i^W - \hat{\boldsymbol{\chi}}_{i-1}^W e^{-\hat{\kappa}^W \Delta t}}{\sqrt{\frac{1-e^{-2\hat{\kappa}^W t}}{2\hat{\kappa}^W}}} \quad \textbf{Eq. 65}$$

$$\hat{\sigma}_\chi^W = \sqrt{Var(\boldsymbol{\pi}^{\boldsymbol{W}})} \quad \textbf{Eq. 66}$$

$$\boldsymbol{\psi} = \frac{\hat{\boldsymbol{\xi}}_i - \hat{\boldsymbol{\xi}}_{i-1} - \hat{\mu}_{\xi,j-1}\Delta t}{\sqrt{\Delta t}} \quad \textbf{Eq. 67}$$

$$\hat{\sigma}_\xi = \sqrt{Var(\boldsymbol{\psi})} \quad \textbf{Eq. 68}$$

$$\hat{\mu}_\xi = \frac{1}{N}\sum_{i=1}^{N}\frac{\hat{\boldsymbol{\xi}}_i - \hat{\boldsymbol{\xi}}_{i-1}}{\Delta t} \qquad \textbf{Eq. 69}$$

$$\begin{aligned} \hat{\rho}_{\chi^B\xi} &= corr[\boldsymbol{\psi}, \boldsymbol{\pi}^{\boldsymbol{B}}] \\ \hat{\rho}_{\chi^W\xi} &= corr[\boldsymbol{\psi}, \boldsymbol{\pi}^{\boldsymbol{W}}] \\ \hat{\rho}_{\chi^B\chi^W} &= corr[\boldsymbol{\pi}^{\boldsymbol{B}}, \boldsymbol{\pi}^{\boldsymbol{W}}] \end{aligned} \qquad \textbf{Eq. 70}$$

The correlation factor $\hat{\rho}_{\chi^B\chi^W}$ is needed when calculating the Cholesky decomposition for simulation purposes, but isn't needed when valuing the futures contracts.

We now have all parameters necessary to make joint simulations of price development for Brent and WTI.

8.4 Expressions for Futures Prices

The formulas for calculating forward prices are identical to the expression for independent products:

$$\begin{aligned} \ln\left(F^B_{t+\Delta t,t}\right) &= \ln\left(E^*\left[S^B_{t+\Delta t}\right]\right) \\ &= E^*\left[\ln\left(S^B_{t+\Delta t}\right)\right] + \frac{1}{2}Var^*\left[\ln\left(S^B_{t+\Delta t}\right)\right] \\ &= \chi^B_t e^{-\kappa^B\Delta t} + \xi^B_t + A^B(\Delta t) \end{aligned} \qquad \textbf{Eq. 71}$$

where

$$\begin{aligned} A^B(\Delta t) = &-\left(1 - e^{-\kappa^B\Delta t}\right)\frac{\lambda^B_\chi}{\kappa^B} + \mu^*_\xi\Delta t \\ &+\frac{1}{2}\left\{\frac{1-e^{-2\kappa^B\Delta t}}{2\kappa^B}\sigma^{B\,2}_\chi + \sigma_\xi{}^2\Delta t \right. \\ &\left. +2(1 - e^{-\kappa^B\Delta t})\frac{\rho_{\chi^B\xi}\sigma^B_\chi\sigma_\xi}{\kappa^B}\right\} \end{aligned}$$

$$\begin{aligned} \ln\left(F^W_{t+\Delta t,t}\right) &= \ln\left(E^*\left[S^W_{t+\Delta t}\right]\right) \\ &= E^*\left[\ln\left(S^W_{t+\Delta t}\right)\right] + \frac{1}{2}Var^*\left[\ln\left(S^W_{t+\Delta t}\right)\right] \\ &= \chi^W_t e^{-\kappa^W\Delta t} + \xi^W_t + A^W(\Delta t) \end{aligned} \qquad \textbf{Eq. 72}$$

where

$$\begin{aligned} A^W(\Delta t) = &-\left(1 - e^{-\kappa^W\Delta t}\right)\frac{\lambda^W_\chi}{\kappa^W} + \mu^*_\xi\Delta t \\ &+\frac{1}{2}\left\{\frac{1-e^{-2\kappa^W\Delta t}}{2\kappa^W}\sigma^{W\,2}_\chi + \sigma_\xi{}^2\Delta t \right. \\ &\left. +2(1 - e^{-\kappa^W\Delta t})\frac{\rho_{\chi^W\xi}\sigma^W_\chi\sigma_\xi}{\kappa^W}\right\} \end{aligned}$$

If we let $\Delta t \to \infty$ we get the expressions for the long-term expected spot prices. Doing this yields

$$\ln\left(E\left[S^B_{t+\Delta t}\right]\right)|_{\Delta t\to\infty} = \xi^B_t + \frac{\sigma^{B\,2}_\chi}{4\kappa^B} + \frac{\rho_{\chi^B\xi}\sigma^B_\chi\sigma_\xi}{\kappa^B} + (\mu_\xi + \frac{\sigma^2_\xi}{2})\Delta t \qquad \textbf{Eq. 73}$$

$$\ln\left(E\left[S^W_{t+\Delta t}\right]\right)|_{\Delta t\to\infty} = \xi^W_t + \frac{\sigma^{W\,2}_\chi}{4\kappa^W} + \frac{\rho_{\chi^W\xi}\sigma^W_\chi\sigma_\xi}{\kappa^W} + (\mu_\xi + \frac{\sigma^2_\xi}{2})\Delta t \qquad \textbf{Eq. 74}$$

$$\ln\left(E^*\left[S^B_{t+\Delta t}\right]\right)|_{\Delta t\to\infty} = \xi^B_t + \frac{{\sigma^B_\chi}^2}{4\kappa^B} + \frac{\rho_{\chi^B\xi}\sigma^B_\chi\sigma_\xi}{\kappa^B} - \frac{\lambda^B_\chi}{\kappa^B} + (\mu^*_\xi + \frac{\sigma^2_\xi}{2})\Delta t \quad \textbf{Eq. 75}$$

$$\ln\left(E^*\left[S^W_{t+\Delta t}\right]\right)|_{\Delta t\to\infty} = \xi^W_t + \frac{{\sigma^W_\chi}^2}{4\kappa^W} + \frac{\rho_{\chi^W\xi}\sigma^W_\chi\sigma_\xi}{\kappa^W} - \frac{\lambda^W_\chi}{\kappa^W} + (\mu^*_\xi + \frac{\sigma^2_\xi}{2})\Delta t \quad \textbf{Eq. 76}$$

8.5 Results from Calibrating the Joint Model

From calibrating the joint model we get the following summary:

Table 5 Calibration summary, joint model (results from independent calibrations in stippled table).

	Brent	WTI	Brent (independent)	WTI (independent)
Data sample	[765x50]		[784x50]	[765x50]
κ	0,45484	0,55284	0,47180	0,65584
λ_χ	-0,06985	-0,11106	-0,07652	-0,08126
μ_ξ*	-0,03731	-0,03731	-0,04002	-0,03401
c_ξ		-0,07892		
μ_ξ	0,06264	0,06264	0,05629	0,09251
σ_χ	0,24196	0,28057	0,23687	0,27730
σ_ξ	0,24403	0,24403	0,24222	0,25367
$\rho_{\chi\xi}$	0,14824	0,19378	0,19596	0,19918
$\rho\chi_B\chi_W$	0,84827		0,84943	
$\rho\xi_B\xi_W$	1,00000		0,96178	
$\rho\chi_B\xi_W$	0,14824		0,25313	
$\rho\chi_W\xi_B$	0,19378		0,19322	
SSE	1,20520	2,30113	0,95625	2,00027
$SSE_{Brent+WTI}$	3,50633		2,95652	

Table 6 Calibration constraints, joint model.

	Lower limit	Upper limit
χ	-6	6
ξ	0	6
κ	1x10^-7	4
$\lambda\chi$	-4	4
μ_ξ*	-4	4
c_ξ	-4	4

Not surprisingly, we see that the SSE has increased for both products (compare Table 5 with Table 3). This is reasonable since we have added the constraints given by Eq. 60 and the shared μ_ξ*.

The estimated parameters are of the same order as for the independent calibrations. When comparing the plots of estimated state variables in Figure 17 to those of the independent calibrations (Figure 8 and Figure 9), we see that the Brent long term variable has about the same values for the joint and the independent model.

Figure 17 Estimated state variables of the joint model.

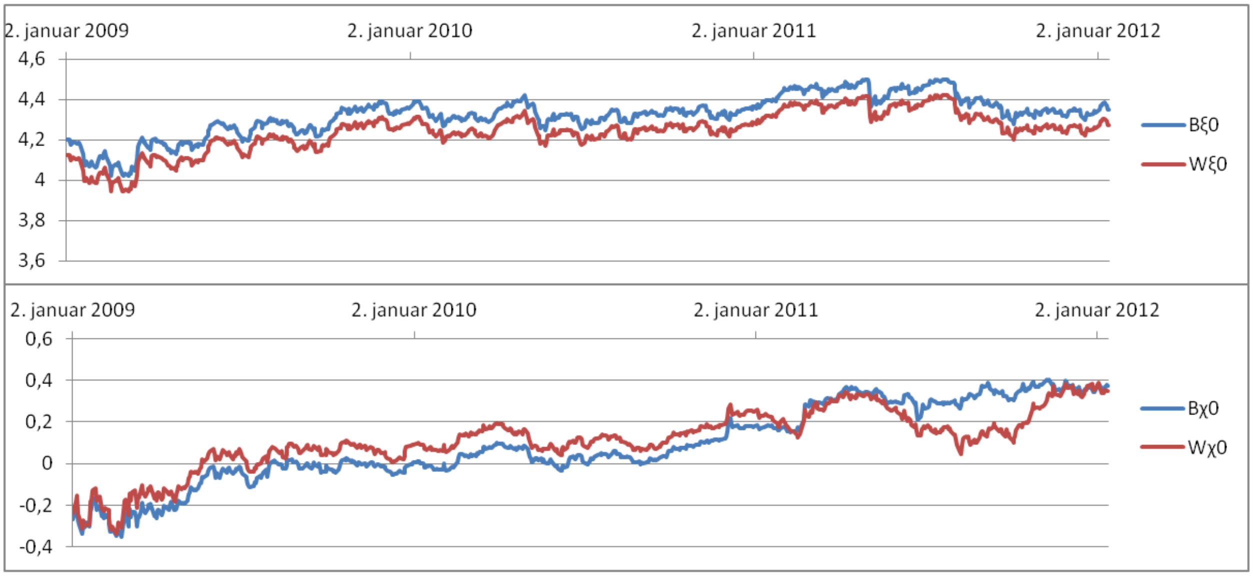

However, when taking a closer look it seems that the Brent variables are about the same while the WTI variables have changed; a lowering of the long-term variable, offset by an increase in the short-term variable, has occurred. This also corresponds to the fact that the kappa value (rate of mean reversion) for WTI has decreased; the short-term variable of WTI has taken on the role of a compensator in order to adjust for the lowered long-term variable. It has to do so because of the rather large (in absolute terms) value of c_ξ . The average deviation of Brent and WTI long-term variables from the independent calibrations is 0.021, however the estimated c_ξ value implies a constant gap of 0.079. Why this apparent illogical estimate of c_ξ appears hasn't been investigated more closely, but it is plausible to think of the reason being a redundancy of factors in the model. As the average gap from independent calibrations is very small, maybe the best thing would just to fix c_ξ as equal to zero (or some other constant like the average gap from independent calibrations).

Figure 18 Comparison of state variables estimated by independent and joint calibrations.

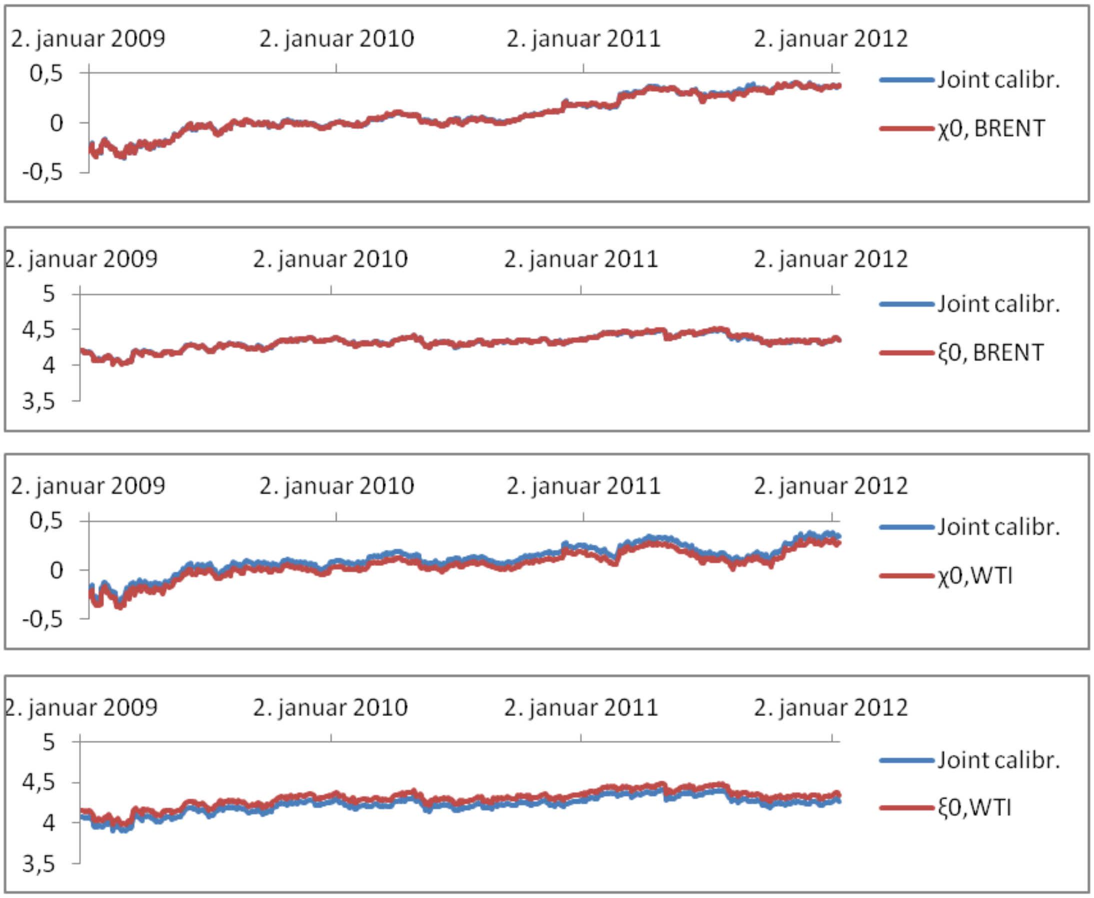

8.6 Simulating Using the Calibrated Joint Model

The estimated correlations between the products can be utilized for simulating. We perform simulations based on the estimated parameters given in Table 5, with the following simulation settings:

Table 7 Simulation settings, joint model.

NoOfDates	365	
NoOfSim	2400	
SimStartDate	13. January 2012	
StartValues	BRENT	WTI
χ_0	0,374883135	0,34840935
ξ_0	4,35129893	4,27237701

The starting values for the simulation are the estimated state variables for the 13th of January 2012. Unfortunately, due to memory restrictions on the computer used for simulations, it wasn't possible to do more than 2400 simulations in one operation. Doing more simulations would of course be possible on a better machine or with a more efficient simulation procedure.

The simulation gives a large amount of data, which needs to be presented in a tangible manner. We choose to show the maximum and minimum values from the simulation, along with the 5th, 50th (median) and 95th percentile. The average and expected prices (both "true" and risk-neutral) are also shown.

Figure 19 Simulated Brent spot prices using joint model. Parameters given in Table 5 and Table 7.

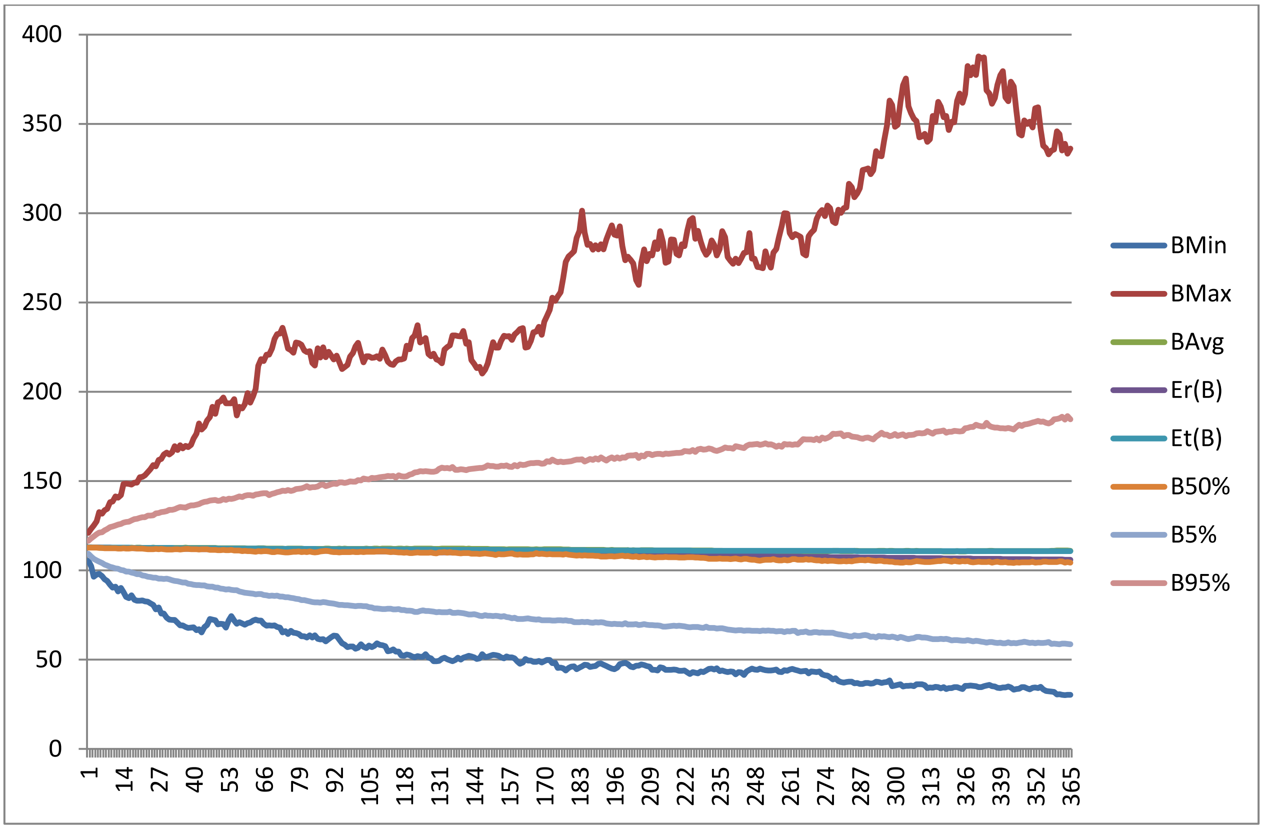

Above, the summary of simulated Brent spot prices is shown. As we see, prices are expected to be fairly stable. The 5th and 95th percentile gives us reasonable limits for how much the price can wander in each direction. The distance downwards from the median price to the 5th percentile price is smaller than the distance upwards to the 95th percentile price. The same holds for the min and max values. This corresponds to the fact that the prices aren't normally distributed, but log-normally distributed. We can also see this by just examining the expectation, median and average of the simulation.

From Figure 20 we note that there's a deviation between the average and the median. This is due to the log-normal distribution of the spot price. The average lies above the median since the distribution is skewed towards the higher values. So in most events, the spot price will be below the average, but in return upward deviations are larger than downward deviations.

Figure 20 Expected values along with simulated median and average Brent spot prices (as shown in Figure 19).

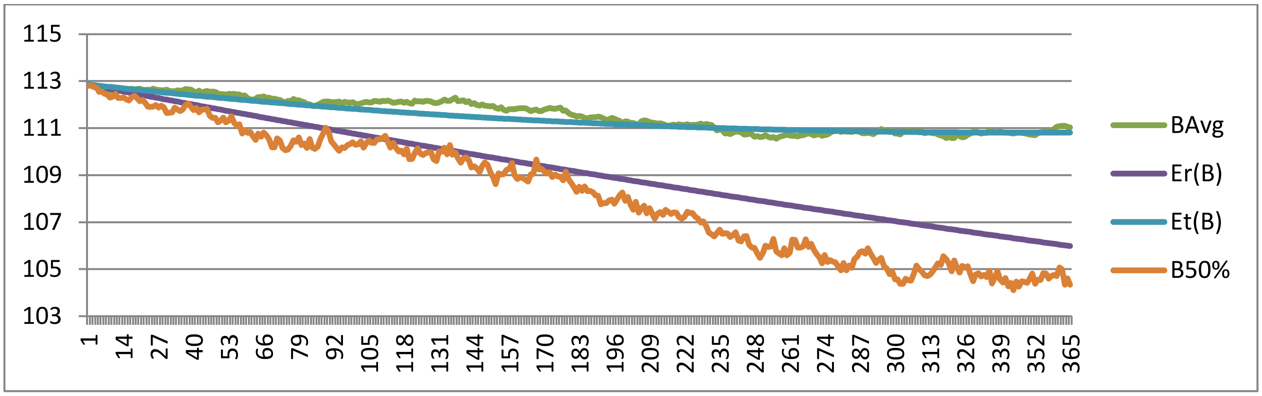

Expected values are calculated using Eq. 20 (Et = expected values of the true process) and Eq. 32 (Er = expected values according to risk-neutral measure).

The spot price is given by the sum of the short- and long-term variable. As the two variables "in reality" evolve according to the true process, we expect the spot price to follow the expected value of the true process (Et). The simulation results coincide with this assumption. The risk-neutral expected spot price (Er), which gives the forward curve, is lower than the average spot price. This is due to the risk-aversion of the buyers in the market.

From the drift parameter (μ_ξ) of Table 5 we would expect a positive growth of spot prices for the true process. However, as the simulation was started at positive values for the short-term deviations (Table 7), the phasing out of the latter results in an expected decrease in spot prices for the short term. However, as we see from the graph, at around a year after the start date it seems like the curve is about to turn into positive growth.

The results for WTI are, as expected, similar to the Brent results. Below is a plot of expected values, median and average for WTI from the simulation.

Figure 21 Expected values along with simulated median and average WTI spot prices.

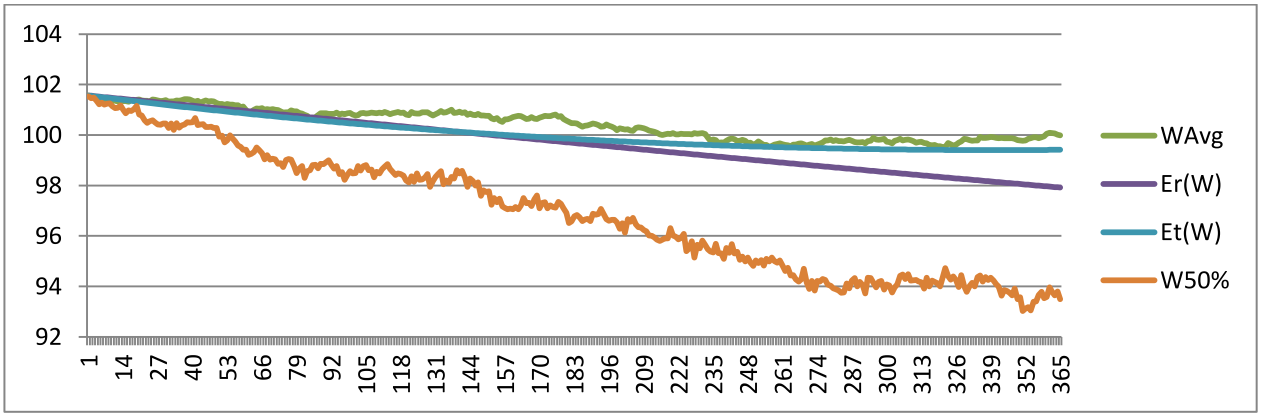

Due to the chosen start variables WTI starts lower than Brent, and the average stays below the Brent price for all dates. This is what we expect from historical data, which is the basis for the parameters of the simulation. The simulation yields the following summary plot of differences (delta) between the Brent and WTI spot prices:

Figure 22 Difference between simulated Brent and WTI prices.

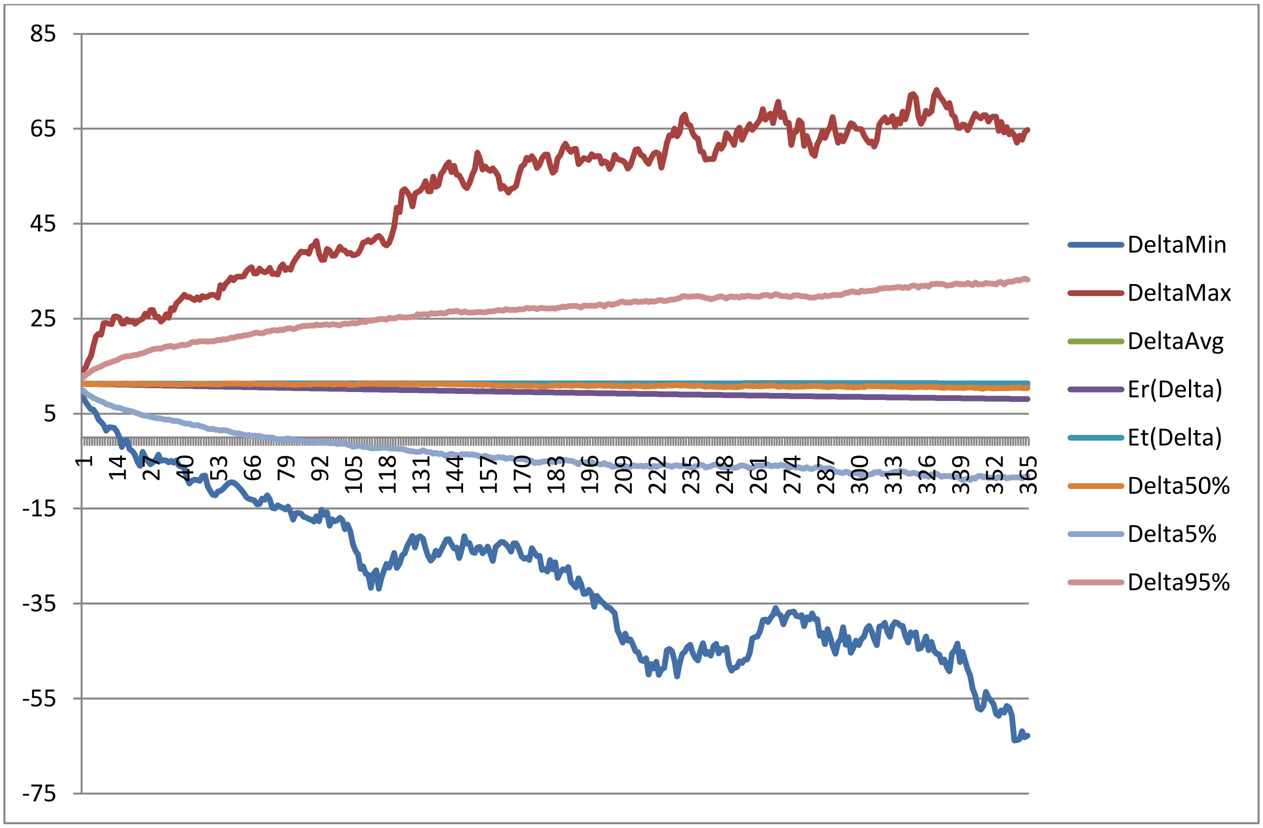

The plot tells us that the difference is expected to be positive (Brent prices higher than WTI), but the opposite may occur. “Zooming in” on the average, median and expected values we get a picture like this:

Figure 23 Average and median delta along with delta of expected values.

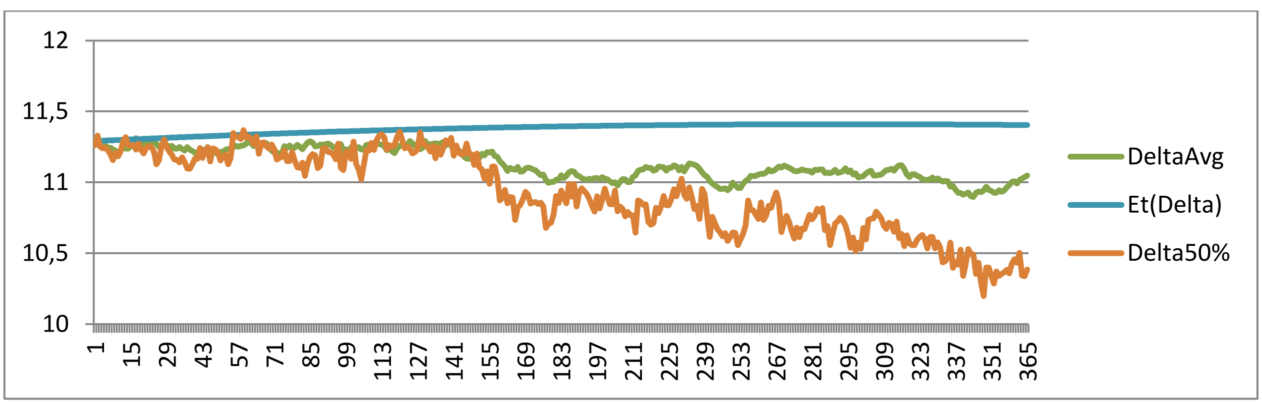

Applying the law that E(X+Y) = E(X) + E(Y), we calculate Et(Delta) as the difference between the expected prices of Brent and WTI. Simulated deltas are calculated as the simulated Brent price minus the WTI price from the same simulation. The expected difference (delta) is relatively stable just below $11.5 one year from the simulation start date. The average of simulated differences should follow the expected difference, which it does pretty well. We anticipate an even closer fit of the curves had we raised the number of simulations.

8.7 A Possible Application for the Joint Model: the Spread Contract

The strength and innovation of the joint model compared to the standard (one-product) Schwartz-Smith model is that we are able to simulate how Brent and WTI develops together. This makes it suited for estimating the expected price differences and also the uncertainty (variance) related to the spread. Below, we will give an example of a derivative which settles according to the Brent/WTI spread, and show how the joint model can serve as a decision basis for a person considering entering a position in the derivative contract.

Say there is a Spread Contract which settles according to the Brent/WTI spread at given dates. For example, let's say the contract settles on the 1st of each month, according to the spot price difference between Brent and WTI (spot price of Brent minus spot price of WTI). Let's further assume that each contract represents the value of 1,000 barrels of crude oil.

Imagine today is January 13th 2012 (the start date for the simulations above), and you want to make a bet about the Brent/WTI price differential on December 1st. There are 323 days between these two dates. Reading from Figure 23 we get an expected[35] spread at December 1st of $11.41, while the 5th and 95th percentile of the spread distribution is -$7.55 and $31.94, respectively.

Further, let's assume that the price at which this contract trades for December 1st is $11.20. This means that you, as a speculator with your calibrated model, regard the contract as underpriced. Therefore, you want to buy 100 contracts (equaling 100,000 barrels). You might end up buying the contract from another speculator who wants to gamble on the spread going down. When you buy the contract, you agree to pay the seller $11.20 per barrel underlying the contract, while you receive the spread occurring at December 1st.

With the information you possess, you expect to make a profit of ($11.41-$11.20) x 100,000 = $21,000. However, there is a 5% probability that your profit will be (-$7.55-$11.20) x 100,000 = -$1,875,000. In return, the upper 95th percentile suggests there is a 5% probability that you may earn ($31.94-$11.20) x 100,000 = $2,074,000.

Thus, there is a possibility of losing money on the deal, but you are expected to earn $21,000 and there is also a possibility of much greater gains.

---

[35] We use the true process since there is no obvious reason to assume that the risk premium should "benefit" one crude price at the expense of the other. In other words, speculators might as well speculate on the spread going up as speculating on it going down; it's not easy to define a hedger and a risk taker in this trade.

Chapter 9 Introducing a Second Short-Term Variable

When modeling price development, we desire to explain great amounts of data by a small number of factors. Herein lies a trade-off between how precisely we want to model the data and how many factors we will allow. We are searching for a model with as few as possible factors that still explains our dataset to a satisfying degree.

Our desire to explain as much as possible of observed data is a reason why we would want to include another factor in our model. Another reason is a desire to get a variable designated to represent the deviations between two products in the joint Brent/WTI model we have developed.

In the model proposed in Chapter 8, the values of the short-term factor are representing the difference between Brent and WTI. But the short-term factors of Brent and WTI are highly correlated. We would like there to be a factor with a small degree of correlation between Brent and WTI, such that this factor solely would represent deviations between the two products.

In this chapter we propose a three-factor model for two correlated products. We show how it can be calibrated and used for simulations, and present results from both calibration and simulation. Differences compared to the results for the two-factor joint model are commented on. For instance, the three-factor model succeeds in explaining more of the data variations. The results also show that the three-factor model is able to isolate the deviations between Brent and WTI by assigning them to the second short-term variable.

101 re-estimations of the model parameters are performed in order to get a measure of the uncertainty related to parameter estimates of the joint three-factor model.

Different assumptions underlying the two-factor and the three-factor models results in differences in the forecasting of prices. Differences specifically materialize for the development of WTI. These differences are discussed and explored in great detailed. A preliminary conclusion is that the three-factor model's assumption of identical equilibrium prices for Brent and WTI is the most reasonable.

As we try to forecast prices into the distant future, the predictions become more and more uncertain. Also, a feature of the proposed models is an expected Brent/WTI price gap that widens as the expected price level increases (which it does due to the drift rate of equilibrium prices). We discuss limitations to the forecasting horizon of the model, and conclude that this widening gap is no problem for reasonable forecasting horizons. However, the reason for the widening gap is described in detail and a solution to mend the problem is proposed.

9.1 Proposing a Three-Factor Model

Motivated by the desires to better describe observed data and isolate the differences between Brent and WTI, we introduce a second short-term factor to the standard Schwartz-Smith model. As a result, we now have the following three processes:

$$\mathrm{d}\chi_t = -\kappa\chi_t \mathrm{dt} + \sigma_\chi dz_\chi \qquad \textbf{Eq. 77}$$

$$d\chi_{2,t} = -\kappa_2 \chi_{2,t} dt + \sigma_{\chi_2} dz_{\chi_2} \qquad \textbf{Eq. 78}$$

$$d\xi_t = \mu_\xi dt + \sigma_\xi dz_\xi \qquad \textbf{Eq. 79}$$

The log spot price is then given by

$$s_t = \ln(S_t) = \chi_t + \chi_{2,t} + \xi_t \qquad \textbf{Eq. 80}$$

The time-discretisation of the second short-term factor will be equivalent to that of the first factor, and therefore for the expected spot prices we just have "double sets" of the terms for short-term factors. Thus, we have

$$E[\ln(S_t)] = [\chi_{t-\Delta t} e^{-\kappa\Delta t} + \chi_{2,t-\Delta t} e^{-\kappa_2 \Delta t} + \xi_{t-\Delta t} + \mu_\xi \Delta t] \qquad \textbf{Eq. 81}$$

$$\begin{aligned} E^*[\ln(S_t)] = [&\chi_{t-\Delta t} e^{-\kappa\Delta t} + \chi_{2,t-\Delta t} e^{-\kappa_2 \Delta t} + \xi_{t-\Delta t} - \left(1 - e^{-\kappa\Delta t}\right)\frac{\lambda_\chi}{\kappa} \\ &- \left(1 - e^{-\kappa_2\Delta t}\right)\frac{\lambda_{\chi_2}}{\kappa_2} + \mu_\xi^* \Delta t] \end{aligned} \qquad \textbf{Eq. 82}$$

The variance of the future spot prices needs some more caution, as we now have introduced two more pairwise correlations. The total variance of $\ln(S_t)$ becomes

$$\begin{aligned} \mathrm{Var}^*[\ln(S_t)] &= Var(\chi_t) + Var(\chi_{2,t}) + Var(\xi_t) + 2Cov(\chi_t, \chi_{2,t}) \\ &\quad + 2Cov(\chi_t, \xi_t) + 2Cov(\chi_{2,t}, \xi_t) \\ &= \frac{1 - e^{-2\kappa\Delta t}}{2\kappa}\sigma_\chi^{\ 2} + \frac{1 - e^{-2\kappa_2\Delta t}}{2\kappa_2}\sigma_{\chi_2}^{\ 2} + \sigma_\xi^{\ 2}\Delta t \\ &\quad + 2\left(1 - e^{-\kappa\Delta t}\right)\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} + 2\left(1 - e^{-\kappa_2\Delta t}\right)\frac{\rho_{\chi_2\xi}\sigma_{\chi_2}\sigma_\xi}{\kappa_2} \\ &\quad + 2\sqrt{(1 - e^{-2\kappa\Delta t})\frac{\sigma_\chi^2}{2\kappa}}\sqrt{(1 - e^{-2\kappa_2\Delta t})\frac{\sigma_{\chi_2}^2}{2\kappa_2}}\rho_{\chi_2\chi} \end{aligned} \qquad \textbf{Eq. 83}$$

where

$$\begin{aligned} Cov(\chi_t, \chi_{2,t}) &= \sqrt{Var(\chi_t)}\sqrt{Var(\chi_{2,t})}\rho_{\chi_2\chi} \\ &= \sqrt{(1 - e^{-2\kappa\Delta t})\frac{\sigma_\chi^2}{2\kappa}}\sqrt{(1 - e^{-2\kappa_2\Delta t})\frac{\sigma_{\chi_2}^2}{2\kappa_2}}\rho_{\chi_2\chi} \end{aligned} \qquad \textbf{Eq. 84}$$

From this we can solve for the futures price:

$$\begin{aligned} \ln\left(F_{t+\Delta t,t}\right) &= \ln(E^*[S_{t+\Delta t}]) \\ &= E^*[\ln(S_{t+\Delta t})] + \frac{1}{2}Var^*[\ln(S_{t+\Delta t})] \\ &= \chi_{t-\Delta t} e^{-\kappa\Delta t} + \chi_{2,t-\Delta t} e^{-\kappa_2\Delta t} + \xi_{t-\Delta t} + A(\Delta t) \end{aligned} \qquad \textbf{Eq. 85}$$

where

$$A(\Delta t) = -\left(1 - e^{-\kappa\Delta t}\right)\frac{\lambda_\chi}{\kappa} - \left(1 - e^{-\kappa_2\Delta t}\right)\frac{\lambda_{\chi_2}}{\kappa_2} + \mu_\xi^*\Delta t + \frac{1}{2}\{\frac{1-e^{-2\kappa\Delta t}}{2\kappa}{\sigma_\chi}^2 + \frac{1-e^{-2\kappa_2\Delta t}}{2\kappa_2}{\sigma_{\chi_2}}^2 + {\sigma_\xi}^2\Delta t + 2\left(1 - e^{-\kappa\Delta t}\right)\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} + 2\left(1 - e^{-\kappa_2\Delta t}\right)\frac{\rho_{\chi_2\xi}\sigma_{\chi_2}\sigma_\xi}{\kappa_2} + 2\sqrt{\left(1 - e^{-2\kappa\Delta t}\right)\frac{\sigma_\chi^2}{2\kappa}}\sqrt{\left(1 - e^{-2\kappa_2\Delta t}\right)\frac{\sigma_{\chi_2}^2}{2\kappa_2}}\rho_{\chi_2\chi}\}$$

The expression for the true process expected future prices is:

$$\ln(E[S_{t+\Delta t}]) = E[\ln(S_{t+\Delta t})] + \frac{1}{2}Var[\ln(S_{t+\Delta t})] = \chi_{t-\Delta t}e^{-\kappa\Delta t} + \chi_{2,t-\Delta t}e^{-\kappa_2\Delta t} + \xi_{t-\Delta t} + A(\Delta t) \qquad \textbf{Eq. 86}$$

where

$$A(\Delta t) = \mu_\xi\Delta t + \frac{1}{2}\{\frac{1-e^{-2\kappa\Delta t}}{2\kappa}{\sigma_\chi}^2 + \frac{1-e^{-2\kappa_2\Delta t}}{2\kappa_2}{\sigma_{\chi_2}}^2 + {\sigma_\xi}^2\Delta t + 2\left(1 - e^{-\kappa\Delta t}\right)\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} + 2\left(1 - e^{-\kappa_2\Delta t}\right)\frac{\rho_{\chi_2\xi}\sigma_{\chi_2}\sigma_\xi}{\kappa_2} + 2\sqrt{\left(1 - e^{-2\kappa\Delta t}\right)\frac{\sigma_\chi^2}{2\kappa}}\sqrt{\left(1 - e^{-2\kappa_2\Delta t}\right)\frac{\sigma_{\chi_2}^2}{2\kappa_2}}\rho_{\chi_2\chi}\}$$

One of the two reasons we desired this three-factor model, was the ability to let one of the short-term factors explain the deviations between Brent and WTI. We do this by adding a restriction to the first short-term factor, so that

$$\chi_t^W = \chi_t^B + c_\chi \qquad \textbf{Eq. 87}$$

We also restrict the parameters κ and λ_χ to be the same for both products. The long-term variables are also tied to each other as for the two-factor joint model. This way, the only difference between the two products is the second short-term variable $\chi_{2,t}$ and its associated parameters (κ_2, $\lambda_{\chi 2}$, $\sigma_{\chi 2}$, $\rho_{\chi 2\xi}$, $\rho_{\chi 2\chi}$).

When performing simulations with this joint three-factor model, we need making four correlated, random draws, and we need a correlation matrix consisting of $\frac{1}{2}n(n-1) = 6$ pairwise correlations.

9.2 Calibrating the Joint Three-Factor Model

Before doing simulations using the three-factor model, we need to calibrate it. We calibrate the joint model where the long-term variable and one of the short-term variables of Brent and WTI are tied together. The model is calibrated the same way as the two-factor model, but we choose to fix both c_ξ and c_χ at zero. This is done to reduce the number of variables and avoid possible redundancy among factors. Setting the constants equal to zero also ensures that all of the difference between Brent and WTI is covered by the second short-term variable and its associated parameters.

We use the same data set as for the joint two-factor model. As guess values[36] for the parameters we choose the estimated parameters of the two-factor model. A complete overview of the guess values are given in Table 8 below:

Table 8 Guess values for calibrating the joint three-factor model.

	B	W
κ	0,45484	0,45484
κ_2	1	5
λ_χ	-0,0699	-0,0699
$\lambda_{\chi 2}$	0	0
μ_ξ^*	-0,0373	-0,0373
c_ξ		0
c_χ		0
μ_ξ	0,06264	0,06264
σ_χ	0,24196	0,24196
$\sigma_{\chi 2}$	0	0
σ_ξ	0,24403	0,24403
$\rho_{\chi 2\chi}$	0	0
$\rho_{\chi 1\xi}$	0,14824	0,14824
$\rho_{\chi 2\xi}$	0	0
$\rho_{\chi 2B\chi 2W}$	0	

The calibration yields the following parameter estimates (results for two-factor joint model in stippled table):

[36] See paragraph 1, page 11.

Table 9 Calibration summary, joint three-factor model (results for two-factor joint model in stippled table).

	B	W
κ	0,54359	0,54359
κ_2	0,35560	6,71655
λ_χ	-0,05259	-0,05259
$\lambda_{\chi 2}$	-0,01147	0,00081
μ_ξ^*	-0,03344	-0,03344
c_ξ		0
c_χ		0
μ_ξ	0,06541	0,06541
σ_χ	0,24214	0,24214
$\sigma_{\chi 2}$	0,12721	0,27229
σ_ξ	0,24506	0,24506
$\rho_{\chi 2\chi}$	-0,27527	0,25526
$\rho_{\chi 1\xi}$	0,29146	0,29146
$\rho_{\chi 2\xi}$	-0,16042	-0,10942
$\rho_{\chi 2B\chi 2W}$	0,14192	
SSE	0,99605	0,70500
$SSE_{BRENT+WTI}$	1,70105	

	B	W
κ	0,45484	0,55284
λ_χ	-0,06985	-0,11106
μ_ξ^*	-0,03731	-0,03731
c_ξ		-0,07892
μ_ξ	0,06264	0,06264
σ_χ	0,24196	0,28057
σ_ξ	0,24403	0,24403
$\rho_{\chi\xi}$	0,14824	0,19378
$\rho\chi_B\chi_W$	0,84827	
SSE	1,2052	2,30113
$SSE_{BRENT+WTI}$	3,50633	

Table 10 Calibration constraints, joint model.

	Lower limit	Upper limit
χ	-6	6
X_2	-6	6
ξ	0	6
κ	1E-07	4
κ_2	1E-07	20
λ_χ	-4	4
$\lambda_{\chi 2}$	-4	4
μ_ξ^*	-4	4

From reading the SSE values, we see that the joint three-factor model explains the variations in the data set better than the joint two-factor model does. Therefore, it might provide a good alternative to the two-factor model, although it is more complicated and involving more factors. This makes it a bit less intuitive than the two-factor model.

If we compare the SSE values of the joint three-factor model to the independent calibrations of the Schwartz-Smith model (Table 3), we perceive that the SSE value of Brent actually has increased slightly for the joint three-factor model (from 0.956 to 0.996). However, this undesired effect is more than offset by the fact that SSE for WTI has been more than halved (going down from 2.000 to 0.705). This adjustment to better fit WTI data can be seen from plotting time series of observed 1st month prices together with 1st month prices implied by the model. Comparing Figure 24 to Figure 5 and Figure 6 we see that the fit for Brent is about the same, while the fit for WTI has become much better. Particularly, the erratic short-term movements of WTI in the beginning of 2009 have been captured by the very short-lived second short-term variable.

Figure 24 Plots of observed 1st month prices vs. model implied prices (log scale).

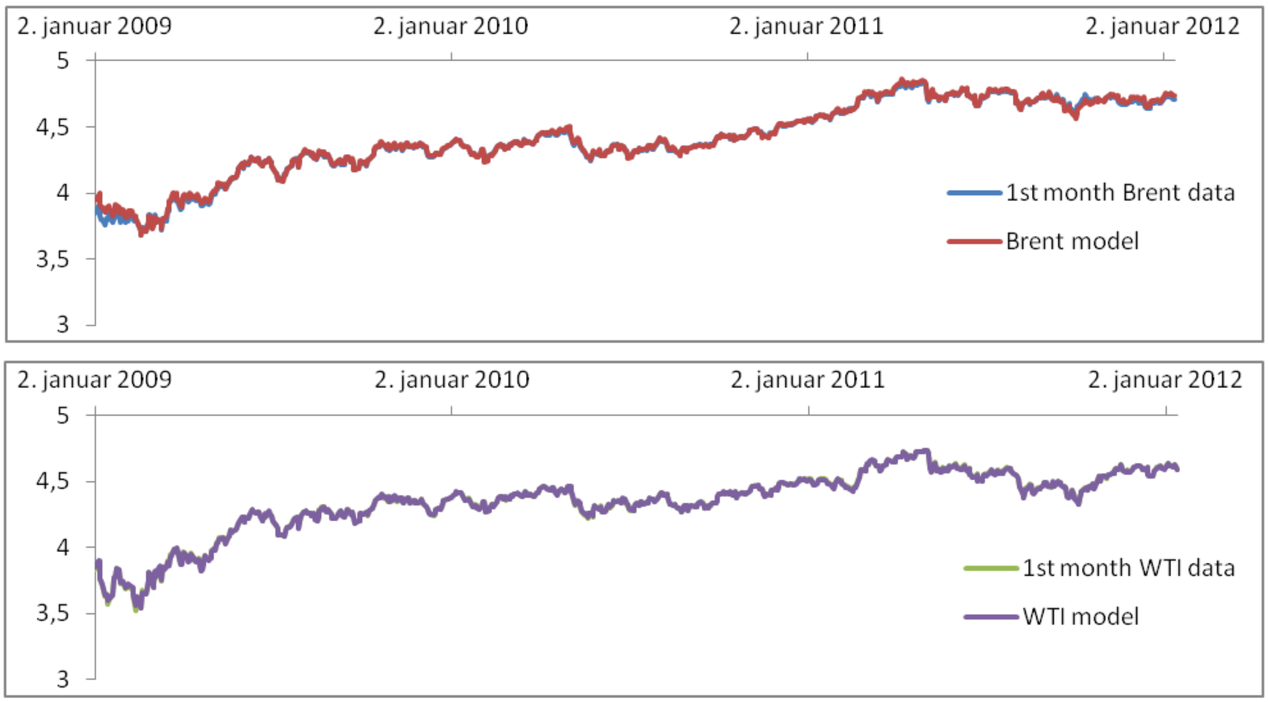

The correlation factor between the second short-term variables of Brent and WTI is estimated at 0.14. This is satisfyingly close to zero, and indicates an almost negligible correlation between the second short-term variables. It appears that we have managed to isolate a variable which explains the deviations between Brent and WTI, while the long-term and the first short-term variable represent the joint price development process.

Plots of the estimate state variables show that the second short-term variable takes on values close to zero when the prices of the two crudes are close to each other. During periods of deviations, like in 2011, the variable takes on different values to explain the price gap.

Figure 25 Estimated state variables of the joint three-factor model.

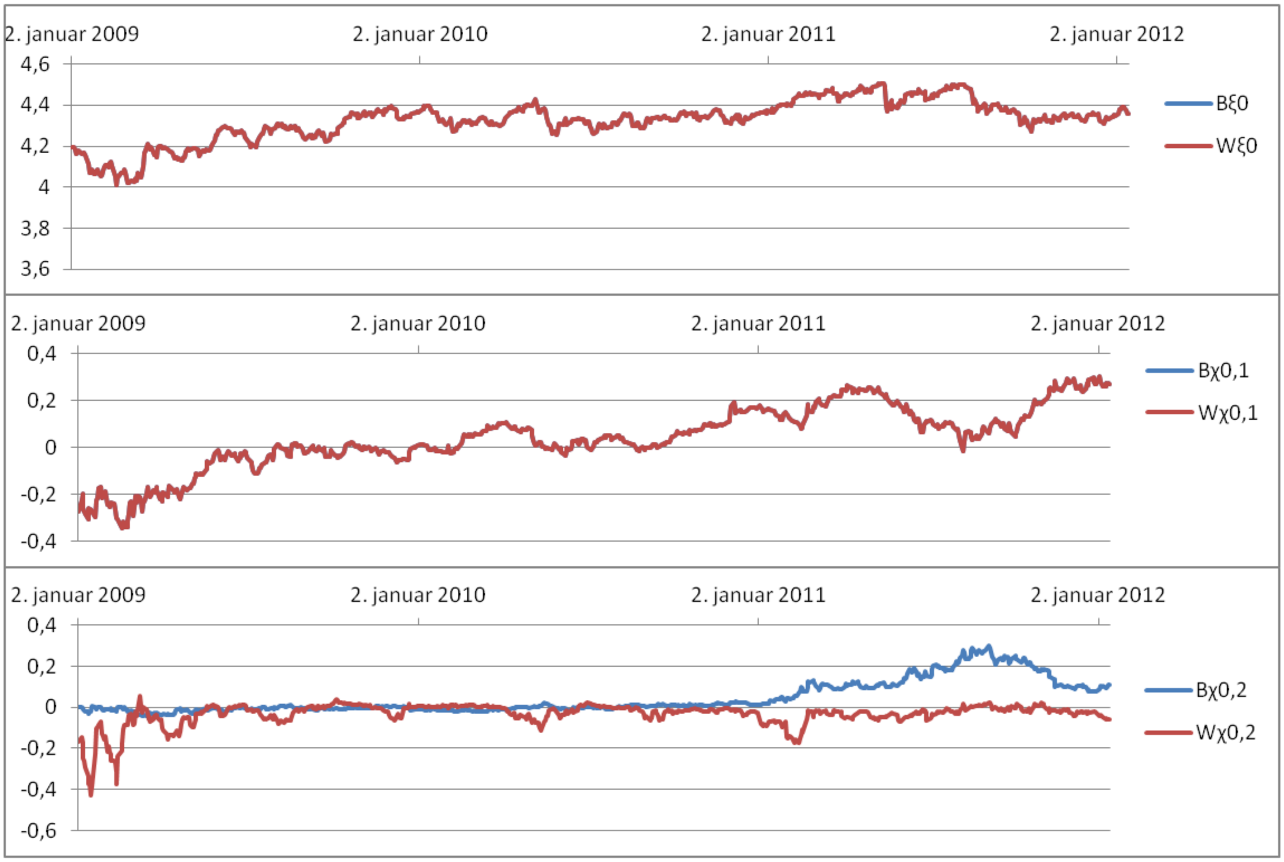

Comparing the estimated state variables with those of the Schwartz-Smith model and the joint two-factor model, we see that the development of the long-term and first short-term variable looks very much alike. However, there are of course some differences as we have restricted both the long-term and first short-term variable to be equal for Brent and WTI. All price deviations are assigned to the second short-term variable.

The restriction that c_ξ shall be equal to zero has made the long-term variable of WTI turn out a bit higher for the joint three-factor calibration than for the independent calibrations. The long-term variable of Brent seems to be estimated at almost exactly the same values for the joint three-factor calibration as for the independent case.

The slightly higher estimated values of the WTI long-term variable results in a corresponding decrease for the short-term variable. Apart from this, the estimates of the WTI short-term variable very well resemble the estimates from the independent calibrations. As the first short-term variable is restricted to be equal for Brent and WTI, this means that the short-term variable of Brent deviates from the independent calibrations during 2011 when the short-term prices of Brent and WTI deviate much. The deviations between Brent and WTI are captured by the second short-term variable, which during 2011 takes on positive values for Brent and remains close to zero for WTI.

From the parameter estimates, the settling of the kappa value close to the value for the independent WTI calibration corresponds to the observation that the first short-term variable adhere to the independent WTI values. The second kappa parameter is very different for the two products, with Brent having quite similar reversion rate for the first and second short-term variable (it's almost like we could merge the two variables into one by addition). Thus, as the second short-term variable is inactive for during most of the assessed time period, it comes into play during 2011 covering the gap between Brent and WTI prices. The WTI second short-term variable, on the other hand, has very fast mean reversion and only very temporary price deviations that the first two factors aren't able to capture.

Figure 26 Comparison of state variables estimated by independent and joint three-factor calibrations.

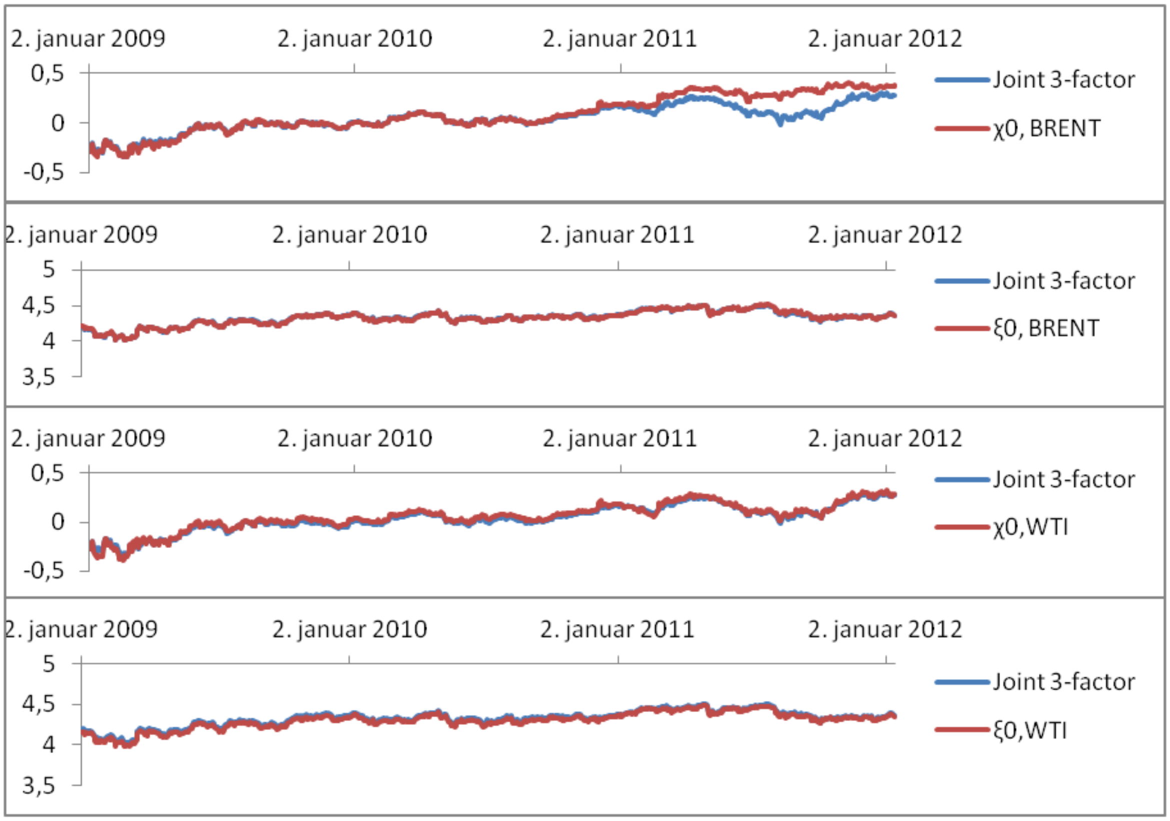

9.2.1 Re-Estimating the Model Parameters

It would be interesting to get a measure of how reliable the parameter estimates of the three-factor model is, because we have a lot of parameters and therefore we have a possibility of redundancy among them. We would also like to know how stable our calibration technique is: does it work every time, or was it just a coincidence when we tried it and it worked?

To investigate this, we perform re-estimation of the parameters. This is done by simulating price development on basis of our parameters in Table 9. From the simulation we get data sets of forward prices, which we use for calibrating the model. What we want to explore is whether or not the procedure is able to recover the parameters underlying the produced data set; namely the parameters of Table 9.

When simulating using the three-factor model, we use the estimated state variables for the last day of our historical data as starting values. The following simulation settings are applied:

Table 11 Simulation settings applied for re-estimation.

NoOfDates	765	
NoOfSim	200	
SimStartDate	13. januar 2012	
StartValues	BRENT	WTI
χ_0	0,270480292	0,270480292
$\chi_{0,2}$	0,108633783	-0,06030279
ξ_0	4,357087429	4,357087429

The algorithm used for re-estimation can be described briefly like this:

1. Use input parameters (TrueValues) and start values given in Table 9 and Table 11 to simulate datasets consisting of 765 days, each day with forward prices of both Brent and WTI for 50 consecutive months.
2. Use the simulated dataset of Brent and WTI forward prices to calibrate the model again. Guess values for calibration are drawn from Table 8.
3. Repeat steps 1 and 2 as many times as possible to get as many parameter estimates as you can. We ended up doing this 101 times.

Re-estimation yields the following descriptive statistics for the estimation routine:

Table 12 Descriptive statistics for iterative estimation routine.

		Min	1st Qu.	Median	Mean	3rd Qu.	Max	TrueValues
BRENT	κ	0,5158	0,5350	0,5394	0,5382	0,5416	0,5517	0,5436
	κ2	0,3536	0,3556	0,3559	0,3561	0,3566	0,3597	0,3556
	λχ	-0,0896	-0,0576	-0,0478	-0,0464	-0,0338	-0,0089	-0,0526
	λχ2	-0,0231	-0,0176	-0,0150	-0,0151	-0,0132	-0,0054	-0,0115
	μξ*	-0,0373	-0,0341	-0,0332	-0,0332	-0,0323	-0,0292	-0,0334
	μξ	-0,3402	-0,0499	0,0875	0,0683	0,2088	0,4300	0,0654
	σχ	0,2261	0,2367	0,2414	0,2410	0,2463	0,2554	0,2421
	σχ2	0,1207	0,1257	0,1275	0,1276	0,1291	0,1405	0,1272
	σξ	0,2284	0,2409	0,2446	0,2445	0,2480	0,2624	0,2451
	ρχ1χ2	-0,3290	-0,2950	-0,2761	-0,2755	-0,2549	-0,1957	-0,2753
	ρχ1ξ	0,1878	0,2699	0,2895	0,2921	0,3119	0,3687	0,2915
	ρχ2ξ	-0,2346	-0,1827	-0,1622	-0,1570	-0,1339	-0,0638	-0,1604
WTI	κ	0,5158	0,5350	0,5394	0,5382	0,5416	0,5517	0,5436
	κ2	6,3045	6,5522	6,6250	6,6235	6,6854	7,0199	6,7165
	λχ	-0,0896	-0,0576	-0,0478	-0,0464	-0,0338	-0,0089	-0,0526
	λχ2	-0,2039	-0,0971	-0,0693	-0,0665	-0,0309	0,0751	0,0008
	μξ*	-0,0373	-0,0341	-0,0332	-0,0332	-0,0323	-0,0292	-0,0334
	μξ	-0,3402	-0,0499	0,0875	0,0683	0,2088	0,4300	0,0654
	σχ	0,2261	0,2367	0,2414	0,2410	0,2463	0,2554	0,2421
	σχ2	0,2560	0,2652	0,2699	0,2707	0,2756	0,2921	0,2723
	σξ	0,2284	0,2409	0,2446	0,2445	0,2480	0,2624	0,2451
	ρχ1χ2	0,1832	0,2341	0,2585	0,2580	0,2801	0,3537	0,2553
	ρχ1ξ	0,1878	0,2699	0,2895	0,2921	0,3119	0,3687	0,2915
	ρχ2ξ	-0,1890	-0,1302	-0,1064	-0,1060	-0,0826	-0,0177	-0,1094
	c_ξ	0	0	0	0	0	0	0
	c_χ	0	0	0	0	0	0	0
	ρBW	0,0601	0,1189	0,1405	0,1440	0,1689	0,2462	0,1419

To get a clearer picture of how large the deviations in the estimated parameters are, we measure the percentual deviations from the true values for each parameter. This means that, for an equal value of deviations, percentual deviations will be larger for parameters with small true values (in absolute terms). Percentual deviations from the true values are given below:

Table 13 Percentual deviations from true values of parameter estimates.

		Min	1st Qu.	Median	Mean	3rd Qu.	Max
BRENT	κ	-5,1 %	-1,6 %	-0,8 %	-1,0 %	-0,4 %	1,5 %
	κ2	-0,6 %	0,0 %	0,1 %	0,1 %	0,3 %	1,1 %
	λχ	-70,4 %	-9,5 %	9,0 %	11,8 %	35,7 %	83,2 %
	λχ2	-101,0 %	-53,6 %	-30,6 %	-31,2 %	-15,1 %	53,1 %
	μξ*	-11,6 %	-2,0 %	0,8 %	0,7 %	3,5 %	12,7 %
	μξ	-620,0 %	-176,3 %	33,8 %	4,4 %	219,2 %	557,4 %
	σχ	-6,6 %	-2,2 %	-0,3 %	-0,5 %	1,7 %	5,5 %
	σχ2	-5,1 %	-1,1 %	0,2 %	0,3 %	1,5 %	10,5 %
	σξ	-6,8 %	-1,7 %	-0,2 %	-0,2 %	1,2 %	7,1 %
	ρχ1χ2	-19,5 %	-7,2 %	-0,3 %	-0,1 %	7,4 %	28,9 %
	ρχ1ξ	-35,6 %	-7,4 %	-0,7 %	0,2 %	7,0 %	26,5 %
	ρχ2ξ	-46,2 %	-13,9 %	-1,1 %	2,1 %	16,5 %	60,2 %
WTI	κ	-5,1 %	-1,6 %	-0,8 %	-1,0 %	-0,4 %	1,5 %
	κ2	-6,1 %	-2,4 %	-1,4 %	-1,4 %	-0,5 %	4,5 %
	λχ	-70,4 %	-9,5 %	9,0 %	11,8 %	35,7 %	83,2 %
	λχ2	-25352,6 %	-12128,9 %	-8682,7 %	-8335,7 %	-3925,6 %	9197,1 %
	μξ*	-11,6 %	-2,0 %	0,8 %	0,7 %	3,5 %	12,7 %
	μξ	-620,0 %	-176,3 %	33,8 %	4,4 %	219,2 %	557,4 %
	σχ	-6,6 %	-2,2 %	-0,3 %	-0,5 %	1,7 %	5,5 %
	σχ2	-6,0 %	-2,6 %	-0,9 %	-0,6 %	1,2 %	7,3 %
	σξ	-6,8 %	-1,7 %	-0,2 %	-0,2 %	1,2 %	7,1 %
	ρχ1χ2	-28,2 %	-8,3 %	1,3 %	1,1 %	9,7 %	38,6 %
	ρχ1ξ	-35,6 %	-7,4 %	-0,7 %	0,2 %	7,0 %	26,5 %
	ρχ2ξ	-72,7 %	-19,0 %	2,8 %	3,1 %	24,5 %	83,8 %
	c_ξ						
	c_χ						
	ρBW	-57,7 %	-16,2 %	-1,0 %	1,5 %	19,0 %	73,5 %

There is a clear pattern indicating that risk premium estimates (given by λ_χ, $\lambda_{\chi 2}$ and μ_ξ) are very uncertain. The correlation factors also seem hard to estimate precisely. To address this problem for the long-term risk premium and the correlation factors one might increase the number of simulated dates, but it would probably demand a very large increase to get satisfactorily results. As mentioned earlier, an obvious solution to the problem with the short-term risk premium hasn't been found.

Apart from the mentioned problems, which also occur for the standard Schwartz-Smith model, the calibration procedure seems to work in a satisfactorily manner also for the joint three-factor model.

9.3 Simulation Results Using the Joint Three-Factor Model

As with the two-factor joint model, we can use the three-factor model to perform simulations about future development of Brent and WTI. Also for the three-factor model, we simulate prices one year ahead, using the following simulation settings:

Table 14 Simulation settings, joint three-factor model.

NoOfDates	365	
NoOfSim	2250	
SimStartDate	13. januar 2012	
StartValues	BRENT	WTI
χ_0	0,270480292	0,270480292
$\chi_{0,2}$	0,108633783	-0,06030279
ξ_0	4,357087429	4,357087429

Due to the three-factor model being a bit more complex, the computer wasn't able to perform as many simulations in one operation as for the two-factor. Therefore, number of simulations was reduced from 2400 to 2250.

Results from simulating one year ahead with the three-factor model are quite similar to results from the two-factor model. Below, we present simulation summary plots for Brent and WTI.

Figure 27 Simulated Brent spot prices using joint three-factor model.

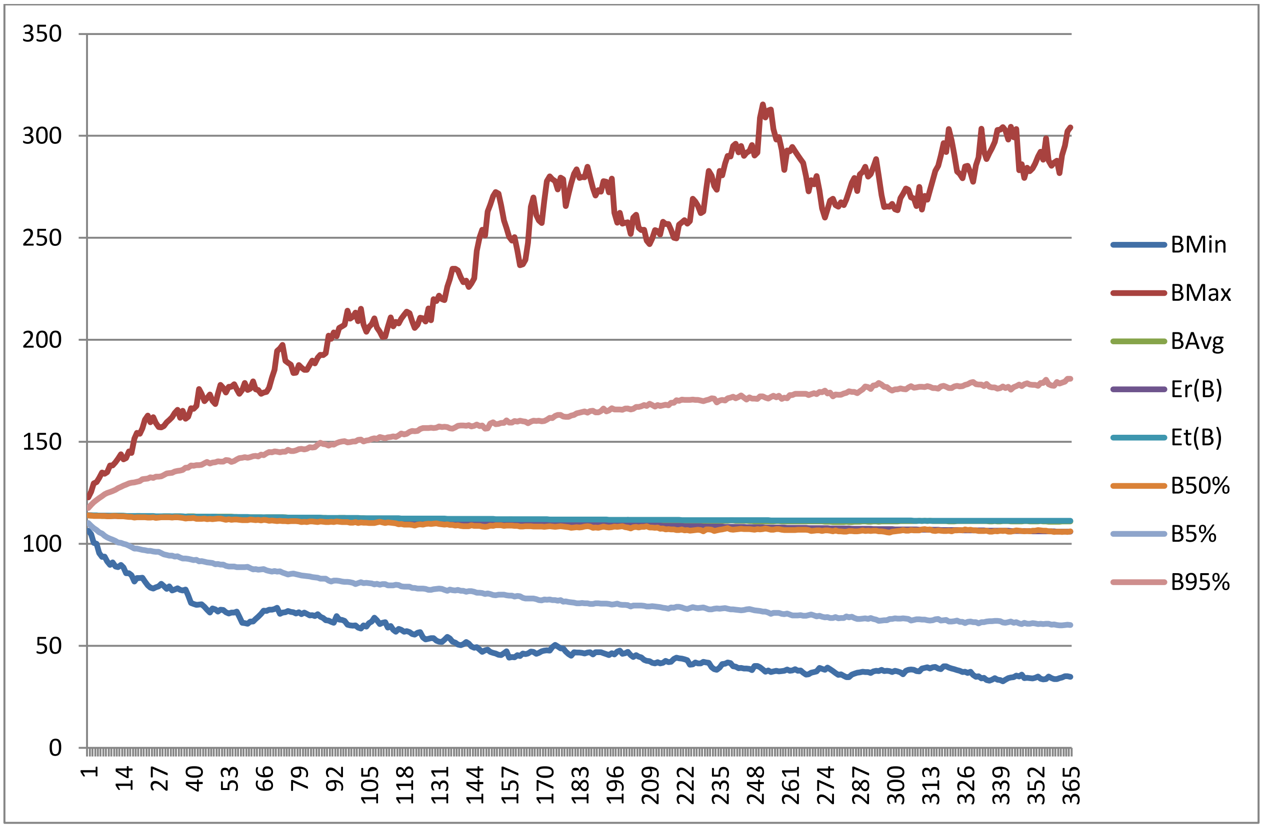

Figure 28 Expected values along with simulated median and average Brent spot prices (as shown in Figure 27)

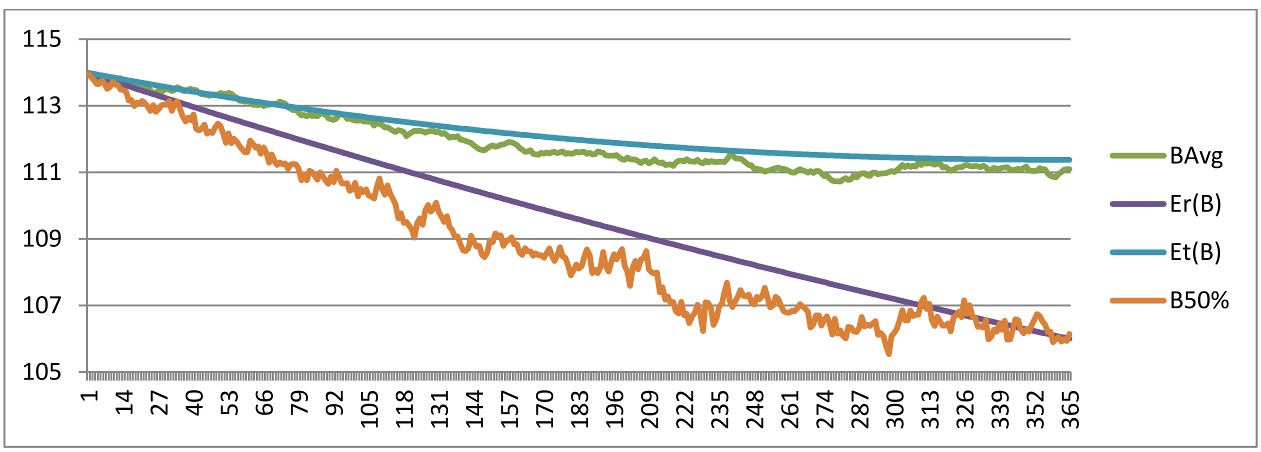

The simulation results for Brent reveal the same main features as for the two-factor simulation (Figure 17 and Figure 18). As for the two-factor simulation, the Brent spot price is expected to end up close to $111 per barrel at the end of the first year. When reviewing the 5% and 95% percentiles of simulated outcomes, the results for the two-factor and three-factor models are in agreement. It seems that the two models have about the same measure of price risk.

The three-factor curve has a slightly higher starting point for the Brent spot price, which is due to a difference in the estimated state variables for January 13th. When reviewing observed data, we find that the two-factor has a better estimate for this particular day. However, for WTI, the three-factor model fits better to observed data.

For January 13th, the two-factor model gives a 1st month price[37] of $112.47 for Brent and $101.39 for WTI. The three-factor model gives $113.53 and $97.87, respectively. The observed values, on the other hand, are $110.44 and $98.70. Hence, the two-factor model yields the closest result for Brent, while the three-factor fits best to WTI.

The differences in estimated 1st month prices result in a remarkable difference in price gap estimates. While the two-factor model implies a price gap at January 13th of about $11, the three-factor model suggests a gap of nearly $16. The observed data, however, reveals Brent prices lying close to $12 above WTI. For this particular day, the three-factor model over-estimates the Brent price, while it under-estimates WTI. The two-factor model over-estimates both, with the error being at about the same size for both crudes.

It's time to have a look at the simulation results for WTI:

Figure 29 Expected values along with simulated median and average WTI spot prices (joint three-factor model).

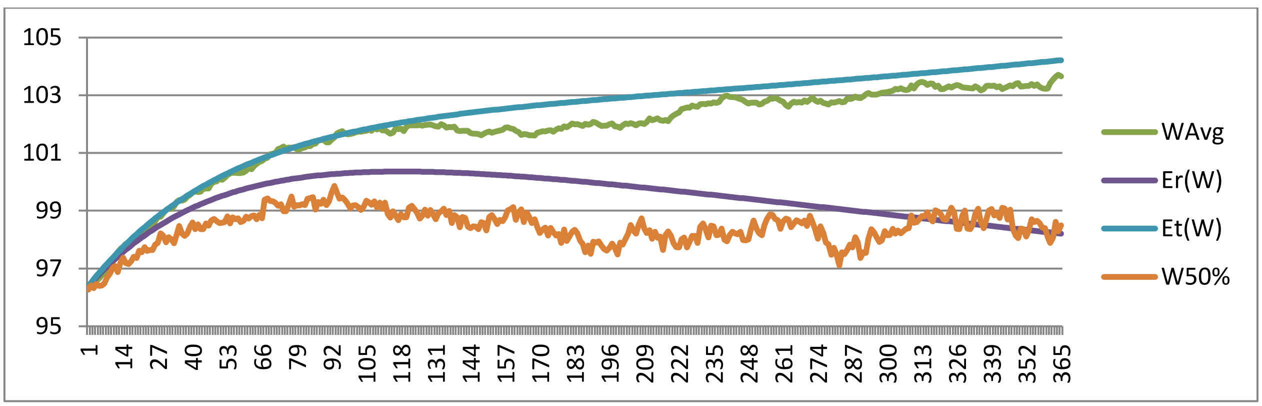

Figure 29 reveals that the (true process) development of WTI is a bit different[38] for the two models (see Figure 21 to compare with the two-factor model). While the two-factor model indicates a curvature of WTI similar to that of Brent, the three-factor model exhibits a WTI curve growing from day one. This difference reflects the different assumptions underlying the two models. The two-factor model assumes some constant gap between the long-term variable of the two products, resulting in price developments that resemble each other but are separated by this constant gap. The three-factor model, on the other hand, assumes the constant gap to be zero, implying that the prices are expected to revert towards each other.

---

[37] We have to use 1st month futures prices for "reality checks", as we don't have observed spot prices.
[38] Please note that we have "zoomed in" on the relevant y-axis interval on the graphs, so that differences might seem more significant than when seen "from distance".

As discussed in relation to the calibration of the two-factor joint model, the value of the constant gap seemed rather illogical as it was much larger than the average price gap between Brent and WTI during the time period of our data sample. Also, from other sources such as Gue (2011) and Fattouh (2011) we know that Brent historically has been traded at a premium to WTI. Therefore, it seems that the three-factor model's assumption of equal equilibrium prices is more realistic than the two-factor's constant gap, c_ξ. Implementing the assumption of equal equilibrium prices can easily be done by restricting c_ξ to be zero.

A more thorough investigation of differences between WTI predictions in the two-factor and three-factor model is done below. But first, let's just review the simulation results for the price differential of Brent and WTI.

Figure 30 Difference between simulated Brent and WTI prices (joint three-factor model).

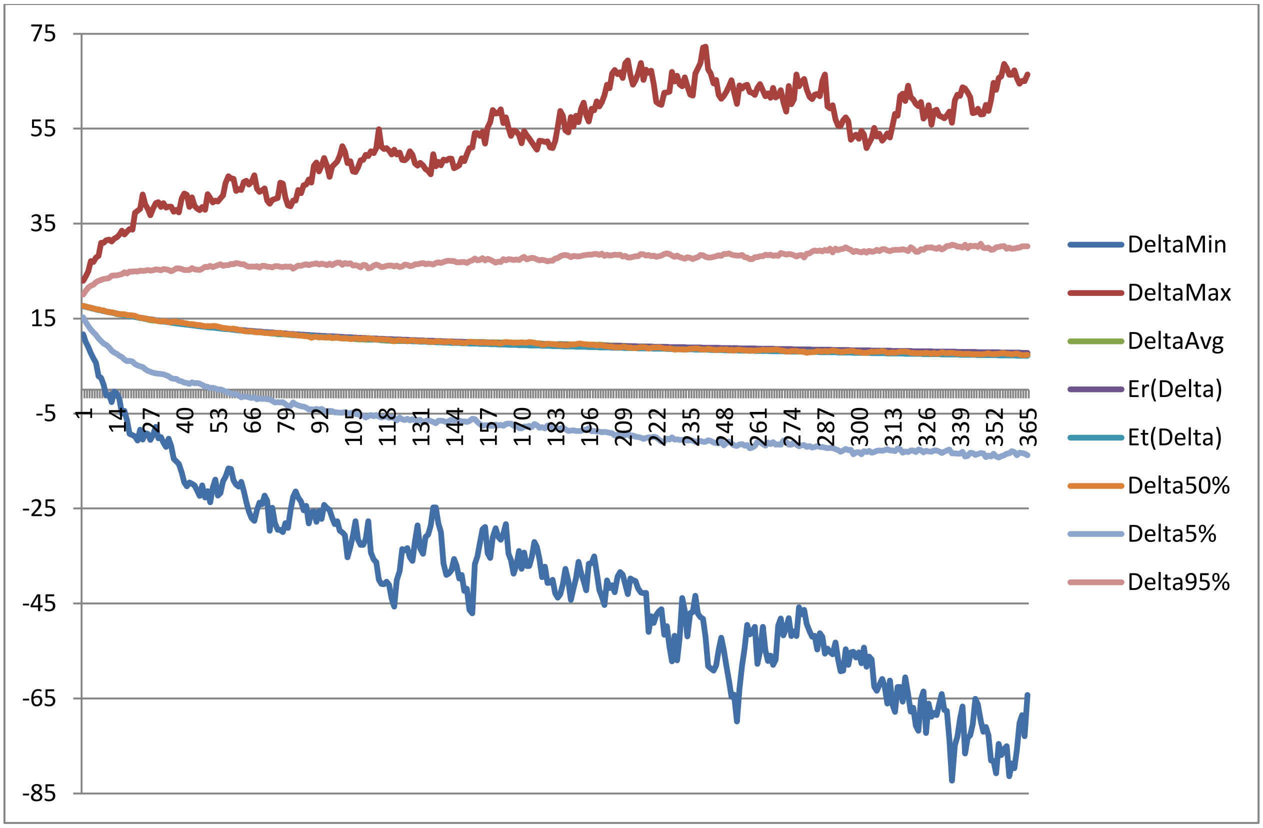

Figure 31 Average and median delta along with delta of expected values (three-factor model).

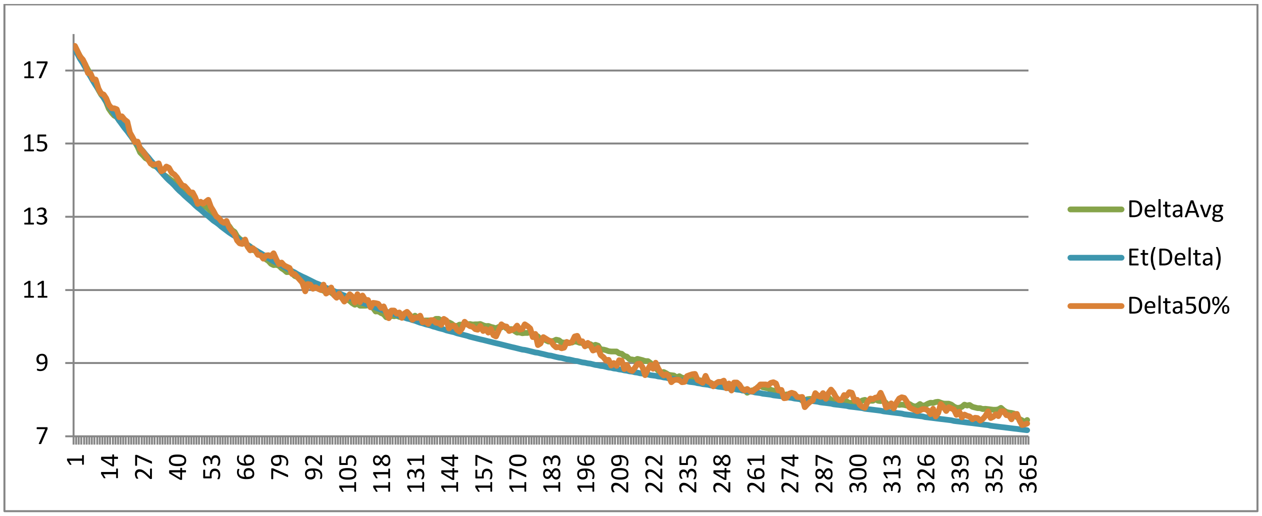

As discussed above, the start value for the price differential is unequal for the two-factor and the three-factor models due to differences in the estimated start variables. However, the most interesting feature of the graph isn't at what value it starts, but how it develops in time.

The results for simulated price differentials are clearly influenced by the expected growth in WTI prices to reach Brent level. Following the assumptions of the model, the price gap of Brent and WTI is expected to grow smaller and smaller as both spot prices revert towards the same long-term equilibrium. Thus, the plots are in accordance with the assumptions of our model.

9.3.1 The Spread Contract Revisited

As the three-factor model gives different results for the Brent/WTI spread than the joint two-factor model, it would be interesting to re-assess the Spread Contract introduced at the end of chapter 8.

Reading from Figure 30 we get an expected spread at December 1st of \$7.57, while the 5th and 95th percentile of the spread distribution is -\$13.11 and \$29.10, respectively. The tree-factor model generally predicts higher WTI prices relative to Brent (yielding a lower expected spread), and it also seems that the variance of the spread distribution has increased slightly compared to the two-factor model.

Consider the case where the price for December 1st spread still is \$11.20. Using the two-factor model, you expect a profit on buying such a contract. However, the three-factor model suggests you will only receive \$7.57 per barrel on December 1st, resulting in a negative profit (loss) of (\$7.57-\$11.20) = -\$3.63 per barrel of the contract. Buying 100 contracts, that would equal a loss of \$363,000. Hence, you would never enter the deal on these terms with the three-factor model as decision basis. In order for you to buy the contract, the market prices would have to move close to (and preferably below) \$7.57.

9.3.2 Why Do the Models Predict Different WTI Prices?

As discussed earlier, the two models exhibit different development of WTI prices; the two-factor model predicts prices at a constant discount to Brent, while the three-factor model seems to predict convergence between the prices. In this section, we will see why this is happening.

To further understand what's going on as the simulated WTI prices grow towards Brent, a more thorough mathematical analysis of how the expected price develops is performed.

First, in order to make it easier to visualize the development, we present plots isolating the expected spot prices of Brent and WTI. The plots show expected development implied by both the two-factor and the three-factor model. In addition to the true process (Figure 32), we also present the numerical solution for the risk-neutral expected development (Figure 33). Both processes are treated in the following discussion.

Figure 32 Comparison of expected spot price development (true process) between two-factor (Et(W) and Et(B)) and three-factor (3_B and 3_W) models.

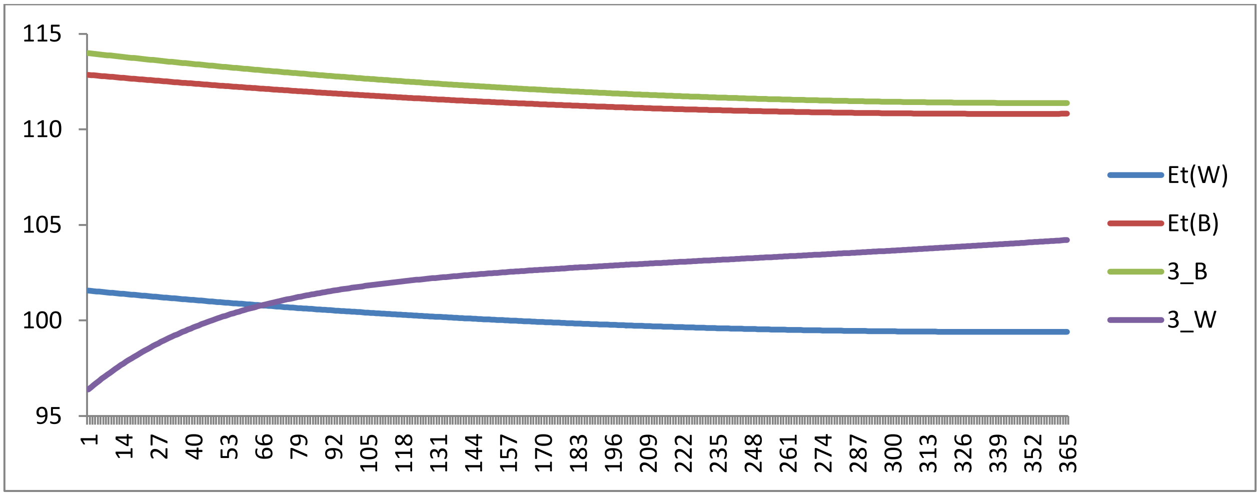

Figure 33 Comparison of expected spot price development (risk-neutral process) between two-factor (Er(W) and Er(B)) and three-factor (3_B and 3_W) models.

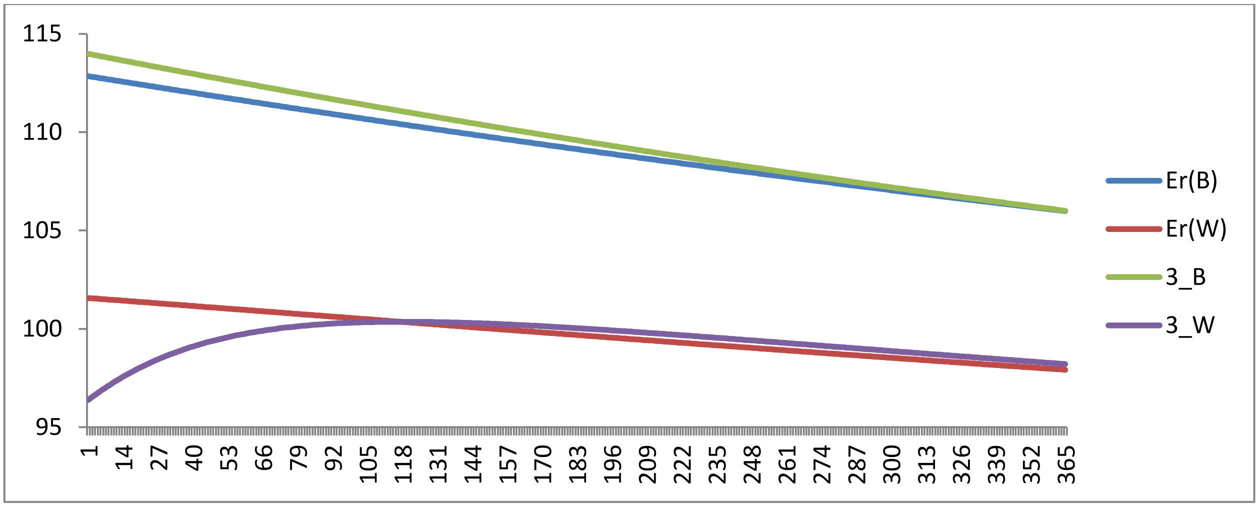

To explain what happens to the expected spot prices of WTI, we split up the expression for the expected spot price. This way, we are able to explain the different factors affecting the development of the expected spot price. We want to provide a mathematical understanding to the sudden growth of expected WTI prices implied by the three-factor simulations.

The (true process) expected spot price is made up of several components (see Eq. 86). The true process future price is $E[S_{t+\Delta t}] = e^{E[\ln(S_{t+\Delta t})]+\frac{1}{2}Var[\ln(S_{t+\Delta t})]}$ where $E[\ln(S_{t+\Delta t})] + \frac{1}{2}Var[\ln(S_{t+\Delta t})]$ is made up of the components presented in Table 15 below:

Table 15 Components of three-factor true process expected spot price.

Long term	1st short term	2nd short term		Drift	"the rest"	
$\xi_{t-\Delta t}$	$+\chi_{t-\Delta t}e^{-\kappa\Delta t}$	$+\chi_{2,t-\Delta t}e^{-\kappa_2\Delta t}$		$+(\mu_\xi + \frac{\sigma_\xi{}^2}{2})\Delta t$	*See footnote*[39]	
Component values for 12 consecutive months, starting at Jan 13th 2012:						
		Brent	WTI		Brent	WTI
4,36	0,27	0,11	-0,06	0,00	0,00	0,00
4,36	0,26	0,11	-0,03	0,01	0,00	0,01
4,36	0,25	0,10	-0,02	0,02	0,01	0,01
4,36	0,24	0,10	-0,01	0,02	0,01	0,01
4,36	0,23	0,10	-0,01	0,03	0,01	0,02
4,36	0,22	0,09	0,00	0,04	0,01	0,02
4,36	0,21	0,09	0,00	0,05	0,02	0,02
4,36	0,20	0,09	0,00	0,06	0,02	0,03
4,36	0,19	0,09	0,00	0,06	0,02	0,03
4,36	0,18	0,08	0,00	0,07	0,02	0,03
4,36	0,17	0,08	0,00	0,08	0,02	0,03
4,36	0,16	0,08	0,00	0,09	0,03	0,03
4,36	0,16	0,08	0,00	0,10	0,03	0,04

In the three-factor model, the long-term component is shared by the two products, as is the drift term and the first short-term component. The difference lies in the second short-term variable and its associated parameters. The differences are therefore displayed in the development of second short term variable and the component named "the rest", where parameters associated with the second short term variable are present. The reason for the expected rapid growth in WTI is the phaseout of the negative second short term variable (which has a very large κ value). Viewed in the perspective of the assumption behind the model, namely that the differences between the products are covered by the second short-term variable, this feature of the model seems logical. The downward slope of WTI by the two-factor model can be explained by the positive short-term start variable (see Table 7) which is phased out as time goes by (phasing out the positive short-term variable results in reduced prices).

39

$$+\frac{1-e^{-2\kappa\Delta t}}{4\kappa}\sigma_\chi{}^2 + \frac{1-e^{-2\kappa_2\Delta t}}{4\kappa_2}\sigma_{\chi_2}{}^2 + \left(1-e^{-\kappa\Delta t}\right)\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} + \left(1-e^{-\kappa_2\Delta t}\right)\frac{\rho_{\chi_2\xi}\sigma_{\chi_2}\sigma_\xi}{\kappa_2} +$$

$$\sqrt{\left(1-e^{-2\kappa\Delta t}\right)\frac{\sigma_\chi^2}{2\kappa}}\sqrt{\left(1-e^{-2\kappa_2\Delta t}\right)\frac{\sigma_{\chi_2}^2}{2\kappa_2}}\rho_{\chi_2\chi}$$

As the effect of the second short-term factor is gone (around month six, according to Table 15), the WTI curve still seems to grow faster than the Brent curve. This can be explained by the positive second short-term start variable of Brent, which is still not outphased due to the lower κ value (discussed above in relation to the calibration results). In addition, “the rest” contributes with larger positive values for WTI than Brent.

Figure 33 shows the risk-neutral expected spot prices, starting at January 13th 2012. Here, the differences are smaller than the true process graphs. This is no surprise, as the risk-neutral expected spot prices are nothing other than the forward curve of Jan 13th 2012. Since both models are calibrated to fit observed forward curves, we expect the risk-neutral expected spot prices of the two- and three-factor model to lie close to each other. The downward turn of the WTI curve is made possible through the reduced drift term (the introduction of a risk premium) due to risk-neutrality. The “the rest” term has also become more alike between the products due to the phase inn of the short-term risk premiums. It seems like the products are expected to have about the same prices when the second short-term variable is phased out. Table 16 below shows the components of the expected spot price under the risk-neutral measure.

Table 16 Components of three-factor risk-neutral process expected spot price.

Long term	**1st short term**	**2nd short term**		**Drift**	**“the rest”**	
$\xi_{t-\Delta t}$	$+\chi_{t-\Delta t}e^{-\kappa\Delta t}$	$+\chi_{2,t-\Delta t}e^{-\kappa_2\Delta t}$		$+(\mu_\xi^* + \frac{\sigma_\xi{}^2}{2})\Delta t$	*See footnote*[40]	
Component values for 12 consecutive months, starting at Jan 13th 2012:						
		Brent	WTI		Brent	WTI
4,36	0,27	0,11	-0,06	0,00	0,00	0,00
4,36	0,26	0,11	-0,03	0,00	0,01	0,01
4,36	0,25	0,10	-0,02	0,00	0,02	0,02
4,36	0,24	0,10	-0,01	0,00	0,02	0,03
4,36	0,23	0,10	-0,01	0,00	0,03	0,03
4,36	0,22	0,09	0,00	0,00	0,04	0,04
4,36	0,21	0,09	0,00	0,00	0,04	0,05
4,36	0,20	0,09	0,00	0,00	0,05	0,05
4,36	0,19	0,09	0,00	0,00	0,06	0,06
4,36	0,18	0,08	0,00	0,00	0,06	0,06
4,36	0,17	0,08	0,00	0,00	0,07	0,07
4,36	0,16	0,08	0,00	0,00	0,07	0,07
4,36	0,16	0,08	0,00	0,00	0,08	0,08

It would be interesting to plot the expected spot prices for a longer time horizon, to see if the models give reasonable expected values. We choose to roll three years into the future. The plot below shows expected true process spot prices of Brent and WTI.

[40] Equal to the expression for the true process plus extra terms $-(1-e^{-\kappa\Delta t})\frac{\lambda_\chi}{\kappa} - (1-e^{-\kappa_2\Delta t})\frac{\lambda_{\chi 2}}{\kappa_2}$

Figure 34 Plot of expected true process Brent and WTI spot prices for two-factor (Et(B) and Et(W)) and three-factor (3_B and 3_W) model.

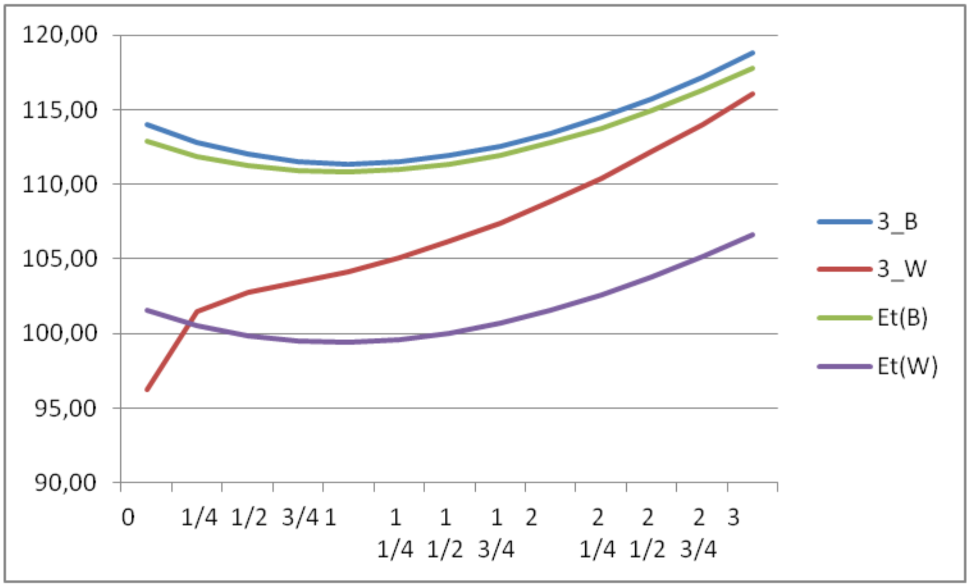

We see that the two models still have pretty similar expectations for Brent, while there is difference in the progress for WTI. The differences between the models are now clearly reflecting the assumptions underlying the models: The two-factor joint model maintains some constant level gap between the two products, while the three-factor joint model wants the gap to revert to approximately[41] zero.

9.4 The Brent/WTI Spread Gets Out of Hand

To really stress the models, we have to make the time horizon long. We want to see if the models give reasonable results with a time horizon of twenty years. The expected values will then develop like this:

[41] We say ”approximately” because there will be some gap between the expected prices of the two products due to differences in the term ”the rest”, stemming from parameters relating to the second short-term variable.

Figure 35 Expected spot prices with twenty-year time horizon.

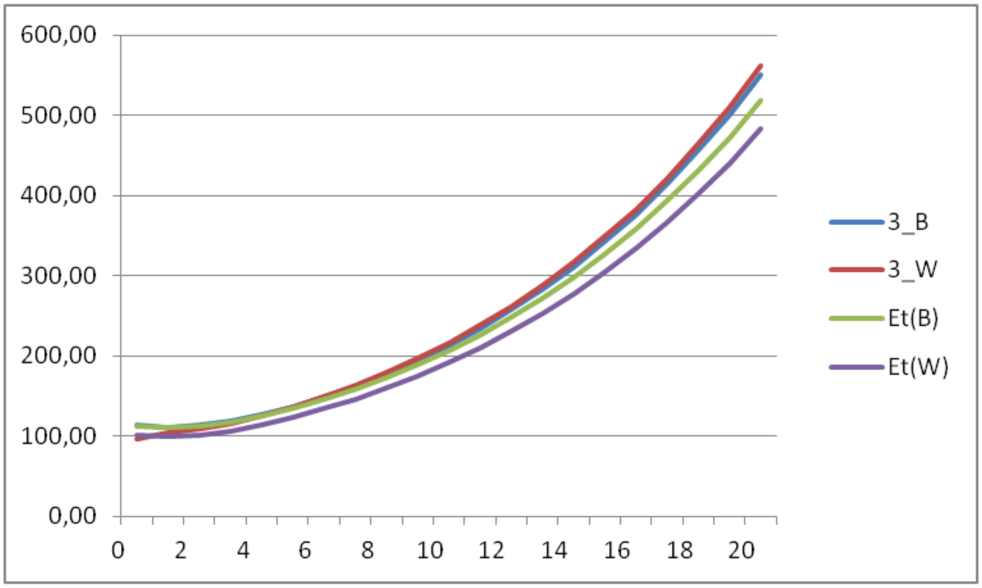

Let's first just state that the values in the long end of the curve are very high, and there is a great uncertainty related to them. We don't know much about oil prices 20 years ahead. Therefore, no-one would make a bet about oil prices that far into the future. We might want to look maximum 5 or 10 years into the future.

However, the plot reveals a weakness in the model; an increasing gap between expected values. A perfect model would exhibit a restricted gap between prices as the time horizon increases, not a growing gap as we see here. What's the reason for this growing gap? A new look at the components of the expected spot price might help us get the answer.

Table 17 Components of true process expected spot prices for twenty-year horizon (three-factor joint model).

Year	$E(S_B)$	$E(S_W)$	$E(S_B\text{-}S_W)$	ξ_0	χ_0	$\chi_{0,2}$		Drift	"the rest"	
						Brent	WTI		Brent	WTI
0	114,00	96,28	17,72	4,36	0,27	0,11	-0,06	0,00	0,00	0,00
1	111,38	104,17	7,21	4,36	0,16	0,08	0,00	0,10	0,03	0,04
2	113,42	108,85	4,57	4,36	0,09	0,05	0,00	0,19	0,04	0,05
3	118,77	116,03	2,74	4,36	0,05	0,04	0,00	0,29	0,04	0,06
4	126,64	125,27	1,37	4,36	0,03	0,03	0,00	0,38	0,05	0,06
5	136,58	136,29	0,30	4,36	0,02	0,02	0,00	0,48	0,05	0,06
6	148,37	148,96	-0,59	4,36	0,01	0,01	0,00	0,57	0,05	0,06
7	161,91	163,25	-1,34	4,36	0,01	0,01	0,00	0,67	0,05	0,06
8	177,19	179,20	-2,01	4,36	0,00	0,01	0,00	0,76	0,05	0,06
9	194,27	196,89	-2,62	4,36	0,00	0,00	0,00	0,86	0,05	0,06
10	213,25	216,45	-3,20	4,36	0,00	0,00	0,00	0,95	0,05	0,06
11	234,25	238,02	-3,77	4,36	0,00	0,00	0,00	1,05	0,05	0,06
12	257,46	261,79	-4,33	4,36	0,00	0,00	0,00	1,15	0,05	0,06
13	283,05	287,96	-4,91	4,36	0,00	0,00	0,00	1,24	0,05	0,06
14	311,26	316,77	-5,51	4,36	0,00	0,00	0,00	1,34	0,05	0,06
15	342,32	348,48	-6,15	4,36	0,00	0,00	0,00	1,43	0,05	0,06
16	376,53	383,36	-6,84	4,36	0,00	0,00	0,00	1,53	0,05	0,06
17	414,17	421,74	-7,57	4,36	0,00	0,00	0,00	1,62	0,05	0,06
18	455,60	463,97	-8,37	4,36	0,00	0,00	0,00	1,72	0,05	0,06
19	501,19	510,43	-9,24	4,36	0,00	0,00	0,00	1,81	0,05	0,06
20	551,36	561,54	-10,19	4,36	0,00	0,00	0,00	1,91	0,05	0,06

The first couple of years, the difference is heavily influenced by the difference in second short-term variable. This is no problem, as we expect short-term shocks to impact the close future. At five years, the difference in the outphased short-term variables is negligible, and expected spot price difference is close to zero. However, the price difference still changes and now gets negative, meaning WTI prices are higher than Brent prices. A continued development is evident also for the two-factor model, where the positive price difference grows larger and larger. In order to understand this continued development, let's look at Table 17 for the three-factor model, from year ten and onwards. The long-term component is constant, effects of short-term shocks are phased out, and the drift term is a shared variable. The "the rest" term is constant[42] for both products. One would imagine that the difference between prices would converge to some constant value. However, it does not. Why is that?

The reason is the expected price being *e* raised to the power of sum of components; *ln*(E[S]) = sum of components. The difference in ln(E[S]) is constant for far time horizons, even though both $\ln(E[S_B])$ and

[42] One might say it has been put to rest.

$\ln(E[S_W])$ change – because they change with the same amount. However, even though the difference in $\ln(E[S])$ is constant, that doesn't mean that the difference in E[S] is constant. We can present it this way: the drift term is the only growing term, all other terms are constant. To make it easy, let's name the drift term ax and all other terms mix together to a constant; b for Brent and w for WTI. ax is the same for both products. Then, even though ax + b – (ax + w) is constant for all values of x (time), $e^{ax+b} - e^{ax+w}$ changes with time. This is easily verified by doing $e^2 - e^1$ and $e^3 - e^2$ on a calculator and find that the results are different even though 2-1 and 3-2 are the same. Luckily, this difference in constant terms (the difference lies in the "the rest" term) is pretty small, and for moderate values of $\ln(E[S])$ the model still gives reasonable results and can be used. But for very long time horizons (when ax + b gets very big), the difference in constant terms will give unrealistic large expected differences in $E(S_B)$ and $E(S_W)$.

9.4.1 Limitations to the Forecasting Horizon

As revealed in the above discussion, both the two-factor and the three-factor model give unrealistic values if we stretch the time horizon too far. From Figure 35 it is evident that the two-factor model yields too big expected price gap between Brent and WTI for the far future, but also the three-factor model has an unsatisfactory development of a continuously growing (negative) gap between Brent and WTI prices. We realize that there is a limitation to the forecasting horizon we can employ.

However, one might argue, we couldn't have trusted *any* model for remote futures. When calibrating the models we have to limit the length of data time series as we want to capture the market fundamentals prevailing at the current point of time. We have only used time series of three years when calibrating the model. Assuming that the market fundamentals will be the same in twenty years time would be naive, therefore we can't use the model to forecast prices that far ahead. In order to give an estimate of prices prevailing that far from now we would have to perform an analysis of expected technological development, demographical conditions, shifts of political power, remaining reserves of crude oil in reservoirs, etc. Nevertheless, the risk of oil prices in say twenty years time is too big to use the model for forecasting.

To sum the discussion up, the model can't be used for very far futures. A reasonable limit would be 10 years maybe. As a reference, WTI futures are listed up to nine years forward[43].

9.5 Attempting to Restrain the Price Gap

An attempt was made to mend the joint model's inadequate feature of diverging expected prices for long time horizons. However, even though the method presented below may serve as a formal solution to the divergence of expected price gaps, we should keep in mind the discussion above about the validity of the models for the distant future. The suggested solution is presented in the following sections.

The persisting growth in expected price difference occurs due to a minor difference in the constant term of the expected log future price. As discussed earlier in this chapter, this constant difference in expected log prices can make significant impact on the difference in expected future prices as the absolute values of the expectations becomes large. However, as the difference in the constants is rather small, only a

[43] http://www.cmegroup.com/trading/energy/crude-oil/light-sweet-crude_contract_specifications.html

minor adjustment to the model should be enough to fix the bug. We would like some small adjustment making the expected price gap constant for very far maturities.

Underlying here is an assumption that the price gap generally doesn't increase as price level itself increases. To check the validity of this assumption, we have to examine historical data. In Figure 36, we present a plot of observed price gaps for different absolute values of the Brent 1st month contract. By reviewing this plot, we can't conclude that the gap increases as price level increases. With reference to Figure 7, the extreme values to the right in Figure 36 don't indicate a correlation between price gaps and price level; rather they can be explained by the bottleneck effects appearing in Cushing, Oklahoma. "Normal situations" prevail in the price interval from $40 to $80, in which there doesn't seem to be any clear pattern in the development of price differentials between Brent and WTI. Therefore, assuming an expected constant price gap seems the most reasonable.

Figure 36 Brent/WTI spot price difference vs. Brent price level.

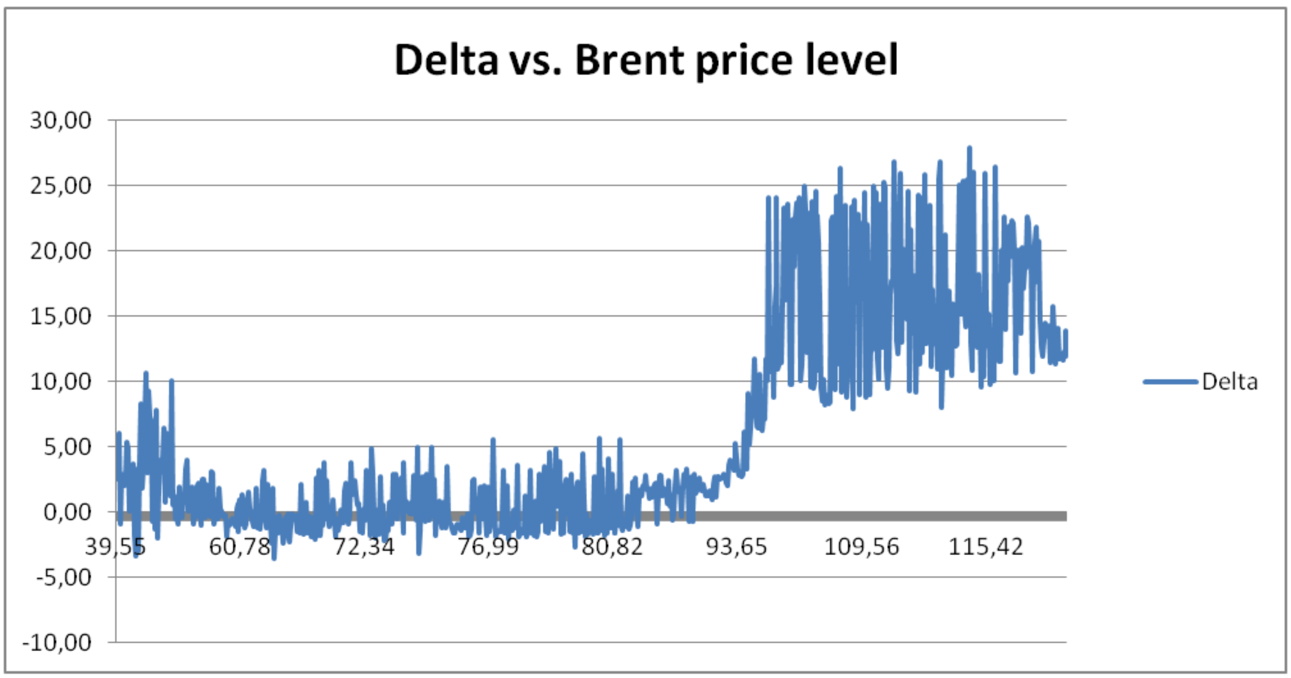

So what we are looking for is some kind of adjustment to the joint model keeping the expected price gap for far maturities constant, which means we need to make $\ln(E[S_B]) - \ln(E[S_W])$ decrease for increasing price level, and vice versa for decreasing price level. This adjustment must lead to a constant price difference $E(S_B) - E(S_W)$. Using our ax + b analogy, a possible solution is introducing an adjustment factor $f(ax,b,w)$ such that $e^{ax+b} - e^{ax+w+f(ax,b,w)} = constant$. Solving for f(ax,b,w) we then get $f(ax,b,c) = \ln(e^{ax+b} - constant) - (ax + w)$. For our application, e^{ax+b} is the expected Brent spot price far ahead in the future (when $\Delta t \rightarrow \infty$). e^{ax+w} is the corresponding WTI value. So, we introduce a new factor *f* to the WTI price development process, which is responsible for keeping the expected WTI price close to the expected Brent price. After presenting a formal expression for factor *f*, let's examine this more in detail for the two-factor joint model.

$$f_{t+\Delta t,t} = \ln\left(E\left(S^B_{t+\Delta t,t}\right)|_{\Delta t \to \infty} - constant\right) - \ln\left(E\left(S^W_{t+\Delta t,t}\right)|_{\Delta t \to \infty}\right) \quad \textbf{Eq. 88}$$

Note that the expected values included in this equation are calculated as if *f* didn't exist, meaning that we first calculate expected values without *f*, use expected values to calculate *f* and lastly include *f* in the calculation of adjusted expected values of WTI. We do not use these adjusted expected values to calculate *f* over again (this is not an iterative process).

For the two-factor model, expected (true process) spot prices when $\Delta t \to \infty$ are given by Eq. 73 and Eq. 74. As Δt increases more and more, the difference between $E(S_B)$ and $E(S_W)$ grows larger and larger (for positive drift terms), according to our recent discussion. This means that *f* has to change as Δt changes, which it will as follows from Eq. 7.

After calculating *f* according to Eq. 7, we add it to the expression for S_W, so that

$$\ln(S^W_t) = \chi^W_t + \xi^W_t + f_t \quad \textbf{Eq. 89}$$

$$E[\ln(S^W_t)] = \chi^W_{t-\Delta t} e^{-\kappa^W \Delta t} + \xi^W_{t-\Delta t} + \mu_\xi \Delta t + f_{t,t-\Delta t} \quad \textbf{Eq. 90}$$

Note that we only do adjustments to the expressions for WTI. Brent still develops according to the standard Schwartz-Smith model. Also note that *f* isn't a stochastic variable, and thus there is no variance associated with it. *f* is a completely deterministic function given by the calculated expectations of Brent and WTI spot prices together with the chosen constant. Therefore, the expression for the variance of log WTI spot prices is still given by Eq. 18.

To understand what value the constant should take on, remember that the constant works as a reversion equilibrium; if expectations excluding *f* indicate a larger difference between Brent and WTI than the constant value, *f* will raise expected WTI price to the level where expected price difference equals the constant. Note that the inclusion of *f* only will give constant price differences for the far future where $\Delta t \to \infty$, meaning where the impact of short-term variations have evaporated. For the near future, short-term variations will still give variations in the expected price gap between Brent and WTI. *f* only serves as a factor that keeps the expected "long-term equilibrium price gap" stable.

Again, we have to decide what we want the "long-term equilibrium price gap" to be. An intuitive answer is that the equilibrium price gap should be the average of all long-term equilibrium price gaps implied by the calibrated joint model. For each day, we calculate the long-term expected price gap using Eq. 73 and Eq. 74. We then take the average of these values and use it as an estimator for the long-term equilibrium price gap.

We can now include *f* in simulating the spot price development. We still simulate state variables using Eq. 56 and Eq. 57 for Brent together with Eq. 59 and Eq. 60 for WTI. Eq. 55 still holds for the calculation of spot price for Brent, but in the expression for spot prices of WTI is replaced by Eq. 89. For each new day, *f* is calculated again based on yesterday's expected values according to Eq. 88. Said another way; $f_{t,t-\Delta t}$ of Eq. 7 is updated for each day *t* based on the expectations of yesterday (t-Δt). This is because the expectations for today is based solely on the latest observed state variables, and there is no dependency whatsoever on the state variables' history before the last observation.

With reference to Eq. 88, Eq. 73 and Eq. 74 we can then give a more formal expression for how f_t is calculated in the simulation process:

$$f_{t,t-\Delta t} = \ln\left(e^{\xi^B_{t-\Delta t} + \frac{\sigma_\chi^{B^2}}{4\kappa^B} + \frac{\rho_{\chi^B\xi}\sigma_\chi^B\sigma_\xi}{\kappa^B} + \left(\mu_\xi + \frac{\sigma_\xi^2}{2}\right)\Delta t} - constant \right) - \xi^W_{t-\Delta t} + \frac{\sigma_\chi^{W^2}}{4\kappa^W} + \frac{\rho_{\chi^W\xi}\sigma_\chi^W\sigma_\xi}{\kappa^W} + (\mu_\xi + \frac{\sigma_\xi^2}{2})\Delta t \qquad \textbf{Eq. 91}$$

f_t works as an adjustment to the time series of spot prices, which develops through the "true process". The forward curve, on the other hand, is a forecast of spot prices following the risk-neutral process. For the model to be consistent, we need to include an adjustment-factor f_t* also along WTI's forward curve. We denote this factor f_t* because it operates under the risk-neutral paradigm. The formula for the forward curve's adjustment factor becomes

$$f^*_{t+\Delta t,t} = \ln\left(E^*\left(S^B_{t+\Delta t,t}\right)|_{\Delta t\to\infty} - constant\right) - \ln\left(E^*\left(S^W_{t+\Delta t,t}\right)|_{\Delta t\to\infty}\right) \qquad \textbf{Eq. 92}$$

where $E^*\left(S^B_{t+\Delta t,t}\right)|_{\Delta t\to\infty}$ is given by Eq. 75 and the expression for $E^*\left(S^W_{t+\Delta t,t}\right)|_{\Delta t\to\infty}$ becomes

$$E^*\left(S^W_{t+\Delta t,t}\right)|_{\Delta t\to\infty} = e^{\xi^W_t + f_t + \frac{\sigma_\chi^{W^2}}{4\kappa^W} + \frac{\rho_{\chi^W\xi}\sigma_\chi^W\sigma_\xi}{\kappa^W} - \frac{\lambda_\chi^W}{\kappa^W} + (\mu^*_\xi + \frac{\sigma_\xi^2}{2})\Delta t} \qquad \textbf{Eq. 93}$$

Note that we include $f_t$in the expression for the before-adjustment risk-neutral expected WTI spot price. We add f_t to all prices on the forward curve of one day, this way it can be interpreted as a general WTI price-level adjuster. $f^*_{t+\Delta t}$ is then an adjuster along the forward curve, but it won't give any great impact as long as we only have maturities up to 50 months. Great deviations between before-adjustment expected prices don't occur before after 5 to 10 years.

After calculating $f^*_{t+\Delta t,t}$, we have to include it in the expression for $E^*\left(S^W_{t+\Delta t,t}\right)$.

The expression for WTI futures prices then becomes:

$$\begin{aligned}\ln\left(F^W_{t+\Delta t,t}\right) &= \ln\left(E^*\left[S^W_{t+\Delta t}\right]\right) \\ &= E^*\left[\ln\left(S^W_{t+\Delta t}\right)\right] + \frac{1}{2}Var^*\left[\ln\left(S^W_{t+\Delta t}\right)\right] \\ &= \chi^W_t e^{-\kappa^W\Delta t} + \xi^W_t + f_t + f^*_{t+\Delta t,t} + A^W(\Delta t)\end{aligned} \qquad \textbf{Eq. 94}$$

where

$$A^W(\Delta t) = -\left(1 - e^{-\kappa^W\Delta t}\right)\frac{\lambda^W_\chi}{\kappa^W} + \mu^*_\xi\Delta t$$

$$+\frac{1}{2}\{\frac{1-e^{-2\kappa^W\Delta t}}{2\kappa^W}\sigma_\chi^{W^2} + \sigma_\xi{}^2\Delta t$$
$$+2(1-e^{-\kappa^W\Delta t})\frac{\rho_{\chi^W\xi}\sigma_\chi^W\sigma_\xi}{\kappa^W}\}$$

9.5.1 Results from Calibrating the Adjusted Model

We can calibrate the model, estimating parameters and state variables that fit the forward curves of WTI (Eq. 94) and Brent (Eq. 71). We then get the following calibration summary:

Table 18 Calibration summary, adjusted two-factor joint model (results for ordinary two-factor joint model in stippled table).

	B	W	B	W
Data sample	[765x50]		[765x50]	
κ	0,438657	0,575944	0,45484	0,55284
λ_χ	-0,05616	-0,10055	-0,06985	-0,11106
μ_ξ*	-0,04044	-0,04044	-0,03731	-0,03731
c_ξ		-0,07803		-0,07892
μ_ξ	0,069964	0,069964	0,06264	0,06264
σ_χ	0,258192	0,312445	0,24196	0,28057
σ_ξ	0,264162	0,264162	0,24403	0,24403
$\rho_{\chi\xi}$	0,192665	0,228419	0,14824	0,19378
$\rho\chi_B\chi_W$	0,839939		0,84827	
SSE	1,351476	2,270491	1,20520	2,30113
$\text{SSE}_{\text{Brent+WTI}}$	3,621967		3,50633	
Constant	5,54674802788032			

Table 19 Calibration constraints, adjusted joint model.

	Lower limit	Upper limit
χ	-6	6
ξ	0	6
κ	1x10^-7	4
$\lambda\chi$	-2	2
μ_ξ*	-4	4
c_ξ	-4	4

It is satisfying to see that the estimated parameters are of about the same order as the parameters of the ordinary joint two-factor model. This is what we expected, since the *f* adjustment is of very small numerical values. The slight increase in SSE values is also no surprise, since the inclusion of the *f* factor serves as a rigid adjustment to the forward curves.

The long-term equilibrium price gap is estimated at $5.55 in favor of Brent. As a comparison, the average price difference of 1st month contracts in our data sample is $5.83. Thus, the value of $5.55 doesn't seem unreasonable on the basis of our data sample; however we know that equilibrium conditions (without bottleneck effects at Cushing) yield price gaps closer to zero. So maybe we should have chosen a value closer to zero as the long-term expected price gap.

9.5.2 Expected Future Prices Implied by the Adjusted Model

The adjusted model including factor *f* gives expected values of WTI approaching a constant distance to the expected values of Brent. For 10000 days the curves of expected values takes on shapes as shown in Figure 37. This graph visualizes what we were looking for when searching for a solution to the problem of growing price gaps; we see that the price gap here is constant when the effect of short-term deviations has faded out.

Figure 37 Expected spot prices using factor *f*.

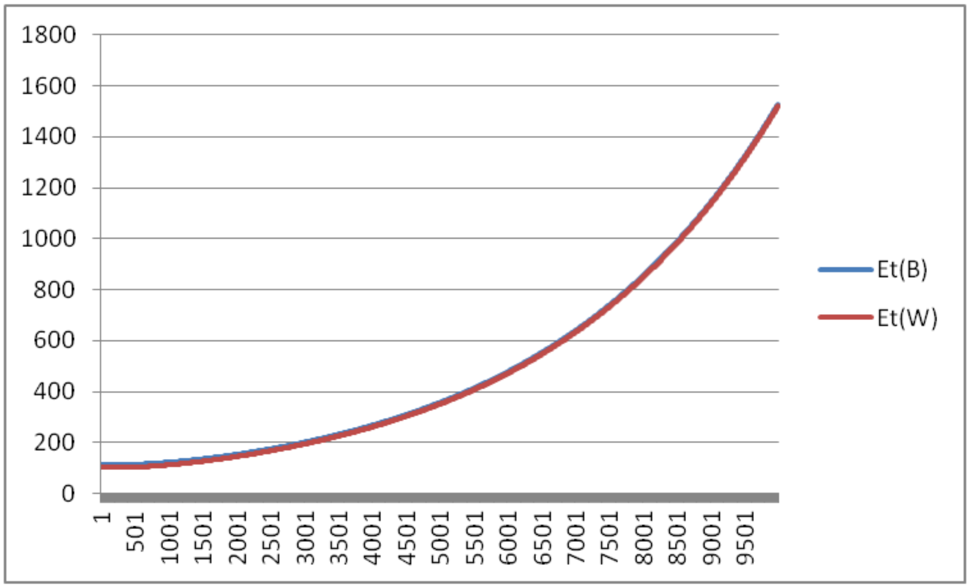

Plotting the expected price spread gives an even clearer picture of the trend indicated by Figure 37; the expected price differential takes some time to settle due to short-term deviations, but eventually it will approach the *constant* value of about $5.55.

Figure 38 Graph of differences in expected Brent and WTI prices, adjusted model.

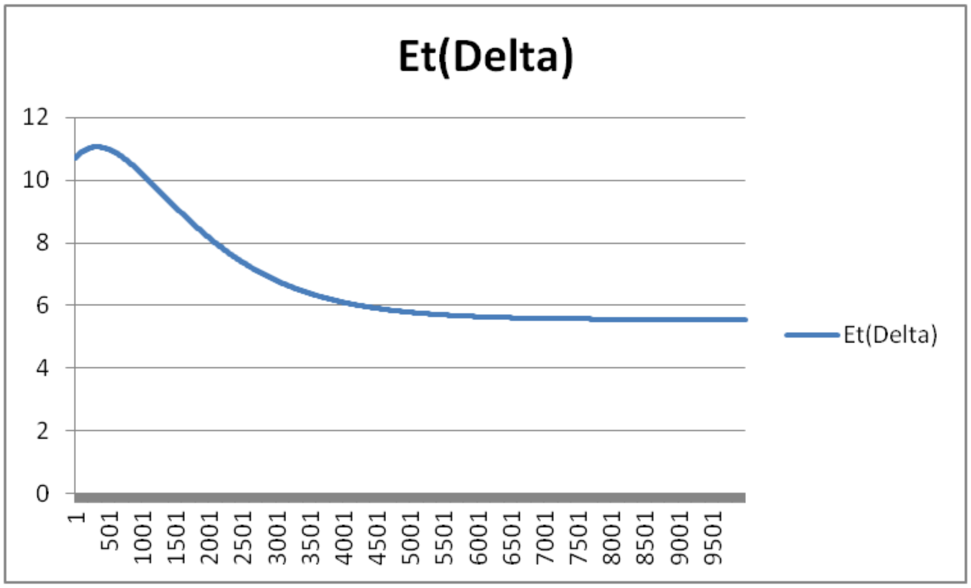

9.5.3 Simulation Results Using the Adjusted Model

We want to show that the adjusted model can be used for simulations. When simulating using the adjusted model, we use the settings given in Table 20. As for previous simulations, the start values are chosen according to the calibration estimates for the start date.

Parameters underlying the simulation are found in Table 18. We use the same time horizon in order to compare the results with result for the ordinary two-factor model.

Table 20 Simulation settings, adjusted joint two-factor model.

NoOfDates	365	
NoOfSim	2350	
SimStartDate	13. januar 2012	
StartValues	BRENT	WTI
χ_0	0,372674081	0,34903568
ξ_0	4,350797708	4,2727661

When comparing the results of Figure 39 to the results of Figure 20 we see that the adjusted model has estimated the Brent spot price on January 13th 2012 slightly below the ordinary model. However, it is likely that the adjusted model's estimate is closer to the true value than the ordinary model's estimate (ref. the discussion about differences in spot price estimates when comparing two-factor and three-factor results). The difference is anyway not big enough to be regarded as a problem.

The adjusted model predicts a higher Brent price after a year than the ordinary model did. The expected price is about $113, while the ordinary model predicted the one-year price to be $111. The higher drift term, affected by increased estimates of μ_ξ and σ_ξ, along with the slower mean reversion of the positive short-term deviation implies faster growth.

Figure 39 Simulated Brent spot prices using adjusted two-factor joint model.

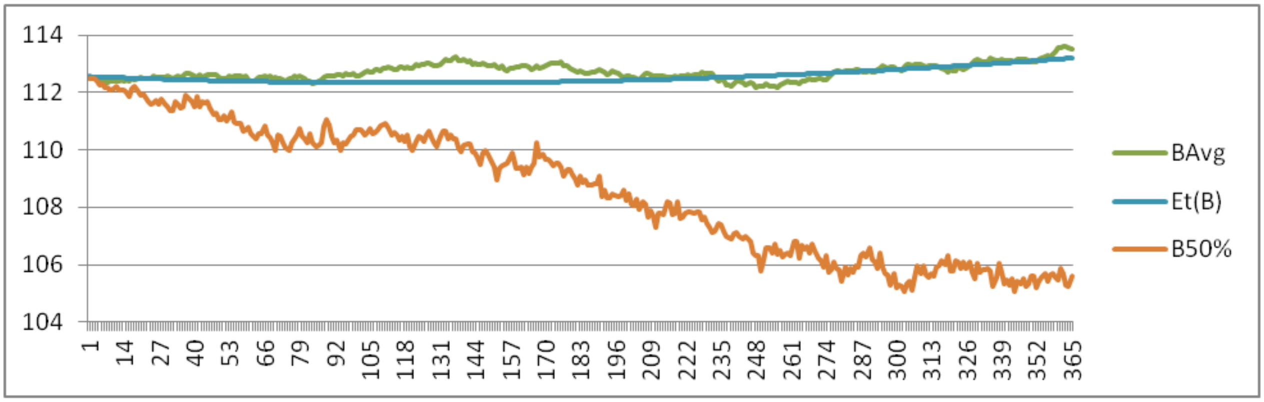

Also for the WTI forecast, the adjusted model predicts faster growth than the ordinary model. However, the growth is not of such a significant character as for the three-factor model (ref. earlier discussion). The adjusted model has maintained the assumption of a constant gap between Brent and WTI prices.

Figure 40 Simulated WTI spot prices using adjusted two-factor joint model.

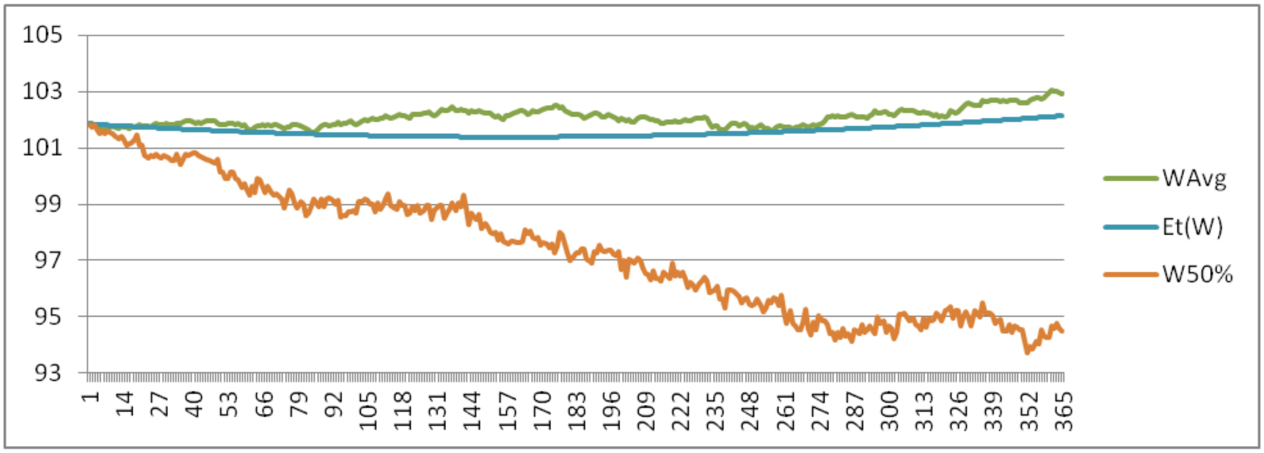

The adjusted model yields slightly higher expected values of both Brent and WTI one year from now. The difference in expected values, on the other hand, is a bit smaller. For the ordinary model expected price gap after a year was close to $11.5, while the adjusted model expects a price gap of $11.0. This is in agreement with the assumption of a slightly lower c_ξ value, the inclusion of the *f* factor and also a WTI short-term variable that has strengthened itself relative to Brent (compare Table 20 to Table 7).

The difference in expected values will, as shown in Figure 38, eventually converge to the value of the *constant* in Eq. 88.

From looking at Figure 23 and Figure 41, we see that the average of simulated price spreads in both cases ends up below the expected price difference. There hasn't been found any explanation to this except the randomness of simulations. For an infinite number of trials, we would expect the average price to converge to the expected price. However, this matter could be a subject for further examination.

Figure 41 Differences in simulated Brent and WTI spot prices using adjusted two-factor joint model.

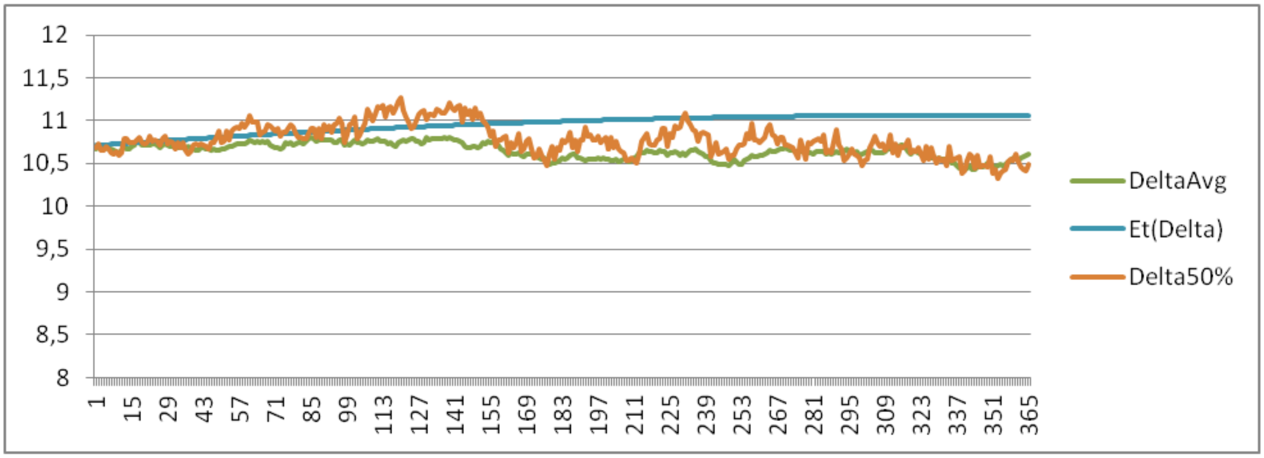

Chapter 10 Summary and conclusion

In this final chapter we present the conclusions of the thesis, and give some thoughts on further work. Main conclusions is that we have been successful in developing a stochastic model for the correlated price development of Brent and WTI, and we have also managed to isolate the deviations in Brent and WTI prices by introducing a second short-term variable to the Schwartz-Smith model.

10.1 Conclusion

Reviewing Brent and WTI futures data show great correlation between Brent and WTI prices. From comparing individual calibrations of the Schwartz and Smith (2000) model on the two products, we conclude that the Brent/WTI price spread principally can be explained by the short-term variable.

In this thesis we have developed a model for price development of two correlated products, using the Schwartz-Smith model as a starting point. We have shown that this model can be calibrated on datasets consisting of observed Brent and WTI futures prices, and used for simultaneously simulating future price development of the two crude oils. A model for the price development of two correlated products could be utilized by companies searching to hedge the risk of cross-product price spreads. It can also be used by investors wanting to speculate on the Brent/WTI spread.

The risk premiums of the Schwartz-Smith model can't be estimated with much accuracy, which is a problem for the forecasting of Brent and WTI prices since we aren't able to obtain reliable estimates of the true process drift parameter. Therefore, great concern and common sense must be applied when using the model for simulations.

The proposed three-factor model gives a better fit to observed data than the original two-factor model. We are also able to designate deviations in Brent and WTI prices to the second short-term variable, which is an appealing feature of the three-factor model.

Re-estimation of the three-factor model parameters confirms the unsatisfying degree of uncertainty in risk premium estimates previously observed for the Schwartz-Smith model.

Simulations performed using the two-factor and the three-factor model shows differences in the WTI price development. This can be shown to be caused by the different assumptions underlying the models. The two-factor model assumes some constant gap between the equilibrium prices of Brent and WTI, whereas the three-factor model assumes identical equilibrium prices. By reviewing historical data, we perceive that the assumptions underlying the three-factor model are the most reasonable. However, these assumptions can easily be implemented into the two-factor model.

As the forecasting horizon becomes long, the models exhibit growing deviations in expected prices of Brent and WTI. This is an undesired feature of the models. We are able to counteract the growing deviations by including an adjustment factor to the WTI process. However, the problem of an increasing expected price gap is only significant for time horizons which are unreasonable to try forecasting using such a model anyway.

10.2 Further work

The problems with estimating the risk premiums of the Schwartz-Smith model undermine its forecasting power. Alternative methods to estimate the risk premiums, for example by looking at exogenous measurements of risk premiums, could add credibility to the model's forecasting power.

The assumption of identical equilibrium prices for Brent and WTI could perhaps be replaced by a constant equilibrium price gap such as that applied for the two-factor model. However, the estimated equilibrium gap was estimated at a rather illogical value when we calibrated the two-factor model. A better way to determine this constant gap could result in a more realistic model than by just setting the gap equal to zero as done for the three-factor model. However, when reviewing the individual calibrations of the Schwartz-Smith model we conclude that the deviations of Brent and WTI equilibrium prices are close to zero anyway.

The problem with divergence of expected Brent and WTI prices for long forecasting horizons could also be a topic for further exploration. The solution presented in this thesis is likely to be one of several possible methods to overcome this problem. Other, more intuitive, adjustments to the joint model could be examined in order to come up with a satisfying solution to the problem.

Also, it is reasonable to believe that the model for correlated price development could be expanded to include more than two products. The most obvious approach would be to continue using Brent as the model's benchmark, and refer the price of new products to it via a level difference c_ξ and a correlated, random variable such as the second short-term variable. Further research on this could end up providing a general stochastic model for prices of *n* correlated products.

Bibliography

Andresen, A., & Sollie, J. M. (2011). *Multi-factor Models and the Risk Premiums: A Simulation Study.* Trondheim.

Brennan, M. J. (1989). *The Price of Convenience and the Valuation of Commodity Contingent Claims.* University of British Columbia.

CME Group. (n.d.). *Light Sweet Crude Oil (WTI) Futures.* Retrieved June 8, 2012, from http://www.cmegroup.com/trading/energy/crude-oil/light-sweet-crude_contract_specifications.html

CME Group. (n.d.). *NYMEX WTI Crude Oil, Natural Gas, Heating Oil, and RBOB Gasoline Futures Daily Settlement Procedure.* Retrieved June 8, 2012, from http://www.cmegroup.com/trading/energy/files/NYMEX_Energy_Futures_Daily_Settlement_Procedure.pdf

Cortazar, G., & Schwartz, E. S. (2003). Implementing a stochastic model for oil futures prices. *Energy Economics 25(3)* , pp. 215-238.

Fattouh, B. (2011). *An Anatomy of the Crude Oil Pricing System.* Oxford Institue for Energy Studies.

Gabillon, J. (1995). Analyzing the Forward Curve. In ENRON, *Managing Energy Price Risk* (pp. 29-38). London: Risk Publications.

Gabillon, J. (1991). *The Term Structures of Oil Futures Prices.* Oxford Institute for Energy Studies.

Gibson, R., & Schwartz, E. S. (1990). Stochastic Convenience Yield and the Pricing of Oil Contingent Claims. *Journal of Finance 45(3)* , pp. 959-976.

Gue, E. (2011, August 28). *What Does The Growing Spread Between WTI and Brent Crude Oil Mean for Energy Holdings?* Retrieved June 7, 2012, from http://seekingalpha.com/article/290273-what-does-the-growing-spread-between-wti-and-brent-crude-oil-mean-for-energy-holdings

Haug, E. G. (2007). *The complete guide to option pricing formulas (2nd ed.).* McGraw-Hill.

ICE Futures Europe. (2008, November). *Exchange Futures for Physical (EFPs) for ICE Brent Futures*. Retrieved June 8, 2012, from https://www.theice.com/publicdocs/futures/ICE_Brent_EFP_Explained.pdf

ICE Futures Europe. (n.d.). *ICE Brent Crude Futures*. Retrieved June 8, 2012, from https://www.theice.com/productguide/ProductSpec.shtml?specId=219#

ICE Futures Europe. (n.d.). *The Brent Index*. Retrieved June 8, 2012, from https://www.theice.com/publicdocs/futures/ICE_Futures_Europe_Brent_Index.pdf

IEA (International Energy Agency). (2011, November 9). *The world is locking itself into an unsustainable energy future which would have far-reaching consequences, IEA warns in its latest World Energy Outlook*.

Retrieved June 7, 2012, from http://www.iea.org/newsroomandevents/pressreleases/2011/november/name,20318,en.html

Iordanova, T. (2007, 18 September). *Introduction To Stationary And Non-Stationary Processes*. Retrieved June 7, 2012, from http://www.investopedia.com/articles/trading/07/stationary.asp#axzz1uB6XNIN8

Lucia, J. J., & Schwartz, E. S. (2002). Electricity Prices and Power Derivatives: Evidence from the Nordic Power Exchange. *Review of Derivatives Research 5(1)* , 5-50.

McDonald, R. L. (2006). *Derivatives Markets (2nd edition).* Pearson Education, Inc.

Norgate Investor Services. (n.d.). *Understanding Futures Market Data*. Retrieved June 08, 2012, from http://www.premiumdata.net/support/understandingfuturesdata.php

Pickrell, E. (2012, January 16). *Lack of pipelines drives price gap between Brent and WTI Index*. Retrieved June 7, 2012, from http://fuelfix.com/blog/2012/01/16/lack-of-pipelines-drives-price-differences-between-brent-and-wti-index/

Platts. (2010). *Derivatives Methodology Guide: Platts Forward Curve-Oil (December).* The McGraw Hill Companies.

Platts. (2011). *Methodology and Specifications Guide: Crude Oil (June).* The McGraw Hill Companies.

Samuelson, P. A. (1965). Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review 6(2)* , 41-49.

Schwartz, E. S. (1997). The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging. *The Journal of Finance 52(3)* , pp. 923-973.

Schwartz, E., & Smith, J. E. (2000). Short-Term Variations and Long-Term Dynamics in Commodity Prices. *Management Science 46(7)* , 893-911.

List of Tables

List of Figures