Norwegian School of Economics
Bergen, Spring, 2019

Mean Reversion and Seasonality in Heating Oil Futures

Ola Vold Rennemo

Supervisor: Jørgen Haug

Master thesis, Finance

NORWEGIAN SCHOOL OF ECONOMICS

Takk til veileder Jørgen Haug for motiverende samtaler og konstruktive innspill.

Takk til Jørgen Vold Rennemo for korrekturlesing.

Takk til venner og familie for at dere viste interesse for prosjektet og stilte ører til disposisjon.

Contents

Introduction 1

Section 1 3

1.1 Heating Oil 3

1.2 NY Harbor ULSD Futures (HO) 8

Section 2 10

2.1 The Commodities Literature 10

2.2 Mathematical Representation of the Model 12

2.3 Nested Models 14

2.4 Futures Prices and Volatility Term Structure of Futures Returns 18

2.5 European Call Option on a Futures Contract 20

2.6 Seasonality Function 21

Section 3 22

3.1 Estimation Method 22

3.2 The Kalman Filter 22

3.3 State Space Representation of a Model Defined By (1), (2) and (3) 25

3.4 Simulation 26

Section 4 29

4.1 The Data 29

Section 5 33

5.1 Estimated Parameters 34

5.2 Likelihood-Ratio Tests 35

5.3 Estimated Seasonality Function 36

5.4 Accuracy 38

5.5 Volatility Term Structure of Futures Returns 41

5.6 One-Factor, Three-Factor and Four-Factor Model 42

Section 6 45

6.1 Practical Implications 45

6.2 Conclusions 49

6.3 Assumptions and Further Works 51

References 52

Appendix 55

Production, Trade and Financial Markets 55

Nested Model Relation Equations 56

Abstract

We examine mean reversion and seasonality in heating oil futures prices using an affine N-factor Gaussian model and NY Harbor ULSD futures. We find strong empirical evidence for mean reversion and seasonality in heating oil futures prices.

Introduction

The valuation and hedging of commodity contingent claims are of great importance to commodity producers, commodity consumers, financial intermediaries, and speculators. Consequently, appropriate modeling and estimation of the stochastic behavior of commodity prices are vital for these market participants' applications of sound investment and risk-management strategies.

In this thesis, we consider an affine N-factor Gaussian model developed by Cortazar and Naranjo (2006), and study its ability to explain the stochastic behavior of heating oil futures prices. In doing so, we determine the parsimonious number of latent unobservable factors to include in the N-factor model, which depends on the existence and scale of mean reverting properties in heating oil's price process under an equivalent martingale measure. Furthermore, we examine seasonality in heating oil futures prices, which depends on the existence and scale of seasonality in heating oil's price process under an equivalent martingale measure. Lastly, we examine some implications of our findings regarding the mean reverting properties and seasonality in heating oil futures prices, utilizing model estimated heating oil futures prices and model estimated prices of a European call option on a heating oil futures contract.

The base model for this thesis is Schwartz and Smith's (2000) two-factor model, which is a two-factor version of Cortazar and Naranjo's (2006) N-factor model. The seasonal adjustment to this base model is modeled as a deterministic continuous trigonometric function with seasonal frequencies, similar to the seasonal adjustment suggested by Sorensen (2002).

The model parameters are estimated using maximum-likelihood techniques and the Kalman filter. We conduct a simulation study of the parameter estimation procedure, following Andresen and Sollie (2013), to highlight potential econometric issues. We find that the models are robust for use in valuation problems and generally not robust for forecasting purposes, confirming Schwartz and Smith's (2000) and Andresen and Sollie's (2013) findings.

In Cortazar and Naranjo's (2006) study of crude oil price modeling, the authors found grounds for using three- or four-factor models, suggesting two or three mean reverting factors in crude oil's price process under an equivalent martingale measure. Furthermore, the authors found that the accuracy of model estimated futures prices given by a two-, three- and four-factor model was comparable for futures maturities between 3 months and 3 years, and that for maturities between 3 and 7 years the two-factor model was substantially less accurate than the three- and four-factor models. Crude oil is a major cost driver for heating oil as heating oil is refined from crude oil. Hence, mean reversion in heating oil's price process under an equivalent martingale measure is probable. Indeed, based on parameter estimates of a two-factor model, fitted on NY Harbor ULSD futures, we find significant mean reversion in heating oil's price process under an equivalent martingale measure. However, we could not find grounds for including more than one mean reverting factor in the model. The maximal time to maturity of the futures contracts used to estimate parameters in this thesis is approximately 3 years, so we believe the difference in maximal contract maturity is the root cause of the difference in our findings regarding the appropriate number of mean reverting factors. Still, we cannot rule out that the difference in findings reflect differences in the two commodities' price process under an equivalent martingale measure.

Heating oil's use as a heating fuel implies that there is an area dependent seasonal distillate fuel oil consumption pattern that is inversely related to the temperature in that area. We find signs of this seasonality in the U.S. distillate fuel oil inventory, and there appears to be a positive correlation between changes in the U.S. distillate fuel oil inventory and changes in the expected development of heating oil prices. Hence, seasonality in heating oil's price process is probable. Indeed, based on likelihood-ratio tests of the two-factor base model and seasonally adjusted two-factor base models, we find that there is statistically significant seasonality in heating oil's price process under an equivalent martingale measure. Furthermore, we find that allowing for more than two inflection points to capture this seasonality provided a statistically significantly better fit. In Roberts and Lin's (2006) study of seasonality in the price process for different commodities the authors found significant seasonality in crude oil's price process under an equivalent martingale measure, though the authors could not find grounds for allowing more than two inflection points in describing this seasonality. The contrast to our results indicates that there are significant differences between the seasonality of the two commodities' price process under an equivalent martingale measure.

The outline for this thesis is as follows. In Section 1, we present an overview of heating oil and NY Harbor ULSD futures. In Section 2, we provide a brief overview of the commodity price model literature, present the chosen model, and derive some of its properties, which are important for parameter estimation and valuation and hedging of commodity contingent claims. In Section 3, the parameter estimation procedure is presented and tested in a controlled environment through simulation. In Section 4, the NY Harbor ULSD futures price panel is presented. In Section 5, estimation results are presented and reviewed. Lastly, in Section 6, we provide examples of practical implications of our findings, conclude, discuss assumptions and questions open for further study.

Section 1

In this section we present an overview of heating oil and NY Harbor ULSD futures, and extract some features of these that will be used for appropriately modeling the heating oil price process.[1]

1.1 Heating Oil

Heating oil is one of the petroleum products produced by refining crude oil (EIA (c), 2019). Heating oil is classified as a distillate alongside other petroleum products with similarities in chemical make-up (EIA (a), 2019). In this thesis, we use the distillate naming convention used by the United States Energy Information Administration (EIA). In EIA's (a) (2019) naming convention, No. 2 fuel oil (heating oil) and No. 2 diesel fuel are considered as different kinds of No. 2 distillate, and No. 2 distillate is one of the distillates considered as distillate fuel oil. No. 2 diesel fuel is mainly used in diesel engines and No. 2 fuel oil (heating oil) is mainly used in atomizing type burners. To better align this naming convention with heating oil futures and to ease reading we use "heating oil" as a substitute for "No. 2 fuel oil (heating oil)".

[1] Hull (2012), defines a futures contract as a contract that obligates the holder to buy or sell an asset at a predetermined delivery price during a specified future time period with a given settlement scheme.

Heating oil is mainly used for commercial and residential space and water heating (EIA (c), 2019). In the United States, the northeastern region accounts for most of the heating oil consumption, and the other regions of the United States typically use Natural Gas for space and water heating (EIA (c), 2019).

There are numerous factors driving fluctuation in heating oil prices, and we will now discuss the most important ones.

Figure 1 shows daily Cushing, OK Crude Oil Future Contract 1 (Dollars per Barrel) prices, daily New York Harbor No. 2 Heating Oil Future Contract 1 (Dollars per Gallon) prices multiplied by 42 and the daily heating oil crack spread from January 2, 2001 to December 12, 2018. The time series are sourced from EIA (2019) which defines Contract 1 as a futures contract specifying the earliest delivery date, often called the front month. As both products have, for each month in a year, a single futures contract with delivery date in that month, the average time to delivery for Contract 1 is approximately 11 business days.[2]

The heating oil crack spread represents a theoretical profit margin for heating oil refineries. It is defined as the price difference between heating oil futures and crude oil futures with the same delivery month and the same unit (CME Group (b), 2019). As 42 gallons is equivalent to one barrel we derive the crack spread by multiplying the heating oil dollars per gallon futures prices by 42 and subtract crude oil dollars per barrel futures prices.

As futures prices tend towards the spot price when the time to maturity goes to zero (Hull, 2012), the futures prices shown in Figure 1 provides a good approximation for their respective spot prices.

[2] We're assuming 250 business days in a year such that average time to delivery is $\left(\frac{250}{12}+1\right)\times\frac{1}{2}\approx 11$ business days.

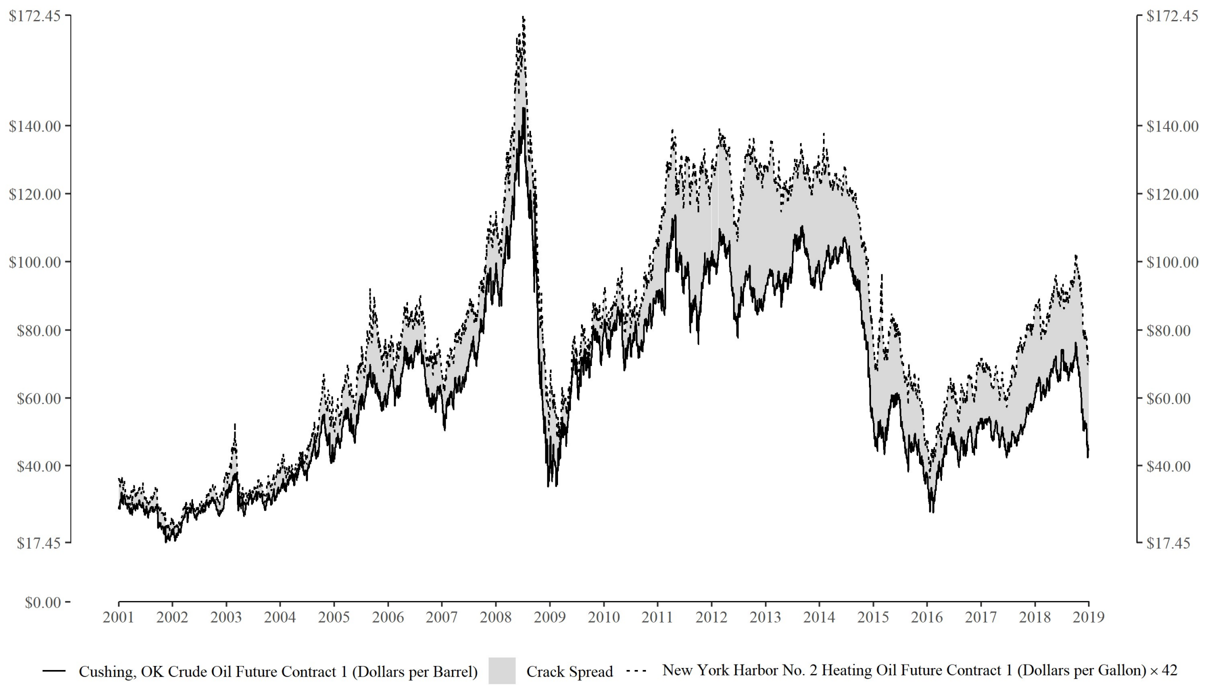

Figure 1. Crude oil front month, heating oil front month and crack spread. Daily Cushing, OK Crude Oil Future Contract 1 (Dollars per Barrel), New York Harbor No. 2 Heating Oil Future Contract 1 (Dollars per Gallon) times 42 and heating oil crack spread.

From Figure 1 we see that heating oil and crude oil prices are very volatile and highly positively correlated. Moreover, since the cost of crude oil is a primary cost driver in heating oil production, we can assert that the price of heating oil generally follows the cost of crude oil. Consequently, the factors that impact worldwide supply and demand for crude oil and expected supply and demand for crude oil will indirectly impact the price of heating oil. As an example, the pre-2008 crude oil price growth was driven by supply factors like production falls due to lack of investments and the Venezuelan oil worker strike, expected supply factors like decline in Saudi spare capacity and concerns about stability in the Middle East and demand factors like rapid economic development in Asia (World Economic Forum, 2016). The price drop in 2008 was caused by the global financial crisis (World Economic Forum, 2016). Moreover, heating oil’s price dependence on the cost of crude oil suggests similarity between heating oil’s price process and crude oil’s price process, and as crude oil’s price process has been extensively studied in the commodities literature, we can find useful guidance regarding appropriate modeling of heating oil’s price process in the commodities literature (see Section 2).

Figure 1 also shows that there have been substantial differences in the yearly average heating oil crack spread from 2001 to 2018. Though it’s less visually apparent from Figure 1, the short-term variability in the heating oil crack spread have also been substantial from 2001 to 2018. The fluctuation in the heating oil crack spread is driven by current and expected changes in supply and demand for heating oil. EIA (b) (2019) segments factors which impact current and expected changes in supply and demand into the categories Consumption, Production, Trade, Balance and Financial Markets. In the following we will define and examine the categories Consumption and Balance (see the Appendix for a brief description of the remaining categories). Statements regarding Consumption and Balance are sourced from EIA (b) (2019) unless otherwise specified.

Consumption is directly linked to demand, and at present distillate fuel oil is the most consumed petroleum product worldwide. Global distillate fuel oil consumption has increased substantially from 1999 to 2016, with an annual continuous growth rate of approximately 2%.

The observed distillate consumption growth is partially a reflection of the robust economic growth several developing countries have shown in this time window. Economic growth in developing countries tends to increase distillate fuel oil demand as economic growth in developing countries is usually centered around growth in the industry rather than services. Conversely, if the global economy were to become more services-oriented one would expect a decline in distillate consumption.

Heating oil's use as a heating fuel drives an area dependent seasonal distillate fuel oil consumption pattern that is inversely related to the temperature in that area. In addition, this relationship can lead to substantial unexpected spikes in consumption caused by deviations from expected weather patterns.

The Balance category deals with inventory management questions induced by the storability of distillates and other petroleum products. Hence, Balance is linked to supply. Essentially, expected price growth encourages inventory builds to satisfy future demand while expected price decline encourages inventory drawdowns to satisfy current demand.

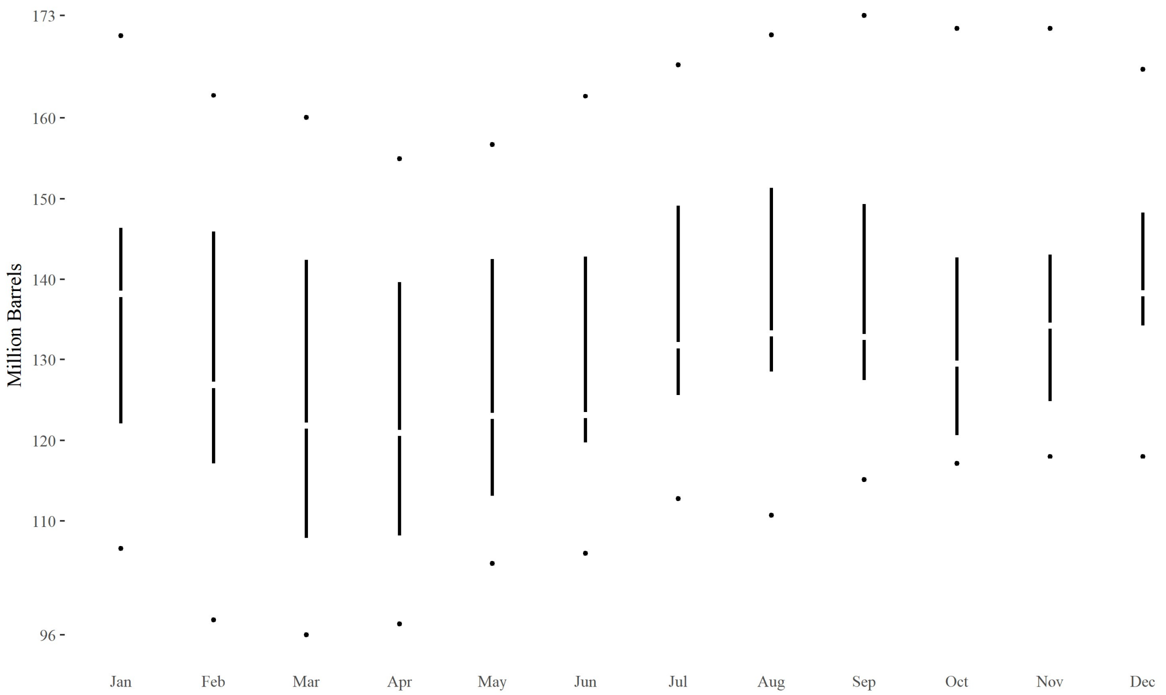

Figure 2. EOM U.S distillate fuel oil inventory. Boxplot of 2000 to 2018 end of month U.S distillate fuel oil inventory measured in million barrels. White points indicate the median, lines indicate the interquartile range and black points indicate the minimum and maximum values.

Figure 2 shows a boxplot of 2000 to 2018 end of month United States distillate fuel oil inventory measured in million barrels. We see signs of the seasonal consumption pattern of heating oil in the inventory drawdowns in the late winter months and inventory buildups prior to the winter months. Furthermore, the interquartile inventory range is relatively narrow in December and wide in March fitting the seasonal consumption pattern of heating oil through increased uncertainty regarding future supply and demand relations as winter progress. There is a clear inventory drawdown tendency in October that is difficult to link to the seasonal consumption pattern of heating oil. However, this does explain the tendency to start builds as early as in July.

Figure 3 shows year-over-year end-of-quarter change in United States distillate fuel oil inventory measured in million barrels, and year-over-year end-of-quarter change in NY Harbor ULSD futures 3rd less 1st-month prices measured in cents from 2001 to 2019. We see that distillate inventory drawdowns and builds is related to the market expectation of future price development.

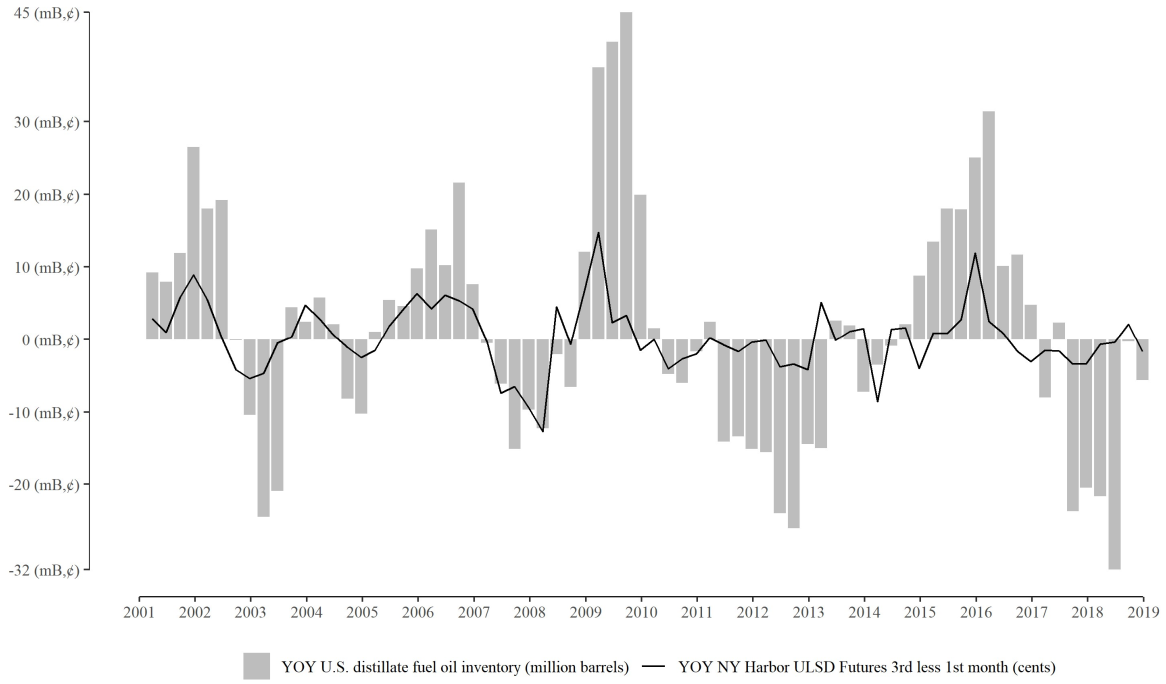

Figure 3. YOY comparison of distillate fuel oil inventory and heating oil futures spread. Year-over-year end-of-quarter change in U.S distillate fuel oil inventory (million barrels) and year-over-year end-of-quarter change in NY Harbor ULSD Futures 3rd less 1st-month price (cents).

Hence, seasonality in heating oil's price process is probable as heating oil have an area-dependent seasonal consumption pattern, and there are signs of seasonality in U.S. Distillate fuel oil inventory and there appears to be a positive correlation between changes in U.S. Distillate fuel oil inventory and changes in the expected development of heating oil prices.

1.2 NY Harbor ULSD Futures (HO)

The heating oil futures used in this thesis is NY Harbor ULSD futures, ticker symbol HO. NY Harbor ULSD futures are New York Mercantile Exchange (NYMEX) listed futures with daily settlement and physical delivery at New York Harbor (NYH) upon expiration (CME Group (a), 2019). The NY Harbor ULSD futures contract unit is 42000 gallons, priced in USD with minimal fluctuation of $0.0001 per gallon (CME Group (a), 2019). At a given point in time different NY Harbor futures contracts are listed, maturing in the current month, one of the remaining months in the current year, one of the months in the three years following the current year and the month following the three years following the current year (CME Group (a), 2019). The futures are settled by CME Group staff based on information gathered from trading activity like prices, volume and calendar spreads, using one set of rules for the closest to maturity contract and a different set of rules for all other contracts (CME Group (c), 2019). As settlement prices calculated by CME Group staff are used for marking to market all account balances, they are widely scrutinized by all market participants and therefore represent a very accurate measure of the true market price (Trolle & Schwartz, 2008).

In May 2013, the New York Mercantile Exchange altered the contract specification for the heating oil futures contracts to the ultra-low sulfur diesel specification (ULSD) (EIA (d), 2019). The ULSD specification requires that the delivered distillate fuel oil contains less than 15 parts per million (ppm) of sulfur (EIA (d), 2019), which is in line with the specification of most diesel fuels (EIA (d), 2019). As a reference, distillate fuel oil with sulfur content up to 2000 ppm have historically been allowed for delivery on heating oil futures contracts (EIA (d), 2019). In general, the switch makes it easier for market participants to hedge their distillate fuel oil investments, unless they specifically want to hedge traditional heating oil, which has a higher sulfur content (EIA (d), 2019).

In 2012, New York imposed a sulfur content limit on distillate fuel oil used for heating purposes and other northeastern states plan to follow suit (EIA (d), 2019). Therefore, though the ticker symbol is HO, which stands for heating oil, and the actual futures contract name contain “Diesel Fuel”, it’s more accurate to think of the underlying product in terms of the broader category distillate fuel oil with a sulfur restriction.

A futures contract’s liquidity is given by how frequently it is traded, volume, and its total number of outstanding positions, open interest (Hull, 2012). In this thesis, we make extensive use of theoretical financial results which assume infinite liquidity, so that high or at the very least decent liquidity is required for the results to be valid/applicable. Figure 4 shows yearly aggregated NY Harbor ULSD futures volume in the time window 2012 to 2018. We see that NY Harbor ULSD futures are frequently traded and that the trading volume has increased from 2013 onwards, possibly due to the contract specification change making the NY Harbor ULSD futures better suited to hedge other low sulfur distillates or a reflection of distillate fuel oil consumption growth. According to the Futures Industry Association, (FIA, 2019), NY Harbor ULSD futures was the 13th most traded future or option in 2018 worldwide. As a reference point, the second and fourth most traded future or option in the world in 2018 were the NYMEX traded WTI Light Sweet Crude Oil futures and Henry Hub Natural Gas futures with 2018 volumes of 306.5 and 114.3 million contracts traded, respectively (FIA, 2019). 2018 end of year open interest figures for WTI Light Sweet Crude Oil futures, Henry Hub Natural Gas futures and NY Harbor ULSD futures were 2, 1.2, and 0.4 million contracts outstanding, respectively (FIA, 2019). Hence, the liquidity of NY Harbor ULSD futures may be classified as high.

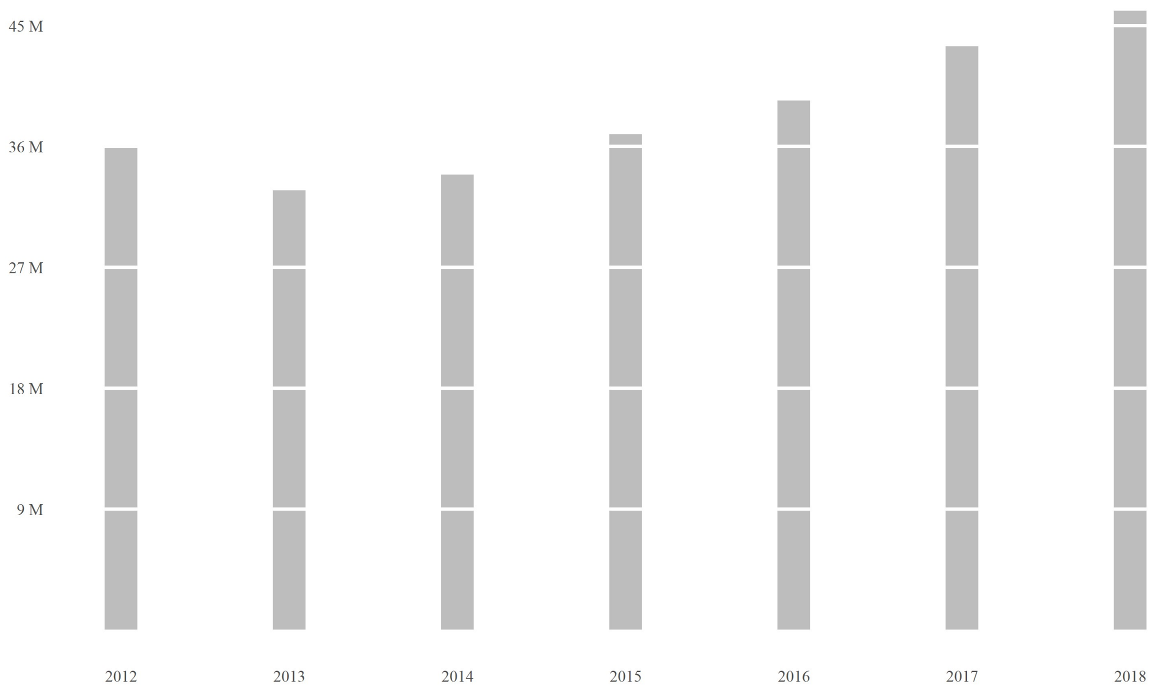

Figure 4. NY Harbor ULSD futures volume. Aggregated yearly NY Harbor ULSD futures volumes, measured in million contracts traded.

Section 2

In this section we provide a brief overview of the commodity price model literature, building on Cortazar and Naranjo's (2006) overview while focusing on the parts relevant for appropriate specification of a heating oil price model. Furthermore, we present our chosen model and derive some of its properties. The derived properties are either required for parameter estimation or properties that are important for valuation and hedging of commodity contingent claims.

2.1 The Commodities Literature

There are many different models of the stochastic process followed by commodity prices proposed in the literature. In general, the models differ in how they specify cost of carry and spot price innovations. The cost of carry summarizes the relationship between futures prices and spot prices and is defined as the storage cost of an asset plus the interest that is paid to finance the asset less the income earned on the asset (Hull, 2012). The income earned on the asset, benefitting the asset holder and not the futures contract owner, is in the commodities literature called convenience yield and is often represented as a dividend yield (Hull, 2012).

A frequently used and highly tractable model of commodity prices is the one-factor geometric Brownian motion model for the commodity spot price with a constant interest rate and convenience yield, proposed by Brennan and Schwartz (1985). This model implies a constant cost of carry and exhibits a constant volatility term structure of futures returns. This latter property contradicts the Samuelson Hypothesis, which predicts increasing futures price return volatility as contract expiration nears (Bessembinder, Coughenour, Seguin, & Smeller, 1996).

Bessembinder, Coughenour, Seguin, and Smeller (1996) showed that the Samuelson Hypothesis should be expected to hold if there is negative covariation between spot returns and changes in the slope of the futures term structure.[3] Furthermore, they argue and show empirically that the hypothesis will hold in markets where spot prices are mean reverting and the reversion is associated with variation in the cost of carry. These properties are typically found in commodity markets where convenience yields exist and display substantial intertemporal variation. Notably, Bessembinder, Coughenour, Seguin, and Smeller (1996) found strong support for the Samuelson Hypothesis in markets for crude oil. Moreover, several studies published after 1996 (see e.g. Schwartz (1997), Schwartz and Smith (2000), Cortazar and Schwartz (2003), Roberts and Lin (2006) and Cortazar and Naranjo (2006)) have found significant mean reversion in crude oil spot prices indicating that Bessembinder, Coughenour, Seguin, and Smeller's (1996) findings are persistent for crude oil markets. As explained in Section 1, the cost of crude oil is a major cost driver for heating oil, and as several studies have found persistent mean reversion in crude oil spot prices, mean reversion in heating oil spot prices is probable.

Several alternative one-factor models of commodity prices that account for mean reversion in commodity prices have been proposed. A remaining issue, common for all one-factor models, is that under the assumption that there is only a single source of uncertainty, it follows that futures prices with different maturities are assumed to be perfectly correlated, which contradicts existing evidence (Cortazar & Naranjo, 2006).

[3] Bessembinder, Coughenour, Seguin, and Smeller (1996) defines "the slope of the futures term structure" as s_t where the value of s_t is given by $F_{t,T} = P_t \times e^{s_t(T-t)}$, where $F_{t,T}$ denotes the futures price at time t for a contract maturing at time T and P_t denotes the spot price of the underlying product of the futures contract at time t.

In later years, commodities price models with multiple sources of uncertainty have been proposed. These models have in previous empirical studies outperformed single-factor models with respect to futures price and volatility term structure fit metrics (see e.g. Schwartz (1997) and Cortazar and Naranjo (2006)).

The model chosen for this thesis is Cortazar and Naranjo's (2006) N-factor Gaussian model, derived from Dai and Singleton's (2000) $A_0(N)$ canonical representation of interest rates models, using Schwartz and Smith's (2000) two-factor model as the base model.[4]

We chose this model because it is tractable for an arbitrary number of risk factors, flexible with regards to mean reverting properties and seasonality, and nests most Gaussian models commonly found in the literature. This latter property is highly desirable as we want our findings to be as independent of the model specification as possible.

2.2 Mathematical Representation of the Model

We now turn to the precise mathematical description of Cortazar and Naranjo's (2006) model. We begin by defining the model under the subjective probability measure (Miltersen, 2003), also called the original probability measure (Miltersen, 2003), the real-world measure (Hull, 2012) and the P-measure (Hull, 2012), often denoted P. The subjective probability measure is the probability measure that investors have and use to calculate expectations, variances, covariances, etc., of future prices and returns (Miltersen, 2003). At a later stage, under "Nested Models", we define the model under an equivalent martingale measure, often, and hence forth denoted Q, which will be used for pricing purposes. We denote Wiener increments under P dw, and Wiener increments under Q dw^Q.

Let

$$ln(S_t) \quad = \quad \mathbf{1}^\top \boldsymbol{x}_t + s(t), \tag{1}$$

where $ln(S_t)$ is the natural logarithm of a spot price at time t, $s(t)$ is a deterministic function of time and

$$\boldsymbol{x}_t \quad = \quad \begin{pmatrix} x_t^1 \\ x_t^2 \\ \vdots \\ x_t^m \end{pmatrix}$$

[4] Cortazar and Naranjo (2006) use "Gaussian" when referring to their own model and equivalent models found in the commodities literature, where Gaussian refers to distributional properties of the model.

is a vector of latent unobservable state variables. The vector of state variables follows the process

$$d\boldsymbol{x}_{\mathrm{t}} \quad = \quad (-\mathrm{K}\boldsymbol{x}_{\mathrm{t}} + \boldsymbol{\beta})dt + \Sigma d\mathbf{w}_{\mathrm{t}}, \tag{2}$$

where

$$\mathrm{K} \quad = \quad \begin{pmatrix} \kappa_1 \equiv 0 & 0 & \cdots & 0 \\ 0 & \kappa_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \kappa_m \end{pmatrix}, \kappa_i > 0 \; \forall \; \mathrm{i} \in \{2,3,\dots,\mathrm{m}\},$$

$$\boldsymbol{\beta} \quad = \quad \begin{pmatrix} \mu \\ 0 \\ \vdots \\ 0 \end{pmatrix},$$

$$\Sigma \quad = \quad \begin{pmatrix} \sigma_1 & 0 & \cdots & 0 \\ 0 & \sigma_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sigma_m \end{pmatrix}, \sigma_i \geq 0 \; \forall \; \mathrm{i} \in \{1,2,\dots,\mathrm{m}\}$$

and

$$d\boldsymbol{w}_t \quad = \quad \begin{pmatrix} dw_t^1 \\ dw_t^2 \\ \vdots \\ dw_t^m \end{pmatrix}.$$

Here $d\boldsymbol{w}_t$ is a vector of correlated Wiener increments such that $d\boldsymbol{w}_t d\boldsymbol{w}_t^{\top} = \Omega dt$,

where

$$\Omega \quad = \quad \begin{pmatrix} 1 & \rho_{21} & \cdots & \rho_{m1} \\ \rho_{21} & 1 & \cdots & \rho_{m2} \\ \vdots & \vdots & \ddots & \vdots \\ \rho_{m1} & \rho_{m2} & \cdots & 1 \end{pmatrix}$$

and $\rho_{ij} \in [-1,1] \; \forall \; \mathrm{i,j} \in \{1,2,\dots,\mathrm{m}\}$ are the instantaneous correlations between state variables i and j.

From this model definition, it follows that the state variables have a multivariate normal distribution.

x_1 is defined as a geometric Brownian motion with drift coefficient μ and diffusion coefficient σ_1. Therefore, x_1 is non-stationary for $\sigma_1 \neq 0$ and $\mu \neq 0$, which introduces the possibility of non-stationarity in the spot price process. There are several ways to remove this non-stationarity possibility, e.g. by setting $\kappa_1 > 0$, $\mu \equiv 0$ and $s(t) = \tilde{\mu} \times t$. However, Cortazar and Naranjo (2006) found no significant difference in comparing model estimates with and without a non-stationarity assumption imposed for crude oil markets. A look at Figure 1 indicates that allowing the possibility of a unit root is appropriate for heating oil markets. Moreover, allowing the possibility of a unit root is in line with traditional models found in the commodities literature (Cortazar & Naranjo, 2006).

The changes in x_1 capture changes in supply or demand, caused by for instance changes in technology and preferences, that are expected to persist (Schwartz & Smith, 2000).

The remaining state variables, $x_2, x_3, \dots, x_m$, follow an Ornstein–Uhlenbeck process that reverts to zero at a mean reversion rate given by $\kappa_2, \kappa_3, \dots, \kappa_m$, respectively. Changes in these state variables capture transitory changes in supply and demand, caused by for instance weather variations and irregular supply disruptions (Schwartz & Smith, 2000). Furthermore, the effect of time to maturity is captured by these state variables (Power & Turvey, 2008). This will become apparent from the closed form formulas for futures and option prices derived at a later stage in Section 2.

Seasonality will in this thesis be modeled through $s(t)$, which we define in the last subsection of Section 2.

2.3 Nested Models

Cortazar and Naranjo's (2006) model is unusual in that it makes no mention of convenience yields. Yet, choosing parameters and assumptions correctly it can be specialized to models equivalent to, "nests", most existing two and three-factor Gaussian models found in the commodities literature, Gibson and Schwartz (1990), Schwartz (1997), Roberts and Fackler (1999), Schwartz & Smith (2000), Cortazar and Schwartz (2003), among others, and can therefore be expressed in terms of the parameters of these models. This nested models property is desirable because it ensures that empirical findings using Cortazar and Naranjo's (2006) model are directly transferable to a wide range of commodity price models found in the literature. In other words, we get tractability and flexibility with respect to mean reverting properties and seasonality without loss of generality or theoretical soundness.

Following Cortazar and Naranjo (2006), the nested model property can be derived using the following series of arguments:

Let $\boldsymbol{y}_t$ be the state vector of a model that can be written on the form:

$$ln(S_t) = \bar{\boldsymbol{h}}^\top \boldsymbol{y}_t + \bar{s}(t)$$

$$d\boldsymbol{y}_t = \left(-\bar{K}\boldsymbol{y}_t + \bar{\boldsymbol{\beta}}\right)dt + \bar{\Sigma} d\bar{\mathbf{w}}_t$$

Where $d\bar{\boldsymbol{w}}_t d\bar{\boldsymbol{w}}_t^\top = \bar{\Omega} dt$, $\bar{s}(t)$ is a deterministic function of time, $\bar{\boldsymbol{h}}$ and $\bar{\boldsymbol{\beta}}$ are vectors and $\bar{K}$ and $\bar{\Sigma}$ are matrices.

Next, consider a model with state vector $\boldsymbol{z}$ with the same dimensions as $\boldsymbol{y}$ defined by an affine transformation of $\boldsymbol{y}$ such that $\boldsymbol{z}_t = L\boldsymbol{y}_t + \boldsymbol{\varphi}(t)$, where the matrix L is invertible and $\boldsymbol{\varphi}(t)$ is a vector of deterministic functions of time with dimensions matching the dimensions of $\boldsymbol{y}_t$. There is a one-to-one relationship between the state vectors $\boldsymbol{z}_t$ and $\boldsymbol{y}_t$. Moreover, it follows that

$$ln(S_t) = \boldsymbol{h}^\top \boldsymbol{z}_t + s(t)$$

$$d\boldsymbol{z}_t = (-K\boldsymbol{z}_t + \boldsymbol{\beta})dt + \Sigma d\mathbf{w}_t$$

where $\boldsymbol{h}^\top = \bar{\boldsymbol{h}}^\top L^{-1}$, $s(t) = \bar{s}(t) - \boldsymbol{h}^\top \boldsymbol{\varphi}(t)$, $K = L\bar{K}L^{-1}$, $\beta = K\boldsymbol{\varphi}(t) + L\bar{\boldsymbol{\beta}} + \frac{d\varphi(t)}{dt}$,[5]

$\Sigma = L\bar{\Sigma}$, $d\boldsymbol{w}_t d\boldsymbol{w}_t{}' = \Omega dt$ and $\Sigma\Omega\Sigma = L\bar{\Sigma}\bar{\Omega}\bar{\Sigma}L'$.

From the analysis of Dai and Singleton (2000) and Cortazar and Naranjo (2006), it follows that for any model that can be expressed like the model containing state vector $\boldsymbol{y}_t$, where K have term wise different eigenvalues, there exists an affine transformation $\boldsymbol{z}_t = L\boldsymbol{y}_t + \boldsymbol{\varphi}(t)$ such that the model containing state vector $\boldsymbol{z}_t$ equals our chosen model containing state vector $\boldsymbol{x}_t$. Hence, any model that can be expressed like the model containing state vector $\boldsymbol{y}_t$, where K have term wise different eigenvalues, is nested in our chosen model.

[5] Cortazar and Naranjo (2006) write this equation as $\boldsymbol{\beta} = K\boldsymbol{\varphi}(t) + L\bar{\boldsymbol{\beta}}$. We belive this is a typo, see Appendix.

Schwartz and Smith (2000) showed how to transform a two-factor version of a model given by (1) and (2) to the two-factor stochastic convenience yield model presented in Schwartz (1997). Next, we show, building on Schwartz and Smith's (2000) derivation, that a two-factor stochastic convenience yield model with time-dependent equilibrium convenience yield level is nested within our chosen model.[6]

Let

$$
\begin{aligned}
ln(S_t) &= \begin{bmatrix} 1 & 0 \end{bmatrix} \boldsymbol{y}_t \\
d\boldsymbol{y}_t &= \left(-\begin{bmatrix} 0 & 1 \\ 0 & \kappa \end{bmatrix} \boldsymbol{y}_t + \begin{bmatrix} \mu - \frac{1}{2}\sigma^2 \\ \kappa\alpha(t) \end{bmatrix} \right) dt + \begin{bmatrix} \sigma_1 & 0 \\ 0 & \sigma_1 \end{bmatrix} d\mathbf{w}_t,
\end{aligned}
$$

where $y_t^{(1)}$ is the natural logarithm of a time t spot price, $y_t^{(2)}$ is the time t convenience yield, $\alpha(t)$ is a deterministic function of time and $d\mathbf{w}_t$ is defined as in (2). By letting

$$
\boldsymbol{x}_t = \begin{bmatrix} 1 & \frac{-1}{\kappa} \\ 0 & \frac{1}{\kappa} \end{bmatrix} \boldsymbol{y}_t + \begin{bmatrix} \frac{1}{2}\sigma_1^2 t + \int \alpha(t)\ dt \\ -e^{-\kappa t} \int \alpha(t)\, e^{\kappa t}\, dt \end{bmatrix}
$$

and applying the transformation relations previously shown, we get the model

$$
\begin{aligned}
ln(S_t) &= \begin{bmatrix} 1 & 1 \end{bmatrix} \boldsymbol{x}_t - \frac{1}{2}\sigma_1^2 t - \int \alpha(t)\ dt + e^{-\kappa t} \int \alpha(t)\, e^{\kappa t}\, dt \\
d\boldsymbol{x}_t &= \left(-\begin{bmatrix} 0 & 0 \\ 0 & \kappa \end{bmatrix} \boldsymbol{x}_t + \begin{bmatrix} \mu \\ 0 \end{bmatrix} \right) dt + \begin{bmatrix} \sqrt{\sigma_1^2 - \frac{2\rho\sigma_1\sigma_2}{\kappa} + \frac{\sigma_2^2}{\kappa^2}} & 0 \\ 0 & \frac{\sigma_2}{\kappa} \end{bmatrix} d\widetilde{\boldsymbol{w}}_t,
\end{aligned}
$$

where $d\widetilde{\mathbf{w}}_t d\widetilde{\mathbf{w}}_t^{\top} = \Omega \mathrm{dt} = \begin{bmatrix} 1 & \bar{\rho} \\ \bar{\rho} & 1 \end{bmatrix} dt$ and $\bar{\rho} = (-\kappa\rho\sigma_1 + \sigma_2)\sqrt{\frac{1}{\kappa^2\sigma_1^2 - 2\kappa\rho\sigma_1\sigma_2 + \sigma_2^2}}$.[7]

---

[6] A two-factor stochastic convenience yield model with time-dependent equilibrium convenience yield level has been used by Roberts and Fackler (1999), Roberts and Lin (2006), among others, to test for seasonality in convenience yield in several different futures markets. Note that the two-factor stochastic convenience yield model with time dependent equilibrium convenience yield level naturally nests the two-factor stochastic convenience yield model with constant convenience yield level found in Schwartz (1997).

[7] The change of Wiener process is derived using that almost surely, we can find a constant a such that $aW = bW_1 + cW_2$ where W denote a Wiener process and a, b and c are constants.

This latter model is a two-factor version of the model given by (1) and (2) expressed with the parameters of a two-factor stochastic convenience yield model with time-dependent equilibrium convenience yield level. As the matrix L is invertible, we conclude that a two-factor stochastic convenience yield model with time dependent equilibrium convenience yield level is equivalent to a model specified by (1) and (2).

A similar exercise can be done with $\boldsymbol{y}$'s process under an equivalent martingale measure. However, carrying out this exercise, an econometric difference between the classical convenience yield models and Cortazar and Naranjo's (2006) model becomes apparent. Utilizing Girsanov's theorem (Hull, 2012) and assuming constant risk premiums, $\boldsymbol{\lambda}$**,** it can be shown that the risk-adjusted dynamics of (2) are

$$d\boldsymbol{x}_{\mathrm{t}} \quad = \quad (-\mathrm{K}\boldsymbol{x}_{\mathrm{t}} + \boldsymbol{\beta} - \boldsymbol{\lambda})dt + \Sigma d\boldsymbol{w}_t^Q \tag{3}$$

where $d\boldsymbol{w}_t^Q$ is a vector of correlated Wiener increments under Q. Furthermore, it can be shown that under Q, the drift of the spot price of a commodity, must be equal to the difference between the continuous risk-free interest rate, r, and the commodity's continuous convenience yield (see e.g. Schwartz (1997) or Hull (2012)), e.g. $y_t^{(2)}$, such that $\boldsymbol{y}'s$ risk adjusted dynamics are

$$d\boldsymbol{y}_t \quad = \quad \left(-\begin{bmatrix} 0 & 1 \\ 0 & \kappa \end{bmatrix}\boldsymbol{y}_t + \begin{bmatrix} r - \frac{1}{2}\sigma_1^2 \\ \kappa\alpha(t) - \lambda_2 \end{bmatrix}\right)dt + \begin{bmatrix} \sigma_1 & 0 \\ 0 & \sigma_1 \end{bmatrix} d\mathbf{w}_t^Q,$$

where $d\boldsymbol{w}_t^Q$ is a vector of correlated Wiener increments under Q.

By counting the number of parameters present in $ln(S_t) = f(\boldsymbol{y}_t,\cdot)$ and $d\boldsymbol{y}_t$ under P and Q, and $ln(S_t) = f(\boldsymbol{x}_t,\cdot)$ and $d\boldsymbol{x}_t$ under P and Q, where $\boldsymbol{y}$ and $\boldsymbol{x}$ have similar dimensions, we find that the model expressed in terms of $\boldsymbol{x}$ is short one parameter compared to the model expressed in terms of $\boldsymbol{y}$. Given the formal equivalence of the two models, this suggest that one of the parameters in the model expressed in terms of $\boldsymbol{y}$ is econometrically redundant. Indeed, Schwartz and Smith (2000) found that the interest rate rate r is not required for specifing spot price dynamics under P or Q, nor for valuing futures contracts or estimating models from futures prices.

2.4 Futures Prices and Volatility Term Structure of Futures Returns

In this subsection we derive the closed form futures price formula in terms of the variables of a model defined by (1), (2) and (3). Based on this futures price formula, we derive the formula for the volatility term structure of futures returns. We use these formulas for parameter estimation, model inspection and to examine some practical implications of our findings.

In an economy where we assume fixed interest rates, the price of a futures contract at time t maturing at time T will be equal to the expected value of the underlying asset of the contract at time T under an equivalent martingale measure (Cox, Ingersoll, & Ross, 1981).

Using this equivalent martingale measure property, we have that

$$F_{t,T} \quad = \quad E_t^Q(S_T),$$

where $F_{t,T}$ is the price of a futures contract at time t maturing at time T where $T - t \geq 0$.

The conditional distribution of S_T within our chosen mode framework is lognormal. Therefore, using well known properties of lognormal variables it follows that

$$E_t^Q(S_T) \quad = \quad exp\left(E_t^Q(ln(S_T)) + \frac{1}{2}Var_t^Q(ln(S_T))\right). \tag{4}$$

For a model given by (1), (2) and (3), this simplifies to

$$E_t^Q(S_T) \quad = \quad exp\left(\mathbf{1}^\top E_t^Q(\boldsymbol{x}_T) \; + \; s(T) \; + \frac{1}{2}(\mathbf{1}^\top Cov_t^Q(\boldsymbol{x}_T)\mathbf{1})\right), \tag{5}$$

where, using properties derived by Schwartz and Smith (2000) and Cortazar and Naranjo (2006),

$$E_t^Q(x_T^i) \quad = \quad \begin{cases} x_t^1 \; + (\mu \; - \; \lambda_1)(T-t), & i = 1 \\ exp(-\kappa_i(T-t))x_t^i \; - \; \dfrac{1 - exp\left(-\kappa_i(T-t)\right)}{\kappa_i}\lambda_i, & i \neq 1 \end{cases}$$

and

$$Cov_t^Q(x_T^i, x_T^j) \quad = \quad \begin{cases} (\sigma_1)^2(T-t), & i = j = 1 \\ \sigma_i\sigma_j\rho_{ij} \; \dfrac{1 - exp\left(-(\kappa_i + \kappa_j)(T-t)\right)}{\kappa_i + \kappa_j}, & i \neq 1 \; or \; j \neq 1 \end{cases}.$$

By applying the natural logarithm to both sides of (5) we get a linear futures price formula in terms of the variables and parameters present in our chosen model,

$$ln(F_{t,T}) \quad = \quad \begin{bmatrix} 1 \\ exp(-\kappa_2(T-t)) \\ \vdots \\ exp(-\kappa_m(T-t)) \end{bmatrix}^{\top} \boldsymbol{x}_t \qquad (6)$$

$$+ s(T) + A(T-t)$$

where

$$A(T-t) \quad = \quad (\mu - \lambda_1)(T-t)$$

$$-\sum_{i=2}^{m} \frac{1-exp(-\kappa_i(T-t))}{\kappa_i}\lambda_i$$

$$+\frac{1}{2}(\mathbf{1}^{\top}\, Cov_t^{Q}(\boldsymbol{x}_T)\mathbf{1}).$$

and $s(T)$ is a deterministic function of time.

Furthermore, by applying Ito's lemma to $ln(F_{t,T})$ given by (6) we have that,

$$d\left(ln(F_{t,T})\right) \quad = \quad \frac{\partial ln(F_{t,T})}{\partial t}dt + \nabla_x ln(F_{t,T})^{\top} d\boldsymbol{x}_t,$$

where $\nabla_x ln(F_{t,T})^{\top} = [exp(-\kappa_1(T-t)) \quad \cdots \quad exp(-\kappa_m(T-t))]$, and so the model variance term structure of futures returns is given by

$$\left(\nabla_x ln(F_{t,T})\right)^{\top} \Sigma\Omega\Sigma \left(\nabla_x ln(F_{t,T})\right) \quad = \quad \sum_{i=1}^{m}\sum_{j=1}^{m} \sigma_i\sigma_j\, \rho_{ij} e^{-(\kappa_i+\kappa_j)(T-t)},\ \kappa_1 \equiv 0. \qquad (7)$$

Hull (2012) defines volatility term structure as the variation of volatility with time to maturity, where volatility is a measure of the uncertainty of the return realized on an asset. We follow the usual convention of measuring uncertainty in standard deviations, so that the that the model volatility term structure of futures returns is given by the square root of the model variance term structure of futures returns defined in (7).

From (7) we see that in our framework the model volatility term structure of futures returns is independent of the state variables. Furthermore, we see that the geometric Brownian motion factor does not induce time to maturity dependence in the model volatility term structure of futures returns, because as $T-t$ approaches 0 the volatility approaches the volatility of $ln(S)$, i.e. the volatility approaches the volatility of the sum of the short-term factors and the long-term factor, and as $T-t$ approaches infinity the volatility approaches the volatility of the geometric Brownian motion factor, i.e. the volatility approaches σ_1, which is the volatility of the long-term factor (Schwartz & Smith, 2000).

2.5 European Call Option on a Futures Contract

In this subsection we derive the closed form formula for the value of a European call option on a futures contract in terms of the variables of a model defined by (1), (2) and (3). We use this formula to examine some practical implications of our findings.

Hull (2012) defines a European call option as an option to buy an asset at a certain price at a certain date. By using (6) and the distributional properties of $\boldsymbol{x}_t$ we can, following Schwartz and Smith (2000), derive a closed form formula for the value of a European call option on a futures contract expiring at time T with strike price K and time t until option maturity, where $T-t \geq 0$.

Let c_0 denote the value of a European call option on a futures contract expiring at time T with strike price K and time t until option maturity where $0 < t \leq T$.

As $F_{t,T}$ is lognormally distributed under Q it follows that

$$c_0 = exp(-rt)E_0^Q\left(\left(F_{t,T}-K\right)^+\right) \quad (8)$$

$$exp(-rt)\left(F_{0,T}N(d) - KN\left(d-\sigma_{\mathrm{t,T}}\right)\right),$$

where $N(\cdot)$ indicates cumulative probabilities for the standard normal distribution,

$$\sigma_{\mathrm{t,T}}^2 = Var_0^Q\left(ln\left(F_{t,T}\right)\right)$$

$$= \begin{bmatrix} exp(-\kappa_1(T-t)) \\ \vdots \\ exp(-\kappa_m(T-t)) \end{bmatrix}^{\mathsf{T}} Cov_0^Q(\boldsymbol{x}_t) \begin{bmatrix} exp(-\kappa_1(T-t)) \\ \vdots \\ exp(-\kappa_m(T-t)) \end{bmatrix}, \kappa_1 \equiv 0$$

and

$$d = \frac{ln\left(\frac{F_{0,T}}{K}\right)}{\sigma_{\mathrm{t,T}}} + \frac{1}{2}\sigma_{\mathrm{t,T}}.$$

2.6 Seasonality Function

We can now define the seasonality function, $s(t)$, with knowledge of how $s(t)$ impacts $F_{t,T}$ and c_0. From (6) we have that s affects the level of $ln(F_{t,T})$. From (8) we have that s impacts c_0 though a percentage adjustment of $F_{t,T}$ and a level adjustment of d scaled by $\sigma_{\mathrm{t,T}}$.

Following Roberts and Fackler (1999), Sorensen (2002), Roberts and Lin (2006) and Power and Turvey (2008), we define $s(t)$ as

$$s(t) = \sum_{i=1}^{I} \gamma_i \cos(2\pi i t) + \gamma_i^* \sin(2\pi i t)$$

where the upper limit of i determines how many inflection points $s(t)$ can have in a year. As s captures price movements that are related to the season, the value of γ_i and γ_i^*, number of inflection points and location of the inflection points is determined by the seasonal patterns in the spot price.

The main benefit of this specification of $s(t)$ is that it provides much flexibility with few parameters. The main potential drawback of this specification of $s(t)$ is that it is continuous and does not differentiate between T and $T + a \in \mathbb{Z} > 0$.

Considering the seasonal consumption pattern of heating oil described in Section 1, general knowledge about how the four seasons in the northern hemisphere develop throughout a year and the pattern seen in Figure 2, we do not expect that $s's$ continuity will pose an issue, i.e. we do not expect several systematic seasonal heating oil price spikes or drops.

Section 3

In this section we present the chosen estimation procedure and test this procedure in a controlled environment utilizing simulation. The aim is to ensure that the codes used in the estimation procedure are working properly and to highlight potential problem areas.

3.1 Estimation Method

In Section 2 we derived a closed-form formula relating the natural logarithm of futures prices to the value of latent unobservable state variables. Therefore, assuming no measurement errors, it is possible to estimate the value of the latent unobservable state variables by inverting the futures price formula relating futures prices to state variable values. Schwartz and Smith (2000) defines measurement errors as errors in the reported futures prices or errors in the model's fit to observed prices. Several empirical studies (see e.g. Schwartz (1997), Schwartz and Smith (2000), Cortazar and Naranjo (2006), among others) have found statistically significant measurement error variability, i.e. the measurement errors cannot be assumed to be constantly equal to zero, using a similar setup to that of this thesis on several different futures markets. Therefore, it is reasonable to assume that some degree of measurement errors will be observed in fitting our chosen model using observed heating oil futures prices.

An estimation methodology that can handle multifactor models, latent unobservable state variables and measurement errors is the Kalman filter. The Kalman filter is an estimation methodology that is widely used in the commodities price modeling literature, see e.g. Schwartz (1997), Roberts and Fackler (1999), Schwartz and Smith (2000), Sorensen (2002), Cortazar and Naranjo (2006), among others.

3.2 The Kalman Filter

The Kalman filter, first introduced by R.E Kalman in 1960, is a set of mathematical equations that provides an efficient computational means to estimate the state of a process such that the mean of the squared error is minimized (Welch & Bishop, 2006). Moreover, the Kalman filter provides a means for estimating the parameters that defines a state space model with unobservable variables through maximum likelihood techniques.

We now describe the filter more precisely, following Koopman and Durbin (1998). Consider the multivariate Gaussian linear state space model with continuous states and discrete time intervals $t \in \{1,2,\dots,n\}$ given by

$$\boldsymbol{y}_t = Z_t\boldsymbol{\alpha}_t + \boldsymbol{\varepsilon}_t \qquad (observation\ equation)$$

$$\boldsymbol{\alpha}_{t+1} = \mathrm{T}_t\boldsymbol{\alpha}_t + R_t\boldsymbol{\eta}_t \qquad (state\ equation)$$

where $\boldsymbol{\varepsilon}_t \sim N(0, H_t)$, $\boldsymbol{\eta}_t \sim N(0, Q_t)$ and $\boldsymbol{\alpha}_1 \sim N(\boldsymbol{a}, P)$ are mutually independent. Here $\boldsymbol{y}_t$ is a $u_t \times 1$ vector of observations, $\boldsymbol{\alpha}_t$ is a $m \times 1$ vector of unobservable states and $\boldsymbol{\varepsilon}_t$ is a vector of disturbances. The state vector follows a Markov process with a $m \times 1$ disturbance vector $\boldsymbol{\eta}_t$. The system matrices, $Z_t, H_t, \mathrm{T}_t, R_t$ and Q_t, have appropriate dimensions and need not be time varying.

The Kalman filter recursions evaluate the mean of the vector of unobservable states conditional on the observations, $\boldsymbol{a}_{t+1} = E(\boldsymbol{\alpha}_{t+1}|\{\boldsymbol{y}_1, \boldsymbol{y}_2, \dots, \boldsymbol{y}_t\})$, usually called the one-step-ahead prediction, and its covariance matrix $P_{t+1} = var(\boldsymbol{\alpha}_{t+1}|\{\boldsymbol{y}_1, \boldsymbol{y}_2, \dots, \boldsymbol{y}_t\})$ for $t \in \{1,2, \dots, n\}$. The Kalman filter recursion equations are:

$$\boldsymbol{v}_t = \boldsymbol{y}_t - Z_t\boldsymbol{\alpha}_t,$$

$$F_t = Z_tP_tZ_t^\top + H_t,$$ [8]

$$K_t = P_tZ_t^\top,$$

$$\boldsymbol{a}_{t+1} = \mathrm{T}_t(\boldsymbol{a}_t + K_tF_t^{-1}\boldsymbol{v}_t),$$

and

$$P_{t+1} = \mathrm{T}_t(P_t - K_tF_t^{-1}K_t^\top)\mathrm{T}_t^\top + R_tQ_tR_t^\top.$$

$\boldsymbol{v}$ is usually called the prediction error and F is its related covariance matrix. The formal proof of these equations can be derived through the properties of the multivariate normal distribution.

The estimated mean of the vector of unobservable states conditional on the observations, $\boldsymbol{a}_{t+1} = E(\boldsymbol{\alpha}_{t+1}|\{\boldsymbol{y}_1, \boldsymbol{y}_2, \dots, \boldsymbol{y}_t\})$, can be used as estimates for the state vector, or as inputs for

[8] Note that F does not denote a futures price in the current subsection.

a smoothing algorithm producing state vector estimates based on the full set of observations, $Y_n = \{\boldsymbol{y}_1, \boldsymbol{y}_2, \dots, \boldsymbol{y}_n\}$.

If the system matrices, $Z_t, H_t, \mathrm{T}_t, R_t$ and Q_t, are dependent upon a vector of unknown parameters, ψ, the log-likelihood for a given ψ can be computed with output from the Kalman filter as

$$log(L) \quad = \quad constant - \frac{1}{2}\sum_{t=1}^{n} log|F_t| + \boldsymbol{v}_t^{\top} F_t^{-1} \boldsymbol{v}_t.$$

This allows for maximum likelihood estimation of ψ through maximization of $log\big(L|\widehat{\psi}\big)$.

To get a better feel for the internal logic of the Kalman filter equations it is useful to step through an iteration. Imagine that we are at time 2 such that that $\boldsymbol{a}_{t-1} \ = \ \boldsymbol{a}$ and $P_{t-1} = \ P$, derived from $\boldsymbol{\alpha}_1 \sim N(\boldsymbol{a}, P)$ assumed known a priori for now. From the state equation we can compute the one-step-ahead prediction, $\widehat{\boldsymbol{a}}_{t\,|\,t\,-\,1} = \mathrm{T}_{t-1}\boldsymbol{a}_{t-1}$, providing an initial estimate for the vector of unobservable state variables at time t, and its covariance matrix $P_{t\,|\,t\,-\,1} \ = \mathrm{T}_{t\,-\,1}P_{t\,-\,1}\mathrm{T}_{t\,-\,1}^{\top} + \ R_{t\,-\,1}Q_{t\,-\,1}R_{t\,-\,1}^{\top}$. Next, we can, using the observation equation and the initial estimate for vector of unobservable state variables at time t, compute the one-step-ahead predicted $\boldsymbol{y}$-values, $\widehat{\boldsymbol{y}}_{t\,|\,t\,-\,1} = Z_t\widehat{\boldsymbol{a}}_{t\,|\,t\,-1}$. As $\boldsymbol{y}$ is observable, we can compute a measure of the accuracy of our predictions and its covariance matrix. This accuracy measure is the one-step-ahead prediction error, computed as $\boldsymbol{v}_t = \boldsymbol{y}_t \ - \ \widehat{\boldsymbol{y}}_{t\,|\,t\,-\,1}$ with covariance matrix $F_t = Z_tP_{t\,|\,t-\,1}Z_t^{\top} + H_t$. Now we have an initial estimate for the vector of unobservable state variables at time t, a measurement of how well this estimate agrees with something observable and their variability measures. Intuitively, all these values contain some information about what a good one-step-ahead prediction for the vector of unobservable state variables should be, so it should be possible to produce a new and improved one-step-ahead prediction for the vector of unobservable state variables at time t using these values. Indeed, in the last step in a single iteration, we calculate a new, updated, estimate for the vector of unobservable state variables at time t using all the previously calculated values, $\widehat{\boldsymbol{a}}_t \ = \widehat{\boldsymbol{a}}_{t\,|\,t\,-\,1} \ + \ P_{t\,|\,t\,-\,1}Z_t^{\top}F_t^{-1}\boldsymbol{v}_t$, and its covariance matrix $P_t \ = \ P_{t\,|\,t\,-\,1} - P_{t\,|\,t\,-\,1}Z_t^{\top}F_t^{-1}Z_tP_{t\,|\,t\,-\,1}^{\top}$.

In walking through an iteration, we see that the Kalman filter requires initial values for $\boldsymbol{a}$ and P, supplied through $\alpha_1 \sim N(\boldsymbol{a}, P)$. If these initial values are unknown, we can make an educated guess based on assumed properties and available data or use numerical techniques like diffuse initialization. Koopman and Durbin (1998) suggest that if the vector of unobservable states contains non-stationary components, which is probable in our case, utilizing diffuse initialization is preferred for some elements of the state vector.

3.3 State Space Representation of a Model Defined By (1), (2) and (3)

We've seen that to implement maximum likelihood parameter estimation for a model defined by (1), (2) and (3) through output from the Kalman filter we need to express the model in state space form.

Using well known solutions to the stochastic differential equations in (2) and relations shown in Section 2 we can cast our model as a multivariate Gaussian linear state space model with continuous states and discrete time intervals $t \in \{t_{1,} t_2 = t_1 + \Delta t, \dots, t_n = t_{n-1} + \Delta t\}$ given by

$$ln(\boldsymbol{F}_{t,\boldsymbol{T}}) = \boldsymbol{b}_t + Z_t \boldsymbol{x}_t + \boldsymbol{\varepsilon}_t \qquad (observation\ equation)$$

$$\boldsymbol{x}_{t+1} = \boldsymbol{c}_t + \mathrm{T_t} \boldsymbol{x}_t + \boldsymbol{\eta}_{\mathrm{t}} \qquad (state\ equation)$$

where $\boldsymbol{\epsilon}_t \sim N(0, H_t)$, $\boldsymbol{\eta}_{\mathrm{t}} \sim N(0, Q_t)$ and $\boldsymbol{x}_{t_1} \sim N(\boldsymbol{a}, P)$ are mutually independent. Furthermore,

$$ln(\boldsymbol{F}_{t,\boldsymbol{T}}) = \begin{bmatrix} ln(F_{t,T_1}) \\ \vdots \\ ln\left(F_{t,T_{u_t}}\right) \end{bmatrix},$$

$$\boldsymbol{b_t} = \begin{bmatrix} s(T_1) + A(T_1 - t) \\ \vdots \\ s(T_{u_t}) + A(T_{u_t} - t) \end{bmatrix},$$

$$Z_t = \begin{bmatrix} 1 & exp(-\kappa_2(T_1 - t)) & \cdots & exp(-\kappa_m(T_1 - t)) \\ \vdots & \vdots & \vdots & \vdots \\ 1 & exp\left(-\kappa_2(T_{u_t} - t)\right) & \cdots & exp\left(-\kappa_m(T_{u_t} - t)\right) \end{bmatrix},$$

$$H_t = [\xi_1^2 \quad \cdots \quad \xi_1^2 \quad \xi_2^2 \quad \cdots \quad \xi_2^2 \quad \xi_3^2 \quad \cdots \quad \xi_3^2] I_{u_t},$$

$$\boldsymbol{c}_t = \begin{bmatrix} \mu\Delta t \\ 0 \\ \vdots \\ 0 \end{bmatrix},$$

$$\mathrm{T_t} = [1 \quad exp(-\kappa_2 \Delta t) \quad \cdots \quad exp(-\kappa_m \Delta t)] I_m,$$

and

$$Q_t \quad = \quad \begin{cases} (\sigma_1)^2\Delta t, & i = j = 1 \\ \sigma_i\sigma_j\rho_{ij}\,\dfrac{1 - exp\big(-(\kappa_i + \kappa_j)\Delta t\big)}{\kappa_i + \kappa_j}, & i \neq 1\; or\; j \neq 1 \end{cases}.$$

In parameterizing the covariance of the measurement errors, H_t, we've opted to follow Schwartz (1997), Schwartz and Smith (2000), Cortazar and Naranjo (2006), among others, and assumed independence. Furthermore, as there are for each date, more than 36 listed NY Harbor ULSD futures contracts with different maturities, the dimensionality of our problem is large, i.e. u_t is large. Therefore, we've opted to follow Cortazar, Kovacevic and Schwartz (2013) and cast H_t as $[\xi_1^2 \quad \cdots \quad \xi_1^2 \quad \xi_2^2 \quad \cdots \quad \xi_2^2 \quad \xi_3^2 \quad \cdots \quad \xi_3^2]I_{u_t}$ using three maturity groups of equal length to reduce the number of estimated parameters.

3.4 Simulation

To create a controlled environment where we can test the Kalman filter and parameter estimation procedure we simulate futures price panels, estimate parameters and review the results. We limit the simulation study to a two-factor version of a model given by (1), (2) and (3), because the simulation study is very demanding in terms of computational resources, and because a simulation study using a two-factor model is likely sufficient for the insights we are after here.

Following Andresen and Sollie (2013), the initial step of the simulation study is to simulate state variable paths for the two-factor model with the following algorithm:

1. Generate $\epsilon_{i,j}$ where $\epsilon \sim N(0, \Sigma\Omega\Sigma)$
2. $\boldsymbol{x}_{t_0} \equiv \boldsymbol{a} \equiv \boldsymbol{0}$
3. $\boldsymbol{for}\; i \;=\; 2\; \boldsymbol{to}\; 2501\; \boldsymbol{do}$
   3.1. $x_{t_i}^{(1)} = x_{t_{i-1}}^{(1)} + \mu\Delta t + \sqrt{\Delta t}\epsilon_{i,1}$
   3.2. $x_{t_i}^{(2)} = x_{t_{i-1}}^{(2)} exp(-\kappa_2\Delta t) + \sqrt{\frac{1-exp(-2\ _2\Delta t)}{2\kappa_2}}\epsilon_{i,2}$
4. $\boldsymbol{end\; for}$

Given the simulated state vectors we evaluate (6) using parameter estimates presented in Schwartz and Smith (2000) and suitable time to maturity to get a price panel $ln(F)$,

$dim(ln(F)) = 2500 \times 20$. The panel of maturities are generated by assuming end of month settlement and no holidays. The first column in $ln(F)$ contains 2500 daily observations of the natural logarithm of simulated futures prices for contracts that are closest to maturity, the second column of $ln(F)$ contains 2500 daily observations of the natural logarithm of simulated futures price for contracts that are second closest to maturity, and so forth. By assuming 52 weeks in a year we get that the average time to maturity for contracts in the first column is $\left(\frac{52\times5}{12}+1\right)\times\frac{1}{2}\approx 11$ business days, in the second column is $\frac{52\times5}{12}+11 \approx 32$ business days, in the third column is $\frac{52\times5}{12}\times 2+11 \approx 54$ business days, and so forth.

Lastly, we add a small measurement error, $\varepsilon\sim N(0,0.001^2)$, to each element in $ln(F)$.

We repeat these steps 100 times creating an array with dimensions $2500 \times 20 \times 100$ such that each slice of the array contains a panel of the natural logarithm of simulated futures prices with characteristics as previously described.

For each slice of the array we estimate a parameter vector, $\hat{\psi}$, by maximizing a function that performs Kalman filtering and returns the log-likelihood for a given $\hat{\psi}$.

We used the Kalman filter implementation created by Helske (2017), which is based on the works of Koopman and Durbin. Helske's (2017) Kalman filter implementation is a univariate version of the Kalman filter presented in "The Kalman Filter" that utilize exact diffuse initialization as described in Koopman and Durbin (1998).

We maximized the log-likelihood by applying an evolutionary search algorithm and quasi-Newton optimization algorithms.[9]

[9] In choosing an optimization procedure we tested all the most promising candidates presented in Mullen (2014) and found that the Genoud algorithm, presented in Mebane and Sekhon (2007), provided the best combination of user control, accuracy and speed. The Genoud algorithm combines a genetic evolution strategy with a quasi-Newton optimizer for local hill climbing. The quasi-Newton optimizer we used was the base R implementation of a LBFGS algorithm. As we did not have an analytical gradient we used Richardson's extrapolation to approximate the gradient numerically. We also tried several different completely derivative free optimizers and found their performance lacking compared to the chosen procedure.

Table I
Summary statistics of parameter estimates based on simulated futures price panels

Summary statistics of parameter estimates based on 100 simulated futures price panels. Each panel of simulated futures price consist of daily observations of 20 futures prices for 2500 business days. The average time to maturity for contracts in the first column of each panel is $\left(\frac{52\times5}{12}+1\right)\times\frac{1}{2}\approx 11$ business days, the average time to maturity for contracts in the second column of each panel is $\frac{52\times5}{12}+11 \approx 32$ business days, the average time to maturity for contracts in the third column of each panel is $\frac{52\times5}{12}\times 2+11 \approx 54$ business days, and so forth. Row Actual reports the true parameter values, row Min. reports the minimum of the estimated parameter values, row 1st and 3rd Qu. reports the first and the third quarter of the estimated parameter values, row Mean reports the mean of the estimated parameter values and row Max. reports the maximum of the estimated parameter values. Negative values are reported in parenthesis.
Note that $\mu^Q = \mu - \lambda_1$.

	μ^Q	μ	κ_2	λ_1	λ_2	ρ_{12}	σ_1	σ_2	ε
Actual	0.016	(0.039)	1.190	(0.055)	0.014	0.189	0.115	0.158	0.001
Min.	0.016	(0.153)	1.188	(0.169)	(0.099)	0.149	0.111	0.154	0.001
1st Qu.	0.016	(0.066)	1.190	(0.082)	(0.024)	0.175	0.114	0.157	0.001
Mean	0.016	(0.041)	1.190	(0.057)	0.019	0.189	0.115	0.158	0.001
3rd Qu.	0.016	(0.014)	1.191	(0.030)	0.055	0.205	0.116	0.159	0.001
Max.	0.017	0.063	1.194	0.047	0.167	0.231	0.120	0.162	0.001

Table I shows the distribution of a set of parameter estimates based on the natural logarithm of 100 simulated futures price panels with dimensions 2500×20, using parameter estimates presented in Schwartz and Smith (2000) as the actual parameter value. We see that for all the parameters the mean estimated value is approximately equal to the actual value, indicating that our estimation procedure is working. Moreover, we see that for $\mu^Q = \mu - \lambda_1, \kappa_2, \rho_{12}, \sigma_1, \sigma_2$ and ε the distribution of estimated values is narrow and that μ, λ_1 and λ_2 are not estimated with much accuracy, comparatively.

The lack of accuracy in the estimates of μ, λ_1 and λ_2 is due to econometric issues. However, Schwartz and Smith (2000) provide an economic explanation which offers valuable insights. As we do not have direct observations of expected future spot prices we cannot accurately determine the expected future spot price location nor its long-term growth rate. Therefore, μ is poorly estimated as it's partially determined by the long-term growth rate of expected future spot prices. Furthermore, λ_1 is poorly estimated as it's partially determined by the difference between the long-term growth rate of expected future spot prices and long-term growth rate of observed futures prices. Lastly, λ_2 is poorly estimated as it's partially determined by the difference between the long-term tangent line intercepts of expected future spot prices and the long-term tangent line intercepts of observed futures prices. Andresen and Sollie (2013) showed that though the indeterminacy is only partial, the increase in parameter estimate accuracy as a function of number of observations is weak.

Note that $\mu^Q = \mu - \lambda_1$ is less affected by this indeterminacy issue as it is solely determined by the long-term growth rate of the observed futures prices.

Thus, using futures data, we can precisely estimate the process for spot prices under an equivalent martingale measure, i.e. under Q, but we cannot precisely estimate the process under P. Consequently, the indeterminacy issue does not affect the robustness of the model for use in valuation problems, but robustness for forecasting purposes is generally lacking. To precisely estimate risk premia and, indirectly, the long-term growth rate of expected future spot prices, either a large time window or observations related to the spot price process under P is required (Schwartz and Smith (2000) and Andresen and Sollie (2013)).

Section 4

In this section we present and review the NY Harbor ULSD futures price panel.

4.1 The Data

The data used to estimate model parameters and test the models consists of daily observations of NY Harbor ULSD futures nominal settlement prices from December 31, 2015 to December 31, 2018, sourced from Thomson Reuters Datastream (2019). We do not adjust the settlement prices for inflation, which is justified by the relatively short time span of the time series, though Casassus and Collin-Dufresne (2005) suggest that mean reversion in real prices is more likely than in nominal prices.

Using the same notation as in Schwartz (1997), we denote the contracts that are closest to maturity F01, the second closest to maturity F02, and so forth. This procedure generates 49 generic futures contracts of whom 38 are continuously observable.

We limited the contract depth to the first 38 contracts because by discarding contract F39 to F49 we get contracts with maturity ranges that are continuously observable, which is preferred given our relatively narrow time window. Moreover, the liquidity pattern painted by open interest and volume figures suggest that for maturities greater than approximately four weeks liquidity reduces as maturity increases. Thus, discarding the tail contracts increases the average sample liquidity.

Of these 38 contracts, we used the even-numbered contracts for parameter estimation, leaving the odd-numbered contracts for within time window out-of-sample statistics. The choice of using even-numbered contracts for parameter estimation was based on the fact that open interest and volume was higher for F02 than F01, and F02 settlement prices are less effected by investors rolling their positions as maturity nears.[10]

Table II shows summary statistics for contract F1 to F38, which is the contract depth for this thesis. From Table II we see that mean futures prices increase with increased time to maturity, indicating that the average market is a normal market, as opposed to an inverted market where the futures price decreases with the maturity of the futures contract (Hull, 2012). Furthermore, we see that the standard error of mean futures prices decreases as time to maturity increases, indicating time to maturity dependence in the volatility term structure of futures returns.

[10] Mean daily volume, mid-month open interest and end-of-month open interest for F01 and F02 was 48396, 75528 and 1430, and 50314, 95780 and 124062, respectively.

Table II
Statistics of heating oil futures contracts

Statistics for daily observations of heating oil futures contracts from December 31, 2015 to December 31, 2018. Settlement prices are in dollars per gallon and maturities are in years. The contract denomination specifies relative closeness to maturity and the approximate number of months to maturity.

Contract	Number of observations	Mean price (SE)	Maturity (SE)
F01	756	1.71 (0.348)	0.040 (0.024)
F02	756	1.71 (0.342)	0.123 (0.025)
F03	756	1.72 (0.336)	0.206 (0.025)
F04	756	1.72 (0.328)	0.290 (0.025)
F05	756	1.73 (0.321)	0.373 (0.025)
F06	756	1.73 (0.313)	0.456 (0.025)
F07	756	1.74 (0.305)	0.539 (0.025)
F08	756	1.74 (0.299)	0.623 (0.025)
F09	756	1.74 (0.293)	0.706 (0.025)
F10	756	1.75 (0.289)	0.789 (0.025)
F11	756	1.75 (0.285)	0.872 (0.025)
F12	756	1.75 (0.282)	0.955 (0.024)
F13	756	1.76 (0.280)	1.039 (0.024)
F14	756	1.76 (0.278)	1.122 (0.024)
F15	756	1.76 (0.275)	1.205 (0.025)
F16	756	1.76 (0.272)	1.289 (0.025)
F17	756	1.76 (0.270)	1.372 (0.025)
F18	756	1.76 (0.267)	1.455 (0.025)
F19	756	1.77 (0.264)	1.538 (0.025)
F20	756	1.77 (0.260)	1.622 (0.025)
F21	756	1.77 (0.256)	1.705 (0.025)
F22	756	1.77 (0.252)	1.788 (0.025)
F23	756	1.77 (0.248)	1.872 (0.025)
F24	756	1.77 (0.243)	1.955 (0.024)
F25	756	1.78 (0.240)	2.039 (0.024)
F26	756	1.78 (0.236)	2.122 (0.024)
F27	756	1.78 (0.231)	2.205 (0.025)
F28	756	1.78 (0.226)	2.288 (0.025)
F29	756	1.78 (0.220)	2.372 (0.025)
F30	756	1.79 (0.215)	2.455 (0.025)
F31	756	1.79 (0.210)	2.538 (0.025)
F32	756	1.79 (0.205)	2.622 (0.025)
F33	756	1.79 (0.202)	2.705 (0.025)
F34	756	1.80 (0.198)	2.788 (0.025)
F35	756	1.80 (0.194)	2.872 (0.024)
F36	756	1.80 (0.191)	2.955 (0.024)
F37	756	1.80 (0.187)	3.038 (0.024)
F38	756	1.80 (0.184)	3.122 (0.024)

Figure 5 shows time to maturity for contract F01, F03, F05, F07 and F09 from December 31, 2015 to December 31, 2018 measured in years. We see that time to maturity remains within a narrow range for each contract. This pattern of time to maturity is representative for all the contracts in the dataset.

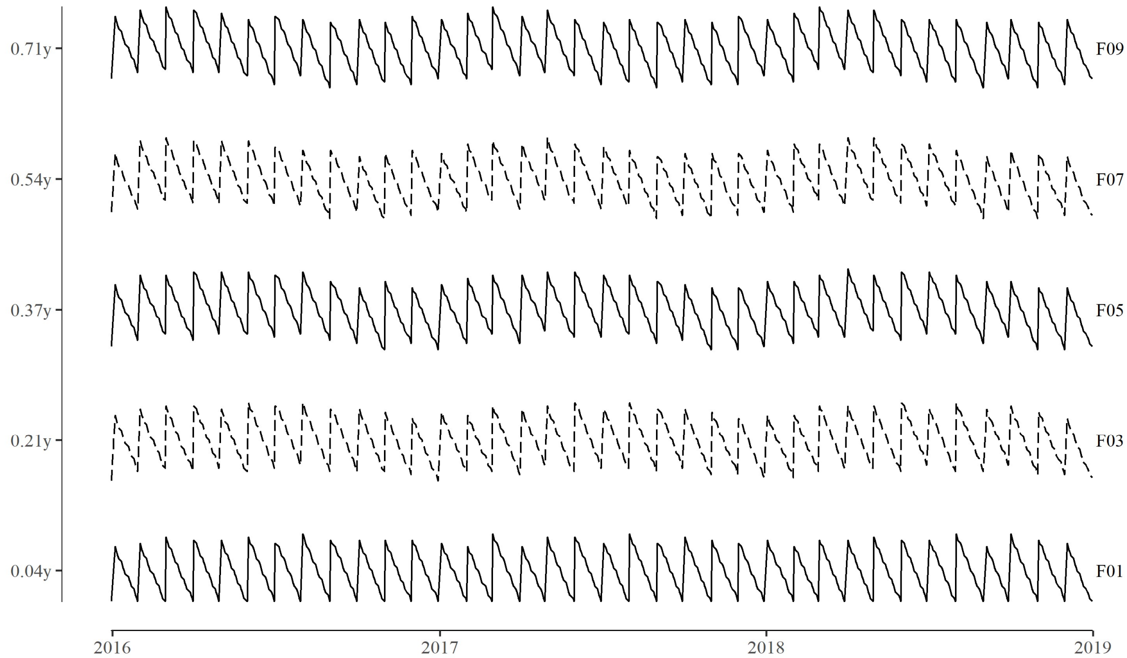

Figure 5. Time to maturity as a function of time. Time to maturity for contract F01, F03, F05, F07 and F09 from December 31, 2015 to December 31, 2018 measured in years (y). The y-axis tick-marks are mean time to maturity measured in years for contract F01, F03, F05, F07 and F09.

Figure 6 shows daily settlement prices for contract F01 and F18 from December 31, 2015 to December 31, 2018 in dollars per gallon. We see a price dip in the lead and tail of the time series with an intermediate growth period. Furthermore, the F18 – F01 spread alternate signs, indicating mean reversion in heating oil prices under an equivalent martingale measure (Casassus & Collin-Dufresne, 2005). Alternating signs in the F18 – F01 spread suggest periods of strong backwardation. Strong backwardation is present if $S_t - F_{t,T} > 0$, and due to convergence $S_t \approx$ F01 (Trolle & Schwartz, 2008).[11] In terms of the classical convenience yield models, strong backwardation implies that the convenience yield of a commodity exceeds the cost of storing the commodity plus the interest forgone (Bessembinder, Coughenour, Seguin, & Smeller, 1996). Intermittent periods of strong backwardation imply

[11] Note that if strong backwardation is present for all maturities the market will also be an inverted market.

mean reversion in one or more of these quantities, convenience yield, cost of storage and interest forgone, a property that is normally assigned to the convenience yield.

When it comes to seasonality, it's difficult to discern any clear sinusoidal pattern from Figure 6. This could be explained by the high volatility, the growth and decline periods or a lack of seasonality.

Figure 6. Futures prices. Contract F01 and F18 daily settlement prices from December 31, 2015 to December 31, 2018 in dollars per gallon.

Section 5

In this section we present and review three different two-factor versions of the model defined by (1), (2) and (3) in Section 2. Model 1 have no seasonal adjustment, Model 2 have a seasonal adjustment with two inflection points and Model 3 have a seasonal adjustment with four inflection points. Furthermore, we examine the one-factor model and models with more than two-factors defined by (1), (2) and (3).

5.1 Estimated Parameters

Table III
Maximum-Likelihood Parameter Estimates

Maximum-likelihood parameter estimates for the two-factor model defined by (1), (2) and (3) in Section 2 for heating oil daily settlement prices from December 31, 2015 to December 31, 2018. Standard errors are estimated as the square root of the diagonal entries of the inverse of the estimated observed Fisher information matrix.

	Model 1		Model 2		Model 3	
Parameter	Estimate	(SE)	Estimate	(SE)	Estimate	(SE)
μ^Q	-0.002	(0.001)	-0.003	(0.001)	-0.003	(0.001)
μ	-0.014	(0.077)	-0.011	(0.080)	-0.013	(0.079)
κ_2	0.330	(0.007)	0.341	(0.007)	0.332	(0.007)
λ_2	0.143	(0.126)	0.144	(0.130)	0.145	(0.131)
ρ_{12}	0.534	(0.044)	0.557	(0.041)	0.541	(0.042)
σ_1	0.134	(0.005)	0.138	(0.005)	0.137	(0.005)
σ_2	0.218	(0.009)	0.224	(0.010)	0.226	(0.010)
γ_1			0.005	(0.000)	0.005	(0.000)
γ_1^*			-0.001	(0.000)	-0.001	(0.000)
γ_2					0.001	(0.000)
γ_2^*					0.001	(0.000)
ξ_1	0.011	(0.000)	0.009	(0.000)	0.009	(0.000)
ξ_2	0.010	(0.000)	0.010	(0.000)	0.010	(0.000)
ξ_3	0.005	(0.000)	0.004	(0.000)	0.004	(0.000)
Log-likelihood	45,610		47,527		47,688	

Table III shows maximum-likelihood parameter estimates for the two-factor model presented in Section 2 applied to heating oil daily settlement prices from December 31, 2015 to December 31, 2018. Standard errors are estimated as the square root of the diagonal entries of the inverse of the estimated observed Fisher information matrix.[12] The observed Fisher information matrix is calculated as the negative hessian of $l(\widehat{\psi}_{ML})$ where l is a log-likelihood function.

Comparing the three models we see high consistency in parameter estimates. Furthermore, we see that μ and λ_2 have standard errors which are large compared to their absolute value and are not significantly different from zero. This is in accord with the indeterminacy issue discussed in Section 3.

[12] Given certain restrictions, it can be shown that $\widehat{\psi}_{ML}$ is asymptotically normally distributed with mean ψ and covariance matrix equal to the inverse Fisher information matrix (see e.g. Hallam (2019)).

The estimated value for $\mu^Q = \mu - \lambda_1$ is barely significantly different from zero in Model 1 and significantly different from zero in Model 2 and Model 3, with an estimated value close to zero. The lack of significance for μ^Q in Model 1 is possibly due to the relatively small sample size compared to the simulation study, or the major shifts in the long-term growth rate of the observed futures prices throughout our sample period, indicated by Figure 6, which puts stress on the constant parameter assumption.

Furthermore, we see that $\kappa_2, \rho_{12}, \sigma_1, \sigma_2, \xi_1, \xi_2$ and ξ_3 are highly significant. κ_2 is estimated to be approximately 0.33. This indicates strong mean reversion where deviations from zero in x_2, the factor that captures the effect of transitory changes in supply and demand, are expected to halve in approximately 2.1 years.[13]

σ_1 and σ_2 are estimated to be approximately 14% and 22%, respectively, reflecting the high volatility in heating oil prices.

The value of the estimated standard deviation of the measurement errors, ξ_1, ξ_2 and ξ_3, reduces as time to maturity increases for Model 1, indicating reduction in errors in the reporting of prices or increase in the model's fit to observed prices as maturity increases. ξ_1 and ξ_3 are lower in Model 2 and Model 3 than in Model 1, indicating that the increase in the model's fit to observed prices caused by the seasonal adjustment is greater for short and long maturity contracts.

5.2 Likelihood-Ratio Tests

As Model 1 and 2 are restricted versions of Model 3, fitted on the same sample, we can use likelihood-ratio tests to determine if the seasonal adjustments are statistically significant.

[13] $2e^{-\kappa t} = 1 \Leftrightarrow t = \frac{ln(2)}{\kappa}, \ \frac{ln(2)}{0.33} \approx 2.1$

Table IV
Likelihood-Ratio Tests

The table shows the results of likelihood-ratio tests for Model 1 and Model 2 ($\gamma_1 = \gamma_1^* = 0$), Model 1 and Model 3 ($\gamma_1 = \gamma_1^* = \gamma_2 = \gamma_2^* = 0$) and Model 2 and Model 3 ($\gamma_2 = \gamma_2^* = 0$). The likelihood-ratios are calculated as -2(log-likelihood of the restricted model less log-likelihood of the unrestricted model). The critical values for the likelihood-ratio tests at the 5% significance level is given by $Prob(\chi^2_{df} \geq c) = 0.05$, where df denotes degrees of freedom and c is the critical value for the test.

Model	Restriction	Likelihood-Ratio	$Prob(\chi^2_{df} \geq c) = 0.05$
(1-2)	$\gamma_1 = \gamma_1^* = 0$	3,835	df = 2, c = 5.99
(1-3)	$\gamma_1 = \gamma_1^* = \gamma_2 = \gamma_2^* = 0$	4,158	df = 4, c = 9.49
(2-3)	$\gamma_2 = \gamma_2^* = 0$	323	df = 2, c = 5.99

Table IV shows likelihood-ratio test for Model 1 and Model 2 ($\gamma_1 = \gamma_1^* = 0$), Model 1 and Model 3 ($\gamma_1 = \gamma_1^* = \gamma_2 = \gamma_2^* = 0$) and Model 2 and Model 3 ($\gamma_2 = \gamma_2^* = 0$) and the 5% significance level for each restriction. We see that γ_1 and γ_1^*, $\gamma_1, \gamma_1^*, \gamma_2$ and γ_2^* and γ_2 and γ_2^* are jointly significant. We conclude that there is statistically significant seasonality in heating oil futures prices, and that more than two inflection points per year provide a statistically significant improvement in describing this seasonality. In Roberts and Lin's (2006) study of seasonality in the price process for different commodities the authors found significant seasonality in crude oil's price process under an equivalent martingale measure, though the authors could not find grounds for allowing more than two inflection points in describing this seasonality. The contrast to our results indicates that there are significant differences between the seasonality of the two commodities' price process under an equivalent martingale measure.

5.3 Estimated Seasonality Function

From (5) we have that

$$F_{t,T} \quad = \quad exp\left(\mathbf{1}^\top E_t^Q(\boldsymbol{x}_T) + \frac{1}{2}(\mathbf{1}^\top Cov_t^Q(\boldsymbol{x}_T)\mathbf{1})\right) \times exp(s(T,\boldsymbol{\gamma})),$$

so that $s(T,\boldsymbol{\gamma})$ can be viewed as a seasonal percentage adjustment. Furthermore, from Table II we see that the natural logarithm of the mean of $F_{t,T}$ is approximately 0.5 so that the proportion of $F_{t,T}$ explained by seasonal variations is approximately $2 \times s(T,\boldsymbol{\gamma})$. Though we think this latter approximation has some value it must be stressed that it's a crude approximation, even more so for negative values of $s(T,\boldsymbol{\gamma})$, and it is less applicable for contracts with longer maturities.

Figure 7 shows $s(T, \gamma_1, \gamma_1^*)$ and $s(T, \gamma_1, \gamma_1^*, \gamma_2, \gamma_2^*)$ where $\boldsymbol{\gamma}$ takes on the estimated parameter values for Model 2 and Model 3, respectively. We see that the seasonal percentage adjustment at a given point in time ranges between approximately -0.6% and 0.6% for both models. By multiplying these ranges by 2 we get a crude approximation of the proportion of $F_{t,T}$ explained by seasonal variation.

Sorensen (2002), using a similar setup, examined seasonality in corn, soybean and wheat futures prices. Naturally, corn, soybean and wheat are not directly comparable with heating oil, but since corn, soybean and wheat are known for having substantial seasonal price dependence the ranges estimated by Sorensen (2002) provide a useful reference point to evaluate the limits of the ranges for heating oil. Sorensen (2002) found that the seasonal percentage adjustment for Corn, Soybean and Wheat ranged between approximately -2.5% and 2.5%. In light of Sorensen's (2002) estimates, the magnitude of the limits of our estimated ranges seems reasonable.

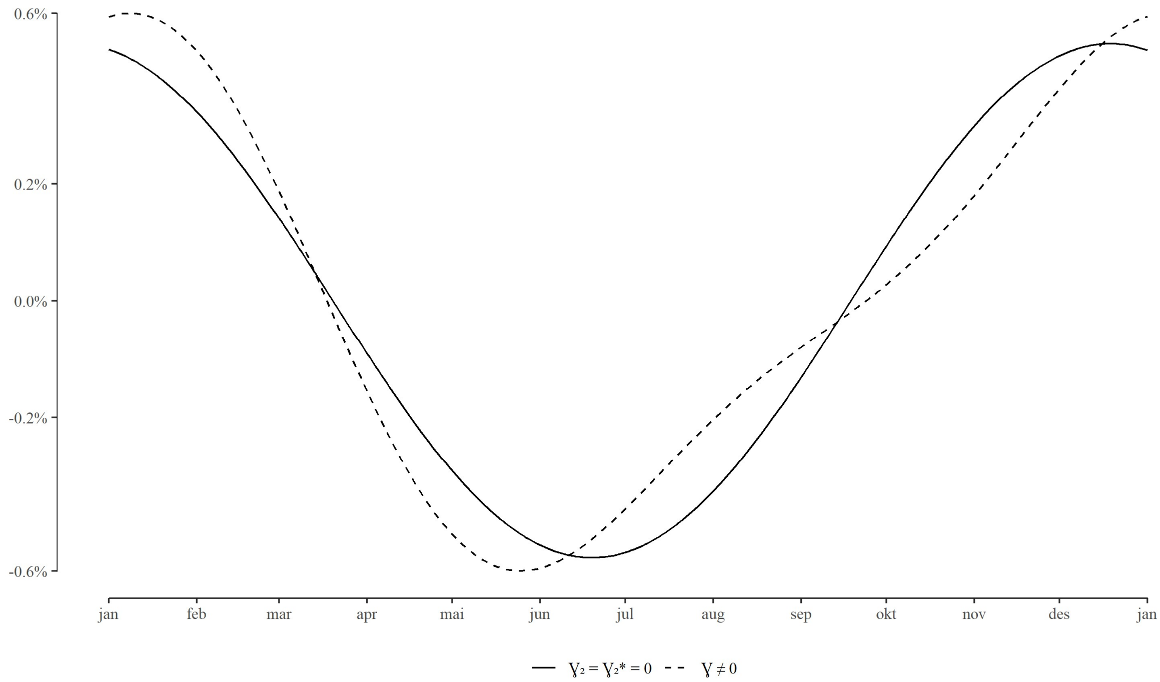

Figure 7. Futures prices seasonality. The solid line is $s(T, \gamma_1, \gamma_1^*)$ with estimated parameters for Model 2 and the dashed line is $s(T, \gamma_1, \gamma_1^*, \gamma_2, \gamma_2^*)$ with estimated parameter for Model 3

In Section 1 we found that United States distillate fuel oil inventory changes appear to be positively correlated with changes in the NY Harbor ULSD 3rd less 1st-month futures spread.

Moreover, we found signs of seasonality in United States distillate fuel oil inventory descriptive statistics. If we compare Figure 2 with Figure 7 we find strong pattern similarity, which is to be expected. Furthermore, we see that $s(T, \gamma_1, \gamma_1^*, \gamma_2, \gamma_2^*)$'s pattern is closer to the pattern of United States distillate fuel oil inventory with reduced ascent around September.

5.4 Accuracy

Table V
In time window dollar accuracy

Root mean squared error (RMSE) in dollars and mean error (ME) in dollars for Model 1, Model 2 and Model 3 from December 31, 2015 to December 31, 2018. The pairs of odd and even contracts displayed in the table represents minimum values, maximum values or turning points for RMSE or ME as a function of time to maturity.

	RMSE in dollars			ME in dollars		
Model	1	2	3	1	2	3
F1	0.0325	0.0286	0.0292	-0.0074	-0.0067	-0.0067
F2	0.0285	0.0243	0.0250	-0.0053	-0.0049	-0.0050
F10	0.0168	0.0107	0.0107	0.0029	0.0032	0.0032
F11	0.0167	0.0113	0.0113	0.0031	0.0033	0.0034
F19	0.0250	0.0233	0.0232	0.0011	0.0008	0.0009
F20	0.0256	0.0245	0.0242	0.0008	0.0007	0.0007
F26	0.0144	0.0137	0.0134	-0.0033	-0.0037	-0.0036
F27	0.0137	0.0117	0.0113	-0.0029	-0.0034	-0.0034
F33	0.0051	0.0043	0.0040	0.0007	0.0009	0.0008
F34	0.0055	0.0040	0.0035	0.0006	0.0008	0.0007
F1-F38	0.0173	0.0153	0.0153	-0.0001	-0.0002	-0.0002

Table V shows Root mean squared error (RMSE) in dollars and mean error (ME) in dollars for Model 1, Model 2 and Model 3 from December 31, 2015 to December 31, 2018.[14] We calculated RMSE and ME for all 38 contracts using Model 1, Model 2 and Model 3 and selected pairs of odd and even contracts to display in Table V. The pairs of odd and even contracts displayed in Table V represents minimum values, maximum values or turning points for RMSE or ME as a function of time to maturity.

From Table V we see that it's difficult to separate Model 2 and Model 3 in terms of monetary accuracy. Furthermore, we see that the seasonally adjusted models are more accurate than the unadjusted model, with less prevalent differences for contracts with time to maturity around

[14] Formulas for RMSE and ME are $RMSE(contract_i) = \sqrt{\frac{1}{n}\sum_t \left(F_{t,T_i|t}^{obs} - F_{t,T_i|t}^{model}\right)^2}$ and $ME(contract_i) = \frac{1}{n}\sum_t \left(F_{t,T_i|t}^{obs} - F_{t,T_i|t}^{model}\right)$, respectively.

20 months. There is no fourth digit difference in full sample accuracy for the seasonally adjusted models and a third digit full sample accuracy difference between the unadjusted model and the seasonally adjusted models.

All the models are less accurate for shorter time to maturity contracts with accuracy increasing toward an approximate time to maturity of 10 months. After the 10th month maturity mark, accuracy decreases towards the 20th month maturity mark, followed by another increase towards peak accuracy with time to maturity at around 33 months. This accuracy pattern is in line with the estimated standard deviation of the measurement errors.

Furthermore, in comparing the accuracy for odd and even contracts we cannot detect out-of-sample accuracy differences that are not better explained by the time to maturity difference between said odd and even numbered contracts.

From the mean error measurements, we see that all the models tend to underestimate shorter time to maturity contract value, overestimate contract values with time to maturity of around 10 months and underestimate contract values for contracts with time to maturity of around 26 months. Overall the models have a fourth digit tendency to underestimate contract value.

Figure 8.A and 8.B shows the term structure of the natural logarithm of NY Harbor ULSD futures prices on November 11, 2016 and January 1, 2018, respectively. The crosses are the observed futures prices, the dashed line represent Model 1 estimated futures prices and the solid line represent Model 3 estimated futures prices. Note that Figure 8.A and 8.B have different y-axis ranges, so one must be careful in comparing distances observed in the figures.

In both Figure 8.A and 8.B we see clear signs of seasonality in heating oil futures prices. Moreover, we see that the seasonally adjusted model, Model 3, captures this seasonality well, though less so for contracts with maturities between 13 and 24 months.

Furthermore, we see that the unadjusted model, Model 1, provides a decent fit. Naturally, it cannot compete with the seasonally adjusted model in terms of accuracy, but it captures both the average slope and location tendency well.

8.A

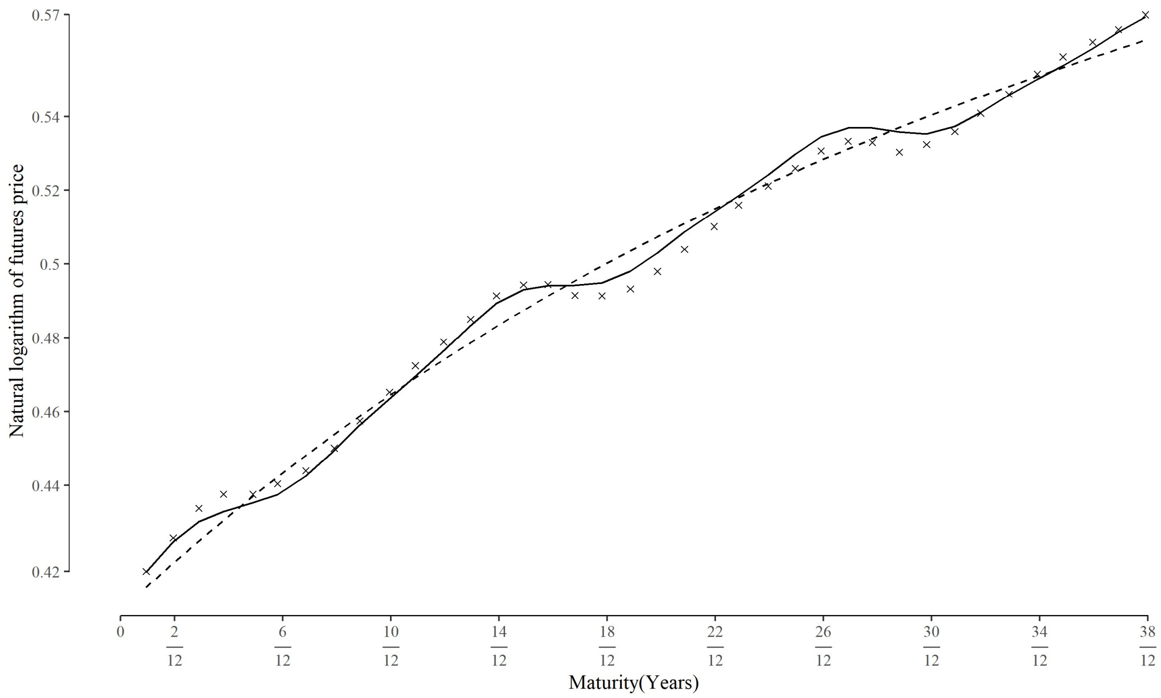

8.B

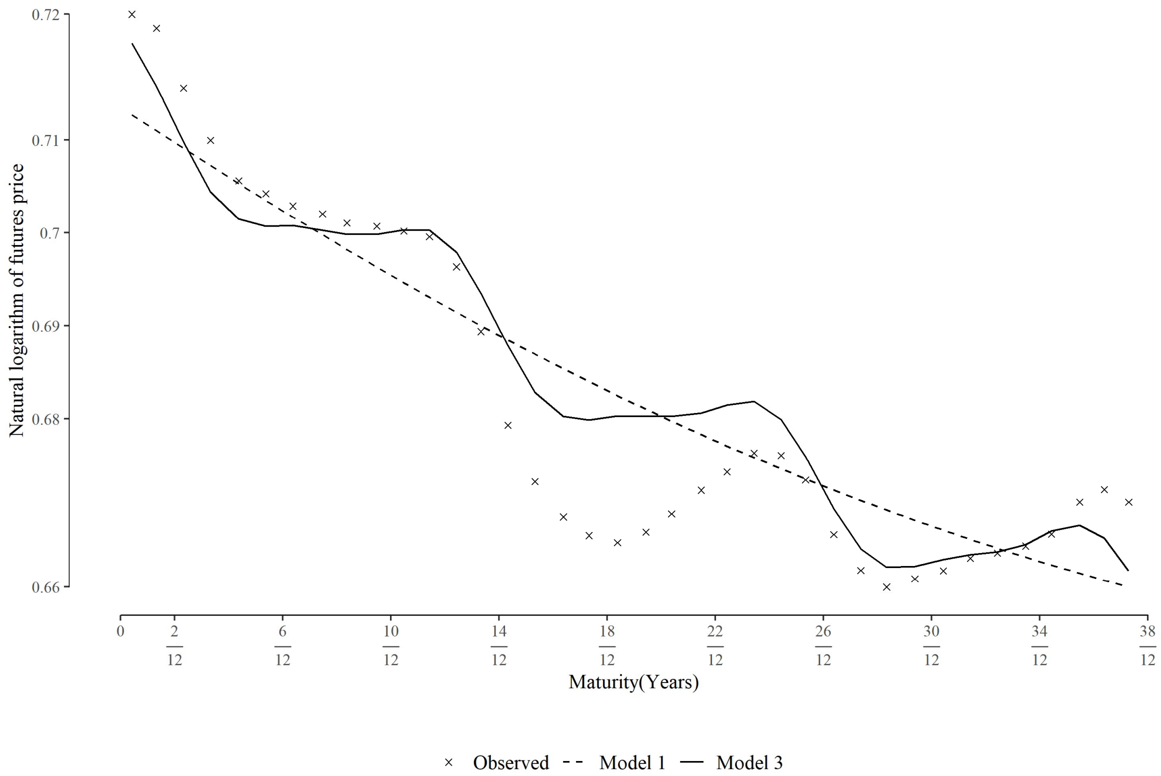

Figure 8. Term structure of the natural logarithm of futures prices. Figure 8.A and 8.B shows the term structure of the natural logarithm of NY Harbor ULSD futures prices on November 11, 2016 and January 1, 2018, respectively. The crosses are the observed futures prices, the dashed line represent futures prices estimated with Model 1 and the solid line represent futures prices estimated with Model 3. Note that Figure 8.A and 8.B have different y-axis ranges.

5.5 Volatility Term Structure of Futures Returns

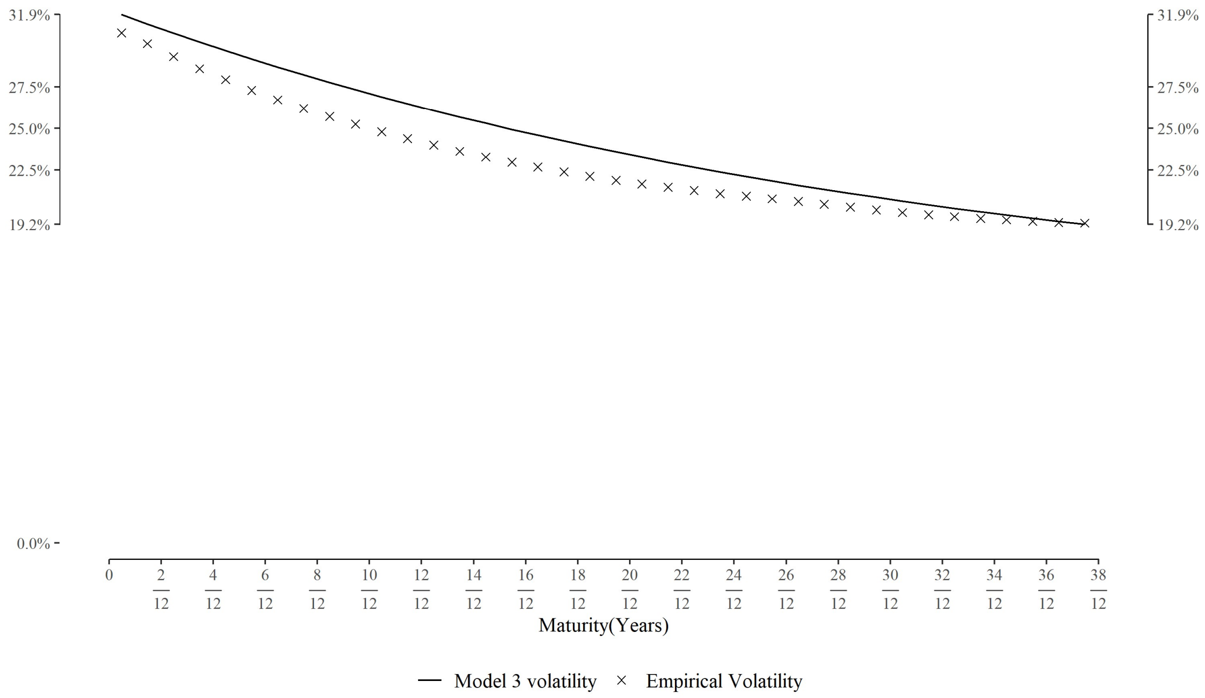

Figure 9. Volatility term structure of futures returns. The solid line shows the theoretical volatility term structure of futures returns given by (7) and the estimated parameters of Model 3. The crosses show the empirical volatility of futures returns based on NY Harbor ULSD futures contracts from December 31, 2015 to December 31, 2018.

Figure 9 shows the theoretical volatility term structure of NY Harbor ULSD futures returns given by (7) and Model 3's parameter estimates, and the empirical volatility of futures returns based on NY Harbor ULSD futures contracts from December 31, 2015 to December 31, 2018. We calculated the empirical volatility of futures returns by first calculating a return series for each individual actual contract in the dataset, i.e. not the generated F# contracts. Next, we mapped these return values to the F# contracts and calculated sample standard deviations of these series using standard techniques and annualized by multiplying with $\sqrt{1/\Delta t}$.

We see that the Samuelson Hypothesis holds for heating oil futures in our time window, which is to be expected as we've established significant mean reversion in heating oil's price process under an equivalent martingale measure. Moreover, we see that the theoretical and empirical volatility term structure of futures returns follow each other closely, though the theoretical volatilities are slightly overstated. Nevertheless, we think the proximity of the estimates is striking as volatility term structure of futures returns is not directly fitted on returns series.

For a one-factor model defined by (1). (2) and (3) the theoretical volatility term structure of futures returns is given by a horizontal line. The uppermost portion of the white space in Figure 9, between 0% and 19.2%, is where we should expect to find the theoretical volatility term structure of futures returns for a one-factor model. Indeed, when we fitted one-factor models we found that that σ_1 for these models was around 19%.

5.6 One-Factor, Three-Factor and Four-Factor Model

In addition to the models presented in Table III we investigated a one-factor, a three-factor and a four-factor model.

The objective function for the one-factor model was well behaved and quickly converged to points on the log-likelihood surface yielding economically sensible parameter estimates. However, the one-factor model could not compete with the two-factor model in terms of accuracy. Indeed, if we look at (6) from Section 2 in terms of a one-factor model we get

$$ln(F_{t,T}) - s(T) \quad = \quad x_t + \left(\mu^Q + \frac{1}{2}\sigma_1^2\right)(T-t).$$

Hence, the one-factor model assumes that the seasonally adjusted natural logarithm of the term structure of futures prices is linear and only allows for parallel shifts in the natural logarithm of the term structure of futures prices as time progresses. In view of Figure 6, Table III and to some extent Table II, we see that these assumptions do not hold in our time window. However, when we relax the fixed parameter assumption, allowing for time dependence in $\left(\mu^Q + \frac{1}{2}\sigma_1^2\right)_t$, the one-factor model performs well in longer subsections of our time window, though the two-factor models presented in Table III (still) provide a better fit.

The objective function for models with more than two factors was less well behaved and converged slowly to points on the log-likelihood surface yielding economically insensible parameter estimates. We believe this was caused by overparameterization as we observed similar behavior when conducting simulated overparameterization tests.[15] However, in the simulated overparameterization tests we got accurate estimates for some of the actual

[15] As a simulated overparameterization test we tried to fit a three-factor model using some of the futures price panels used in the simulation study presented in Section 3 Simulation.

parameters, leading us to believe that sample size was a contributing factor in the problems with parameter estimates using real data.

In order to test the overparameterization hypothesis, we examine the linear properties of the relationship between futures prices and Model 3 for each date in our sample through ordinary least squares (OLS) regression. Furthermore, we examine the trade-off between goodness-of-fit and parsimony using the output from the OLS regressions.

From (6) we have that

$$\underbrace{ln(F_{t,T}) - s(T) - A(T-t)}_{\parallel} = \underbrace{\begin{bmatrix} 1 \\ exp(-\kappa_2(T-t)) \\ \vdots \\ exp(-\kappa_m(T-t)) \end{bmatrix}^{\top} \boldsymbol{x}_t}_{\parallel}$$

$$\underbrace{y_{t,T}}_{\parallel} = \underbrace{\mathrm{x}_t^{(1)} + \boldsymbol{x}_t^{(2,3,\ldots,\mathrm{m})\top} \begin{bmatrix} exp(-\kappa_2(T-t)) \\ \vdots \\ exp(-\kappa_m(T-t)) \end{bmatrix}}_{\parallel}$$

$$y_{t,T} = \alpha_t + \boldsymbol{\beta}_t^{\top} \begin{bmatrix} exp(-\kappa_2(T-t)) \\ \vdots \\ exp(-\kappa_m(T-t)) \end{bmatrix}. \quad (9)$$

Using (9) we regressed $y_{t,T_{1,2,\ldots,u_t}}$ on $exp\left(-\kappa_2\left(T_{(1,2,\ldots,u_t)} - t\right)\right)$ for each date in our sample, using Model 3's estimated parameter values to calculate $exp\left(-\kappa_2\left(T_{(1,2,\ldots,u_t)} - t\right)\right)$, $s(T)$ and $A(T-t)$. Next, we gathered R^2 for each regression and inspected its distribution. We found that the minimum value of R^2 was 0.9, the interquartile range was [0.97,1], the mean was 0.98 and the median was 0.99. In other words, $y_{t,T}$ is approximately linear with respect to $exp\left(-\kappa_2\left(T_{(1,2,\ldots,u_t)} - t\right)\right)$ when $exp\left(-\kappa_2\left(T_{(1,2,\ldots,u_t)} - t\right)\right)$, $s(T)$ and $A(T-t)$ are calculated with Model 3's estimated parameter values.

Next, let's consider adding a second independent variable in (9), using the sample size adjusted Akaike information criterion (AIC) to measure the trade-off between goodness-of-fit and parsimony. AIC for a least squares model is given by

$$AIC = n \times \log\left(\frac{\mathrm{SSR}}{\mathrm{n}}\right) + 2k + \frac{2k(k+1)}{n-k-1},$$

where SSR is the sum of the squared residuals for a least squares model, n is the number of observations and k is the total number of parameters including SSR and the intercept (see e.g. Burnham and Anderson (2004)). Clearly, the first summand of AIC measures goodness-of-fit in terms of prediction error and the last two summands provide a penalty for increased model complexity (so that we seek to minimize AIC).

As $R^2 = 1 - \frac{SSR}{SST} \approx 1$ for 75% of our sample, where SST is the total sum of squares, it follows that SSR is approximately 0 for 75% of our sample, so that adding a second independent variable to (9) would result in an increase in AIC. Hence, the AIC tells us that adding a second independent variable to (9) is ill advised in terms of the trade-off between goodness-of-fit and parsimony.

Therefore, adding more factors to the two-factor models can be considered if $\alpha_t \not\approx x_t^{(1)}$ and $\beta_t \not\approx x_t^{(2)}$, else adding more factors would for 75% of the sample provide a worse trade-off between goodness-of-fit and parsimony. We found that the mean absolute deviation between α_t and $x_t^{(1)}$ and between β_t and $x_t^{(2)}$ was 0.0082 and 0.0112, respectively. Furthermore, we found that for approximately 95% of the sample, $x_t^{(1)}$ and $x_t^{(2)}$ was within the 95% confidence interval of α_t and β_t, respectively.

This result holds regardless of change in s and A caused by adding more factors to the two-factor model. In fact, taking A into account strengthens the result as adding more factors increases the number of parameters in A, which reduces parsimony. This reduction in parsimony is not directly observable in our examination, but can be inferred through the general idea of penalizing the inclusion of more parameters.

In Cortazar and Naranjo's (2006) study of crude oil price modeling the authors found grounds for using three or four-factor models, suggesting two or three mean reverting factors in crude oil's price process under an equivalent martingale measure. Furthermore, the authors found that the accuracy of a two-, three- and four-factor model was comparable for maturities between 3 months and 3 years, and for maturities between 3 and 7 years the two-factor model was substantially less accurate than the three- and four-factor models. The maximal time to maturity of the futures contracts used in this thesis is approximately 3 years, so we believe the difference in maximal contract maturity is the root cause of the difference in our findings regarding the appropriate number of mean reverting factors. However, we cannot rule out that the difference in findings reflect differences in the two commodities' price process under an equivalent martingale measure.

Section 6

In this section we provide examples of practical implications of our findings, conclude, discuss assumptions and questions open for further study.

6.1 Practical Implications

In this subsection we examine some of the implications of our findings regarding mean reversion and seasonality in heating oil futures prices, using the relations derived in Section 2 and two simple business examples to provide context. In particular, we use the relation between futures prices and spot prices under Q given constant interest rates, the closed form solution to the value of a futures contract and the closed form solution to the value of a European call option on a futures contract.

First, let's consider a business example where a heating oil refinery generate profits given by π_t, where π_t is linear with respect to the spot price of heating oil:

$$\pi_t = KS_t - F_c.$$

Here K is maximal output per period, S_t is the heating oil spot price per gallon at time t and F_c is the fixed cost per period.[16] Assuming constant interest rates r, the present value of π_T, where $T > 0$, is given by

$$\begin{aligned} Pv(\pi_T) &= e^{-rT} E_0^Q(KS_T - F_c) \qquad (10) \\ &= e^{-rT}\left(KE_0^Q(S_T) - F_c\right) \\ &= e^{-rT}\left(KF_{0,T} - F_c\right), \end{aligned}$$

where $F_{0,T}$ is the current heating oil futures price for a contract maturing at time T. This result follows straightforwardly from $F_{t,T} = E_t^Q(S_T)$, and equivalent martingale properties (see e.g. Hull (2012)). From (10) we see that changes in $F_{0,T}$ yield a similar change in $Pv(\pi_T)$ scaled by $e^{-r}\ K$. In other words, the error in $Pv(\pi_T)$ caused by a lack of seasonal adjustment of the

[16] This setup is taken from class notes in the course *ECO423 - Principles of Derivatives Pricing and Risk Management*.

estimated futures price is equal to the error in the estimated futures price caused by the lack of seasonal adjustment multiplied by a discount factor multiplied by an output factor. Though the discount factor is important, the major implication here is that the error scales with output. For example, Table V tells us that for a heating oil futures contract with time to maturity at approximately 10 months, the accuracy improvement in $Pv(\pi_{10/12})$ provided by including a seasonal adjustment is approximately \$0.0061$\times e^{-r\times\frac{10}{12}} \times K$. If we let the interest rate r be 0 and output K equal to the contract unit of a NY Harbor ULSD futures contract, 42000 gallons, the accuracy improvement is \$252.6.

Next, let's consider a business example where a firm that utilize heating oil to generate profit has been offered a contract giving them the right to enter into a long position in a heating oil futures contract for a fixed price K at time t, where the heating oil futures contract matures at time T, where $0 < t \leq$ T. The present value of the contract offered to the firm is equal to the value of a European call option on a heating oil futures contract expiring at time T with strike price K and time t until option maturity.

From Table III, we have that for a two-factor model defined by (1), (2) and (3) the shared parameters of the seasonally adjusted models and the unadjusted model are similar. Hence, (8) tells us that the seasonal adjustment mainly affects the value of the contract through changes in estimated $ln(F_{t,T})$ and not the volatility of $ln(F_{t,T})$.

From Figure 10 we see that if we use Model 1 to estimate futures prices, which lacks a seasonal adjustment, we get both positive and negative estimation errors of the present value of the contract, dependent upon the year placement of T. For example, assuming a contract unit of 42000 gallons, the error in present value estimates on November 11, 2016 using Model 1 for a contract with strike price $K = 1.42$, option and futures maturities $t = T \approx 14/12$ and interest rate $r = 0$ is estimated to be $\$0.0126 \times 42000 = \$529.2$. Furthermore, we see that the magnitude of the estimation error is slightly dependent upon moneyness, which is due to differences in the estimated volatility of $ln(F_{t,T})$ using the estimated parameters of Model 1 and Model 3.

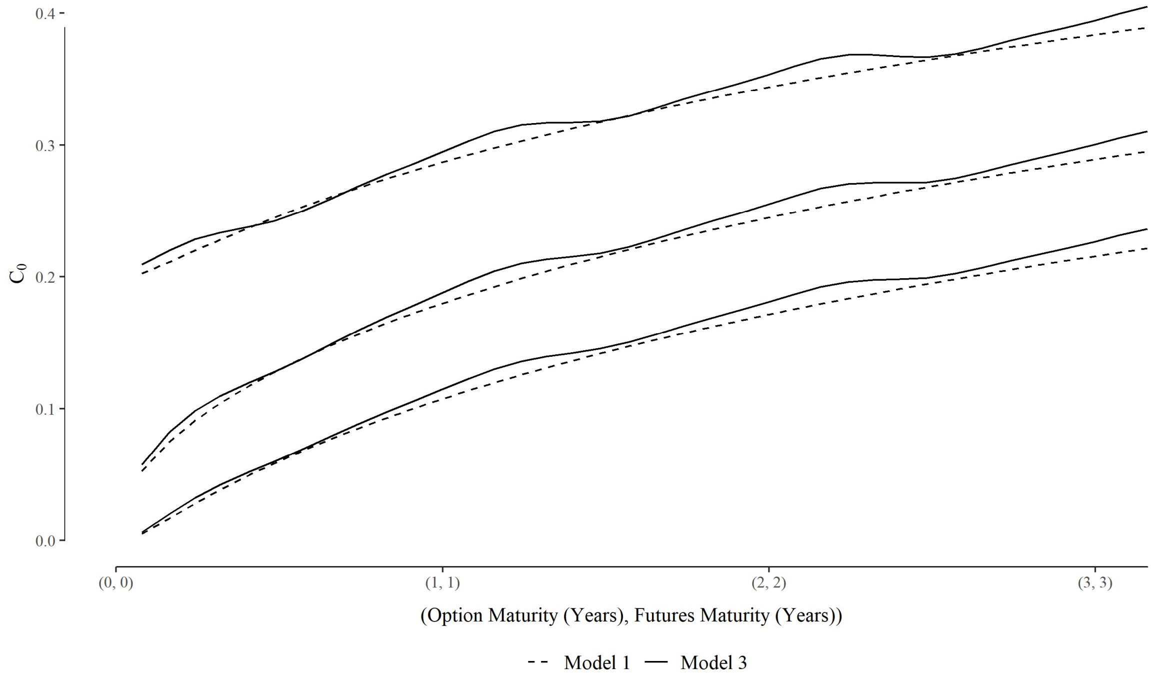

Figure 10. The value of a European call option on a heating oil futures contract at different levels of moneyness, with futures prices estimated by using a seasonally adjusted and an unadjusted model defined by (1), (2) and (3). The value of a European call option on a heating oil futures contract with futures prices estimated by using a seasonally adjusted and an unadjusted model defined by (1), (2) and (3). The solid lines shows option value at different levels of moneyness using a two-factor model with the parameter estimates of Model 3, and the dashed line shows option value at different levels of moneyness using a two-factor model with the parameter estimates of Model 1. $F_{0,T}$ is equal to e raised to the power of the values presented in Figure 8.A for Model 3 and Model 1, the strike prices K are $F_{0,T}^{Model\ 1} - 0.2$, $F_{0,T}^{Model\ 1}$ and $F_{0,T}^{Model\ 1}\ + 0.2$, option maturities are equal to futures maturities and the interest rate r is 0.

As the shared parameters of Model 1 and Model 3 are similar, a comparison of the differences in contract value caused by the differences in the volatility of $ln(F_{t,T})$ holds little interest. However, from the examination of the term structure of volatility of futures returns, we found major differences between theoretical volatilities of futures returns using parameter estimates of a one-factor and a two-factor model defined by (1), (2) and (3). These differences translate to differences in the volatility of $ln(F_{t,T})$, which in turn impacts contract value.

Figure 11 shows the difference between contract value estimates using a two-factor model and a one-factor model, where the estimated futures price is assumed to be equal for the two models. Hence, Figure 11 shows the difference between contract value estimates using a two-factor model and a one-factor model that is caused by differences in the estimated volatility of $ln(F_{t,T})$. We see that the difference between contract value estimates using a two-factor model and a one-factor model decreases with moneyness. Moreover, we see that as the difference between T and t increases, the difference between contract value estimates using a two-factor model and a one-factor model decreases. Intuitively, this makes sense as the additional factor in the two-factor model captures transitory changes in supply and demand of heating oil, which has less impact on the volatility of $ln(F_{t,T})$ as $T - t$ grows large.

From Figure 9 we see that the theoretical volatility of futures returns is slightly overstated using parameter estimates of a two-factor model and that the theoretical volatility of futures returns is substantially understated using parameter estimates of a one-factor model for short maturities. Hence, the contract value will lie somewhere between the solid line and the dashed line in Figure 11, closer to the solid line for small $T - t$. For example, assuming a contract unit of 42000 gallons, option maturity $t = 1$, futures maturity $T = 1.25$, futures price $F_{0,1.25} = 1$, strike price $K = 0.75$ and interest rate $r = 0$, the present value of the contract is estimated to be between $\$0.1695 \times 42000 = \$7119$ and $\$0.1937 \times 42000 = \$8135.4$, closer to $\$8135.4$ than $\$7119$.

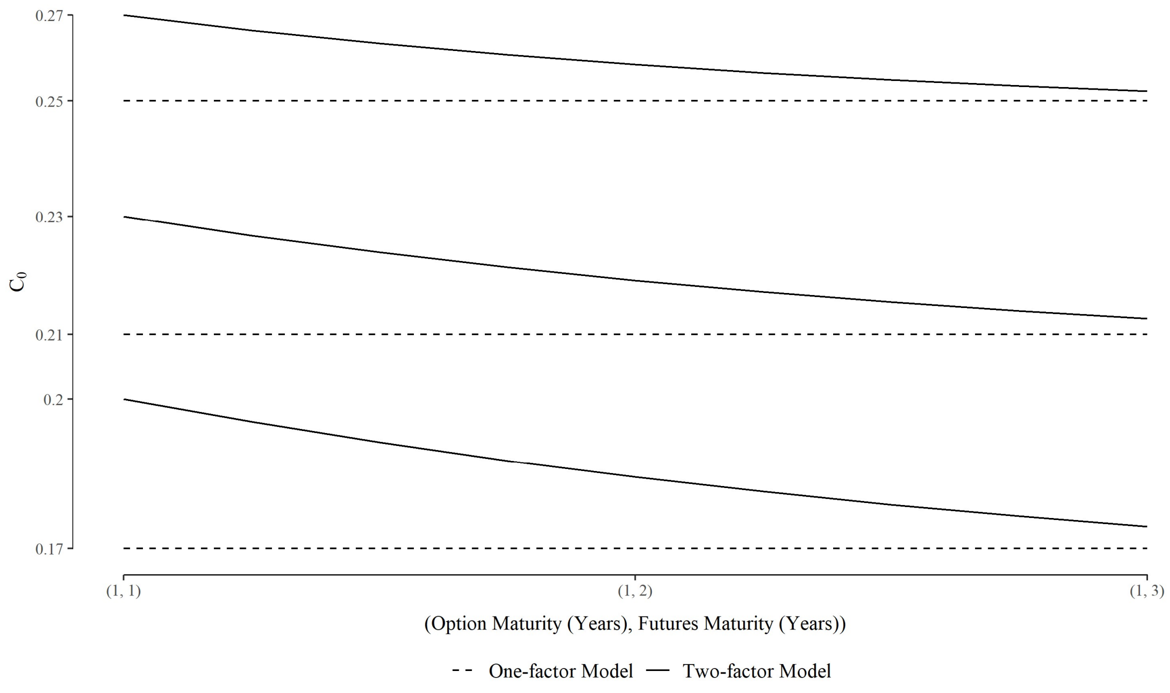

Figure 11. The value of a European call option on a heating oil futures contract at different levels of moneyness, with volatility of the natural logarithm of future prices estimated by using a one-factor model and a two-factor model defined by (1), (2) and (3). The figure shows the value of a European call option on a heating oil futures contract given by (8). The volatility of the natural logarithm of futures prices are estimated by using a model defined by (1), (2) and (3). The solid lines shows option value at different levels of moneyness using a two-factor model with the parameter estimates of Model 3, and the dashed line shows option value at different levels of moneyness using a one-factor model with $\sigma_1 \equiv 19.2\%$. The remaining inputs of (8) are $F_{0,T} = 1$, the strike price K is 0.75, 0.8 and 0.85 and the interest rate r is 0.

6.2 Conclusions

In this thesis, we've considered an affine N-factor Gaussian model developed by Cortazar and Naranjo (2006), and studied its ability to explain the stochastic behavior of heating oil futures prices. In doing so, we've determined the parsimonious number of latent unobservable factors to include in the N-factor model, which depends on the existence and scale of mean reverting properties in heating oil's price process under an equivalent martingale measure. Furthermore, we examined seasonality in heating oil futures prices, which depends on the existence and scale of seasonality in heating oil's price process under an equivalent martingale measure.

The model parameters were estimated using maximum-likelihood techniques and the Kalman filter. We conducted a simulation study of the parameter estimation procedure, following Andresen and Sollie (2013), to highlight potential econometric issues. We found that the models are robust for use in valuation problems and generally not robust for forecasting purposes, confirming Schwartz and Smith's (2000) and Andresen and Sollie's (2013) findings.

The cost of crude oil is a major cost driver for heating oil, and as several previous studies have found mean reversion in crude oil's price process under an equivalent martingale measure mean reversion in heating oil's price process under an equivalent martingale measure is probable. Furthermore, from an examination of two timeseries of heating oil futures prices with different maturities we found that the long-term growth rate of observed heating oil futures prices alternate signs, indicating presence of mean reversion in heating oil's price process under an equivalent martingale measure.

In Cortazar and Naranjo's (2006) study of crude oil price modeling, the authors found grounds for using three- or four-factor models, suggesting two or three mean reverting factors in crude oil's price process under an equivalent martingale measure. Furthermore, the authors found that the accuracy of model estimated futures prices given by a two-, three- and four-factor model was comparable for futures maturities between 3 months and 3 years, and that for maturities between 3 and 7 years the two-factor model was substantially less accurate than the three- and four-factor models. Based on parameter estimates of a two-factor model, fitted on NY Harbor ULSD futures, we found significant mean reversion in heating oil's price process under an equivalent martingale measure. However, we could not find grounds for including more than one mean reverting factor in the model. The maximal time to maturity of the futures contracts used to estimate parameters in this thesis was approximately 3 years, so we believe the difference in maximal contract maturity is the root cause of the difference in our findings regarding the appropriate number of mean reverting factors. Still, we cannot rule out that the difference in findings reflect differences in the two commodities' price process under an equivalent martingale measure.

Heating oil's use as a heating fuel implies that there is an area dependent seasonal distillate fuel oil consumption pattern that is inversely related to the temperature in that area. We found signs of this seasonality in the U.S. distillate fuel oil inventory, and there appears to be a positive correlation between changes in the U.S. distillate fuel oil inventory and changes in the expected development of heating oil prices. Hence, seasonality in heating oil's price process is probable. Indeed, based on likelihood-ratio tests of the two-factor model and seasonally adjusted two-factor models, we found that there is statistically significant seasonality in heating oil's price process under an equivalent martingale measure. Furthermore, we found that allowing for more than two inflection points to capture this seasonality provided a statistically significantly better fit. In Roberts and Lin's (2006) study of seasonality in the price process for different commodities the authors found significant seasonality in crude oil's price process under an equivalent martingale measure, though the authors could not find grounds for allowing more than two inflection points in describing this seasonality. The contrast to our results indicates that there are significant differences between the seasonality of the two commodities' price process under an equivalent martingale measure.

Lastly, we provided two simple business examples to show some simple practical implications of our findings. In the first example, we evaluated the present value of a cash flow that is linear in the price of heating oil and find that errors in estimates of heating oil futures prices, caused by for example an omitted seasonal factor, yield a similar error in estimates of the present value of the cash flow scaled by a discount factor and output. In the second example, we evaluate the value of a contract offer with characteristics equal to a European option on a heating oil futures contract. We show that seasonality in heating oil futures prices affects the present value of such a contract, and that errors caused by omitting the seasonal factor have magnitude and sign dependent upon the year placement of futures maturity. Furthermore, we find that omission of mean reverting properties in heating oil's price process under an equivalent martingale measure impacts the volatility term structure of futures returns, which in turn impacts the estimated present value of the contract. We find that contract value estimates based on a model without mean reverting properties will generally understate the present value of a contract offered on futures with short maturities and will be fairly accurate for a contract offered on futures with longer maturities. The contract value accuracy problems of a model without mean reverting properties are reduced when moneyness increases.

6.3 Assumptions and Further Works

The three main assumptions of this thesis are the no-arbitrage assumption, the functional form of the stochastic process for heating oil prices and the constant parameter assumption.

The no-arbitrage assumption is reasonable considering the high liquidity and the product maturity of NY Harbor ULSD futures.

The functional form assumption can always be debated, though we've shown that the chosen functional form is equal to the functional form utilized in several previous comparable studies of commodities with similar properties to that of heating oil. Moreover, the standard deviation of the measurement errors and accuracy estimates indicate that the functional form is relatively close to the true process.

The constant parameter assumption is known to be a simplification, as it violates previous empirical findings and general knowledge (e.g. interest rates are not constant). For example, Casassus and Collin-Dufresne (2005) found empirical evidence for time varying risk premia which implies mean reversion in commodity prices under the subjective measure. However, we've seen that the constant parameter simplification is less critical for the two-factor model and given this study's time span we do not expect that findings regarding existence of mean reversion or seasonality will change if we relax this assumption.

In future work one could relax some of the assumptions and utilize additional sources of information to pin down the parameter values. As an example, one could use the capital asset pricing model (CAPM) (see e.g. Hull (2012)) β's to implement time varying CAPM restricted risk premia, lifting the constant risk premia assumption and reducing the indeterminacy issue. One could utilize option implied volatility, lifting the constant volatility assumption and providing market guidance for the values of the parameters. Furthermore, there might be room for a more complex seasonality function, i.e. one that takes on different values for T and $T + a \in \mathbb{Z} > 0$.

Lastly, assuming well-functioning markets, one could examine the implications of our findings in terms of different economic quantities such as convenience yield.

References

Andresen, A., & Sollie, J. M. (2013). Multi-factor models and the risk premiums: a simulation study. *Energy Systems 4*, 301–314.

Bessembinder, H., Coughenour, J. F., Seguin, P. J., & Smeller, M. M. (1996). Is There a Term Structure of Futures Volatilities? *The Journal of Derivatives 4(2)*, 45-58.

Brennan, M., & Schwartz, E. S. (1985). Evaluating Natural Resource Investments. *The Journal of Business 58(2)*, 135-157.

Burnham, K. P., & Anderson, D. R. (2004). Multimodel Inference - Understanding AIC and BIC in Model Selection. *Sociological Methods & Research 33(2)*, 261-304.

Casassus, J., & Collin-Dufresne, P. (2005). Stochastic Convenience Yield Implied from Commodity Futures and Interest Rates. *The Journal Of Finance 60(5)*, 2283-2331.

CME Group (a). (2019). *NY Harbor ULSD Futures Contract Specs*. Retrieved from cmegroup.com: https://www.cmegroup.com/trading/energy/refined-products/heating-oil_contract_specifications.html

CME Group (b). (2019). *Conversion Calculator*. Retrieved from cmegroup.com: https://www.cmegroup.com/tools-information/calc_crack.html

CME Group (c). (2019). *NYMEX Heating Oil*. Retrieved from cmegroup.com: https://www.cmegroup.com/confluence/display/EPICSANDBOX/NYMEX+Heating+Oil

Cortazar, G., & Naranjo, L. (2006). An N-Factor Gaussian Model of Oil Futures Prices. *Journal of Futures Markets 26*, 243-268.

Cortazar, G., & Schwartz, E. S. (2003). Implementing a Stochastic Model for Oil Futures Prices. *Energy Economics 25*, 215–238.

Cortazar, G., Kovacevic, I., & Schwartz, E. S. (2013). Commodity And Asset Pricing Models: An Integration (No. w19167). *National Bureau of Economic Research*.

Cox, J. C., Ingersoll, J., & Ross, S. (1981). The relation between forward prices and futures prices. *Journal of Financial Economics 9*, 321-346.

Dai, Q., & Singleton, K. J. (2000). Specification Analysis of Affine Term Structure Models. *The Journal of Finance 55(5)*, 1943-1978.

EIA (a). (2019). *Glossary*. Retrieved from eia.gov: https://www.eia.gov/tools/glossary/

EIA (b). (2019). *What drives petroleum product prices?* Retrieved from eia.gov: https://www.eia.gov/finance/markets/products/

EIA (c). (2019). *Heating Oil*. Retrieved from eia.gov: https://www.eia.gov/energyexplained/index.php?page=heating_oil_home

EIA (d). (2019). *Heating oil futures contract now uses ultra-low sulfur diesel fuel*. Retrieved from eia.gov: https://www.eia.gov/todayinenergy/detail.php?id=11211

FIA. (2019). *FIA releases annual trading statistics showing record ETD volume in 2018*. Retrieved from fia.org: http://fia.org/articles/fia-releases-annual-trading-statistics-showing-record-etd-volume-2018

Gibson, R., & Schwartz, E. S. (1990). Stochastic convenience yield and the pricing of oil contingent claims. *The Journal of Finance 45(3)*, 959-976.

Hallam, A. (2019). *Asymptotic Distribution of Maximum Likelihood Estimators*. Retrieved from iastate.edu: http://www2.econ.iastate.edu/classes/econ671/hallam/documents/Asymptotic_Dist.pdf

Helske, J. (2017). KFAS: Exponential Family State Space Models in R. *Journal Of Statistical Software 78(10)*, 1 - 39.

Hull, J. C. (2012). *Options, futures, and other derivatives 8th ed.* Prentice Hall.

Koopman, S. J., & Durbin, J. (1998). Fast Filtering And Smoothing For Multivariate State Space Models. *Journal of Time Series Analysis 21(3)*, 281-296.

Mebane, W. R., & Sekhon, J. S. (2007). Genetic Optimization Using Derivatives: The rgenoud Package for R. *Journal of Statistical Software 42(11)*, 1-26.

Miltersen, K. R. (2003). Mathematical methods and models in finance.

Mullen, K. M. (2014). Continuous Global Optimization in R. *Journal Of Statistical Software 60(6)*, 1-45.

Power, G. J., & Turvey, C. G. (2008). On Term Structure Models of Commodity Futures Prices and the Kaldor-Working Hypothesis. *Proceedings of the NCCC-134 Conference on Applied Commodity Price Analysis, Forecasting, and Market Risk Management.*

Roberts, M. C., & Lin, C. (2006). A Term Structure Model for Commodity Prices: Does Storability Matter? *2006 NCCC-134 Conference on Applied Commodity Price Analysis, Forecasting, and Market Risk Management.*

Roberts, M., & Fackler, P. L. (1999). A Term Structure Model for Agricultural Futures. *In Selected Paper Presented at AAEA annual meeting*.

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

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

Sorensen, C. (2002). Modeling Seasonality in Agricultural Commodity Futures. *The Journal of Futures Markets 22(5)*, 393–426.

Thomson Reuters Datastream. (2019). *Thomson Reuters Datastream Database*. Retrieved from https://infobase.thomsonreuters.com/infobase/login/?next=/infobase/

Trolle, A. B., & Schwartz, E. S. (2008). Unspanned Stochastic Volatility and the. *Review of Financial Studies 22(11)*, 4423-4461.

Welch, G., & Bishop, G. (2006). *An Introduction to the Kalman Filter.*

World Economic Forum. (2016). *155 years of oil prices - in one chart*. Retrieved from weforum.org: https://www.weforum.org/agenda/2016/12/155-years-of-oil-prices-in-one-chart/

Appendix

Production, Trade and Financial Markets

For completeness, we here briefly discuss the factors impacting the current and expected supply and demand for heating oil that were not discussed in Section 1. Statements regarding Production, Trade and Financial Markets are sourced from EIA (b) (2019).

Production is directly linked to supply. In general, refineries adjust their product mix and volumes according to long-term price and demand signals through long-term investment decisions and other long-term strategy decisions. The more complex refineries can also adjust their product mix and output volume according to short-term price and demand signals.

The refinery's response to price signals is reflected in, among other metrics, the positive correlation between the distillate to gasoline production ratio and the spread between distillate- and gasoline futures prices. The effect of the refinery's response to demand signals is reflected in the aggregated growth or decline in distillate production and the seasonal production adaptations induced by the seasonal consumption pattern of heating oil and other petroleum products.

Trade represents import and export-related factors. In general, the level of free trade impacts the size of the total marketplace and the local market's ability to cover unexpected shortages and surpluses.

Financial Markets represent supply and demand factors related to supply and demand for financial instruments like futures, options and crack spreads that are connected to heating oil. As an example, the rapid price decline in 2008 shown in Figure 1 was mostly due to financial market factors brought forth by the global financial crisis.

Nested Model Relation Equations

In Section 2.3, Nested Models, we presented a set of equations derived by Cortazar and Naranjo (2006) relating the parameters of two models through an affine transformation. We believe there is a typo in one of these equations ($\boldsymbol{\beta} = K\boldsymbol{\varphi}(t) + L\overline{\boldsymbol{\beta}}$) based on the following:

Let

$$\boldsymbol{f}(t, \boldsymbol{y}_t) = \boldsymbol{z}_t \tag{A.1}$$

$$= L\boldsymbol{y}_t + \boldsymbol{\varphi}(t)$$

where L is a time invariant matrix, $\boldsymbol{z}$ and $\boldsymbol{y}$ are vectors of Ito processes and $\boldsymbol{\varphi}(t)$ is a vector of functions of time. As $f_{1,2,\ldots,m}$ is linear in $\boldsymbol{y}$ Ito's lemma states that

$$dz_t^i = \frac{\partial f_i}{\partial t} dt + \nabla_{\boldsymbol{y}} f_{\mathrm{i}}^{\top} d\boldsymbol{y}_t, \tag{A.2}$$

where $\nabla_{\boldsymbol{y}} f_{\mathrm{i}}^{\top}$ is the gradient of f_i with respect to $\boldsymbol{y}$.

Next, let

$$d\boldsymbol{y}_{\mathrm{t}} = \left(-\overline{\mathrm{K}}\boldsymbol{y}_{\mathrm{t}} + \overline{\boldsymbol{\beta}}\right)dt + \overline{\Sigma}\mathrm{d}\overline{\mathbf{w}}_{\mathrm{t}} \tag{A.3}$$

By combining equation A.1, A.2 and A.3, we get

$$d\boldsymbol{z}_t = \frac{\partial \boldsymbol{f}}{\partial t} dt + L d\boldsymbol{y}_t$$

$$= L\left(-\overline{\mathrm{K}}\boldsymbol{y}_{\mathrm{t}} + \overline{\boldsymbol{\beta}} + L^{-1}\frac{\partial \boldsymbol{f}}{\partial t}\right)dt + \mathrm{L}\overline{\Sigma}\mathrm{d}\overline{\mathbf{w}}_{\mathrm{t}}$$

$$= L\left(-\overline{\mathrm{K}}\mathrm{L}^{-1}\left(\boldsymbol{z}_t - \boldsymbol{\varphi}(t)\right) + \overline{\boldsymbol{\beta}} + L^{-1}\frac{\partial \boldsymbol{f}}{\partial t}\right)dt + \mathrm{L}\overline{\Sigma}\mathrm{d}\overline{\mathbf{w}}_{\mathrm{t}}$$

$$= \left(-\mathrm{L}\overline{\mathrm{K}}L^{-1}\boldsymbol{z}_t + \mathrm{L}\overline{\mathrm{K}}L^{-1}\boldsymbol{\varphi}(t) + \mathrm{L}\overline{\boldsymbol{\beta}} + \frac{\partial \boldsymbol{f}}{\partial t}\right)dt + \mathrm{L}\overline{\Sigma}\mathrm{d}\overline{\mathbf{w}}_{\mathrm{t}}$$

$$= (-K\boldsymbol{z}_t + \boldsymbol{\beta}_{\mathrm{t}})dt + \Sigma d\boldsymbol{w}_t,$$

where $K = L\bar{K}L^{-1}$, $\boldsymbol{\beta}_{\mathrm{t}} = K\boldsymbol{\varphi}(t) + L\bar{\boldsymbol{\beta}} + \frac{d\boldsymbol{\varphi}(t)}{dt}$ and $\Sigma = L\bar{\Sigma}$. The change of Wiener process is derived using that almost surely, we can find a constant a such that $aW = bW_1 + cW_2$ where W denote a Wiener process and a, b and c are constants.