Commodity market modeling and physical trading strategies

by

Per Einar S. Ellefsen

Ingénieur de l'Ecole Polytechnique, 2008

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Master of Science in Mechanical Engineering
at the
Massachusetts Institute of Technology

June 2010

Signature of Author: ____

Department of Mechanical Engineering
May 13, 2010

Certified by: ____

Paul D. Sclavounos
Professor of Mechanical Engineering and Naval Architecture
Thesis Supervisor

Accepted by: ____

David E. Hardt
Professor of Mechanical Engineering
Chairman, Department Committee on Graduate Students

Commodity market modeling and physical trading strategies

by

Per Einar S. Ellefsen

Submitted to the Department of Mechanical Engineering on May 13, 2010 in partial fulfillment of the requirements for the degree of Master of Science in Mechanical Engineering

ABSTRACT

Investment and operational decisions involving commodities are taken based on the forward prices of these commodities. These prices are volatile, and a model of their evolution must correctly account for their volatility and correlation term structure. A two-factor model of the forward curve is proposed and calibrated to the crude oil, shipping, natural gas, and heating oil markets. The theoretical properties of this model are explored, with focus on its decomposition into independent factors affecting the level and slope of the forward curve. The two-factor model is then applied to two problems involving commodity prices. An approximate analytical expression for the prices of Asian options is derived and shown to explain the market prices of shipping options. The floating storage trade, which appeared in the oil market in late 2008, is presented as an optimal stopping problem. Using the two-factor model of the forward curve, the value of storing crude oil is derived and analyzed historically. The analytical framework for physical commodity trading that is developed allows for the calculation of expected profits, risks involved, and exposure to the major risk factors. This makes it possible for market participants to analyze such physical trades in advance, creates a decision rule for when to sell the cargo, and allows them to hedge their exposure to the forward curve correctly.

Thesis supervisor: Paul D. Sclavounos
Title: Professor of Mechanical Engineering and Naval Architecture

Table of Contents

1. INTRODUCTION

1. Commodity markets

On July 3rd, 2008, Brent crude oil futures were trading at 146 US dollars per barrel. The TD3 Arabian Gulf – Japan shipping route was quoted at 240 Worldscale, and analysts were predicting crude prices at 200 dollars within the next months. On December 3rd, Brent traded at 45 US dollars per barrel and TD3 at 70 Worldscale, drops of 69 and 71 percent respectively.

The commodity markets are among the most volatile in the world, and their volatility is a source of both profits and risks for the actors involved. In order to manage these risks the physical spot markets have from an early stage been accompanied by forward markets, later transforming into financial derivatives markets. The Chicago Board of Trade introduced exchange-traded futures contracts on agricultural products in 1848, and crude oil was traded forward from its beginnings in the 1860s (Yergin, 2008).

Most modern commodity markets consist of two intertwined markets: the physical and the financial market. The physical – or spot – market is made up of all market participants selling or taking delivery of the commodity product. In the crude oil market these are oil companies, refiners and physical trading companies. Trading in the spot market usually occurs through brokers, matching sellers and buyers of cargoes at specific dates and locations.

The financial commodity market is the market for derivative contracts based on the spot. These derivatives take the form of forwards, futures and options, and are used for risk management by companies involved in the physical market and speculation by other players. Importantly, the derivatives settle against the physical market, thereby linking the two. In some cases the derivatives are physically settled, i.e. the buyer receives the actual commodity. In others, the derivatives settle financially against a spot index published daily based on transactions in the physical market.

The relative volumes of the financial and physical markets depend on the level of development of the derivatives market. As seen in Appendix 1, in 2009, the volume of derivatives (futures and options) traded on crude oil was 303 billion barrels, compared with an annual world production of 33 billion barrels (CIA, 2009), making the derivatives market nine times the size of the physical market. In tanker shipping, the derivatives market traded 304 million tonnes of oil cargo in 2009, compared to 145 million tanker deadweight tonnes traded spot in 2006 (Stopford, 2009), which evaluates the size of the derivatives market to twice the physical market. There is still a large growth potential in the freight derivatives market, which will happen through standardization and changes in the conventions for physical price setting, similar to what has occurred in the oil market since the 1980s.

The linkage between the spot and forward markets for commodities will be the main topic of this thesis, and in particular how the financial market can be used to gain greater insight into physical trading decisions. While the focus will be on crude oil and tanker shipping, we will present results in a general setting and the same principles apply for dry commodities such as coal.

2. Definitions and markets

Definitions

In this thesis we will be considering a commodity market where the commodity is trading at a spot price $S(t)$ on date *t*. This is the market price for delivery as soon as possible, which can be the next day for electricity or during the next month for crude oil.

Associated with this market are forward prices $F(t,T)$ on date *t*. These are the prices in the market for delivery of the commodity at the date *T*, which is in the future. The distinction is often made between forward and future contracts, the latter being more standardized and marked-to-market daily, but we will not make such a distinction here. In the case when the forward is financially settled, a long position in the forward contract entered on date *t* will pay off $S(T)-F(t,T)$ on the settlement date *T*. As long as the physical market is liquid, entering a physical or financial forward contract is therefore equivalent with respect to market risk – both give a fixed purchase price of $F(t,T)$.

The set of forward contracts trading in the market allows us to construct a forward curve $F(t,T)$. Usually the maturities *T* are monthly, but they can be more granular in the short end. We will also index this curve by the time-to-maturity $\tau = T - t$, which is the time to settlement of the forward contract: $f(t,\tau) = F(t,t+\tau)$.

Crude oil market

The crude oil derivatives market is by far the largest commodity market, with a volume of 303 billion barrels traded in 2009 (ICE, 2009 and CME, 2009). It is, however, not a single commodity – the price of crude oil depends on its grade (mainly specific gravity and sulphur content) and location. There are however two reference grades of crude oil: BFOE (Brent, Forties, Oseberg and Ekofisk) in the North Sea and light sweet crude oil at Cushing, Oklahoma in the United States, also known as West Texas Intermediate (WTI). Most grades of oil at other locations are priced at a differential to these marker crudes.

WTI futures contracts trade on the New York Mercantile Exchange (NYMEX) and are physically deliverable into the pipeline system in Cushing, Oklahoma, during the month of the contract. This makes the front-month WTI contract the spot contract of crude oil at that location. However Cushing is inland and only reachable by pipeline, while imported crude oil will generally arrive by tanker at the Louisiana Offshore Oil Port (LOOP) in the Gulf of Mexico. We will therefore also be considering the spot price of Louisiana Light Sweet (LLS) which is a light sweet crude priced at St. James, Louisiana.

BFOE is the most complete crude forward market. The spot price, known as Dated Brent, is assessed daily by Platts from trades during the "Platts trading window", and corresponds to cargoes delivered between 10 and 21 days forward. Starting a month out there are futures contracts trading on the Intercontinental Exchange (ICE) and settling financially against the ICE Brent Index. Between the two, traders can hedge prices during more specific time windows using over the counter Brent Contracts For Difference (CFDs).

Using these different prices, a very precise forward curve can be constructed for BFOE, especially in the short end.

In addition to futures contracts, there is a liquid market for options on these crude oil marker prices. The exchange-traded options are mostly American and deliver a futures contract when exercised.

Tanker shipping market

While oil has been transported on ships since 1861 (Yergin, 2008), its transition from a logistical exercise controlled by the oil majors to a spot market is relatively recent, with 70% of spot chartering in the 1990s versus only 20% in 1973. In the spot market, tankers are chartered for a single voyage (e.g. Sullom Voe – LOOP) through brokers, with all costs included in the price. There are a number of reference routes for dirty and clean tankers, numbered TD1 through TD18 (dirty, i.e. crude) and TC1 through TC11 (clean, i.e. products). At the end of each trading day the Baltic Exchange polls brokers and publishes an assessment of the price level for each of these routes, making up the Baltic Exchange dirty and clean tanker indices. This is the recognized spot price in the tanker market.

Tanker rates are usually published in a unit called Worldscale (WS). This unit, specific to each route, is updated yearly by the Worldscale Association and represents a reference price for a reference tanker on the specific route, in US dollars per deadweight tonne. A spot price of WS100 will then be equal to 100% of this price, while WS150 would be 150% of this price. These tanker rates for voyage charters include all costs, i.e. fuel, port and canal costs. It is useful to back out a TimeCharter Equivalent (TCE) price for the ship, in US dollars per day, corresponding to the daily price of hiring the ship net of these costs. This requires knowing details about the distance covered, ship speed, fuel consumption and fuel prices. Based on this assessment the Baltic Exchange also publishes daily TCE prices for VLCC, Suezmax, Aframax and MR tankers.

The tanker market has also seen the relatively recent development of a Forward Freight Agreement (FFA) market. While BIFFEX[1] futures were traded as early as 1985 they lost popularity and have since been replaced by route-specific FFAs. FFAs settle financially at the end of their contract month on the arithmetic average of the daily values of the underlying Baltic Exchange index during that month. FFAs are traded through brokers, with the International Maritime Exchange (Imarex) having the largest market share in tanker FFAs. Liquidity is concentrated in a few key routes, such as TD3 (VLCC Arabian Gulf – Japan), TD5 (Suezmax West Africa – US Atlantic Coast) and TC2 (MR product tanker Rotterdam – New York). In addition to broker prices, the Baltic Exchange publishes a daily assessment of FFAs obtained by polling brokers.

An important specificity of the shipping market is that the commodity being traded, tonne-miles, is a service, not a physical commodity that can be stored. It is similar in this respect to electricity markets. While it is not impossible to store tonne-miles, it can be done by slow steaming or laying up ships for example, it is more difficult and this lack of inventories induces higher spot price volatility and lower correlations between forward contracts of different tenors.

There is also a nascent market in freight options, spearheaded by Imarex. These are of Asian style and, like the FFAs, settle on the average of a spot index during a month.

[1] Baltic Index Freight Futures Exchange, a Baltic Exchange initiative, existed from 1985 to 2001

3. Motivation

The volatility of commodity prices exposes market actors to considerable risk. All investment decisions involving commodities expose the investors to the forward curve. Such decisions include buying a coal mine, operating a power plant, ordering and canceling a new ship, trading commodities between different locations and writing options on a commodity.

As stressed in Dixit and Pindyck (1994), these decisions should not be taken based solely on forecasts of prices. The 40% yearly volatility of the crude oil spot price will have substantially more impact on investment decisions in crude oil assets than a forecasted growth of 2%.

The year 2008 was a particularly volatile year in the oil market. It was also marked by the transition of the forward market to a steep contango after years of backwardation as the spot price plummeted. At the same time, tanker rates fell 71%. This led to an array of tankers being used as floating storage facilities, anchored up near delivery ports to store the unused crude oil and take advantage of the contango, a phenomenon not seen since 1973. Deciding when to release crude from such a floating storage trade also depends on the forward market and price volatility.

While many such investments and trades are being executed, they are not, in general, evaluated using a proper framework. The correct valuation and operational decision-making for such investments or trades requires the use of a simple and correct model for the commodity forward curves involved. Such a model opens further possibilities of managing the firm's risk correctly and making informed choices about different possibilities.

4. Objectives

The main objective of this thesis is to develop a simple and efficient framework for the optimal physical trading of commodities. Such a framework will allow us to understand and analyze the floating storage trade that appeared in late 2008 and continued into 2009.

The questions we will attempt to answer in this thesis are:

- Can a simple two-factor forward curve model explain the historical volatilities and correlations of traded forward contracts, in different commodity markets?
- What is the consequence on the spot price process for such a two-factor model?
- How should commodity Asian options, as traded in the shipping market, be interpreted, priced, and hedged by market participants? What is the meaning of implied volatility for such options?
- When has the cross-Atlantic crude oil arbitrage window been open? When were there floating storage opportunities in this trade?

- What is the optimal floating storage strategy to follow to maximize profits for the trader? Is there value to keeping exposure to the forward curve by not selling the cargo forward immediately, and how can we understand this value?

- What is the optimal ship routing strategy to follow when a general physical trading problem is considered? When should a ship be re-routed from its initial destination?

5. Methodology and outline

The general framework we will be working under is that of continuous-time financial markets using Itô's stochastic calculus as formulated in Musiela and Rutkowski (2008). Securities prices will generally be assumed to follow diffusions of the type

$$\frac{dS(t)}{S(t)} = \mu(t)dt + \sigma(t)dW(t)$$

where $W(t)$ is a Brownian motion, $\mu(t)$ will be called the instantaneous drift and $\sigma(t)$ the instantaneous volatility of the stochastic process $S(t)$.

This thesis is both theoretical and practical. We present new models and new theoretical results. Each time we present a new model or result, however, we will also present its calibration to market data or historical performance and analyze those results.

In Part 2 we present a two-factor model of commodity forward curves and show that it reproduces the main historical features of the forward curves of four different commodities. We also explore its theoretical properties and reformulate it in terms of mean-reverting factors shocking the constant-maturity forward curve.

Using this parametric stochastic model of the forward curve we are able to derive an approximate evolution of average price contracts such as FFAs in Part 3. This then allows us to find approximate but closed-form formulas for Asian options that take into account the main features of commodity futures: the term structure of prices and the term structure of volatility, as well as relatively short averaging periods. We then compare the prices obtained to market prices of shipping options and find a very good fit to market data.

In Part 4 we use a crude oil forward curve model, data on shipping markets and stochastic dynamic programming to formulate the optimal routing and floating storage problem. Having formulated the optimal stopping problem for trading crude oil across the Atlantic we examine the empirical results of this trade during 2007-2009 and identify its key features: what conditions must be satisfied for it to be interesting, when it performs well and what the origins of the profits are.

2. MARKET MODELING

1. Rationale

The media and commodity market analysts tend to focus on the trends of prices based on expected supply and demand evolution. This is an important task, but commodity markets are volatile and an expected growth of two percent will be dwarfed by a price volatility of forty percent as is the case for crude oil. The market expectations of future supply and demand balances are reflected in the futures markets for the different commodities. Most commodity markets now have liquid forward curves with long maturities and these complete forward curves should be guiding long-dated investment and operational decisions, not only the spot price.

With this in mind, a model of commodity prices needs to provide a realistic model for the evolution of the complete forward curve and the volatility of the different contracts. Such a model can then be used in a variety of applications, such as pricing other derivatives or real assets with operational flexibility.

Such a model must also have a small number of parameters and correspond to reality when calibrated to market prices. With a realistic parametric model, analytical expressions for the prices of options and real assets can be obtained easily, as will be shown in Parts 3 and 4.

2. Existing literature

Early studies of commodity markets have focused on modeling the spot price, as it has been the only observable market price. Following work in equity markets the spot price has been modeled as geometric Brownian motion with constant growth rate, such as Brennan and Schwartz (1985) and Paddock et al (1988) for crude oil. Observing that price-based decisions on the supply or demand side will have a tendency to bring prices back to an equilibrium level, other authors such as Dixit and Pindyck (1994) have favored modeling the spot price as a process mean-reverting to a known and constant mean value. Ådland (2003) develops a mean-reverting spot price model for freight rates, arguing for the use of a spot price model because of the absence of liquidity in the forward market.

These one-factor models of the spot price give a good intuition about the behavior of prices, but fail to capture important effects, most notably transitions of the forward curves from contango to backwardation and the decreasing volatility of futures contracts with respect to maturity. Longstaff, Santa-Clara and Schwartz (1999) detail how failing to account for several factors leads to suboptimal exercise strategies in the swaptions market. In order to account for this Gibson and Schwartz (1990) introduce a mean-reverting stochastic convenience yield. In their model there are thus two factors shocking the forward curve: the spot price, affecting levels, and the convenience yield, affecting slope.

This two-factor model can be reinterpreted in terms of long-term and short-term shocks, such as in Baker, Mayfield and Parsons (1998) and Schwartz and Smith (2000). In this model the spot price is shocked by a mean-reverting short-term factor and a persistent long-term factor.

These models are all spot price models: they seek to explain the behavior of the spot price, which is traditionally the observable and most liquid price. They then price futures from this process by introducing a market price of risk and arbitrage-free pricing, and derive the process for the forward curve. The converse approach consists in taking the complete forward curve as the primary process. Miltersen and Schwartz (1998), Clewlow and Strickland (2000) and Sclavounos and Ellefsen (2009) develop such a model inspired by the multi-factor Heath, Jarrow and Morton (1992) model for the term structure of interest rates. It consists in decomposing the covariance matrix of the forward curve into a small number of orthogonal principal components. The spot price process is then derived as the front price of the forward curve.

It is this approach that we will adopt, but we will make parametric hypotheses about the principal component shapes and calibrate these to the covariance matrices.

3. Exploratory data analysis

In order to get an idea of the main features of the commodity forward markets we will begin by an analysis on the historical prices of different commodities.

Spot price

In Figure 1 we present the spot price of different commodities over recent time periods. In many markets, such as crude oil, this spot price is understood to be the price of the front-month futures contract with physical delivery. In other markets, such as shipping, the spot price is an index compiled daily using spot fixings from different brokers, on which the financial futures contracts settle.

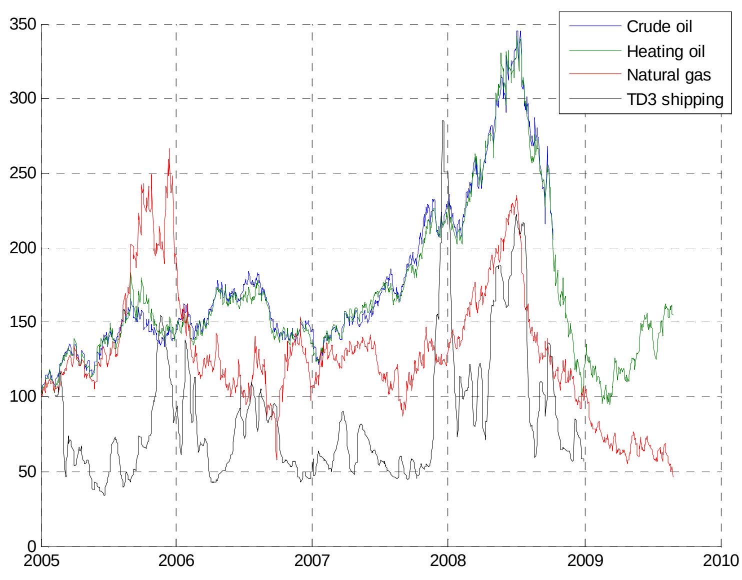

Figure 1. Daily spot prices of crude oil, heating oil, natural gas and the TD3 shipping route from January 2005 to 2009. Index = 100 on January 1st, 2005.

Forward curves

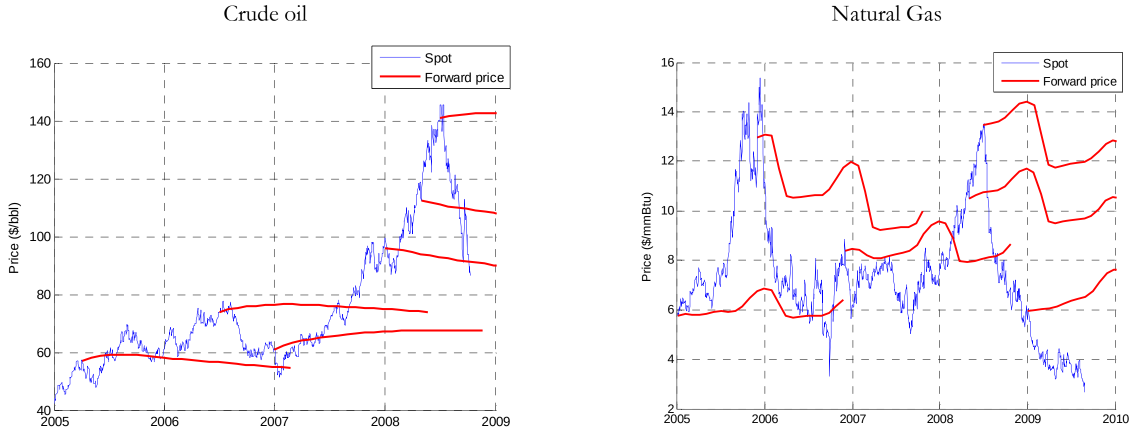

Figure 2. Spot and forward prices of crude oil and natural gas on different dates

Figure 2 presents the forward curves of crude oil and natural gas on different dates. We observe that the level of the forward curves shifts with the spot price and that the curves transition between contango and backwardation. Furthermore, the forward curves for natural gas have a seasonal pattern embedded in them.

Volatility term structure

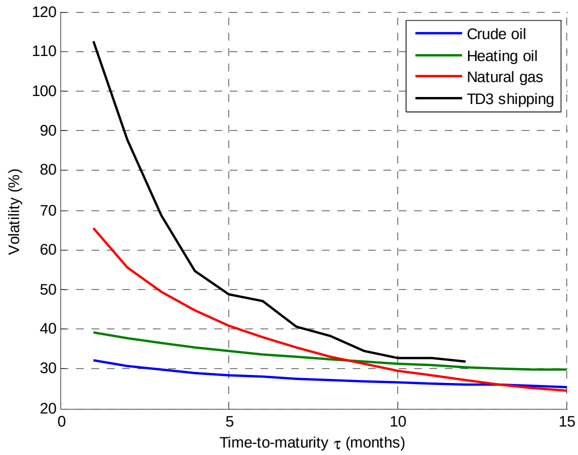

Figure 3. Volatility term structure of futures contracts on crude oil, heating oil, natural gas and TD3 shipping

Figure 3 presents the term structure of volatilities for crude oil, heating oil, natural gas and TD3 shipping. This is the historical volatility of the contracts with fixed time to maturity. We observe that for all these commodities, the volatility of near-term contracts is higher than the volatility of contracts further out on the curve, consistently with Samuelson's (1965) hypothesis. This is an important feature of commodity markets and happens because they are more inelastic in the short run than in the long run. If the tanker market is saturated it is impossible to add new ships within a month, but new ships can be built to accommodate the increasing demand in the next years.

Correlation structure

The forward prices for a given commodity do not move independently. Observing the correlation matrix of contracts with different maturities quantifies the relationship between these movements. Figure 4 presents the correlation matrix for crude oil contracts. The correlation matrix shows a strong correlation between different contracts, with an 84% correlation between the front-month and 60-month contract. However the correlation between the front-month contract and other contracts decays more rapidly than the correlation between the 60-month contract and neighboring contracts.

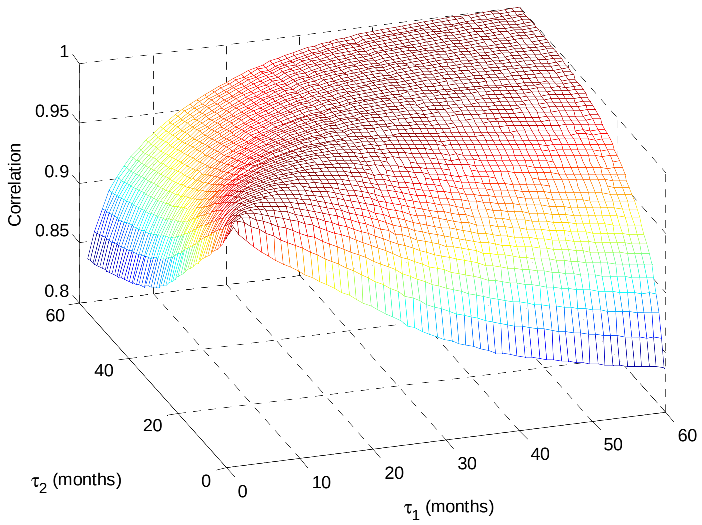

Figure 4. Correlation structure of crude oil futures

Principal components analysis

To get more insight into the structure of the co-movements of the forward prices we can perform a principal components analysis (PCA) of the price series. This consists in finding the eigenvalues and eigenvectors of the covariance matrix. The eigenvalues can be interpreted as the volatilities of each of the factors and the eigenvectors as the weights with which the principal components shock the forward curve.

We present results of a PCA of the crude oil market in Figure 5. As can be seen from these results, the dominant factor is the first factor, which accounts for 96.9% of the variance. This factor is the parallel shift factor, shifting forward prices in the same direction.

The second factor, explaining 2.8% of the variance, affects the slope of the forward curve by shocking the front end and long end of the forward curve with different signs. This accounts for transitions from contango to backwardation. The third factor affects the convexity of the forward curve by shocking the front and long ends positively and the middle of the curve negatively.

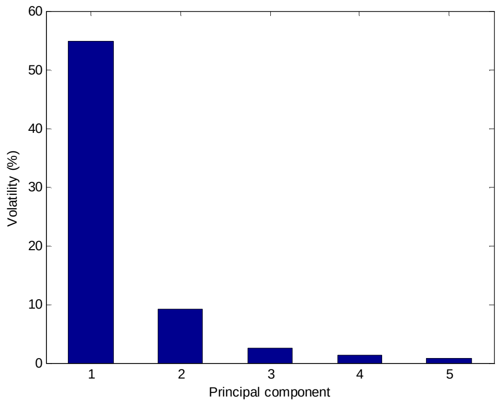

Volatility of the first five principal components

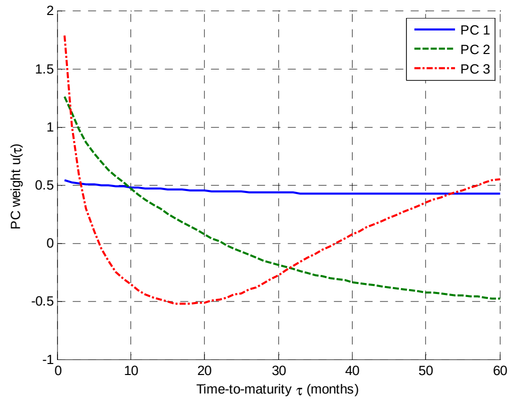

Principal component weights

Figure 5. Volatilities and weights for the first principal components of the crude oil market

4. Two-factor model of commodity futures

Consider a commodity forward market where we on each date t observe a forward curve $F(t,T)$ settling on the spot price $S(t)$ at date T. $S(t)$ could represent the spot price of some tradable commodity at time t (e.g. a specific grade of crude oil at a specific location), or the daily published value of an index. If a long forward position is entered at date t, it will receive the difference $S(T)-F(t,T)$ at date T.

Absence of arbitrage tells us (Musiela and Rutkowski, 2004) that under the risk-neutral measure,

$$\begin{aligned} &E_t^*[B(t,T)(S(T)-F(t,T))]=0 \\ &F(t,T)=E_t^*[S(T)] \end{aligned} \tag{2.1}$$

where $B(t,T)$ the time t price of the zero coupon bond that is matures at time T. The forward price of S at time t is the expectation of the spot price at time T, under the risk-neutral measure and given the information at time t. In some markets where the spot is storable, such as equities or currencies, there is a tight arbitrage enforcing the relationship between spot and forward prices. In markets where storage is limited, such as crude oil or shipping, the forward price is determined by supply and demand. It is not our goal here to impose a parametric model for the shape of the initial forward curve, which we take as given, but to give a model of its future stochastic evolution.

Following Baker, Mayfield and Parsons (1998) and Schwartz and Smith (2000), we suggest a two-factor model for the stochastic evolution of the forward curve. We present this model as a forward curve model rather than a spot price model, considering that the commodity derivative markets are generally more liquid than their physical counterparts, and contain more information to calibrate on than the spot price.

We suggest the following two-factor model for the forward curve under the risk-neutral measure:

$$\begin{aligned} \frac{dF(t,T)}{F(t,T)} &= \sigma_S e^{-\alpha(T-t)} dW_S(t) + \sigma_L dW_L(t) \\ dW_S dW_L &= \rho dt \end{aligned} \tag{2.2}$$

This is a four-parameter model and as we will show the parameters can be interpreted as follows:

- σ_S is the volatility of short-term shocks to the forward curve,
- σ_L is the volatility of long-term shocks,
- α is the mean-reversion speed, quantifying how fast short-term shocks dissipate,
- ρ is the correlation between short-term and long-term shocks.

Covariance and correlation

This model implies a covariance matrix between contracts that can be calculated as a function of the parameters

- Covariance matrix:

$$\begin{aligned} \Sigma_t(T_1,T_2) &= \frac{1}{dt}\text{Cov}\left(\frac{dF(t,T_1)}{F(t,T_1)}, \frac{dF(t,T_2)}{F(t,T_2)}\right) \\ &= (\sigma_S e^{-\alpha(T_1-t)} + \rho\sigma_L)(\sigma_S e^{-\alpha(T_2-t)} + \rho\sigma_L) + (1-\rho^2)\sigma_L^2 \\ &= \Sigma(\tau_1, \tau_2) \end{aligned} \tag{2.3}$$

  where $\tau_k = T_k - t$ is the time to maturity of the contract

- Futures instantaneous volatility function:

$$\sigma_{inst}(t,T) = \sqrt{(\sigma_S e^{-\alpha(T-t)} + \rho\sigma_L)^2 + (1-\rho^2){\sigma_L}^2} = \sigma_{inst}(\tau) \tag{2.4}$$

- Spot volatility:

$$\sigma_0 = \sqrt{(\sigma_S + \rho\sigma_L)^2 + (1-\rho^2)\sigma_L^2} \tag{2.5}$$

- Correlation matrix:

$$
\begin{aligned}
\rho_t(T_1,T_2) &= \text{Corr}\left(\frac{dF(t,T_1)}{F(t,T_1)},\frac{dF(t,T_2)}{F(t,T_2)}\right) = \frac{\Sigma(T_1,T_2)}{\sigma_{inst}(t,T_1)\sigma_{inst}(t,T_2)} \\
&= \frac{(\sigma_S e^{-\alpha(T_1-t)}+\rho\sigma_L)(\sigma_S e^{-\alpha(T_2-t)}+\rho\sigma_L)+(1-\rho^2)\sigma_L{}^2}{\left[(\sigma_S e^{-\alpha(T_1-t)}+\rho\sigma_L)^2+(1-\rho^2)\sigma_L^2\right]^{1/2}\left[(\sigma_S e^{-\alpha(T_2-t)}+\rho\sigma_L)^2+(1-\rho^2)\sigma_L^2\right]^{1/2}} \\
&= \rho(\tau_1,\tau_2)
\end{aligned}
\tag{2.6}
$$

All these quantities depend only on the time-to-maturities $\tau = T - t$ of the contracts involved, and not on time *t*.

Implied spot price process

In Appendix 2 we show that the spot price model consistent with this forward curve model is:

$$
\begin{aligned}
d\log S(t) &= \alpha(\mu(t)-\log S(t))dt+\sigma_S dW_1(t)+\sigma_L dW_2(t) \\
d\mu(t) &= m(t)dt+\sigma_L dW_2(t)
\end{aligned}
\tag{2.7}
$$

i.e. the spot price is mean-reverting to a stochastic mean. This is equivalent to the Schwartz and Smith (2000) model which can be rewritten as

$$
\begin{aligned}
d\log S_t &= \kappa\left(\frac{\mu_\xi}{\kappa}+\xi_t-\log S_t\right)dt+\sigma_\chi dz_\chi+\sigma_\xi dz_\xi \\
d\xi_t &= \mu_\xi dt+\sigma_\xi dz_\xi
\end{aligned}
\tag{2.8}
$$

From equation (2.7) we can see that α can be interpreted as the speed of mean-reversion and σ_L as the volatility of the long-term shocks.

5. Principal components analysis

As discussed in Sclavounos and Ellefsen (2009), futures markets can be analyzed and modeled in a non-parametric way through principal components analysis (PCA) of the covariance matrix, leading to a multi-factor Heath-Jarrow-Morton model of the form

$$
\frac{dF(t,T)}{F(t,T)} = \sum_{k=1}^{d}\sigma_k(t,T)dW_k, \qquad dW_k dW_l = \delta_{kl}dt \tag{2.9}
$$

Given the parametric model presented here, we can perform a PCA of the model's covariance matrix and deduce the shape of its principal components. This will allow us to reformulate the model in terms of independent factors that can be interpreted in terms of their actions on the forward curve.

In the continuous setting we perform the Karhunen-Loève decomposition of the process following Basilevsky (1994). Let $f(t,\tau) = F(t, t+\tau)$ be the constant-maturity forward with time-to-maturity τ. We want to decompose its evolution into:

$$\frac{df(t,\tau)}{f(t,\tau)} = \mu(t,\tau)dt + \sum_{k=1}^{\infty} \sqrt{\lambda_k} u_k(\tau) dz_k \tag{2.10}$$

Where:

- The z_k are independent Brownian motions
- The functions u_k are the eigenvectors of the covariance matrix $\Sigma(\tau_1, \tau_2)$ with associated eigenvalues λ_k : for some arbitrary maximal tenor $\tau_{\max}$,

$$\begin{aligned} &\int_0^{\tau_{\max}} \Sigma(\tau_1,\tau_2) u_k(\tau_2) d\tau_2 = \lambda_k u_k(\tau_1) \\ &\int_0^{\tau_{\max}} u_k^2(\tau) d\tau = 1 \end{aligned} \tag{2.11}$$

We solve this eigenvector problem analytically in Appendix 3, and show that there are only two distinct functions u_k (because it is a two-factor model), and they can be written in the form

$$u_k(\tau) = A_k e^{-\alpha\tau} + B_k \tag{2.12}$$

where (A_k, B_k) and λ_k are solutions of the two-dimensional eigenvalue problem

$$\lambda_k \begin{bmatrix} A_k \\ B_k \end{bmatrix} = \begin{bmatrix} \int_0^{\tau_{\max}} (\sigma_S^2 e^{-\alpha\tau_2} + \rho\sigma_S\sigma_L) e^{-\alpha\tau_2} d\tau_2 & \int_0^{\tau_{\max}} (\sigma_S^2 e^{-\alpha\tau_2} + \rho\sigma_S\sigma_L) d\tau_2 \\ \int_0^{\tau_{\max}} (\rho\sigma_S\sigma_L e^{-\alpha\tau_2} + \sigma_L^2) e^{-\alpha\tau_2} d\tau_2 & \int_0^{\tau_{\max}} (\rho\sigma_S\sigma_L e^{-\alpha\tau_2} + \sigma_L^2) d\tau_2 \end{bmatrix} \begin{bmatrix} A_k \\ B_k \end{bmatrix} \tag{2.13}$$

The volatility of factor k is then related to λ_k by $\sigma_k = \sqrt{\lambda_k}$. The shape of the eigenfunctions is given in Figure 6 in the case of crude oil futures. We can notice that u_1 corresponds to parallel shifts of the forward curve, whereas u_2 corresponds to tilts. This is consistent with the two first factors observed doing a PCA of the historical covariance matrix (Sclavounos and Ellefsen, 2009).

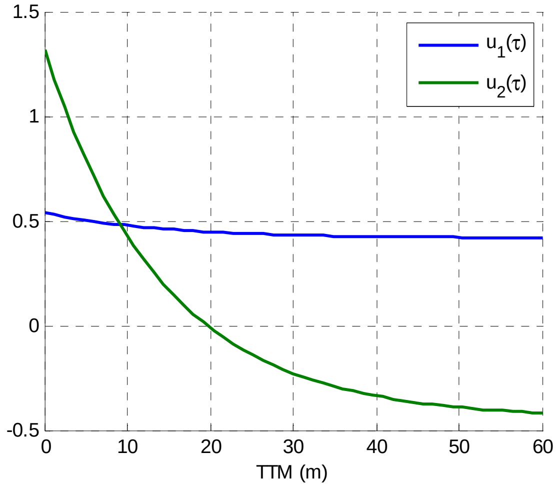

Figure 6. Shape of the eigenfunctions $u_1(\tau)$ and $u_2(\tau)$ for $\sigma_S = 18.1\%$, $\sigma_L = 23.3\%$, $\alpha = 0.842\,\text{yr}^{-1}$, $\rho = 0.195$ and $\tau_{\max} = 5\,\text{years}$

This allows us to reformulate the evolution of the individual forward contract expiring at date *T*, in the risk-neutral measure:

$$\frac{dF(t,T)}{F(t,T)} = \sigma_1 u_1(T-t)dz_1(t) + \sigma_2 u_2(T-t)dz_2(t) \tag{2.14}$$

Constant-maturity forward curve

In Appendix 4 we show how this translates to the evolution of constant-maturity forward curve. We show that the constant-maturity futures price $f(t,\tau) = F(t,t+\tau)$ can be written as:

$$\log f(t,\tau) = \log F(0,t+\tau) + \psi_1(t,\tau) + \psi_2(t,\tau) + g_1(t) + g_2(t) + u_1(\tau) f_1(t) + u_2(\tau) f_2(t) \tag{2.15}$$

where:

$$\begin{aligned} df_k(t) &= -\alpha_k f_k(t)dt + \sigma_k dW_k(t) \\ dg_k(t) &= B_k \alpha_k f_k(t)dt \\ d\psi_k(t,\tau) &= -\frac{1}{2}\sigma_k^2 u_k^2(t+\tau)dt \end{aligned} \tag{2.16}$$

Thereby we have decomposed the forward curve's shape at time t into

- Its initial shape $F(0,t+\tau)$, which under the risk-neutral measure is also its expected shape
- A deterministic risk-neutral drift $\psi_1(t,\tau) + \psi_2(t,\tau)$ ensuring that $E_t^*[F(t,T)] = F(0,T)$

- A stochastic drift $g_1(t)+g_2(t)$, independent of the maturity τ
- Two independent mean-reverting factors $f_1(t)$ and $f_2(t)$ (volatilities σ_k and mean-reversion speeds α_k), giving rise to a parallel shift and a tilt, according to the shape of the factor weights $u_k(\tau)$

The spot price process $S(t)$ is given by the zero time-to-maturity price $f(t,0)$.

This allows us to express the evolution of the forward curve as the result of shocks from two independent mean-reverting factor values. f_1, the parallel shift factor, affects the average level of the forward curve and is the dominant factor. f_2, the tilt factor, affects the slope of the forward curve, as seen in Figure 7.

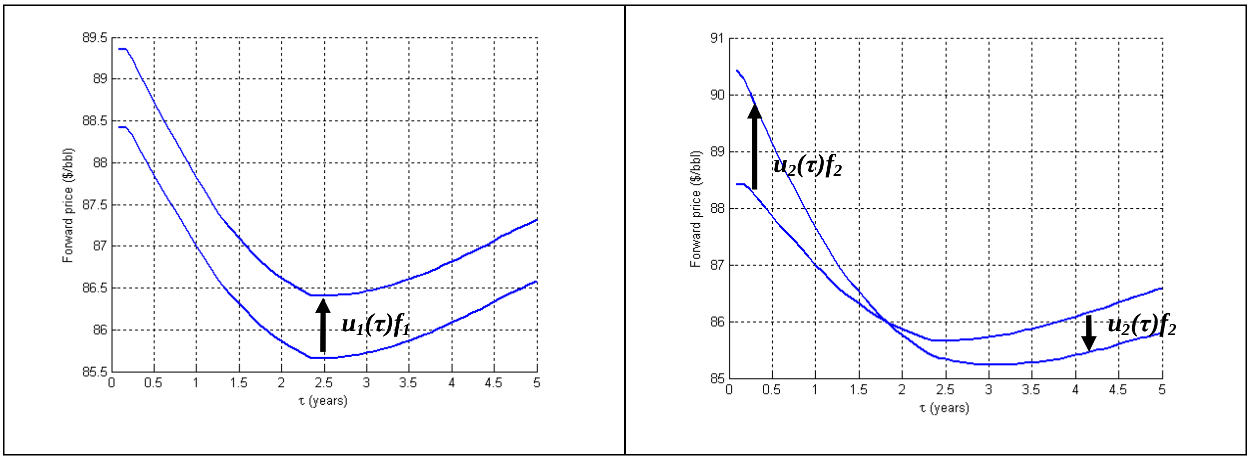

Figure 7. Effect of a positive parallel shift (left) and tilt (right) on the forward curve

6. Forward curve seasonality

A number of commodities have seasonal prices. It appears because demand or supply is seasonal, and inventories are not sufficient to smooth this seasonality out over the year. Examples of seasonal commodities are heating oil and natural gas (winter heating demand), gasoline (summer driving season and different volatility requirements during summer and winter) and agricultural products (seasonal supply). This seasonality in spot prices is reflected in the forward prices because of market expectations.

The difficulty when analyzing such forward prices is that the seasonality masks the underlying shifts in level and tilt that we are interested in. When considering such a seasonal commodity, the forward curve can be decomposed into a trend component and a seasonal component:

$$\log F(t,T) = \log F_T(t,T) + \log F_S(t,T) \tag{2.17}$$

The trend component $F_T(t,T)$ represents the underlying non-seasonal forward curve, whereas the seasonal component $F_S(t,T)$, which for a given t is 1-year-periodic in T, represents the seasonal aspects of the curve.

Pilipovic (2007) suggests a functional form that we have successfully applied to the natural gas and heating oil markets. In Section 2.10 we will show that this form is also consistent with the static shape of our three-factor model. For the trend component,

$$\log F_T(t,T) = (A_1 e^{-\alpha_1(T-t)} + B_1) f_1 + (A_2 e^{-\alpha_2(T-t)} + B_2) f_2 + (A_3 e^{-2\alpha_3(T-t)} + B_3 e^{-\alpha_3(T-t)} + C_3) f_3 \qquad (2.18)$$

This functional form is flexible enough to reproduce the shapes of the underlying forward curve. For the seasonal component, we use sinusoidal seasonality with two harmonics (time must be measured in years)

$$\log F_S(t,T) = a_1 \cos(2\pi(T-t)) + b_1 \sin(2\pi(T-t)) + a_2 \cos(4\pi(T-t)) + b_2 \sin(4\pi(T-t)) \qquad (2.19)$$

The only test of this model is how good the fit to the forward curve is. We find the parameters by least-squares minimization for each day in the data set.

A selection of forward curves is presented in Figure 8. While the fit is not perfect, the trend component seems to correctly capture the underlying trend, and that is what we are ultimately interested in. We then use this trend as the new forward curve, and carry out the rest of the calibration procedure on it.

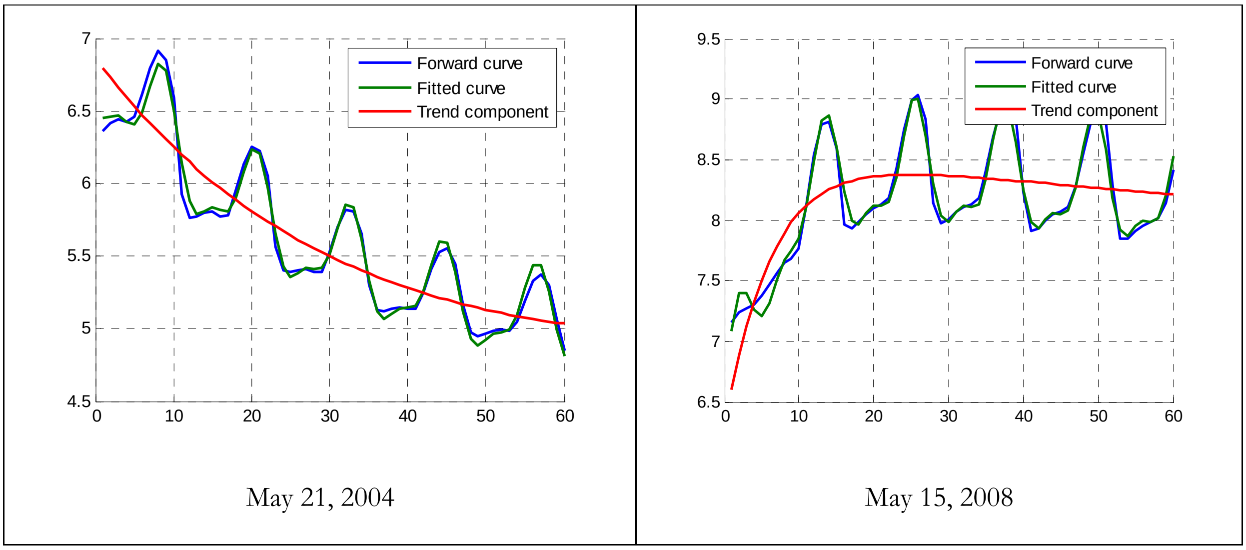

Figure 8. Fitted Natural Gas forward curves on different dates

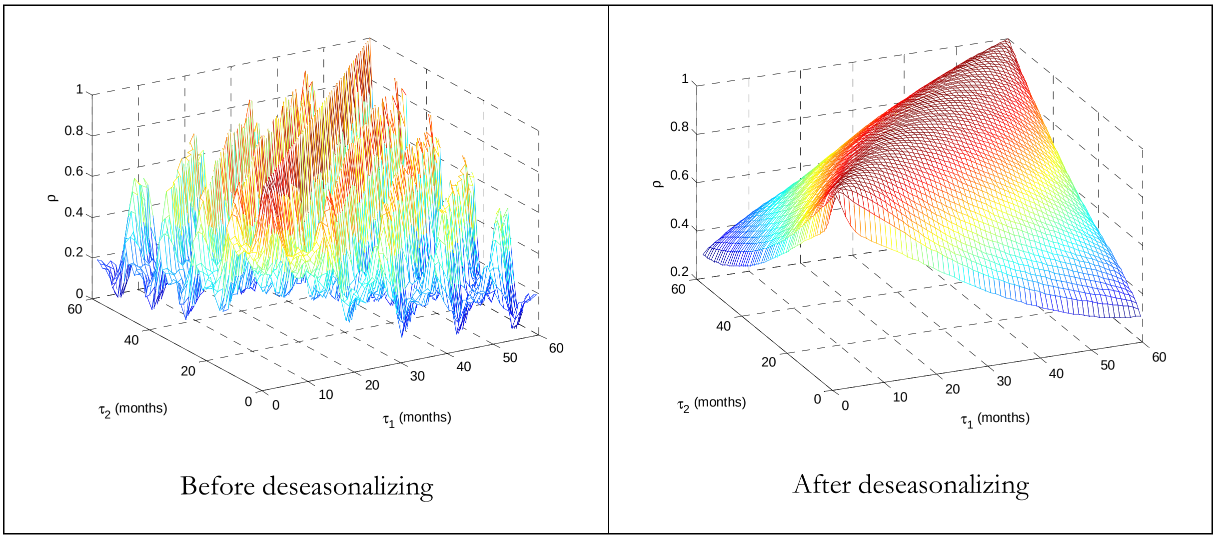

Figure 9. Correlation surfaces of Natural Gas futures before and after deseasonalizing the curves

Figure 9 shows the effect of the procedure on the correlation surface. The results of calibrating the two-factor model to this new correlation surface are discussed below.

7. Market calibration

To be successful the model needs to correctly reproduce the volatilities and instantaneous correlations of the traded instruments. We show that there is a good fit to the crude oil, tanker shipping, natural gas and heating oil markets.

Method 1: Least squares fit of the covariance matrix

In order to calibrate the model, we perform the following steps:

1. If the forward curve is seasonal (such as natural gas, gasoline, heating oil), deseasonalize it using the technique described above (Section 2.6), and keep only the non-seasonal part $F_T(t,T_j)$
2. From the available set of contract prices $F(t,T_j)$, construct constant-maturity prices $f(t,\tau_j)$ by linear interpolation using

$$\log f(t,\tau_j) = \frac{(t+\tau_j - T_j)\log F(t,T_{j+1}) + (T_{j+1} - t - \tau_j)\log F(t,T_j)}{T_{j+1} - T_j}, \quad T_j < t + \tau_j < T_{j+1} \qquad (2.20)$$

3. From observations of $f(t,\tau_j)$ at dates $t_1,\ldots,t_{M+1}$, construct logarithmic returns net of roll yield and their mean value

$$R(t_i,\tau_j)=\log\left(\frac{f(t_i,\tau_j)}{f(t_{i-1},\tau_j)}\right)-(t_i-t_{i-1})\frac{\partial\log f(t_{i-1},\tau_j)}{\partial\tau},\qquad \bar{R}(\tau_j)=\frac{1}{M}\sum_{i=1}^{M}R(t_i,\tau_j) \tag{2.21}$$

4. Calculate the historical covariance matrix

$$\tilde{\Sigma}(\tau_j,\tau_k)=\frac{1}{M}\sum_{i=1}^{M}(R(t_i,\tau_j)-\bar{R}(\tau_j))(R(t_i,\tau_k)-\bar{R}(\tau_k)) \tag{2.22}$$

5. Find the parameters $\sigma_S,\sigma_L,\alpha,\rho$ that minimize the squared error:

$$\min_{\sigma_S,\sigma_L,\alpha,\rho}\sum_{j,k=1}^{N}\left[\Sigma_{\sigma_S,\sigma_L,\alpha,\rho}(\tau_j,\tau_k)-\tilde{\Sigma}(\tau_j,\tau_k)\right]^2 \tag{2.23}$$

The results are presented in Table 1 for the crude oil, shipping, natural gas and heating oil markets. The results indicate that a satisfactory fit to the volatility term structure and correlation surface can be obtained using the two-factor model presented here. The best calibration results are obtained for crude oil futures, which is arguably the most liquid market of the four. It is also interesting to note the differences between the values obtained. The short-term volatility of shipping futures is extremely high, at 143%, reflecting the high spot price volatility, but its long-term volatility is comparable to the other markets, at 28.7%.

Table 1. Calibration results for different commodity markets

<table>
<tr><th></th><th>Crude oil (Nymex WTI)</th><th>Tanker shipping (Imarex TD3)</th></tr>
<tr><td>Data</td><td>Period: April 2005 – October 2008
Contracts: NYMEX WTI futures
Frequency: daily
Source: Thomson Datastream</td><td>Period: January 2005 – March 2009
Contracts: Imarex TD3 futures
Frequency: weekly
Source: Imarex</td></tr>
<tr><td>Parameters</td><td>

σ_S	σ_L	α	ϱ
18.1%	23.3%	0.842	0.195

</td><td>

σ_S	σ_L	α	ϱ
143%	28.7%	3.32	-0.01

</td></tr>
<tr><td>Volatility</td><td>

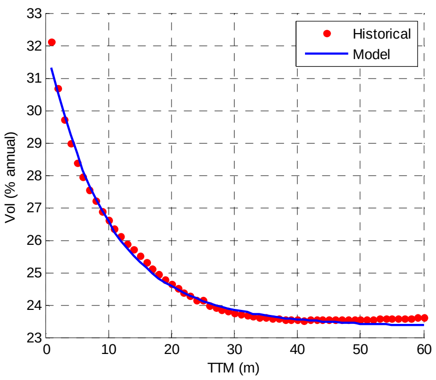

Volatility term structure of crude oil futures

</td><td>

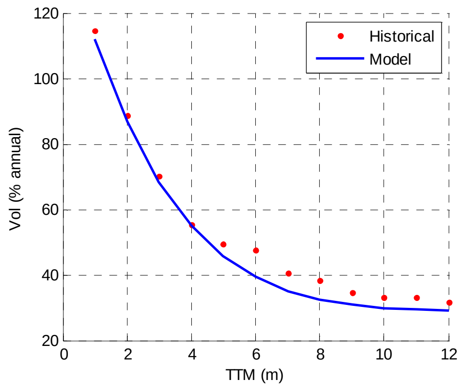

Volatility term structure of TD3 futures

</td></tr>
<tr><td>Correlation</td><td>

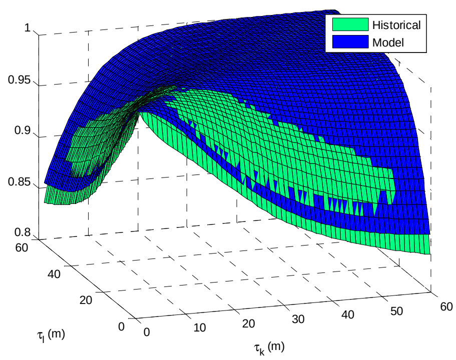

Correlation surface of crude oil futures

</td><td>

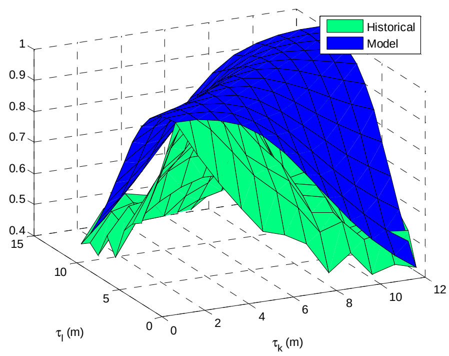

Correlation surface of TD3 futures

</td></tr>
</table>

<table>
<tr><td>Principal Components</td><td>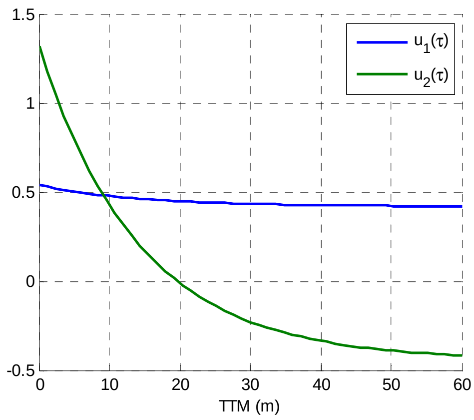

Model principal components for crude oil</td><td>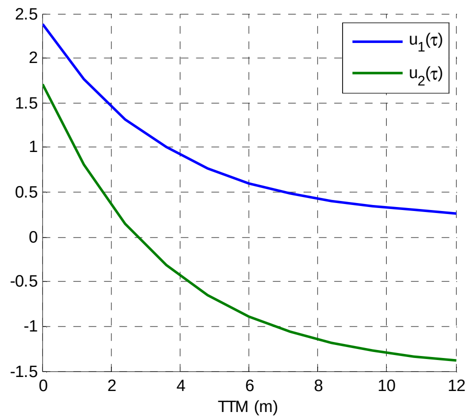

Model principal components for TD3 futures</td></tr>
</table>

<table>
<tr><td></td><td>Natural Gas (Nymex Henry Hub)</td><td>Heating oil (Nymex New York Harbor)</td></tr>
<tr><td>Data</td><td>Period: October 2002 – August 2009
Contracts: NYMEX NG futures
Frequency: daily
Source: Reuters</td><td>Period: May 2002 – August 2009
Contracts: NYMEX HO futures
Frequency: daily
Source: Reuters</td></tr>
<tr><td>Parameters</td><td><table><tr><td>σ_S</td><td>σ_L</td><td>α</td><td>ϱ</td></tr><tr><td>53%</td><td>17.3%</td><td>0.762</td><td>-0.172</td></tr></table></td><td><table><tr><td>σ_S</td><td>σ_L</td><td>α</td><td>ϱ</td></tr><tr><td>27.6%</td><td>26.4%</td><td>1.386</td><td>0.228</td></tr></table></td></tr>
<tr><td>Volatility</td><td>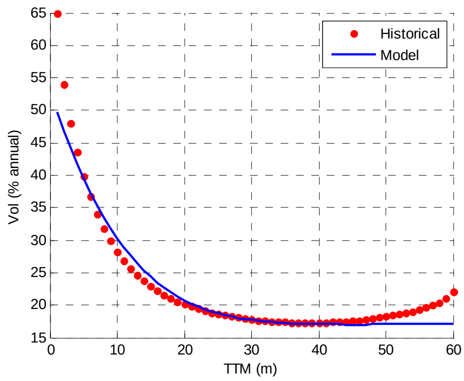

Volatility term structure of natural gas futures</td><td>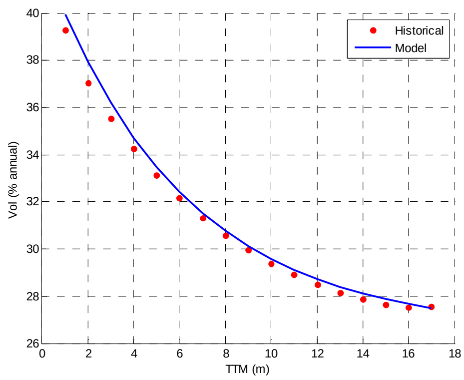

Volatility term structure of heating oil futures</td></tr>
</table>

Correlation

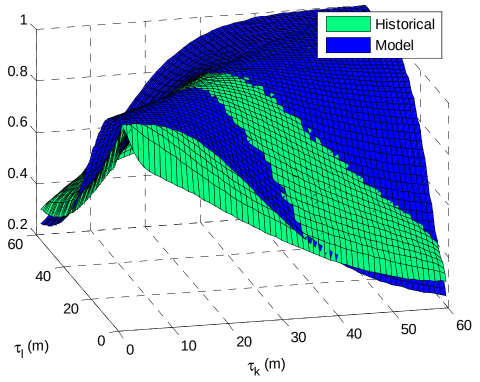

Correlation surface of natural gas futures

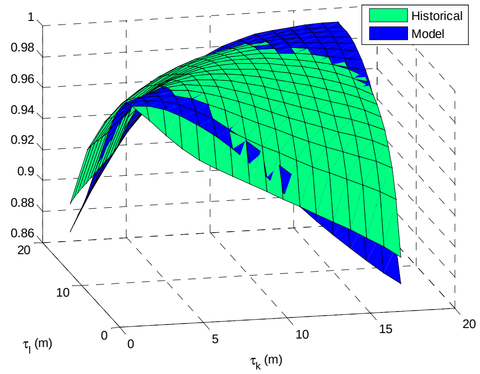

Correlation surface of heating oil futures

Principal Components

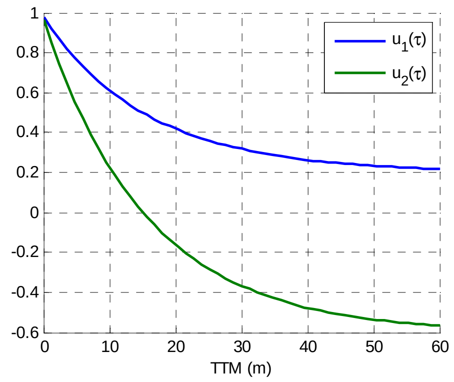

Model principal components for natural gas

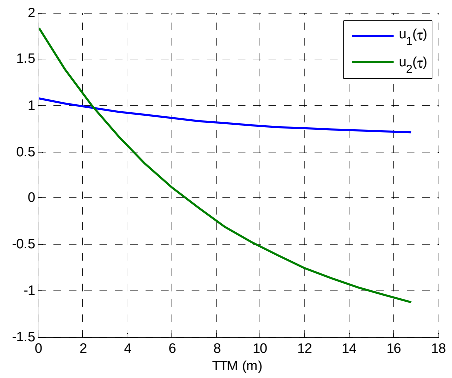

Model principal components for heating oil

Method 2: Calibration of the individual factors

The first two principal components have a simple expression in this model, and can be used for calibration. The method is the same as above, but we replace steps 4 and 5 with:

4'. Calculate the PCA of the historical covariance matrix and extract the first two factor loadings $\tilde{u}_1(\tau_j)$, $\tilde{u}_2(\tau_j)$

5'. Calibrate the exponential functional form on each of the factors by least squares:

$$\min_{A_k, B_k, \alpha_k} \sum_{j=1}^{N} \left[\tilde{u}_k(\tau_j) - (A_k e^{-\alpha_k \tau_j} + B_k) \right]^2 \tag{2.24}$$

We present the results of this method for crude oil futures in Table 2 and Figure 10.

Table 2. Principal component parameters for crude oil futures, using two calibration methods

	Principal Component 1		**Principal Component 2**	
	Method 1	**Method 2**	**Method 1**	**Method 2**
σ	54.91 %	54.91 %	9.53 %	9.50 %
A	0.1218	0.1205	1.7639	1.7385
B	0.4177	0.4189	-0.4435	-0.5148
α	0.8422	0.8713	0.8422	0.6707

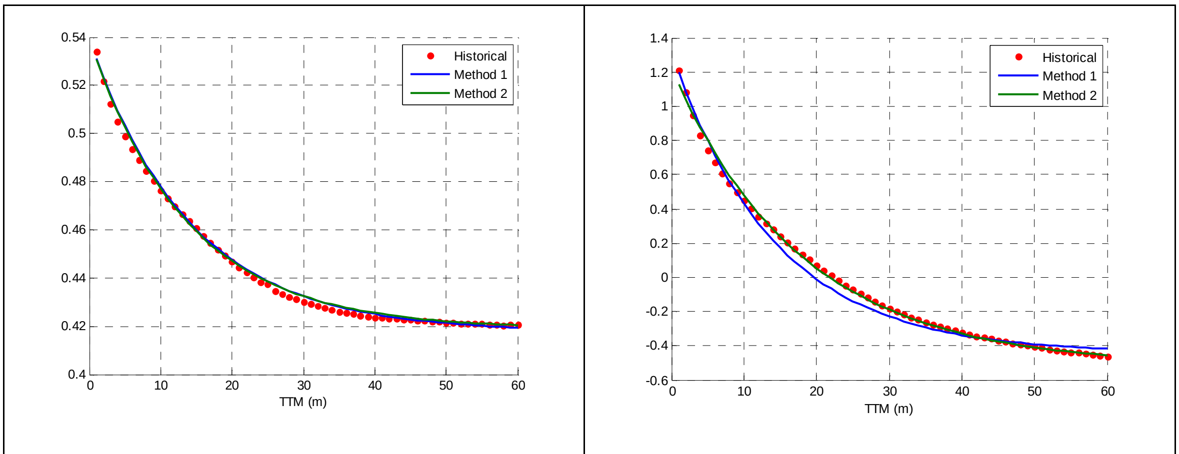

Principal Component 1 | **Principal Component 2**

Figure 10. Fit of the shape of the two principal components using the two different calibration methods

We see that the two methods give very close results, except that the second method allows for a different value of α which gives a slightly better fit to the second principal component. It should be noted that Method 2 adds one extra free parameter by allowing α_1 and α_2 to be different.

8. Forward risk premia – from the risk neutral to the objective measure

The present model has been formulated under the risk-neutral measure. The prices evolve under the real measure. The change of measure from the risk-neutral to the real measure involves introducing a risk-premium λ_k for each of the Brownian motions W_k. We assume this risk premium to be constant.

$$dW_k \rightarrow dW_k + \lambda_k dt \tag{2.25}$$

This will affect the factor processes $f_k(t)$ and $g_k(t)$ studied in Section 2.5:

$$\begin{aligned} f_k(t) &= \sigma_k \int_0^t e^{-\alpha_k(t-s)}(dW_k(s) + \lambda_k ds) \\ df_k(t) &= \sigma_k dW_k(t) + \sigma_k \lambda_k dt + \sigma_k \int_0^t (-\alpha_k e^{-\alpha_k(t-s)})(dW_k(s) + \lambda_k ds) \\ df_k(t) &= (\sigma_k \lambda_k - \alpha_k f_k)dt + \sigma_k dW_k(t) = (\mu_k - \alpha_k f_k)dt + \sigma_k dW_k(t) \end{aligned} \tag{2.26}$$

And for the drift process $g_k(t)$:

$$\begin{aligned} g_k(t) &= B_k \sigma_k \int_0^t (1 - e^{-\alpha_k(t-s)})(dW_k(s) + \lambda_k ds) \\ dg_k(t) &= B_k \sigma_k \left[dW_k(t) + \lambda_k dt - \left(dW_k(t) + \lambda_k dt + \int_0^t (-\alpha_k) e^{-\alpha_k(t-s)}(dW_k(s) + \lambda_k ds) dt \right) \right] \\ dg_k(t) &= B_k \alpha_k f_k(t) dt \end{aligned} \tag{2.27}$$

Hence the factor process $f_k(t)$ follows an Ornstein-Uhlenbeck process mean-reverting to $\mu_k = \sigma_k \lambda_k / \alpha_k$ instead of 0. The definition of $g(t)$ does not change. We let $\mu_k = \sigma_k \lambda_k$ be the drift term for the factor *k*.

The stochastic evolution of the forward price with tenor *T* can then be written as

$$\frac{dF(t,T)}{F(t,T)} = (\mu_1 u_1(T-t) + \mu_2 u_2(T-t))dt + \sigma_1 u_1(T-t) dW_1 + \sigma_2 u_2(T-t) dW_2 \tag{2.28}$$

These results show how to incorporate drifts of the forward curve into the model. These can be based on historical evidence of drifts in prices or subjective evaluations of the expected future prices. This allows valuation models of physical assets to take into account forecasts of future price evolution. Financial derivatives, however, will be valued under the risk-neutral measure.

9. Extension to three factors

The model we have considered is sufficient to reproduce the volatility and correlation term structures of most forward markets. However, it only allows for certain movements of the forward curve, i.e. parallel shifts and tilts. As shown previously in the Principal Components Analysis, the forward curve does have other movements, and the third principal component is generally understood to correspond to changes in curvature. Certain strategies, such as a butterfly trade, are especially sensible to this kind of change.

We suggest modeling the third principal component as

$$u(\tau) = Ae^{-2\alpha\tau} + Be^{-\alpha\tau} + C \qquad (2.29)$$

As is shown in Figure 11 it gives a good fit to the third principal component calculated from a historical covariance matrix. With parameters A and C positive and B negative the function $u(\tau)$ will take positive values for small times-to-maturity, negative values for intermediate τ, and then positive values again, thereby affecting the convexity of the curve.

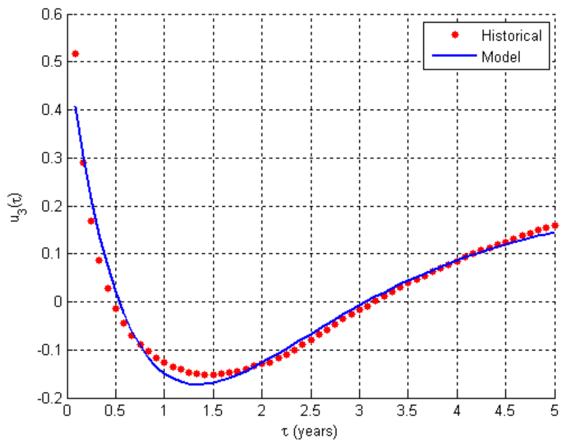

Figure 11. Fit of the parametric third PC to the third PC from the covariance matrix (crude oil)

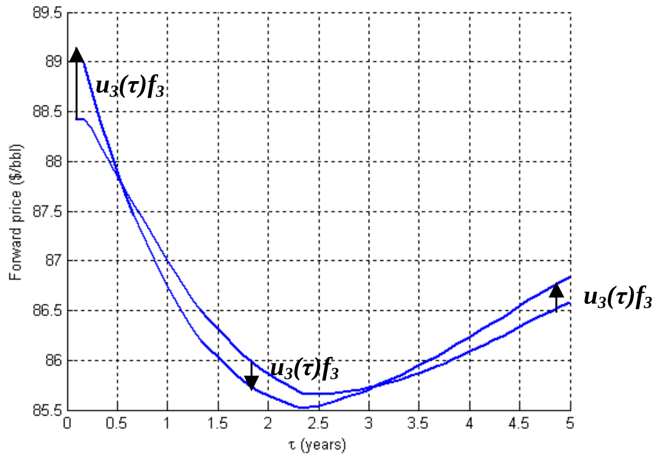

Figure 12. Effect on the forward curve of a positive shock from the third principal component

In order to study its interpretation we will consider its effect on the constant-maturity forward curve, as we did in Section 2.4 for the first two components. In Appendix 5 we show that the constant-maturity forward curve can be written as

$$\begin{aligned}\log f(t,\tau) = \log F(0,t+\tau) + \sum_{k=1}^{2}(\psi_k(t,\tau) + g_k(t) + u_k(\tau) f_k(t)) \\ + \psi_3(t,\tau) + B_3 e^{-\alpha_3 \tau} g_3(t) + C_3 h_3(t) + u_3(\tau) f_3(t)\end{aligned} \quad (2.30)$$

where

$$\begin{aligned} & df_3(t) = -2\alpha_3 f_3(t) + \sigma_3 dW_3(t) && \text{(Ornstein-Uhlenbeck process)} \\ & dg_3(t) = \alpha_3(f_3(t) - g_3(t))dt && g_3(t) = \alpha_3 \int_0^t e^{-\alpha_3(t-s)} f_3(s)ds \\ & dh_3(t) = 2\alpha_3 f(t)dt && h_3(t) = 2\alpha_3 \int_0^t f_3(s)ds \end{aligned} \quad (2.31)$$

The process $f_3(t)$ is an Ornstein-Uhlenbeck process mean-reverting to zero with mean-reversion speed $2\alpha_3$ and volatility σ_3. The processes $g_3(t)$ and $h_3(t)$ are stochastic drifts – integrals of $f_3(t)$ with different weights.

10. Model of the static forward curve

While the starting point of our modeling is that the initial forward curve $F(0,T)$ is given, there are situations where one would want to model this curve with a small number of parameters. Using the factor model presented in this part, we can express the possible shapes of the curve when starting from an initial forward curve:

$$\log f(t,\tau) = \log F(0,t+\tau) + \sum_{k=1}^{N} \psi_k(t,\tau) + g_k(t) + u_k(\tau) f_k(t) \tag{2.32}$$

If we assume the initial forward curve to be flat, $F(0,\tau) = F$, this formulation simplifies to:

$$\log f(t,\tau) = A(t) + \sum_{k=1}^{N} \psi_k(t,\tau) + u_k(\tau) f_k(t) \tag{2.33}$$

Where $A(t)$ is a time-dependent scalar not depending on time-to-maturity τ and

$$\psi_k(t,\tau) = -\frac{1}{2}\int_0^t \sigma_k^2(s,t+\tau)ds \tag{2.34}$$

This gives the possible shapes that can be taken by the forward curve given an initially flat curve. We can further simplify this by remarking that the first factor, the parallel shift factor, has a function $u_1(\tau)$ that is almost constant, such that the constant term can be merged into the first factor value. Thereby the forward curve can be written as

$$f(t,\tau) = \exp\left(\sum_{k=1}^{N} \psi_k(t,\tau) + u_k(\tau) f_k(t)\right) \tag{2.35}$$

Hence it can be described by $N+1$ state variables: $t, f_1, ..., f_N$. Their initial values can be calibrated on the initial forward curve by calculating

$$f_k(0) = \int_0^{\tau_{\max}} u_k(\tau) \log\left(F(0,\tau)\right) d\tau \tag{2.36}$$

This formulation also allows us to relate the average level of the forward curve and the first factor value by forming the geometric average weighted by $u_1(\tau)$:

$$\bar{F}(t) = \exp\left(\int_0^{\tau_{\max}} w(\tau) \log f(t,\tau) d\tau\right), \quad w(\tau) = \frac{u_1(\tau)}{\int_0^{\tau_{\max}} u_1(\tau) d\tau} \tag{2.37}$$

such that

$$\bar{F}(t) = \exp\left[\frac{f_1(\tau) + \sum_{k=1}^{N} \int_0^{\tau_{\max}} u_1(\tau)\psi_k(t,\tau)d\tau}{\int_0^{\tau_{\max}} u_1(\tau)d\tau}\right] \tag{2.38}$$

We can also examine the initial slope of the curve, that we will use to determine if the curve is in backwardation or contango:

$$\left.\frac{\partial f}{\partial \tau}\right|_{\tau=0} = f(t,0)\left.\frac{\partial \log f}{\partial \tau}\right|_{\tau=0} \tag{2.39}$$

and

$$\frac{\partial \log f}{\partial \tau} = \sum_{k=1}^{N} \frac{\partial \psi_k}{\partial \tau} + \frac{\partial u_k}{\partial \tau} f_k(t) \tag{2.40}$$

If we are considering a two-factor model, the first factor is almost flat such that its derivative is zero. In that case the only contribution comes from the second factor:

$$\frac{\partial \log f}{\partial \tau} = \frac{\partial \psi_2}{\partial \tau} + \frac{\partial u_2}{\partial \tau} f_2(t) \tag{2.41}$$

such that the initial slope of the forward curve is

$$\left.\frac{\partial f}{\partial \tau}\right|_{\tau=0} = f(t,0)\left[\left.\frac{\partial \psi_2}{\partial \tau}\right|_{\tau=0} + \left.\frac{\partial u_2}{\partial \tau}\right|_{\tau=0} f_2(t)\right] \tag{2.42}$$

Thus the value of $f_2(t)$ determines the slope of the forward curve.

11. Applications of the market model

Derivatives pricing

The main application of stochastic models of forward curves is in derivatives pricing. The stochastic model that we have derived and calibrated allows for simple pricing of paper derivatives depending on the volatility of prices, such as European or Asian options written on the forward or spot price. In Part 3 we will derive analytical prices of commodity Asian options using the two-factor model derived here.

Real asset valuation and operation

There are a number of physical assets whose value depends on commodity prices and forward curves. Oil or gas reservoirs are a simple example, but more complex assets such as refineries, power plants or oil in transit depend on these prices in a more complex way. Their value depends not only on the spot price but on the complete forward curve, and operational decisions should be made taking into account the possible future evolutions of the complete curve.

The value of such an asset can be written as $V(t, f(\tau))$ where $f(\tau)$ is the current forward curve. If a two-factor model such as the one in this thesis is adopted, $f(\tau)$ is a function of the initial forward curve $F_0(\tau)$, time *t* and the factor values f_1 and f_2, such that the value can be written

$$V(t, f(\tau)) = V(t, f_1, f_2) \tag{2.43}$$

The stochastic evolution of this value function can then be derived, using Ito's formula and the independence of the factors, as

$$\begin{aligned} dV &= \frac{\partial V}{\partial t}dt + \frac{\partial V}{\partial f_1}df_1 + \frac{\partial V}{\partial f_2}df_2 + \frac{1}{2}\frac{\partial^2 V}{\partial f_1^2}df_1^2 + \frac{1}{2}\frac{\partial^2 V}{\partial f_2^2}df_2^2 \\ &= \left[\frac{\partial V}{\partial t} + (\mu_1 - \alpha_1 f_1)\frac{\partial V}{\partial f_1} + (\mu_2 - \alpha_2 f_2)\frac{\partial V}{\partial f_2} + \frac{1}{2}\sigma_1^2\frac{\partial^2 V}{\partial f_1^2} + \frac{1}{2}\sigma_2^2\frac{\partial^2 V}{\partial f_2^2}\right]dt + \frac{\partial V}{\partial f_1}\sigma_1 dW_1 + \frac{\partial V}{\partial f_2}\sigma_2 dW_2 \end{aligned} \tag{2.44}$$

Associated with appropriate boundary conditions this allows for the calculation of the value of the real asset and the hedging of its value using the factors. In Part 4 we present the results of this methodology for a physical crude oil trade involving the shipment and possibly storage of crude oil.

Risk evaluation

Once a portfolio of paper and real assets has been valued, the risk of the portfolio can be evaluated using the market model presented here. We assume that given a forward curve $f(\tau)$ and a date *t* the portfolio has a value $V(t, f(\tau))$. If we assume a two-factor model this value can be re-written as $V(t, f_1, f_2)$ and its stochastic evolution as

$$dV = \mu(t, f_1, f_2)dt + \frac{\partial V}{\partial f_1}\sigma_1 dW_1 + \frac{\partial V}{\partial f_2}\sigma_2 dW_2 \tag{2.45}$$

Thereby

$$V(t) = \int_0^t \mu(s, f_1(s), f_2(s))ds + \int_0^t \frac{\partial V}{\partial f_1}\sigma_1 dW_1(s) + \int_0^t \frac{\partial V}{\partial f_2}\sigma_2 dW_2(s) \tag{2.46}$$

Hence the expected value of *V* at a horizon *t* is

$$E[V(t)] = \int_0^t \mu(s, f_1(s), f_2(s))ds \tag{2.47}$$

and its standard deviation

$$\text{Std}[V(t)] = \left(E\left[\int_0^t \left(\frac{\partial V}{\partial f_1} \right)^2 \sigma_1^2 ds + \int_0^t \left(\frac{\partial V}{\partial f_2} \right)^2 \sigma_2^2 ds \right] \right)^{1/2} \tag{2.48}$$

These values can be calculated if the value of V as a function of the factors and time is known explicitly. Alternatively Monte Carlo simulation can be used, using the two independent processes f_1 and f_2, to estimate the complete distribution of V at the horizon time t. This Monte Carlo simulation will only require the simulation of two independent stochastic variables and not of each forward price separately.

This information about the distribution of the portfolio value can be used to evaluate the risk of the position and calculate risk measures such as value-at-risk.

Hedging

As seen above a portfolio that depends on forward prices has, according to the two-factor model, a stochastic evolution that can be written

$$\begin{aligned} dV &= \mu(t, f_1, f_2)dt + \frac{\partial V}{\partial f_1}\sigma_1 dW_1 + \frac{\partial V}{\partial f_2}\sigma_2 dW_2 \\ &= \mu(t, f_1, f_2)dt + \delta_1\sigma_1 dW_1 + \delta_2\sigma_2 dW_2 \end{aligned} \tag{2.49}$$

In order to hedge the risk related to factor k the portfolio must be complemented with a position of $-\delta_k$ in the factor k. In that case the hedged portfolio $\tilde{V}$ has the stochastic evolution

$$d\tilde{V} = (\mu(t, f_1, f_2) - \delta_k(\mu_k - \alpha_k f_k))dt + \sum_{j \neq k} \delta_j \sigma_j dW_j \tag{2.50}$$

Such a position in a specific factor can only be established with the traded futures $F(t,T_j)$, $j = 1,\dots,N$. The future with tenor T has the instantaneous evolution

$$dF(t,T) = F(t,T)\left(\sigma_1 u_1(T-t)dW_1 + \sigma_2 u_2(T-t)dW_2\right) \tag{2.51}$$

Consider a portfolio with w_j contracts $F(t,T_j)$, such that

$$dW_k \cdot dV = \left(\sum_{j=1}^N w_j F(t,T_j) u_k(T_j - t) \right) \sigma_k dt \tag{2.52}$$

For this portfolio to hedge the factor f_k while being unaffected by the factors $f_l, l \neq k$, the following equations must be satisfied:

$$\begin{aligned}
&\sum_{j=1}^{N} w_j F(t,T_j) u_k(T_j - t) = 1 \\
&\sum_{j=1}^{N} w_j F(t,T_j) u_l(T_j - t) = 0, \quad l \neq k
\end{aligned} \tag{2.53}$$

If the number of factors is smaller than *N* there are several solutions to the equations. If there are only two factors, this can be accomplished using two distinct contracts $F(t,T_1)$ and $F(t,T_2)$. To hedge factor 1:

$$\left.\begin{aligned}
w_1 F(t,T_1) u_1(T_1 - t) + w_2 F(t,T_2) u_1(T_2 - t) = 1 \\
w_1 F(t,T_1) u_2(T_1 - t) + w_2 F(t,T_2) u_2(T_2 - t) = 0
\end{aligned}\right\} \Rightarrow \left\{\begin{aligned}
w_1 &= \frac{u_2(T_2 - t)}{F(t,T_1)} \frac{1}{D(t,T_1,T_2)} \\
w_2 &= -\frac{u_2(T_1 - t)}{F(t,T_2)} \frac{1}{D(t,T_1,T_2)}
\end{aligned}\right. \tag{2.54}$$

where $D(t,T_1,T_2) = u_1(T_1 - t) u_2(T_2 - t) - u_2(T_1 - t) u_1(T_2 - t)$.

The portfolio with w_1 contracts expiring at T_1 and w_2 contracts expiring at T_2 replicates the stochastic part of the factor f_1. Similarly, the portfolio replicating the stochastic part of f_2 with these contracts is

$$\begin{aligned}
w_1 &= -\frac{u_1(T_2 - t)}{F(t,T_1)} \frac{1}{D(t,T_1,T_2)} \\
w_2 &= \frac{u_1(T_1 - t)}{F(t,T_2)} \frac{1}{D(t,T_1,T_2)}
\end{aligned} \tag{2.55}$$

It should be noted that using only two contracts makes the hedge very sensitive to these two contracts. If a continuous forward curve $F(t,T)$ is available, a hedge of f_k can be formed using all the contracts if the following conditions are satisfied:

$$\begin{aligned}
&\int_0^{\tau_{\max}} w(\tau) F(t,t+\tau) u_k(\tau) d\tau = 1 \\
&\int_0^{\tau_{\max}} w(\tau) F(t,t+\tau) u_l(\tau) d\tau = 0 \quad l \neq k
\end{aligned} \tag{2.56}$$

A solution to this equation is then:

$$w(\tau) = \frac{u_k(\tau)}{F(t,t+\tau)} \tag{2.57}$$

3. ASIAN OPTIONS ON COMMODITIES

1. Definitions and markets

Most liquid commodity futures traded on exchanges settle on a specific day. For example, Brent futures trading on the InterContinental Exchange settle on the ICE Brent index price on the day following the last trading day of the futures contract. Futures with physical delivery, such as NYMEX WTI futures, do not have cash settlement but the options trading on them settle on their value on a specific day.

In the case of freight derivatives the spot indices, published daily by the Baltic Exchange, are not considered liquid enough to be used for derivatives settlement. Given that there are relatively few spot transactions on a particular day, a big market participant might be able to manipulate the market to his favor over a period of a couple of days. To avoid this, the forward contracts settle on the average spot price over a month. This structure can also be found in over-the-counter swaps in other markets, such as crude oil or metals.

Given a set of settlement dates $T_1,...,T_N$ (generally the trading days of a given month), the settlement price of the average contract settling on these dates will be

$$F_A(T_N;T_1,...,T_N) = \frac{1}{N}\sum_{k=1}^{N} S(T_k) \tag{3.1}$$

This settlement price is also used for settling Asian options written on the same commodity. For example, the payoff of an Asian call option with strike *K* settling on the spot fixings on the dates $T_1,...,T_N$ is

$$C(T_N,K;T_1,...,T_N) = \max\left[\frac{1}{N}\sum_{k=1}^{N} S(T_k) - K, 0\right] \tag{3.2}$$

Asian options are very common in commodities – indeed they first appeared through commodity-linked bonds (Carr et al, 2008). They are popular not only because they avoid the problems of market manipulation as detailed above, but also because they are less expensive than their European counterparts.

The Asian options we will consider are arithmetic average options with European exercise. Given a set of fixing dates $T_1,...,T_N$, the option will pay off at date T_N the value

$$\begin{aligned} C(T_N,K;T_1,...,T_N) &= \max\left[\frac{1}{N}\sum_{k=1}^{N} S(T_k) - K, 0\right] \quad \text{(for a call)} \\ P(T_N,K;T_1,...,T_N) &= \max\left[K - \frac{1}{N}\sum_{k=1}^{N} S(T_k), 0\right] \quad \text{(for a put)} \end{aligned} \tag{3.3}$$

2. Literature on Asian options

The existing literature on Asian options focuses on Asian options written on stock or foreign exchange rates. In this case the main effect of the averaging is in reducing the standard deviation of the payoff function. However, the distribution of the average of log-normal variables is not log-normal, and this is the main obstacle to pricing Asian options using the standard Black-Scholes framework.

To tackle this, several techniques have been developed. Monte Carlo simulation can be used, such as in Kemna and Vorst (1990), Haykov (1993) and Joy et al. (1996). A partial differential equation depending on the spot price and the observed average price can be derived and solved numerically: Dewynne and Wilmott (1995) and Rogers and Shi (1995).

Geman and Yor (1992) derive a semi-analytical expression for a spot price following geometric Brownian motion. Turnbull and Wakeman (1991) and Levy (1992) derive approximate expressions by matching the moments of a log-normal distribution with the moments of the average price distribution.

A closed form expression is derived in Geman and Yor (1992) for a spot price following geometric Brownian motion. Approximate expressions have been obtained by Turnbull and Wakeman (1991) and Levy (1992). Haug (2006) presents these and other approximations for Asian options on futures. Koekebakker, Ådland and Sødal (2007) find an approximate expression for the Asian options trading in shipping, assuming the spot price follows geometric Brownian motion. Koekebakker and Ollmar (2005) use a one-factor forward curve model with time-varying volatility and derive an approximate process for the shipping forward freight agreement.

A major issue in using these formulas for commodity futures options is that they assume geometric Brownian motion for the spot price, which is not consistent with a multi-factor model with mean-reverting factors. They also ignore the existence of a forward curve which gives the risk-neutral expectations of the spot price.

3. Approximate formulas under the two-factor model

For option pricing we will work in the risk-neutral measure. The two-factor model of the forward curve is, as formulated in Part 2,

$$\frac{dF(t,T)}{F(t,T)} = \sigma_S e^{-\alpha(T-t)} dW_S + \sigma_L dW_L, \qquad dW_S dW_L = \rho dt \tag{3.4}$$

Consider an Asian forward contract F_A settling on the average of the daily contracts $F(\cdot,T_1),...,F(\cdot,T_N)$. The average price contract satisfies:

$$\begin{aligned} F_A(t) &= \frac{1}{N}\sum_{k=1}^{N} F(t,T_k) \\ dF_A(t) &= \frac{1}{N}\sum_{k=1}^{N} dF(t,T_k) \end{aligned} \tag{3.5}$$

where $\sigma_k(\tau)=0$ when $\tau<0$ (the contract has already settled, so its price is fixed). Then

$$\frac{dF_A(t)}{F_A(t)} = \frac{\sum_{k=1}^{N}\sigma_S e^{-\alpha(T_k-t)}F(t,T_k)}{\sum_{k=1}^{N}F(t,T_k)}dW_S + \frac{\sum_{k=1}^{N}\sigma_L F(t,T_k)}{\sum_{k=1}^{N}F(t,T_k)}dW_L \tag{3.6}$$

Assumption 1: we assume the forward curve to be flat through the settlement period of the Asian contract: $F(t,T_k)=F_A(t)$ (all the daily contracts have the same price)

If we make Assumption 1 then the above equation simplifies to the lognormal evolution

$$\frac{dF_A(t)}{F_A(t)} = \left(\frac{1}{N}\sum_{k=1}^{N}\sigma_S e^{-\alpha(T_k-t)}\right)dW_S + \sigma_L dW_L \tag{3.7}$$

That is, the factor volatilities of F_A are the average of the factor volatilities of the individual contracts.

Assumption 2: Let us assume that the fixing dates are equally distributed: $T_k = T_N-(N-k)h$. Denote by $c=T_N-T_1$ the contract length. Then for $t<T_1$ (pre-settlement):

$$\begin{aligned} \frac{dF_A(t)}{F_A(t)} &= \frac{\sigma_S e^{-\alpha(T_N-t)}}{N}\sum_{k=0}^{N-1}e^{\alpha kh}dW_S + \sigma_L dW_L \\ &= \sigma_S e^{-\alpha(T_N-t)}\frac{e^{\alpha Nh}-1}{N(e^{\alpha h}-1)}dW_S + \sigma_L dW_L \\ &= \sigma_S e^{-\alpha(T_N-t)}\frac{e^{\alpha\frac{N}{N-1}(T_N-T_1)}-1}{N(e^{\alpha(T_N-T_1)/(N-1)}-1)}dW_S + \sigma_L dW_L \end{aligned} \tag{3.8}$$

Assumption 3: The number of fixing dates N is large enough that $N(e^{\alpha(T_N-T_1)/(N-1)}-1)\approx \alpha c$. Then

$$\frac{dF_A(t)}{F_A(t)} \approx \sigma_S e^{-\alpha(T_N-t)}\frac{e^{\alpha c}-1}{\alpha c}dW_S + \sigma_L dW_L \tag{3.9}$$

Dynamics inside the settlement period

An essential issue of Asian contracts is what happens inside the settlement period. As the contract enters the settlement period, its constituent prices are progressively discovered, and the uncertainty on its price at expiration diminishes. For example, the day before expiration, the only uncertainty is on the last price's evolution during one day, which only contributes a small part to the average contract price.

Furthermore, as part of the contract is priced, our Assumption 1 of a flat term structure $F(t,T_k) = F_A(t)$ becomes wrong. Indeed on date $t, T_M \le t < T_{M+1}$, the spot prices $S(T_1),...,S(T_M)$ have been observed and they will not be equal to $F_A(t)$. We will therefore consider the observed average $A(t)$ and the adjusted average contract price $F_A^{'}(t)$:

$$
\begin{aligned}
A(t) &= \frac{1}{M}\sum_{k=1}^{M} S(t_k), \qquad T_m \le t < T_{m+1} \\
F_A^{'}(t) &= F_A(t) - \frac{M}{N}A(t) = \frac{1}{N}\sum_{k=M+1}^{N} F(t,T_k)
\end{aligned}
\tag{3.10}
$$

If we assume a flat term structure ahead, $F(t,T_k) = \dfrac{N}{N-M}F_A^{'}(t)$ for $k \ge m+1$ then

$$
\begin{aligned}
dF_A^{'}(t) &= \frac{1}{N}\sum_{k=M+1}^{N} dF(t,T_k) \\
&= \frac{1}{N}\sum_{k=M+1}^{N} \frac{N}{N-M} F_A^{'}(t)(\sigma_S e^{-\alpha(T_k-t)}dW_S + \sigma_L dW_L) \\
\frac{dF_A^{'}(t)}{F_A^{'}(t)} &= \left(\frac{1}{N-M}\sum_{k=M+1}^{N} \sigma_S e^{-\alpha(T_k-t)}\right)dW_S + \sigma_L dW_L
\end{aligned}
\tag{3.11}
$$

The adjusted average contract price $F_A^{'}(t)$ has an approximately lognormal evolution with volatilities equal to the average of the volatilities. This, however, is only valid for $T_M \le t < T_{M+1}$ and assuming a flat term structure. The exact evolution of $F_A(t)$ from $t = T_1$ to $t = T_N$ is complex, and an exact derivation would have to follow the lines of Geman and Yor (1993).

We do however get a good idea of the result by assuming that the already settled prices $S(T_k)$ and the daily forward prices $F(t,T_k)$ are all equal, in which case the average contract follows the evolution:

$$
\frac{dF_A(t)}{F_A(t)} = \left(\frac{1}{N}\sum_{k=1}^{N}\sigma_1(T_k - t)\right)dW_1 + \left(\frac{1}{N}\sum_{k=1}^{N}\sigma_2(T_k - t)\right)dW_2 \tag{3.12}
$$

By definition $\sigma_k(\tau) = 0$ for $\tau < 0$. In that case the average contract still has a lognormal distribution inside the settlement period, but with volatility decaying to 0 as shown in Figure 13.

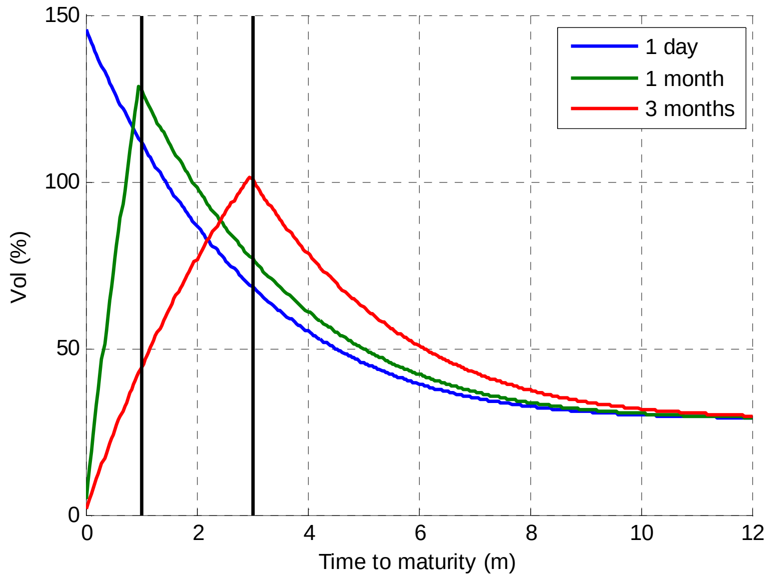

Figure 13. Instantaneous volatility of the Asian contract as a function of time to maturity, for different contract lengths

We want to quantify this contract's stochastic evolution until maturity at T_N. Consider a date *t* such that $T_{M+1} \le T_L \le t < T_{L+1}$. Then:

$$\begin{aligned} \frac{dF_A^{'}(t)}{F_A^{'}(t)} &= \frac{\sigma_S e^{-\alpha(T_N - t)}}{N-M} \sum_{k=L+1}^{N} e^{\alpha(N-k)h} dW_S + \frac{N-L}{N-M} \sigma_L dW_L \\ &\approx \sigma_S e^{-\alpha(T_N-t)} \frac{e^{\alpha(T_N-t)}-1}{\alpha c_M^{'}} dW_S + \frac{T_N - t}{c_M^{'}} \sigma_L dW_L \qquad \text{(Assumption 3)} \end{aligned} \tag{3.13}$$

where $c_M^{'} = T_N - T_M$ is the residual contract period.

Pricing Asian options

As discussed previously, in some markets the Asian forward contract/swap is traded in the market, and also settles on the average of the spot:

$$F_A(T_N; T_1, \ldots, T_N) = \frac{1}{N} \sum_{k=1}^{N} S(T_k) \tag{3.14}$$

Therefore we can rewrite the payoff of the option as

$$\max\left(\frac{1}{N}\sum_{k=1}^{N}S(T_k)-K,0\right)=\max\left(F_A(T_N,T_1,...,T_N)-K,0\right) \tag{3.15}$$

This payoff is a European call option written on the forward contract F_A, which simplifies the problem considerably. We will denote by

$$C(t,F_A(t,T_1,...,T_N),K) \tag{3.16}$$

the price, at time *t*, of a call option settling on the average of the spot prices at times $T_1,...,T_N$, with strike *K*. Similarly $P(t,F_A(t,T_1,...,T_N),K)$ denotes the price at time *t* of a put.

Put-Call parity: because the options are standard European options written on a swap, put-call parity holds in the form

$$C(t,F_A(t;T_1,...,T_N),K)-P(t,F_A(t;T_1,...,T_N),K)=B(t,T_N)(F_A(t;T_1,...,T_N)-K) \tag{3.17}$$

where $B(t,T_N)$ is the discount factor. We will henceforth focus the discussion on calls.

Black's formula under time-dependent volatility

We will begin by establishing Black's (1976) formula for futures options under deterministic time-dependent volatility. We will follow the derivation of Musiela and Rutkowski (2008).

Consider a futures process $F(t,T)$ with time-varying volatility:

$$\frac{dF(t,T)}{F(t,T)}=\mu(t,T)dt+\sigma(t,T)dW_t \tag{3.18}$$

We consider the futures option settling at date *T* on $F(T,T)$ with strike *K*:

$$C(T,F)=(F-K)^+ \tag{3.19}$$

Consider a self-financing futures strategy $\phi=(g(F_t,t),h(F_t,t))$. Since the replicating portfolio ϕ is assumed to be self-financing, the wealth process $V(t,F,\phi)$, which equals

$$V(t,F_t,\phi)=h(F_t,t)B_t=C(t,F_t) \tag{3.20}$$

satisfies

$$\begin{aligned} dV &= g(F_t,t)dF_t+h(F_t,t)dB_t \\ &= rV(t,F_t)dt+\mu(t,T)F_tg(F_t,t)dt+\sigma(t,T)F_tg(F_t,t)dW_t \end{aligned} \tag{3.21}$$

If we assume that the function *C* is sufficiently smooth, we find by Ito's lemma that

$$dC(t,F_t)=\left(\frac{\partial C}{\partial t}+\mu(t,T)F_t\frac{\partial C}{\partial F}+\frac{1}{2}\sigma^2(t,T)F_t^2\frac{\partial^2 C}{\partial F^2}\right)dt+\sigma F_t\frac{\partial C}{\partial F}dW_t \tag{3.22}$$

Equating the values of *V* and *C* we find that we must have

$$g(F_t,t)=\frac{\partial C}{\partial F}(t,F_t) \tag{3.23}$$

And

$$\begin{aligned}&\frac{\partial C}{\partial t}+\frac{1}{2}\sigma^2(t,T)F^2\frac{\partial^2 C}{\partial F^2}-rC=0\\&C(T,F)=(F-K)^+\end{aligned} \tag{3.24}$$

The solution of this partial differential equation is

$$C(t,F(t,T))=B(t,T)\left[F(t,T)N(d_1)-KN(d_2)\right] \tag{3.25}$$

where

$$d_1=\frac{\ln(F(t,T)/K)+\frac{1}{2}(T-t)\sigma_{Black}(t,T)^2}{\sigma_{Black}(t,T)\sqrt{T-t}},\quad d_2=d_1-\sigma_{Black}(t,T)\sqrt{T-t} \tag{3.26}$$

$$\sigma_{Black}(t,T)^2=\frac{1}{T-t}\int_t^T\sigma(s,T)^2ds \tag{3.27}$$

Hence the formula for a call written on a futures contract with time-varying volatility is the same as Black's (1976) formula for a futures option, except that the volatility is replaced with the root-mean-square of the instantaneous volatilities during the period.

Instantaneous volatilities

Let us now calculate the instantaneous volatilities of $F_A(t)$ pre- and in-settlement, based on the stochastic evolution derived in equations (3.9) and (3.13). For $t<T_1$

$$\sigma_A(t,T_N)=\sqrt{\left(\sigma_S e^{-\alpha(T_N-t)}\frac{e^{\alpha c}-1}{\alpha c}+\rho\sigma_L\right)^2+(1-\rho^2)\sigma_L^2} \tag{3.28}$$

and in-settlement, for $T_1\le t\le T_N$:

$$\sigma'_A(t,T_N)=\sqrt{\left(\sigma_S\frac{1-e^{-\alpha(T_N-t)}}{\alpha c'_M}+\rho\frac{T_N-t}{c'_M}\sigma_L\right)^2+(1-\rho^2)\left(\frac{T_N-t}{c'_M}\right)^2\sigma_L^2} \tag{3.29}$$

Black volatilities

Given these instantaneous volatility functions, we can calculate the Black volatility, i.e. the standard deviation of the contract price at expiration:

$$\sigma_{Black}(t,T)^2=\frac{1}{T}\int_t^T\sigma_A^2(s,T)ds \tag{3.30}$$

In-settlement

Consider the case when dates $T_1,...,T_M$ have been priced, such that we are in fact considering the adjusted contract

$$F'_A(t)=F_A(t)-\frac{1}{N}\sum_{k=1}^{M}S(T_k)=\frac{1}{N}\sum_{k=M+1}^{N}F(t,T_k) \tag{3.31}$$

In Appendix 6 we show that its Black volatility is:

Such that the square of the Black volatility is given by:

$$\begin{aligned}\sigma_{Black}(t,T)^2&=\frac{1}{T-t}\int_t^T\sigma_A^2(s,T)ds\\&=\frac{\sigma_S^2}{\alpha^2c_M'^2}\left(1-\frac{2}{\alpha}\frac{1-e^{-\alpha(T-t)}}{T-t}+\frac{1}{2\alpha}\frac{1-e^{-2\alpha(T-t)}}{T-t}\right)+\cdots\\&\quad\frac{2\rho\sigma_S\sigma_L}{\alpha c_M'^2}\left(\frac{T-t}{2}+\frac{e^{-\alpha(T-t)}}{\alpha}-\frac{1}{\alpha^2}\frac{1-e^{-\alpha(T-t)}}{T-t}\right)+\cdots\\&\quad\frac{1}{3}\frac{(T-t)^2}{c_M'^2}\sigma_L^2\end{aligned} \tag{3.32}$$

In the case when $\alpha c \ll 1$ this simplifies to

$$\sigma_{Black}(t,T)\approx\frac{1}{\sqrt{3}}\sqrt{\sigma_S^2+2\rho\sigma_S\sigma_L+\sigma_L^2} \tag{3.33}$$

Pre-settlement

In Appendix 6 we show that

$$\sigma_{Black}^2(t,T) = \frac{1}{T-t}\int_t^T \sigma_A^2(s,T)ds = \frac{1}{T-t}\int_t^{T_1} \sigma_A^2(s,T)ds + \frac{T-T_1}{T-t}\sigma_{Black}^2(T_1,T) \quad (3.34)$$

where

$$\int_t^{T_1} \sigma_A^2(s,T)ds = \sigma_S^2\left(\frac{e^{\alpha c}-1}{\alpha c}\right)^2 \frac{e^{-2\alpha(T-T_1)}-e^{-2\alpha(T-t)}}{2\alpha} + 2\rho\sigma_S\sigma_L\frac{e^{\alpha c}-1}{\alpha c}\frac{e^{-\alpha(T-T_1)}-e^{-\alpha(T-t)}}{\alpha} + (T_1-t)\sigma_L^2 \quad (3.35)$$

and $\sigma_{Black}^2(T_1,T)$ is given by equation (3.32). If $\alpha c \ll 1$,

$$\sigma_{Black}(t,T)^2 \approx \sigma_S^2\frac{c}{T-t}\left[\frac{1-e^{-2\alpha(T_1-t)}}{2\alpha c}+\frac{1}{3}\right] + 2\rho\sigma_S\sigma_L\frac{c}{T-t}\left[\frac{1-e^{-\alpha(T_1-t)}}{\alpha c}+\frac{1}{3}\right] + \sigma_L^2\left[1-\frac{2}{3}\frac{c}{T-t}\right] \quad (3.36)$$

Option price

Once the Black volatility is known, pricing futures options is simply a matter of applying the equation (3.26) to the process $F_A(t,T)$. If the option is pre-settlement,

$$\begin{aligned} C(t,F_A(t,T),K) &= B(t,T)\left[F_A(t,T)N(d_1) - KN(d_2)\right] \\ P(t,F_A(t,T),K) &= B(t,T)\left[KN(-d_2) - F_A(t,T)N(-d_1)\right] \end{aligned} \quad (3.37)$$

where

$$d_1 = \frac{\log(F_A(t,T)/K) + \frac{1}{2}(T-t)\sigma_{Black}(t,T)^2}{\sigma_{Black}(t,T)\sqrt{T-t}}, \quad d_2 = d_1 - \sigma_{Black}(t,T)\sqrt{T-t} \quad (3.38)$$

If it is in-settlement, introduce the already priced average.

$$A(t) = \frac{1}{M}\sum_{k=1}^{M} S(t_k), \qquad T_M \le t < T_{M+1} \quad (3.39)$$

The payoff at expiration can be rewritten as

$$\max\left(F_A(T)-K,0\right) = \max\left(F_A'(T) - \left(K - \frac{M}{N}A(t)\right),0\right) \approx \max\left(F_A'(T) - \left(K - \frac{t-T_1}{c}A(t)\right),0\right) \quad (3.40)$$

such that the option can be considered to be written on the log-normal process $F_A'(t)$ with an adjusted strike $K - A(t)M/N$. If $K - A(t)M/N \ge 0$. The prices of the options are therefore:

$$C(t,F_A(t,T),K)=B(t,T)\left[\left(F_A(t,T)-\frac{t-T_1}{c}A(t)\right)N(d_1)-\left(K-\frac{t-T_1}{c}A(t)\right)N(d_2)\right]$$
$$P(t,F_A(t,T),K)=B(t,T)\left[\left(K-\frac{t-T_1}{c}A(t)\right)N(-d_2)-\left(F_A(t,T)-\frac{t-T_1}{c}A(t)\right)N(-d_1)\right] \quad (3.41)$$

where:

$$d_1=\frac{\log\left(\dfrac{F_A(t,T)-\frac{t-T_1}{c}A(t)}{K-\frac{t-T_1}{c}A(t)}\right)+\frac{1}{2}(T-t)\sigma'_{Black}(t,T)^2}{\sigma'_{Black}(t,T)\sqrt{T-t}}, \quad d_2=d_1-\sigma'_{Black}(t,T)\sqrt{T-t} \quad (3.42)$$

and $\sigma'_{Black}(t,T)$ is given by equation (3.32).

If $K-A(t)M/N<0$, the average over the contract period will always be larger than the strike, such that the call option will always be exercised and the put option will never be exercised. Thus

$$C(t,F_A(t,T),K)=B(t,T)(F_A(t,T)-K)$$
$$P(t,F_A(t,T),K)=0 \quad (3.43)$$

4. Comparison to other Asian option models and market prices

We examine two other Asian option models that are commonly used in the freight options market: the Levy (1992) approximation and the Koekebakker, Adland and Sodal (2007) formula. These models both assume that the spot price follows geometric Brownian motion under the risk-neutral measure:

$$\frac{dS(t)}{S(t)}=\lambda dt+\sigma dW_t \quad (3.44)$$

Let $T_k=T_1+(k-1)h$ be the settlement period dates and define the observed running average $A(t)$ as in equation (3.10) above. Also let $M(t)=A(T_N)-A(t)\frac{m}{N}$ and

$$v(t)=\sqrt{\ln E^*\left[M(t)^2\right]-2\ln E^*\left[M(t)\right]}$$
$$d_1=\frac{\frac{1}{2}\ln E^*\left[M(t)^2\right]-\ln\left[K-A(t)\frac{m}{N}\right]}{v(t)}, \quad d_2=d_1-v(t) \quad (3.45)$$

Then the Levy (1992) approximation is, for $T_m \le t < T_{m+1}$,

$$
\begin{aligned}
C(S(t),A(t),t) &= e^{-r\tau}\left(E^*\left[M(t)\right]N(d_1) - \left[K - A(t)\frac{m}{N}\right]N(d_2)\right) \\
P(S(t),A(t),t) &= e^{-r\tau}\left(E^*\left[M(t)\right]\left[N(d_1)-1\right] - \left[K - A(t)\frac{m}{N}\right]\left[N(d_2)-1\right]\right)
\end{aligned}
\tag{3.46}
$$

Expressions for $E^*[M(t)]$ and $E^*[M(t)^2]$ are given in the appendix of the original article.

Koekebakker, Ådland and Sødal (2007) arrive at an option price along the following lines:

1. Assume geometric Brownian motion for the spot under the risk-neutral measure
2. Derive what the average forward price should be, consistent with this spot process

$$F_A(t;T_1,\dots,T_N) = \frac{1}{N}\sum_{i=1}^{N} E_t^*[S(T_i)] = S_t \frac{e^{\lambda(T_N-t)}}{N}\frac{1-e^{-\lambda N\Delta}}{1-e^{-\lambda\Delta}} \tag{3.47}$$

3. Derive the forward's approximate evolution: $\dfrac{dF_A(t;T_1,\dots,T_N)}{F_A(t;T_1,\dots,T_N)} = \sigma(t)dW_t$ where

$$\sigma(t) = \begin{cases} \sigma & t \le T_1 \\ \sigma \dfrac{\frac{S_t}{N}\sum_{i=M+1}^{N} e^{\lambda(T_i-t)}}{F(t,T_1,T_N)} & T_M < t < T_{M+1} \end{cases} \approx \begin{cases} \sigma & t \le T_1 \\ \sigma \dfrac{N-M}{N} & T_M < t < T_{M+1} \end{cases} \tag{3.48}$$

4. Use Black (1976) with time-varying volatility to price the option

The result is as follows: for $t < T_1$ (pre-settlement) let

$$\sigma_F = \sigma\sqrt{(T_1-t) + R(N)(T_N-T_1)} \tag{3.49}$$

$$R(N) = \frac{1-\dfrac{3}{2N}+\dfrac{1}{2N^2}}{3-\dfrac{3}{N}} \tag{3.50}$$

For $t = T_M$ (in settlement)

$$\sigma_F = \frac{\sigma}{\sqrt{3}}\left[(T_N - T_M)\left(1-\frac{M-1}{N}\right)\left(1-\frac{M}{N}+\frac{1}{2N}\right)\right]^{1/2} \tag{3.51}$$

Then

$$\begin{aligned} C(t, F_A(t;T_1,...,T_N), K) &= B(t,T_N)\left[F_A(t;T_1,...,T_N)N(d_1) - KN(d_2)\right] \\ P(t, F_A(t;T_1,...,T_N), K) &= B(t,T_N)\left[KN(-d_2) - F_A(t;T_1,...,T_N)N(-d_1)\right] \end{aligned} \tag{3.52}$$

where $B(t,T_N)$ is the discount factor and

$$d_1 = \frac{\ln\left(\frac{F_A(t;T_1,...,T_N)}{K}\right) + \frac{1}{2}\sigma_F^2}{\sigma_F}, \quad d_2 = d_1 - \sigma_F \tag{3.53}$$

We now evaluate option premia based on these models. The volatility inputs to the models are evaluated based on the historical volatilities estimated from time series of prices of shipping futures contracts. We compare these to the option premia quoted by Imarex in their weekly Imarex Freight Options reports for the TD3 route.

The main parameters are given in Table 3..

Table 3. Parameters on December 8, 2008

Date	December 8, 2008
3-month LIBOR	2.19%
Spot price (WS)	91.73
Observed average spot 12/1 – 12/5 (WS)	75.88

The key input to each of the models is the volatility. We use

- For Levy and KAS, the historical volatility of the spot, using weekly log returns over the year 2008. This is evaluated to 148.4%.
- We calibrate the two-factor forward curve model to the estimated historical covariance matrix for the year 2008.

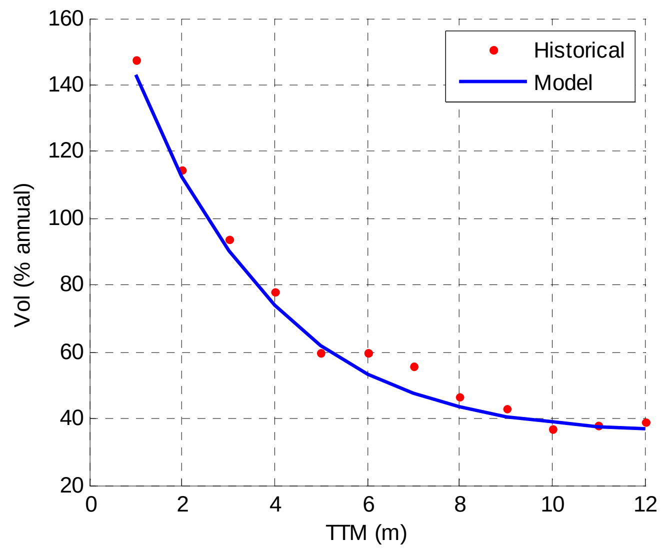

Parameters:

$\sigma_S = 172.4\%$

$\sigma_L = 34.8\%$

$\alpha = 3.245$

$\rho = 0.21$

Figure 14. Two-factor model fitted to historical volatilities

Table 4. Model Call option ATM prices on December 8, 2008 compared to market

Contract		FFA	Imarex Premium ATM	Model option premium ATM (WS)		
Year	**Period**	**WS**	**WS**	**Levy**	**KAS**	**Two-factor**
2008	Dec	81	5.6	4.91	5.45	6.76
2009	Jan	59	8.8	10.69	10.56	11.07
2009	Feb	56	9.2	13.60	13.51	12.87
2009	Mar	46	8.2	13.34	13.28	11.52
2009	Apr	45	8.2	14.98	14.93	11.93
2009	Q2	45	8.5	16.48	16.44	12.27
2009	Q3	44	9.3	19.82	19.79	12.65
2009	Q4	48	11	24.66	24.63	14.18
2010	CAL	78	16.6	48.26	48.23	24.10

The results are listed in Table 4. We note that the Levy and KAS prices are very similar. This is not surprising as they are based on the same model and parameters for the underlying, only with different ways of approximating the option premium.

The option premia obtained with the term structure of volatility from Figure 14 are lower than the Levy and KAS premia. This is obvious from Figure 15: the difference between KAS and the present study is the volatility input to the Black formula - σ_F and σ_{Black} respectively – which are the root-mean-square of the time-dependent instantaneous volatility $\sigma(t)$. Since the historical volatility of futures with increasing tenors is lower than that of the spot, the call premium from the present study will be lower than the KAS premium.

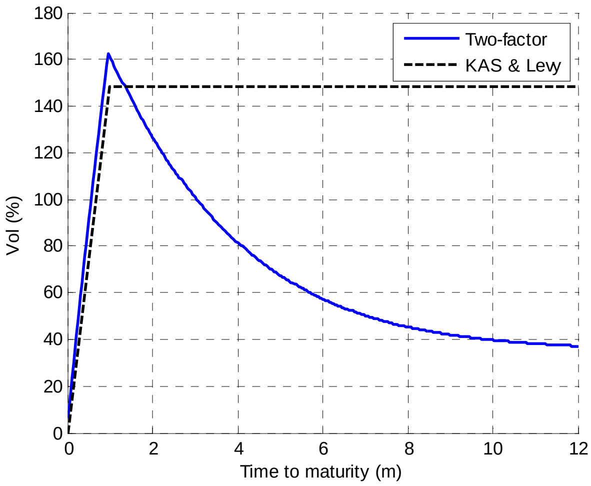

Figure 15. Asian volatility term structure for the two-factor model and the KAS & Levy models

Comparing these premia to market prices quoted by Imarex, the market prices are lower than the prices obtained from the present model supplied with the historical volatilities – namely the market implies a lower future volatility than the historical volatility over 2008. Yet, 2008 was a particularly volatile year in the freight market, giving a high historical volatility, whereas the market might be expecting 2009 to be calmer.

Implied volatilities – calibration of the models to market prices

Volatility is the most important parameter of an option pricing formula. Considering that the other parameters are observable, a quoted market price for an option implies a value for this parameter. It is therefore important to be able to back out this parameter from the options premia observed in the market, resulting in the implied volatility. This follows by solving the equation

$$C(t, F_A(t,T), T, K, r, \sigma) = C_{market}(t, T, K) \tag{3.54}$$

for σ or possibly several parameters that enter the definition of σ.

Levy (1992)

The formula (3.46) has a complicated dependence in σ and a numerical technique must be used to back out the volatility of the spot from the option price. As this is a single parameter, to each option price $C(T,K)$

there corresponds an implied volatility $\sigma_{Levy}(T,K)$. The volatility of the spot and the options implied volatility are not the same because the market doesn't follow the assumptions of the model. In particular, there is a term structure of volatility which is not consistent with the geometric Brownian motion model for the spot price. Imarex quotes these implied volatilities in their Freight Options reports.

Koekebakker, Adland and Sødal (2007)

The formula (3.52) is just Black (1976) with a tweaked volatility input. Extracting implied volatilities from the Black formula is standard, giving rise to $\sigma_{Black}(T,K)$. The implied volatility of the spot is then

$$\sigma_{KAS}(T,K) = \sigma_{Black}(T,K)\sqrt{\frac{T-t}{g(t,T_1,\ldots,T_N)}} \tag{3.55}$$

where, consistently with equations (3.49) and (3.51),

$$g(t,T_1,\ldots,T_N) = \begin{cases} T_1 - t + R(N)(T_N - T_1) & t < T_1 \\ \frac{1}{3}(T_N - T_M)\left(1-\frac{M-1}{N}\right)\left(1-\frac{M}{N}+\frac{1}{2N}\right) & T_1 \le t = T_M \end{cases} \tag{3.56}$$

As with the Levy model, different implied volatilities are obtained for each maturity and strike.

Two-factor model

The input to the Black formula is the time dependent volatility $\sigma_A(t,T)$ modeled according to (3.28) and (3.29). The parameters in the model need to be estimated from the term structure of options prices. Rebonato (2002) discusses in depth the calibration of the LIBOR market model to traded options in the interest rate markets, and much of his discussion applies here. $\sigma_{Black}(T)$ for at-the-money options is first obtained from options market prices, and the parameters of the model are estimated by nonlinear least squares:

$$\min_{\sigma_S,\sigma_L,\alpha,\rho} \sum_T \left[\sigma_{Black}^{Model}(T,\sigma_S,\sigma_L,\alpha,\rho) - \sigma_{Black}(T)\right]^2 \tag{3.57}$$

where the summation is over all available liquid option maturities. It should be noted that unlike the Levy and KAS models which estimate one implied volatility per options contract, the present model estimates four parameters from all liquid options prices by a nonlinear least squares technique, which is more parsimonious, but can lead to inaccuracies if the model is not suitable.

Note on calibrating a multi-factor model to implied volatilities

The two-factor model that we present in this article is intended to reproduce not only the term structure of volatilities but also the correlation surface between the contracts. When pricing vanilla options, however, only the term structure of volatilities matters and the correlations are irrelevant. Thus, when calibrating the four parameters of the model to implied volatilities alone one cannot expect to correctly reproduce the correlation structure.

However, considering that we have observed historical estimates of the parameter ρ to be close to 0, a simple solution consists in fixing this parameter to 0 and calibrating $\sigma_S, \sigma_L, \alpha$ to the implied volatility term structure. As long as this produces a good fit, vanilla options will be priced correctly.

Results

Based on the prices published in the Imarex Freight Options report on Dec 8, 2008, we extract implied parameters for the different models. The results are presented in Table 5. We fit the two-factor model (2.2) to the market prices using the procedure described above, and the result of the optimization is displayed in Figure 16. We can see that the two-factor model gives a very good fit to the option market prices.

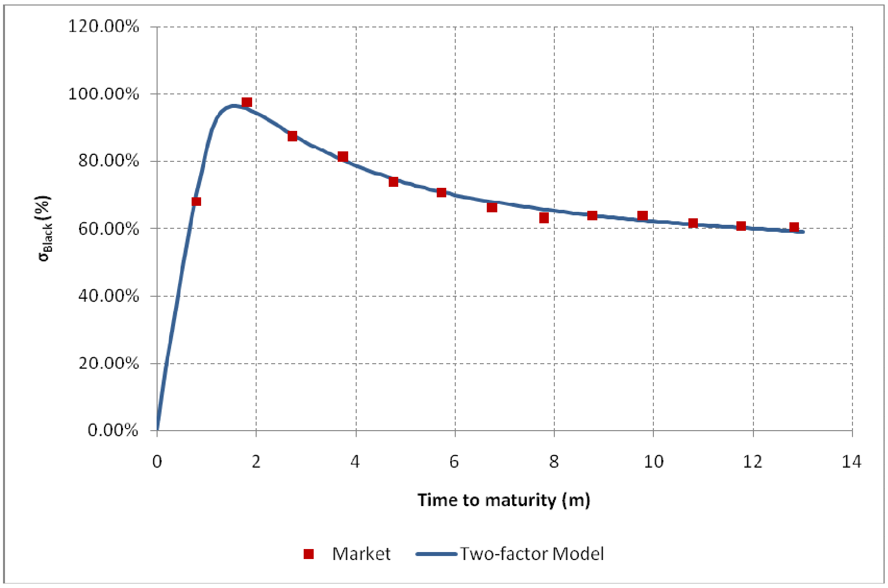

Figure 16. Black implied volatilities from market quoted options and calibrated two-factor model. Market prices are from the Imarex Freight Options report on December 8, 2008

Table 5. Implied volatilities for the different models on December 8, 2008

Contract		Imarex Premium	Implied volatilities					
Year	Month	WS	Imarex report	Levy $\sigma_{Levy}(T)$	KAS $\sigma_{KAS}(T)$	Black $\sigma_{Black}(T)$	Two-factor	
2008	Dec	5.6	165%	169.20%	152.53%	67.96%	σ_S	177.25%
2009	Jan	8.8	125%	121.88%	123.32%	97.43%	σ_L	47.68%
2009	Feb	9.2	100%	99.60%	100.23%	87.51%	α	8.7
2009	Mar	8.2	90%	89.92%	90.31%	81.30%	ϱ	0
2009	Apr	8.2	80%	79.52%	79.76%	73.86%		
2009	May	8.6	75%	74.98%	75.16%	70.79%		
2009	Jun	8.7	70%	69.60%	69.74%	66.14%		
2009	Jul	8.7	66%	65.98%	66.10%	63.03%		
2009	Aug	9.3	66%	66.27%	66.38%	63.76%		
2009	Sep	9.8	66%	66.05%	66.14%	63.81%		
2009	Oct	10.8	63%	63.44%	63.52%	61.49%		
2009	Nov	11.1	62%	62.46%	62.53%	60.70%		
2009	Dec	11.5	62%	62.10%	62.16%	60.42%		

5. Hedging of Asian options

Greeks of Asian options

When writing options, the seller may be interested in delta-hedging his portfolio with the underlying to construct a delta-neutral position. The question is what position to take in the underlying to hedge the option: the delta. The hedge ratio changes with the price of the underlying, time and volatility and must be adjusted regularly, leading to dynamic hedging strategies.

When hedging a number of questions must be addressed: what is the underlying? What instruments are we going to hedge with? The options are written on spot, but in shipping the spot is ton-miles that are not be possible to buy and hold, nor short, since it is a service. In shipping the forward contracts that are trading in the market, the Forward Freight Agreements, settle on the average of the spot, and should therefore be used as hedging instruments.

The Greeks follow upon differentiation of the Black formula. Let d_1 and d_2 be defined as in Section 3.2, then we have for a call option:

$$\begin{aligned}
\Delta &= \frac{\partial C}{\partial F_A} = e^{-r(T-t)} N(d_1) \\
\Gamma &= \frac{\partial^2 C}{\partial F_A{}^2} = e^{-r(T-t)} \frac{n(d_1)}{F_A \sigma_B \sqrt{T-t}} \\
\theta &= \frac{dC}{dt} = \frac{\partial C}{\partial t} + \frac{\partial \sigma_B}{\partial t} \frac{\partial C}{\partial \sigma_B} = rC - e^{-r(T-t)} \frac{F_A \sigma_B}{2\sqrt{T-t}} n(d_1) + \frac{\partial \sigma_B}{\partial t} Vega \\
Vega &= \frac{\partial C}{\partial \sigma_B} = F_A e^{-r(T-t)} \sqrt{T-t}\, n(d_1) \\
\rho &= \frac{\partial C}{\partial r} = -(T-t)C
\end{aligned} \tag{3.58}$$

Because the volatility is time dependent, the theta θ of the option also depends on the temporal variation of the Black volatility:

$$(T-t)\sigma_B^2(t,T) = \int_t^T \sigma^2(s,T)ds \tag{3.59}$$

Differentiating this with respect to t and rearranging we get

$$\frac{\partial \sigma_B(t,T)}{\partial t} = \frac{\sigma_B^2(t,T) - \sigma^2(t,T)}{2(T-t)\sigma_B(t,T)} \tag{3.60}$$

which can be calculated using the formulas in Section 3.

Delta-hedging an option position

When a forward contract with the same settlement period as the Asian option is available, that contract should be used to delta-hedge the option position to avoid basis risk. The position to be taken in this contract is then given by the previous formula. The position to take in the contract $F_A(t,T)$ to hedge a short call position settling on the same period is

$$\Delta_C = e^{-r(T-t)} N(d_1) \tag{3.61}$$

In other markets such a contract is not available. In the crude oil market, for example, the futures contracts settle on a single date while Asian options will settle on the trading days within a month. Luckily, these contracts are highly correlated and we can quantify the required number of contracts using the two-factor model.

The instantaneous evolution of the call price is

$$\begin{aligned} dC &= \frac{\partial C}{\partial t} dt + \frac{\partial C}{\partial F_A} dF_A \\ &= \theta dt + \Delta_{F_A} F_A(t,T)\left(\sigma_S^A(t,T) dW_S + \sigma_L^A(t,T) dW_L\right) \end{aligned} \tag{3.62}$$

Where $F_A(t,T) = F_A(t;T_1,...,T_N)$ pre-settlement and $F_A(t,T) = F_A^{'}(t;T_{M+1},...,T_N)$ in-settlement, and

$$\begin{aligned} \sigma_S^A(t,T) &= \begin{cases} \sigma_S e^{-\alpha(T-t)} \dfrac{e^{\alpha c}-1}{\alpha c} & t < T_1 \\ \sigma_S e^{-\alpha(T-t)} \dfrac{e^{\alpha(T-t)}-1}{\alpha c_M^{'}} & T_M \leq t < T_{M+1} \end{cases} \\ \sigma_L^A(t,T) &= \begin{cases} \sigma_L & t < T_1 \\ \dfrac{T-t}{c_M^{'}} \sigma_L & T_M \leq t < T_{M+1} \end{cases} \end{aligned} \tag{3.63}$$

A daily contract settling on the date *T* has the stochastic evolution

$$\frac{dF(t,T)}{F(t,T)} = \sigma_S e^{-\alpha(T-t)} dW_S + \sigma_L dW_L \tag{3.64}$$

We can use two of these contracts to hedge the call. Assume we take a position w_1 in contract $F(t,T_1)$ and w_2 in contract $F(t,T_2)$, then the hedge of a short call must satisfy:

$$\begin{cases} w_1 F(t,T_1)\sigma_S e^{-\alpha(T_1-t)} + w_2 F(t,T_2)\sigma_S e^{-\alpha(T_2-t)} = \Delta_{F_A} F_A(t,T)\sigma_S^A(t,T) \\ (w_1 F(t,T_1) + w_2 F(t,T_2))\sigma_L = \Delta_{F_A} F_A(t,T)\sigma_L^A(t,T) \end{cases} \tag{3.65}$$

which yields

$$\begin{aligned} w_1 &= \Delta_{F_A} \frac{F_A(t,T)}{F(t,T_1)} \frac{\dfrac{\sigma_S^A(t,T)}{\sigma_S} - \dfrac{\sigma_L^A(t,T)}{\sigma_L} e^{-\alpha(T_2-t)}}{e^{-\alpha(T_1-t)} - e^{-\alpha(T_2-t)}} \\ w_2 &= \Delta_{F_A} \frac{F_A(t,T)}{F(t,T_2)} \frac{\dfrac{\sigma_S^A(t,T)}{\sigma_S} - \dfrac{\sigma_L^A(t,T)}{\sigma_L} e^{-\alpha(T_1-t)}}{e^{-\alpha(T_2-t)} - e^{-\alpha(T_1-t)}} \end{aligned} \tag{3.66}$$

6. Dependence of the Asian option price on the parameters

Carr, Ewald and Xiao (2008) establish that in the Black-Scholes framework, the premium of an arithmetic average Asian call option written on stock increases with volatility. They also show that this is not a trivial result and does not hold outside the Black-Scholes assumption, for example using a binomial model.

In the model presented here there is not a single volatility parameter but four parameters governing the term structure of volatility. We will study the dependence of the option premium on these four parameters. This has an important impact on option risk management, given that the implied term structure of volatility can change stochastically over time, thereby affecting prices.

The dependence of the option premium on the volatility parameters is through σ_{Black}, therefore we can write

$$\frac{\partial C}{\partial \sigma_S} = \frac{\partial C}{\partial \sigma_B}\frac{\partial \sigma_B}{\partial \sigma_S}, \quad \frac{\partial C}{\partial \sigma_L} = \frac{\partial C}{\partial \sigma_B}\frac{\partial \sigma_B}{\partial \sigma_L}, \ldots \tag{3.67}$$

We calculate the sensitivity of the Black volatility on the four parameters, in the simplified case when $\alpha c \ll 1$. In-settlement, for $T_M \le t < T_{M+1}$,

$$\sigma_B(t,T) \approx \frac{1}{\sqrt{3}}\sqrt{\sigma_S^2 + 2\rho\sigma_S\sigma_L + \sigma_L^2} \tag{3.68}$$

$$\frac{\partial \sigma_B}{\partial \sigma_S} = \frac{\sigma_S + \rho\sigma_L}{3\sigma_B}, \quad \frac{\partial \sigma_B}{\partial \sigma_L} = \frac{\sigma_L + \rho\sigma_S}{3\sigma_B}, \quad \frac{\partial \sigma_B}{\partial \alpha} = 0, \quad \frac{\partial \sigma_B}{\partial \rho} = \frac{\sigma_S\sigma_L}{3\sigma_B} \tag{3.69}$$

And pre-settlement, $t < T_1$,

$$\sigma_B(t,T)^2 \approx \sigma_S^2 \frac{c}{T-t}\left[\frac{1-e^{-2\alpha(T_1-t)}}{2\alpha c} + \frac{1}{3}\right] + 2\rho\sigma_S\sigma_L \frac{c}{T-t}\left[\frac{1-e^{-\alpha(T_1-t)}}{\alpha c} + \frac{1}{3}\right] + \sigma_L^2\left[1 - \frac{2}{3}\frac{c}{T-t}\right] \tag{3.70}$$

$$\begin{aligned}
\frac{\partial \sigma_B}{\partial \sigma_S} &= \frac{\sigma_S}{\sigma_B}\frac{c}{T-t}\left(\frac{1-e^{-2\alpha(T_1-t)}}{2\alpha c} + \frac{1}{3}\right) + \frac{\rho\sigma_L}{\sigma_B}\frac{c}{T-t}\left(\frac{1-e^{-\alpha(T_1-t)}}{\alpha c} + \frac{1}{3}\right) \\
\frac{\partial \sigma_B}{\partial \sigma_L} &= \frac{\rho\sigma_S}{\sigma_B}\frac{c}{T-t}\left[\frac{1-e^{-\alpha(T_1-t)}}{\alpha c} + \frac{1}{3}\right] + \frac{\sigma_L}{\sigma_B}\left[\frac{T_1-t}{c} + \frac{1}{3}\right] \\
\frac{\partial \sigma_B}{\partial \alpha} &= \frac{\sigma_S}{\sigma_B}\frac{c}{2\alpha(T-t)}\left[\sigma_S\left[\left(\frac{T_1-t}{c} + \frac{1}{2}\right)e^{-2\alpha(T_1-t)} - 1\right] + \rho\sigma_L\left[\left(\frac{T_1-t}{c} + \frac{1}{2}\right)e^{-\alpha(T_1-t)} - 1\right]\right] \\
\frac{\partial \sigma_B}{\partial \rho} &= \frac{\sigma_S\sigma_L}{\sigma_B}\frac{c}{T-t}\left[\frac{1-e^{-\alpha(T_1-t)}}{\alpha c} + \frac{1}{3}\right]
\end{aligned} \tag{3.71}$$

Let us examine the sign of these quantities. In-settlement, we have, on the condition that $\sigma_S, \rho, \sigma_L \geq 0$, all the derivatives are non-negative. For the pre-settlement values,

$$\frac{\partial \sigma_B}{\partial \sigma_S} = \frac{c}{T-t}\left(\frac{\sigma_S}{\sigma_B} g(2\alpha) + \frac{\rho \sigma_L}{\sigma_B} g(\alpha) \right) \tag{3.72}$$

where $g(\alpha) = \left(1 - e^{-\alpha(T_1 - t)}\right) / (\alpha c) + 1/3$ is decreasing in α, such that

$$\frac{\partial \sigma_B}{\partial \sigma_S} \geq \frac{c}{T-t} g(2\alpha) \frac{\sigma_S + \rho \sigma_L}{\sigma_B} \geq 0 \tag{3.73}$$

For σ_L,

$$\frac{\partial \sigma_B}{\partial \sigma_S} = \frac{c}{T-t}\left(\frac{\rho \sigma_S}{\sigma_B} g(\alpha) + \frac{\sigma_L}{\sigma_B} g(0) \right) \geq \frac{c}{T-t} g(\alpha) \frac{\rho \sigma_S + \sigma_L}{\sigma_B} \geq 0 \tag{3.74}$$

For α we can see that $\sigma_{Black}(t,T)^2$ is decreasing with α. The Black volatility is increasing in ρ without any conditions on the volatilities.

Hence we have proven that, if we assume $\sigma_S \geq 0, \rho \geq 0, \sigma_L \geq 0$, the Black volatility and the call option premium are increasing in the parameters σ_S, σ_L and ρ, and decreasing in α.

4. THE FLOATING STORAGE TRADE

1. Introduction

After the collapse of oil and shipping prices in mid 2008, floating storage, i.e. storing crude oil or products in idle tankers, became a viable opportunity. This was made possible by a steep contango of the crude oil forward curve combined with low freight rates.

This is only one example of what international oil trading consists of: if a price difference – in time or space – is higher than the shipping and capital cost involved, then there is an arbitrage opportunity that can be exploited. Two very basic examples are:

- Shipping crude oil from Nigeria to the US
- Storing crude oil in storage tanks in Cushing, OK when the WTI forward curve is in contango.

Often several opportunities present themselves to an oil trader, and the volatility of the associated forward curves makes it possible that these opportunities could evolve during the voyage. For example, a ship leaving Nigeria with crude oil has the option of going either to the United States or to Europe, and the trader doesn't necessarily have to make that choice immediately – he can choose to stay on a northward course in the mid-Atlantic and defer the choice of destination port to a later date, when one option will be significantly more interesting than the other. The ship can also choose to speed up or slow down to control its fuel consumption and arrive at an optimal date.

When choosing to keep his options open, the oil trader is exposed to movements in the forward curves. The existence of liquid futures and options markets at several key locations makes it possible to hedge this exposure partly, thereby reducing risks.

Our aim is to create a framework for analyzing such trading strategies and derive the optimal route that a ship should follow. This framework can then be used to evaluate expected return and risk beforehand, to find the optimal route that the ship should follow, and to derive hedge ratios to hedge the exposure to the dominant risk factors.

We will concentrate on a simple problem: the cross Atlantic crude oil arbitrage with possibility of floating storage, and present the results for this. We will then proceed to generalize the framework to a general optimal trading problem.

2. The floating storage problem

We consider a ship at location X at time t. The characteristics of the ship and the shipping route are given in Table 6.

The problem we are considering is as follows:

1. At date $t = 0$ the decision is made to load the tanker with a cargo at the spot price S_0

2. The cargo is loaded at date τ_{load}.

3. The ship then sails at the constant speed *u* to the destination port.

4. Upon arrival, and until the date when the cargo is actually delivered, the ship is anchored at the destination port

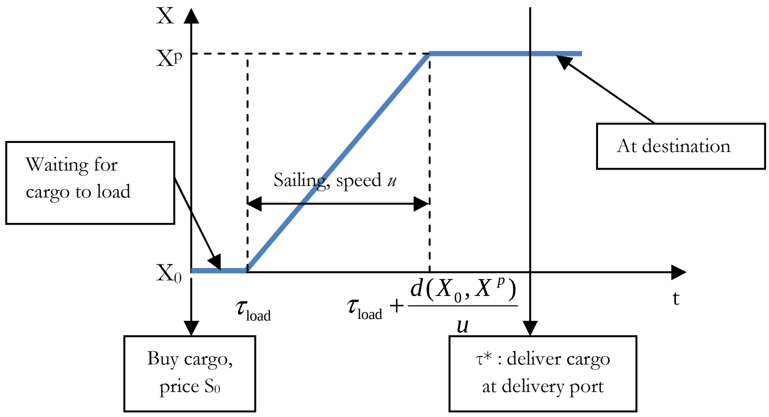

The daily cost paid for the shipping is then, at date *t*,

$$g(t)=\begin{cases}0 & t\le\tau_{load}\\ H+B\cdot FC(u) & \tau_{load}\le t\le\tau_{load}+d(X_0,X^P)/u\\ H+B\cdot FC_a & t\ge\tau_{load}+d(X_0,X^P)/u\end{cases}\qquad(4.1)$$

The cargo is not necessarily sold into the market immediately. At any location *X* and time *t* the decision can be made to sell the cargo for delivery τ days forward, at the price $F(t,\tau)$.

Exercise profit

We define the exercise profit $\Omega(t,F(\tau))$ as the profit that can be earned on the cargo if the ship is at location *X(t)* and the forward curve is $F(\tau)$, by *committing* to a specific delivery price sometime in the future and sailing to deliver at that time. In effect, the trader gives up the possibility of changing delivery time.

At exercise, one chooses a time-to-delivery τ. For one choice of this parameter the profit is

$$\begin{aligned}&\tau_{sail}=d(X(t),X^P)/u\\ &\omega(\tau)=F(\tau)-S_0-H\tau-B\cdot(FC(u)\tau_{sail}+FC_a(\tau-\tau_{sail}))-BH,\quad \tau\ge\tau_{sail}\end{aligned}\qquad(4.2)$$

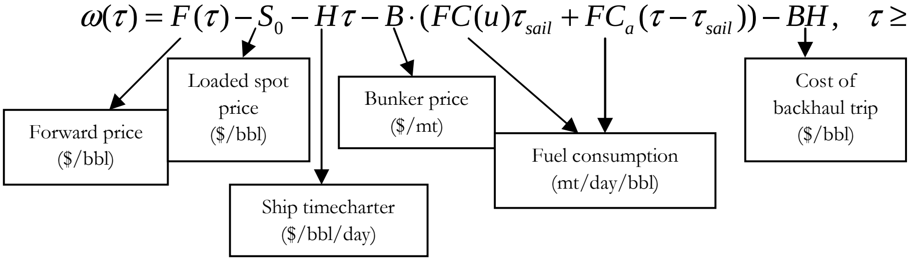

The exercise profit consists in maximizing ω over all possible times-to-maturity:

$$\Omega(t, F(\tau)) = \max_{\tau \geq \tau_{sail}} \omega(\tau) \tag{4.3}$$

Table 6. Ship and route parameters

	Unit	Typical value		Unit	Typical value
Location *X*	Nm		**Ship**[1]		
Time *t*	days		Type		Very Large Crude Carrier (VLCC)
			Cargo size	mt	270 000 mt
Route		Sullom Voe – LOOP	DWT	mt	300 000 mt
Distance *d*	Nm	4535 Nm	Speed *u*	knots	15 knots
Cargo		Brent	Fuel consumption sailing: FC(u) anchor: FC_a	mt/day mt/day	87.5 mt/day (laden) 74 mt/day (ballast) 85 mt/day (pumping) 15 mt/day (anchor)
Barrel factor	bbl/mt	7.578 bbl/mt	Timecharter price *H*	USD/day	VLCC average timecharter equivalent (Baltic Exchange)
Loading port	X_0	Sullom Voe	**Delivery port**	X_P	LOOP
Loading price S_0	USD/bbl	Dated Brent 10-21 days	Delivery price $F(t,\tau)$	USD/bbl	LLS forward curve
Loading delay τ_{load}	Days	15 days			
IFO price *B*	USD/mt	Fuel Oil 3.5% CIF NWE (Platts)			

It should be noted that if the location of the ship X is the loading port, then Ω is the arbitrage profit from that port and if it is positive, the arbitrage is said to be open. Furthermore, if the loading and destination ports are the same and Ω is positive, then there is a floating storage opportunity at that port and Ω is the profit that can be made from it.

This profit is also riskless – at least market-wise – considering that the profit is locked in by selling the cargo forward.

Valuing the expected profit of the voyage

When the tanker is loaded, the arbitrage profit Ω can be locked in without risk. However, the large number of optionalities available to the trader throughout the voyage means that the expected profit is sometimes higher.

[1] Typical values correspond to the modern double-hull VLCC from Clarksons (2009)

Let us define $V(t, F(\tau))$ as the expected profit from the cargo when the tanker is at the location $X(t)$ and the forward curve is given by $F(\tau)$. This is a real option value as detailed in Dixit and Pindyck (1994).

When the ship is at a location $X(t)$, and the maximum exposure time T has not been exceeded, the trader has two choices:

- either "exercise" and sell the cargo forward, thereby earning the exercise profit $\Omega(t, F(\tau))$
- or choose to continue speculating during a time dt without exercising. The expected profit is:

$$V_C(t, F(\tau)) = E_{F(\tau)}\left[V(t+dt, F(\tau)+dF(\tau))\right] - g(t)dt \tag{4.4}$$

This gives the continuation value V_C. The forward curves are evolved during the time period dt using the two-factor model from Part 2.

Hence the expected profit at location $X(t)$ given the forward curve $F(\tau)$ is

$$V(t, F(\tau)) = \max\left[\Omega(t, F(\tau)), \quad V_C(t, F(\tau))\right] \tag{4.5}$$

We notice that we always have $V \geq \Omega$ because the possibility of obtaining Ω is included in V. If the maximal exposure time $t = T$ is reached, the cargo must be sold and $V(t, F(\tau)) = \Omega(t, F(\tau))$.

This value function V contains all the information needed to evaluate and run the physical trade:

- the value $V(0, F_0(\tau))$ is the expected profit from following the optimal trading strategy
- the a priori risk of the strategy and its exposure to the principal risk factors can be evaluated through V, as seen in Section 3.5.
- at a date t, given the forward curve $F(\tau)$, compare the exercise value $\Omega(t, F(\tau))$ and the continuation value $V_C(t, F(\tau))$.
    - If $\Omega \geq V_C$ then the delivery of the cargo should be specified. The optimal date at which to deliver it is obtained from the calculation of Ω
    - If $V_C > \Omega$ then the ship should continue sailing or anchoring without specifying when delivery will take place.

Simplification in the case of a two-factor model

If the dynamics of the forward curve are described by a simple two-factor model as described in Part 2, the forward curve $F(\tau)$ can be expressed in terms of the factor values f_1 and f_2 and the initial forward curve $F_0(\tau)$:

$$\log F_t(\tau) = \log F_0(t+\tau) + \psi(t,\tau) + u_1(\tau) f_1(t) + u_2(\tau) f_2(t) \tag{4.6}$$

Therefore the functions Ω and V only depend on the values of f_1, f_2 and time *t* since the beginning of the trade:

$$\begin{aligned} \Omega(t, F(\tau)) &= \Omega(t, f_1, f_2) \\ V(t, F(\tau)) &= V(t, f_1, f_2) \end{aligned} \tag{4.7}$$

Optimal stopping formulation

Determining *V* can alternatively be seen as an optimal stopping problem. The value function can equivalently be written as

$$V(t, f_1, f_2) = \max_{\tau \in ST_{[t,T]}} E\left[-\int_t^{\tau} g(s)ds + \Omega(\tau, f_1(\tau), f_2(\tau)) \mid f_1(t) = f_1, f_2(t) = f_2 \right] \tag{4.8}$$

where $ST_{[t,T]}$ is the set of all stopping times in $[t,T]$. The optimal stopping time corresponds to the time when V becomes equal Ω, i.e.

$$\tau^* = \min\{ s \le T, V(s, f_1(s), f_2(s)) = \Omega(s, f_1(s), f_2(s) \} \tag{4.9}$$

and the initial expected profit from the trade is

$$V = E\left[-\int_0^{\tau^*} g(s)ds + \Omega(\tau^*, f_1, f_2) \right] \tag{4.10}$$

3. Solution methods

Solving an optimal stopping problem such as the one that has been formulated for the floating storage trade is akin to calculating the value of an American option. A number of numerical methods have been suggested to this effect and we will review some of them here.

Dynamic programming

The conceptually simplest method of solving an American option problem is by dynamic programming. Discretizing time into dates $t_0 = 0, t_1, ..., t_N = T$, the value at date t_i can be written as

$$\begin{aligned} &V(t_N, f_1, f_2) = \Omega(t_N, f_1, f_2) \\ &V_C(t_i, f_1, f_2) = E\left[V(t_{i+1}, f_1 + \Delta f_1, f_2 + \Delta f_2)\right] - g(t_i)\Delta t \\ &V(t_i, f_1, f_2) = \max\left[\Omega(t_i, f_1, f_2), \quad V_C(t_i, f_1, f_2)\right] \end{aligned} \tag{4.11}$$

The expectation in the calculation of $V_C(t_i, f_1, f_2)$ is calculated using the transition probabilities of f_1 and f_2. When a binomial distribution is assumed this yields the binomial tree method for American options, as detailed in Clewlow and Strickland (1998). Using the two-factor model presented here we can evaluate it using transition probabilities.

The factor value space is discretized into a $N_1 \times N_2$ rectangular grid: $(f_1^1,..., f_{N_1}^1) \times (f_1^2,..., f_{N_2}^2)$. The expectation is evaluated numerically using the previously calculated values of $V(t_{j+1}, f_k^1, f_l^2)$ and the joint probability density of (df^1, df^2):

$$\begin{aligned} &E_{f_k^1, f_l^2}\left[V(t_{j+1}, f_k^1 + df^1, f_l^2 + df^2)\right] \approx \frac{1}{NF} \sum_{k'=1}^{N_1} \sum_{l'=1}^{N_2} p_{k'l'} V(t_{j+1}, f_{k'}^1, f_{l'}^2) \\ &p_{k'l'} = \frac{1}{2\pi\sigma_1\sigma_2\Delta t} \exp\left[-\frac{(f_{k'}^1 - f_k^1 - (\mu_1 - \alpha_1 f_k^1)\Delta t)^2}{2\sigma_1^2\Delta t} - \frac{(f_{l'}^2 - f_l^2 - (\mu_2 - \alpha_2 f_l^2)\Delta t)^2}{2\sigma_2^2\Delta t}\right] \\ &NF = \sum_{k',l'} p_{k'l'} \end{aligned} \tag{4.12}$$

Partial Differential Equation

In continuous time, the dynamic programming formulation for V combined with Ito's formula yields a partial differential equation for V. Let us assume that $V_C > \Omega$, i.e. we are in the continuation region, such that $V = V_C$. We develop V using Ito's formula:

$$\begin{aligned} V(t+dt, f_1+df_1, f_2+df_2) &= V + \frac{\partial V}{\partial t}dt + \frac{\partial V}{\partial f_1}df_1 + \frac{\partial V}{\partial f_2}df_2 + \frac{1}{2}\sum_{1\le i,j\le 2} df_i \frac{\partial^2 V}{\partial f_i \partial f_j} df_j \\ &= V + \left[\frac{\partial V}{\partial t} + (\mu_1 - \alpha_1 f_1)\frac{\partial V}{\partial f_1} + (\mu_2 - \alpha_2 f_2)\frac{\partial V}{\partial f_2} + \frac{1}{2}\sigma_1^2 \frac{\partial^2 V}{\partial f_1^2} + \frac{1}{2}\sigma_2^2 \frac{\partial^2 V}{\partial f_2^2}\right]dt \\ &\quad + \frac{\partial V}{\partial f_1}\sigma_1 dW_1 + \frac{\partial V}{\partial f_2}\sigma_2 dW_2 \end{aligned} \tag{4.13}$$

Using the definition of V and V_C we find that

$$V(t, f_1, f_2) = V_C(t, f_1, f_2) = E\left[V(t+dt, f_1+df_1, f_2+df_2)\right] - g(t)dt \tag{4.14}$$

$$\frac{\partial V}{\partial t}+(\mu_1-\alpha_1 f_1)\frac{\partial V}{\partial f_1}+(\mu_2-\alpha_2 f_2)\frac{\partial V}{\partial f_2}+\frac{1}{2}\sigma_1^2\frac{\partial^2 V}{\partial f_1^2}+\frac{1}{2}\sigma_2^2\frac{\partial^2 V}{\partial f_2^2}=g(t) \tag{4.15}$$

This is valid for (t, f_1, f_2) in the continuation region such that $V(t, f_1, f_2) > \Omega(t, f_1, f_2)$. For (t, f_1, f_2) in the exercise region we have $V(t, f_1, f_2) = \Omega(t, f_1, f_2)$. If we define $S^*(t)$ as the continuation region at time *t*, this entails the following boundary conditions on *V*

$$\begin{aligned} &V(t, f_1, f_2)=\Omega(t, f_1, f_2), \quad \forall (f_1, f_2)\in \partial S^*(t) \quad \text{(continuity)} \\ &\nabla V(t, f_1, f_2)=\nabla\Omega(t, f_1, f_2), \quad \forall (f_1, f_2)\in \partial S^*(t) \quad \text{(smooth pasting)} \end{aligned} \tag{4.16}$$

Using these equations for *V* we can calculate *V* using finite differences or a semi-analytical formulation.

Semi-analytical solution

Following Albanese and Campolieti (2006), the partial differential equation for *V* can be solved in closed form if we assume the boundary of the continuation region to be known.

In Appendix 7 we show that the solution of (4.15) can be written as

$$\begin{aligned} V(t, f) &= \int_{\mathbb{R}^2} p(f',T; f,t)(\Omega(T, f')-G(t,T))df' + \int_t^T \int_{\mathbb{R}^2\setminus S^*(s)} p(f',s; f,t)(-\psi(s, f'))df'ds \\ &= V_{eur}(t, f)+V_{early}(t, f) \end{aligned} \tag{4.17}$$

where *T* is the maximum exposure time, and

$$\begin{aligned} &\psi(t, f)=\begin{cases} 0 & f\in S^*(\tau) \\ \dfrac{\partial\Omega}{\partial\tau}+L\Omega & f\notin S^*(\tau) \end{cases} \\ &\mathrm{L}\,\Omega=(\mu_1-\alpha_1 f_1)\frac{\partial\Omega}{\partial f_1}+(\mu_2-\alpha_2 f_2)\frac{\partial\Omega}{\partial f_2}+\frac{1}{2}\sigma_1^2\frac{\partial^2\Omega}{\partial f_1^2}+\frac{1}{2}\sigma_2^2\frac{\partial^2\Omega}{\partial f_2^2}-g(t) \end{aligned} \tag{4.18}$$

The continuation region $S^*(t)$ is defined as

$$S^*(t)=\left\{(f_1, f_2), V(t, f_1, f_2)>\Omega(t, f_1, f_2)\right\} \tag{4.19}$$

The boundary $\partial S^*(t)$ of this domain has to be determined for each date *t*. We write it as a function of the second factor

$$\partial S^*(t)=\left\{(f_1^*(t, f_2), f_2), f_2\in\mathbb{R}\right\} \tag{4.20}$$

such that the equation to be solved by $f_1^*(t, f_2)$ is

$$\begin{aligned} V(t, f_1^*(\tau, f_2), f_2) &= \int_{\mathbb{R}^2} p(f_1', f_2', T; f_1^*(t, f_2), f_2, t)(\Omega(T, f') - G(t,T))df' \\ &\quad + \int_t^T \int_{-\infty}^{+\infty} \int_{f_1^*(s,f_2)}^{+\infty} p_1(f_1', f_2', s; f_1^*(t, f_2), f_2, t)(-\psi(s, f'))df' ds \\ &= \Omega(t, f_1^*(t, f_2), f_2) \end{aligned} \tag{4.21}$$

In Appendix 7 we detail the numerical procedure used to find this exercise boundary, which can be found recursively beginning at $t = T$.

Monte Carlo simulation methods

The methods discussed above, while suitable for a two-factor model, become impractical if the number of factors is higher, for example if several ports are being considered. In this case a Monte Carlo method should be employed. Monte Carlo methods are not ideally suited to American option problems, because of their backward-recursion properties. However, Longstaff and Schwartz (2001) suggest a least-squares Monte Carlo method with projection of the value function onto a small basis, allowing for efficient pricing of American options. A similar method could be employed in this case.

4. Analytical properties of the solution

We will examine some of the properties of the expected profit function V using the analytical expression obtained above. We will decompose the solution as follows:

$$\begin{aligned} V(0, f) &= \int_{\mathbb{R}^2} p(f', T; f, 0)(\Omega(T, f') - G(0,T))df' + \int_0^T \int_{\mathbb{R}^2 \setminus S^*(s)} p(f', s; f, 0)(-\psi(s, f'))df' ds \\ &= \Omega(0, f) + EP_{\text{drift}} + EP_{\text{convexity}} + EP_{\text{early}} \end{aligned} \tag{4.22}$$

where the excess profit components EP_{drift} , $EP_{\text{convexity}}$ and EP_{early} are defined as

$$\begin{aligned} EP_{\text{drift}} &= \Omega(T, Ef_1(T), Ef_2(T)) - G(T) - \Omega(0, f_1, f_2) \\ EP_{\text{convexity}} &= V_{\text{eur}}(T, f_1, f_2) - (\Omega(T, Ef_1(T), Ef_2(T)) - G(T)) \\ EP_{\text{early}} &= V_{\text{early}}(T, f_1, f_2) \end{aligned} \tag{4.23}$$

Model parameter dependence

We want to examine the dependence of the value function on the model parameters α_k, σ_k and μ_k.

Let us consider first the drift component EP_{drift}. Its value does not depend on the volatilities of the factors. Remembering that

$$Ef_k(T) = f_k e^{-\alpha_k T} + \frac{\mu_k}{\alpha_k}\left(1 - e^{-\alpha_k T}\right) \tag{4.24}$$

we can derive its dependence on the drift and on the model parameters:

$$\begin{aligned}
\frac{\partial EP_{\text{drift}}}{\partial \mu_k} &= \frac{1-e^{-\alpha_k T}}{\alpha_k}\frac{\partial \Omega}{\partial f_k}(T, Ef_1(T), Ef_2(T)) \\
\frac{\partial EP_{\text{drift}}}{\partial \alpha_k} &= \frac{\partial Ef_k(T)}{\partial \alpha_k}\frac{\partial \Omega}{\partial f_k}(T, Ef_1(T), Ef_2(T)) \\
\frac{\partial EP_{\text{drift}}}{\partial \sigma_k} &= 0
\end{aligned} \tag{4.25}$$

The convexity component can be written as

$$\begin{aligned}
EP_{\text{convexity}} &= \int_{\mathbb{R}^2} p(f',T; f,0)(\Omega(T, f') - \Omega(T, Ef_1(T), Ef_2(T)))df' \\
&= \int_{\mathbb{R}^2} p(f_1' - Ef_1(T), f_2' - Ef_2(T))\left(\Omega(T, f_1', f_2') - \Omega(T, Ef_1(T), Ef_2(T))\right)df' \\
&= \int_{\mathbb{R}^2} p(f_1', f_2')\left(\Omega(T, Ef_1(T) + f_1', Ef_2(T) + f_2') - \Omega(T, Ef_1(T), Ef_2(T))\right)df'
\end{aligned} \tag{4.26}$$

(in the last line we change variables from f' to $f' + Ef(T)$). We want to show that this convexity premium does not depend on the expected value of the factors.

$$\frac{\partial EP_{\text{convexity}}}{\partial Ef_k} = \int_{\mathbb{R}^2} p(f_1', f_2')\left(\frac{\partial \Omega}{\partial f_k}(T, Ef_1(T) + f_1', Ef_2(T) + f_2') - \frac{\partial \Omega}{\partial f_k}(T, Ef_1(T), Ef_2(T))\right)df' \tag{4.27}$$

This term is zero if Ω is at most quadratic in the factors. In the general case we can write

$$\frac{\partial \Omega}{\partial f_k}(Ef + f') - \frac{\partial \Omega}{\partial f_k}(Ef) = f' \cdot \nabla\left(\frac{\partial \Omega}{\partial f_k}\right) + O(f'^2) \tag{4.28}$$

such that:

$$\left|\frac{\partial EP_{\text{convexity}}}{\partial Ef_k}\right| \leq C\int_{\mathbb{R}^2} p(f_1^{'}, f_2^{'})(f_1^{'2} + f_2^{'2})df^{'} = C\left(\sigma_1^2\frac{1-e^{-\alpha_1 T}}{\alpha_1} + \sigma_2^2\frac{1-e^{-\alpha_2 T}}{\alpha_2}\right) \tag{4.29}$$

For sufficiently small values of T this term is small, of the order $O(\sigma_1^2 T)$. Hence we have established that the excess profits coming from convexity are indeed independent of the expected values of the factors, and therefore also of the drifts.

Early exercise in a backwardated market

As we will see in Section 4.6, it is generally optimal to sell the cargo immediately when the forward curve net of freight is in backwardation. We will show this result here.

We assume that the loading and delivery port are the same, such that the trade is purely a floating storage trade. The forward curve net of freight is in net backwardation if

$$F_0(\tau) - (H + FC_a)\tau < F_0(0) \tag{4.30}$$

In this case it is optimal to sell the cargo spot, such that

$$\Omega(t) - \int_0^t g(s)ds - \Omega(0) = F_0(t) - (H + FC_a)t - F_0(0) < 0 \tag{4.31}$$

The initial expected excess profit is, as seen in equation (4.10),

$$V(0,0,0) = E\left[-\int_0^{\tau^*} g(s)ds + \Omega(\tau^*, f_1, f_2)\right] \tag{4.32}$$

such that

$$V - \Omega = E\left[F_0(\tau^*) - (H + FC_a)\tau^* - F_0(0)\right] + E\left[F_0(\tau^*)\left(\exp(u_1(0)f_1(\tau^*) + u_2(0)f_2(\tau^*)) - 1\right)\right] \tag{4.33}$$

Conditional on $\tau^* = \tau$ (independent of the values of f_1 and f_2, we then find that given the distributions of $f_1(\tau)$ and $f_2(\tau)$,

$$\begin{aligned} &E\left[F_0(\tau^*)\left(\exp(u_1(0)f_1(\tau^*) + u_2(0)f_2(\tau^*)) - 1\right)\right] = \\ &\quad F_0(\tau^*)\left(\exp\left(u_1(0)E[f_1(\tau^*)] + u_2(0)E[f_2(\tau^*)] + \frac{u_1(0)^2\sigma_1^2(1-e^{-\alpha_1\tau^*})}{2\alpha_1} + \frac{u_2(0)^2\sigma_2^2(1-e^{-\alpha_2\tau^*})}{2\alpha_2}\right) - 1\right) \end{aligned} \tag{4.34}$$

We now introduce the initial slope of the curve $a = \left.\frac{\partial F_0}{\partial \tau}\right|_{\tau=0}$ and consider only small values of τ^*, then:

$$V - \Omega = F_0(0)\left[\frac{a-g}{F_0(0)} + u_1(0)\mu_1 + u_2(0)\mu_2 + \frac{1}{2}(u_1(0)^2\sigma_1^2 + u_2(0)\sigma_2^2)\right]\tau^* + O(\tau^{*2}) \qquad (4.35)$$

Hence, if

$$\frac{a-g}{F_0(0)} + u_1(0)\mu_1 + u_2(0)\mu_2 + \frac{1}{2}(u_1(0)^2\sigma_1^2 + u_2(0)\sigma_2^2) < 0 \qquad (4.36)$$

there is no value to exercising later, such that $V = \Omega$. The other parameters being fixed, this can always be achieved for a strong enough net backwardation.

5. Profit and risk

The calculation of the functions V and Ω defines a physical trading strategy that can be applied in practice. In order to assess how interesting this strategy is, we would like to assess a priori its expected return and risk. Furthermore we would like to assess the dependence of the profits on the different risk factors, to define a financial hedging strategy using futures or options.

Expected and realized profit on the trading strategy

As discussed above, the values $V(t=0, f_1=0, f_2=0)$ and $\Omega(t=0, f_1=0, f_2=0)$ are respectively the expected profit and arbitrage profit that can be obtained on the initial date. These numbers are expressed in US dollars per barrel ($/bbl) for crude oil or US dollars per gallon ($/gal) for gasoil.

When the trading strategy is executed the realized profit is not necessarily equal to the expected profit, given that the distribution of forward curves is stochastic. The realized profit of the trip is, in the simple case of a single port,

$$W = \int_0^{t^*} -g(t)dt + \Omega(t^*, f_1(t^*), f_2(t^*)) \qquad (4.37)$$

where t^* is the exercise date. This can be calculated a posteriori to get the realized profit. But seen at $t = 0$ this is a random variable with a certain distribution. Its expected value is *V*:

$$V(t=0, f_1^0, f_2^0) = E\left[W(\tau^*, f_1(\tau^*), f_2(\tau^*)) \mid f_1(t=0) = f_1^0, f_2(t=0) = f_2^0\right] \qquad (4.38)$$

where τ^* is the optimal stopping time, which is a random variable depending on the realized values of f_1 and f_2.

Expected risk and Sharpe ratio

There is no market risk tied to the arbitrage profit *Ω*, because the cargo is sold forward and the profit is fixed at the moment the decision is taken. However there is financial risk tied to the physical trading strategy with expected profit V: the forward curves will change before the decision to deliver the cargo is taken. This risk is reflected in the distribution of the realized profit W. We have seen that this distribution is centered on V:

$$V(t=0, f_1^0, f_2^0) = E\left[W(\tau^*, f_1(\tau^*), f_2(\tau^*)) \mid f_1(t=0) = f_1^0, f_2(t=0) = f_2^0\right] \tag{4.39}$$

Furthermore, at exercise,

$$W(t^*, f_1, f_2) = \int_0^{t^*} -g(t)dt + \Omega(t^*, f_1, f_2) = \int_0^{t^*} -g(t)dt + V(t^*, f_1, f_2) \tag{4.40}$$

We define the process U representing the expected profit and loss (P&L) on the trade at time t by

$$U(t, f_1, f_2) = \int_0^{t} -g(s)ds + V(t, f_1, f_2) \tag{4.41}$$

The value of this process at exercise equals W, the realized P&L of the trade. To find its distribution we differentiate U using Ito's formula:

$$\begin{aligned} dU &= \left(-g(t) + \frac{\partial V}{\partial t} + (\mu_1 - \alpha_1 f_1)\frac{\partial V}{\partial f_1} + (\mu_2 - \alpha_2 f_2)\frac{\partial V}{\partial f_2} + \frac{1}{2}\sigma_1^2 \frac{\partial^2 V}{\partial f_1^2} + \frac{1}{2}\sigma_2^2 \frac{\partial^2 V}{\partial f_2^2}\right)dt + \frac{\partial V}{\partial f_1}\sigma_1 dW_1 + \frac{\partial V}{\partial f_2}\sigma_2 dW_2 \\ &= \frac{\partial U}{\partial f_1}\sigma_1 dW_1 + \frac{\partial U}{\partial f_2}\sigma_2 dW_2 \end{aligned} \tag{4.42}$$

We find that the process U has zero drift. This reflects the fact that V was correctly priced initially. The instantaneous volatility of U over a time period dt is

$$\sigma_U = \left[\left(\sigma_1 \frac{\partial U}{\partial f_1}\right)^2 + \left(\sigma_2 \frac{\partial U}{\partial f_2}\right)^2\right]^{1/2} = \left[\left(\sigma_1 \Delta_1\right)^2 + \left(\sigma_2 \Delta_2\right)^2\right]^{1/2} \tag{4.43}$$

where

$$\Delta_k = \frac{\partial V}{\partial f_k} \tag{4.44}$$

is the delta of the value function with respect to factor k.

Furthermore, the distribution of U given the stopping time τ^* is

$$U(\tau^*, f_1, f_2) = U(0, f_1^0, f_2^0) + \int_0^{\tau^*} \frac{\partial U}{\partial f_1}(t, f_1(t), f_2(t))\sigma_1 dW_1(t) + \int_0^{\tau^*} \frac{\partial U}{\partial f_2}(t, f_1(t), f_2(t))\sigma_2 dW_2(t) \tag{4.45}$$

We can calculate the first moments of U:

$$E[U(\tau^*, f_1, f_2)] = U(0, f_1^0, f_2^0) = V(0, f_1^0, f_2^0) \tag{4.46}$$

$$\mathrm{Var}\left[U(\tau^*, f_1, f_2)\right] = \sigma_1^2 E\left[\int_0^{\tau^*}\left(\frac{\partial V}{\partial f_1}(t, f_1(t), f_2(t))\right)^2 dt\right] + \sigma_2^2 E\left[\int_0^{\tau^*}\left(\frac{\partial V}{\partial f_2}(t, f_1(t), f_2(t))\right)^2 dt\right] \tag{4.47}$$

The variance depends on the stopping time and can best be evaluated through Monte Carlo simulation, simulating the paths of (f_1, f_2) and using the value function already calculated.

If we make the simplifying assumption that the deltas of the value function are constant, the variance can be approximated as:

$$\mathrm{Var}[W] = \mathrm{Var}\left[U(\tau^*, f_1, f_2)\right] \approx \left(\sigma_1^2\Delta_1^2 + \sigma_2^2\Delta_2^2\right)E[\tau^*] \tag{4.48}$$

Alternatively, the complete distribution of *W* can be evaluated using Monte Carlo simulation. It should be noted that this Monte Carlo simulation is simpler than the least squares Monte Carlo technique used for finding the optimal stopping time.

If we know the expected profit and the standard deviation of the realized profit, we can calculate the annualized Sharpe ratio of the strategy a priori:

$$\begin{aligned} SR &= \frac{\text{Expected profit}}{\text{Std. deviation}} \\ &= \frac{1}{E[\tau^*]^{1/2}} \frac{V}{\left(\sigma_1^2 E\left[\int_0^{\tau^*}\left(\frac{\partial V}{\partial f_1}(t, f_1(t), f_2(t))\right)^2 dt\right] + \sigma_2^2 E\left[\int_0^{\tau^*}\left(\frac{\partial V}{\partial f_2}(t, f_1(t), f_2(t))\right)^2 dt\right]\right)^{1/2}} \\ &\approx \frac{V}{\left(\sigma_1^2\Delta_1^2 + \sigma_2^2\Delta_2^2\right)^{1/2} E[\tau^*]} \end{aligned} \tag{4.49}$$

6. Results

Arbitrage results

In this section we present the results from the calculation of the function Ω at different dates. This function, evaluated at trade initiation time (*t = 0*), gives the arbitrage profit that can be obtained from the shape of the forward curve at the current date. By studying its dependence on the factor values f_1 and f_2, we can also evaluate its dependence on the level and slope of the curve.

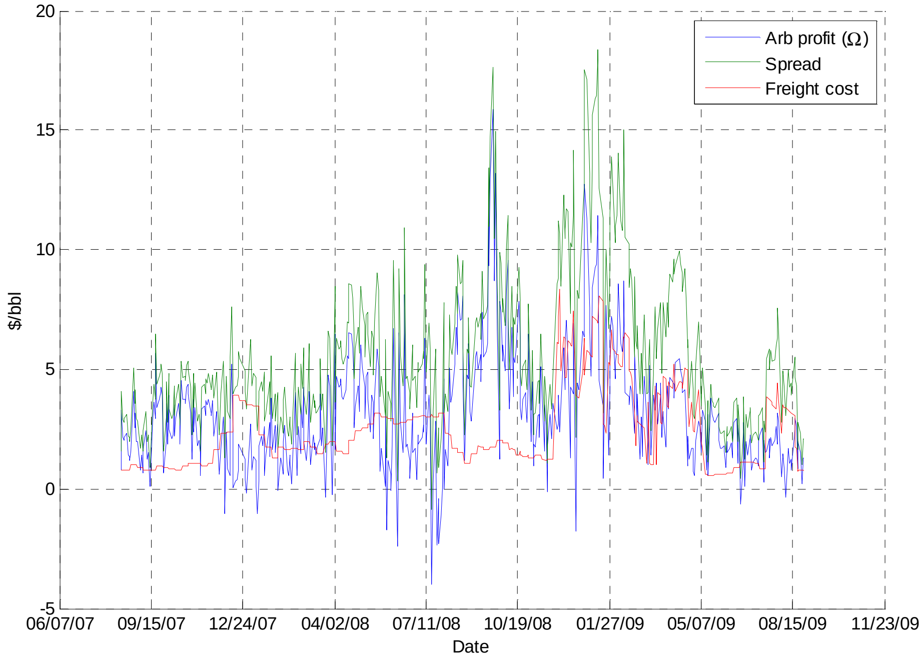

Figure 17. Arbitrage profit per barrel on the Sullom Voe-LOOP route

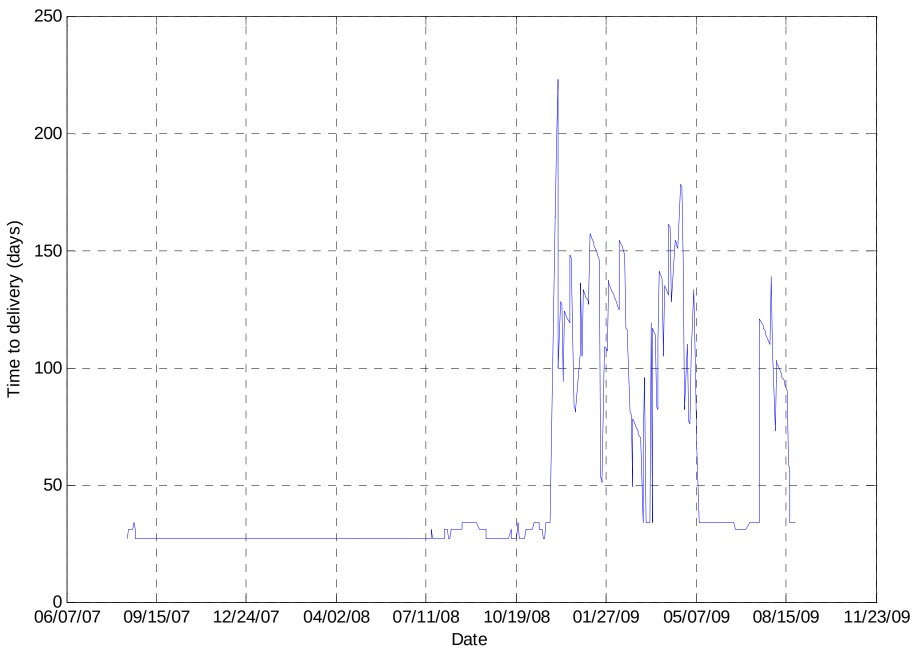

Figure 18. Time to delivery for static arbitrage on the Sullom Voe – LOOP route

Figure 19. Spot prices of BFOE @ Sullom Voe (blue) and LLS @ LOOP (green)

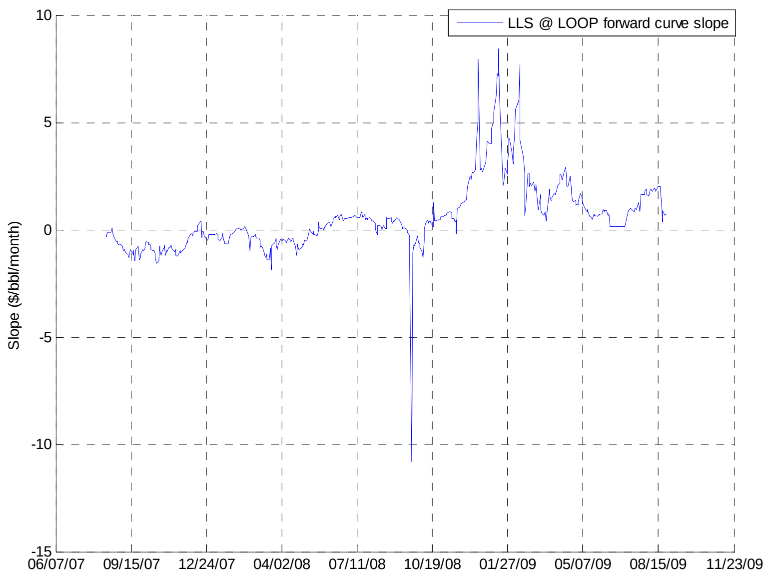

Figure 20. Slope of the LLS forward curve, in US Dollars per barrel per month, measured on the front two month contracts

On each day in the sample period (August 2007 to October 2009) we calculate the arbitrage profit Ω that can be obtained from the observed forward curve and freight prices on that date. The arbitrage profit is calculated as $\Omega = \max_{\tau \geq \tau_{sail}} \omega(\tau)$ where

$$\begin{aligned} &\tau_{sail} = d(X_0, X^P)/u \\ &\omega(\tau) = F(\tau) - S_0 - H\tau - B\cdot(FC(u)\tau_{sail} + FC_a(\tau - \tau_{sail})) - BH, \quad \tau \geq \tau_{sail} \end{aligned} \tag{4.50}$$

The arbitrage profit is obtained by buying BFOE crude at Sullom Voe on the trade initiation date and delivering it into LOOP in the optimal time τ^*, sailing at speed *u* to get there, and anchoring up for a time $\tau - \tau_{sail}$ to wait for delivery. The cargo is bought at the spot price S_0 and sold forward at the price $F(\tau^*)$.

We decompose this arbitrage profit into a geographical spread $\mathrm{Sp}(\tau) = F(\tau) - S_0$ and a freight cost Fc,

$$\mathrm{Fc}(\tau) = H\tau + B\cdot(FC(u)\tau_{sail} + FC_a(\tau - \tau_{sail})) + BH \tag{4.51}$$

The physical arbitrage is said to be *open* when $\Omega > 0$, i.e. $\mathrm{Sp}(\tau^*) > \mathrm{Fc}(\tau^*)$: the spread that can be earned on the crude oil is higher than the cost of transportation and storage. In the opposite case it is said to be closed. When the arbitrage window is open it is profitable to ship a cargo of oil on the considered route.

We present the results for the Sullom Voe-LOOP route in Figure 17. The geographical spread is shown in green, the freight cost in red, and the net arbitrage profit Ω in blue.

We observe that the arbitrage window is open during large parts of the time period under consideration. In Figure 18 we show the value of τ^*, the time between the current date and the optimal date to exercise the option of delivering the cargo. We observe that the large profits from late 2008 and early 2009 came from the large opportunities in floating storage created by a steep contango and low timecharter rates.

In Figure 19 we show the spot prices of crude oil at the loading and delivery ports. In Figure 20 we show the slope of the forward curve at the delivery port. Figure 19 confirms what makes this physical arbitrage possible: the spread in spot prices between European and American crude. However, the arbitrage profits seem to be uncorrelated with the general level of crude prices. This stems from the fact that international crude prices largely move together, partly because of such arbitrage activity. There does, however, seem to be some relation between the forward curve slope and the arbitrage profit. This is witnessed in Figure 21 where we regress the arbitrage profit on the forward curve slope. The relationship is stronger for a steeper contango.

In Figure 22 we present data collected from different research reports on the actual crude oil in floating storage worldwide alongside the optimal time to delivery for the arbitrage trade. We see that there is a substantial increase in the amount of crude oil stored at sea starting in October 2008. This coincides with the appearance of floating storage opportunities according to our model. Furthermore, the short disappearance of floating storage opportunities according to our model in June 2009 was accompanied by a clear downward trend in the number of tankers storing crude in the Goldman Sachs and Gibson Research data.

The same analysis can be performed for different markets and different routes. As a point of comparison we present the results for the gasoil arbitrage between Europe and the United States. The product being traded is No. 2 fuel oil, also known as gasoil or heating oil. The loading port is the Amsterdam-Rotterdam-Antwerp (ARA) region, Europe's major refining hub. The destination is New York harbor (NYH), which is the main delivery point of refined products on the east coast of the United States. Details of the route and cargo are presented in Appendix 8.

The arbitrage profit, decomposed as described above, is presented in Figure 23. The optimal delivery time is presented in Figure 24 and these results should be compared to the spot prices in Figure 25 and the forward curve slope in Figure 26.

We note that the arbitrage window is open less frequently than was the case for crude oil and the spread has been negative on occasions, making the inverse arbitrage (U.S. to Europe) interesting. However, there have been significant floating storage opportunities since the end of 2008 as witnessed in Figure 24, and these profits have been very interesting: around 20 cents per gallon for a gallon costing less than 2 dollars.

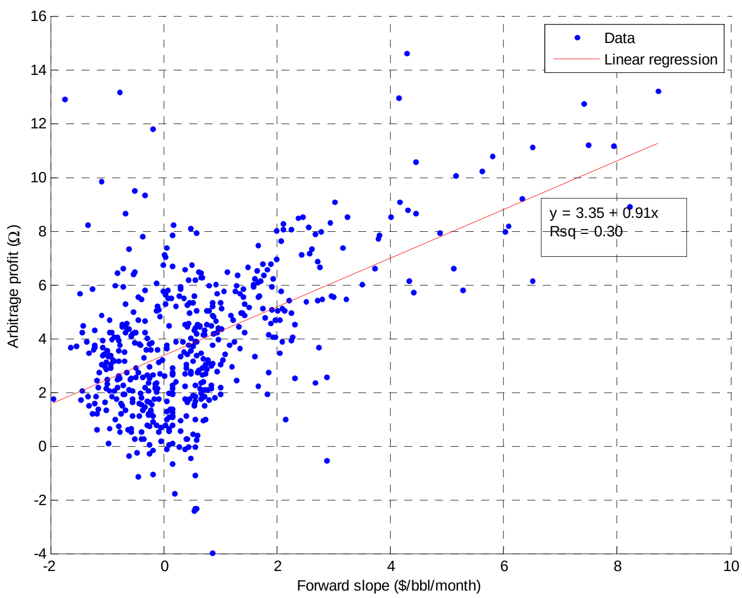

Figure 21. Relationship between the forward curve slope at the delivery port (LLS @ LOOP) and the arbitrage profit

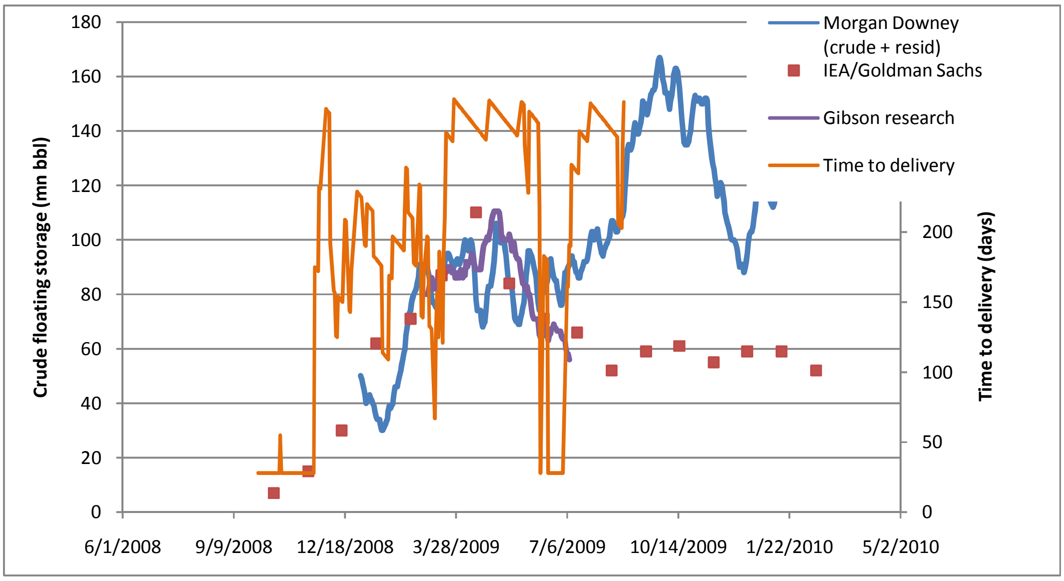

Figure 22. Crude oil in floating storage worldwide (left axis) and optimal time to delivery of the floating storage trade (right axis).

Sources: IEA/Goldman Sachs Global ECS Research, Gibson Research, Morgan Downey

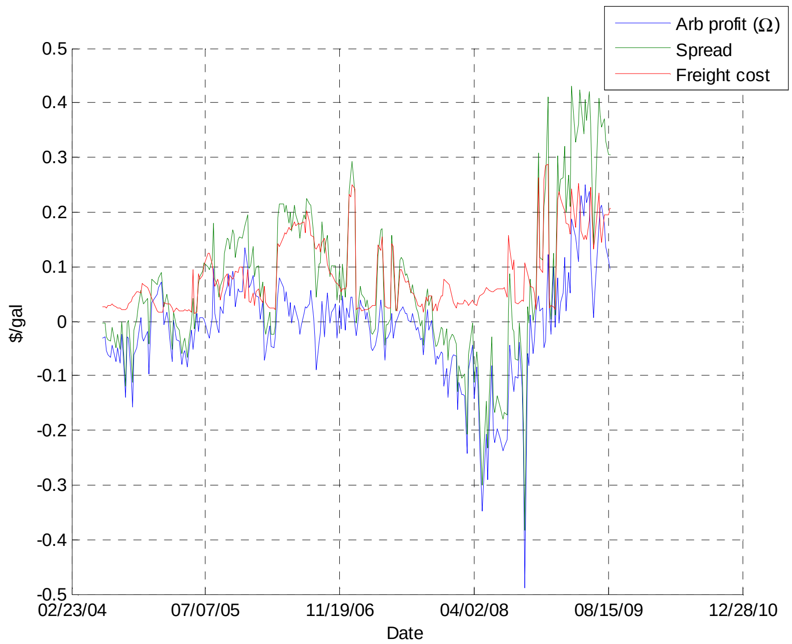

Figure 23. Arbitrage profit per gallon for gasoil trade between ARA and NYH

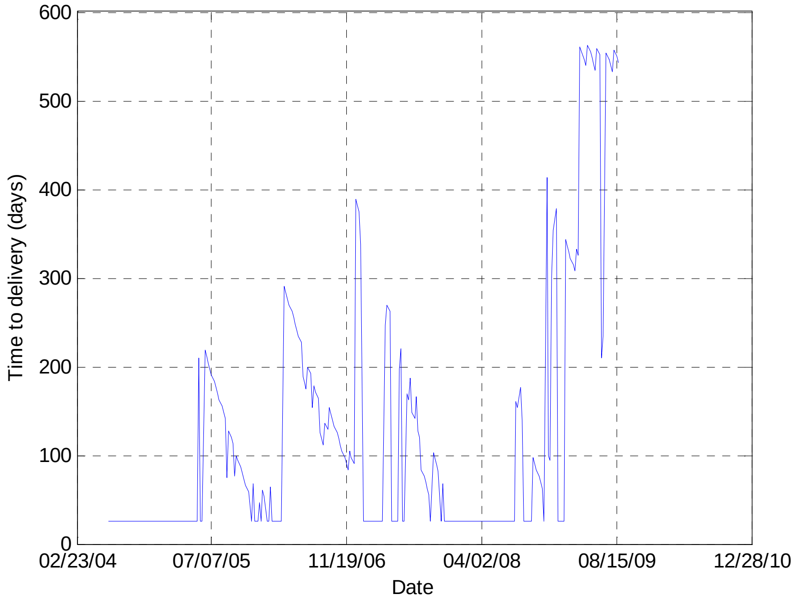

Figure 24. Time to delivery for arbitrage on the ARA-NYH route

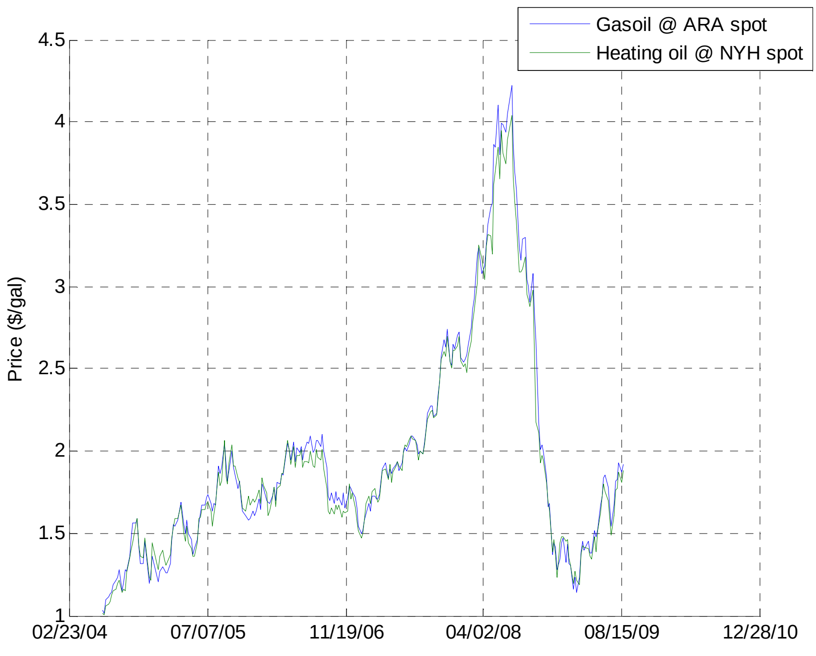

Figure 25. Spot prices of No. 2 fuel oil at Amsterdam-Rotterdam-Antwerp (blue) and New York harbor (green), in US Dollars per gallon

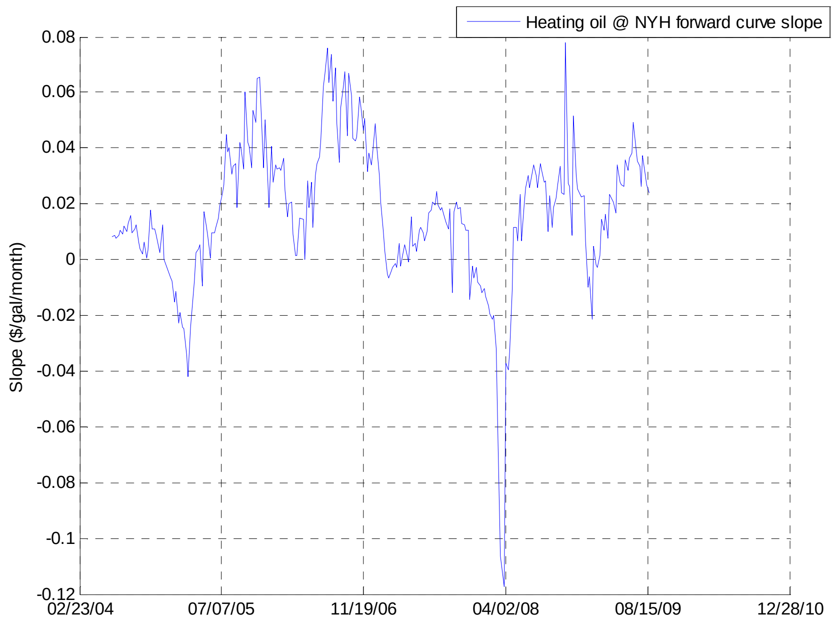

Figure 26. Slope of the heating oil forward curve at New York Harbor, in US Dollars per gallon per month

Expected and excess profits

The results presented for Ω were *static* arbitrage results. We now consider the optimal trading strategy presented in Section 4.2. This trading strategy yields a value function V which is the expected profit of the physical trading strategy.

These results are obtained using the semi-analytical formulation presented in Section 4.3, calculating the exercise boundary numerically as described in Appendix 7. The two-factor model used is calibrated on the crude oil futures market as described in Section 2.4, and we make the assumption that drifts are zero: the expected spot price is therefore equal to the forward price. Furthermore, trades are limited to a maximal exposure time T equal to 100 days.

We study the shape of Ω and V with the initial date set to December 18, 2008. As seen in Figure 27 the forward curve on that date was in contango. We plot Ω and V as functions of the factor values f_1 and f_2 at trade initiation.

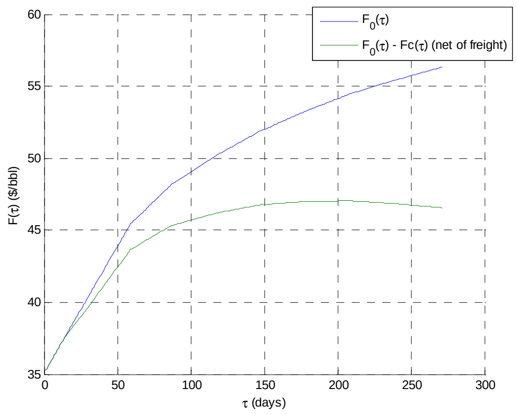

Figure 27. Forward curve and forward curve net of freight on December 18, 2008

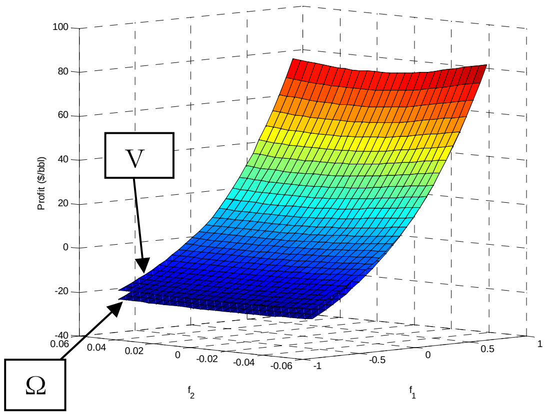

Figure 28. Ω and V as functions of the factors f_1 and f_2 at trade initiation (December 18, 2008)

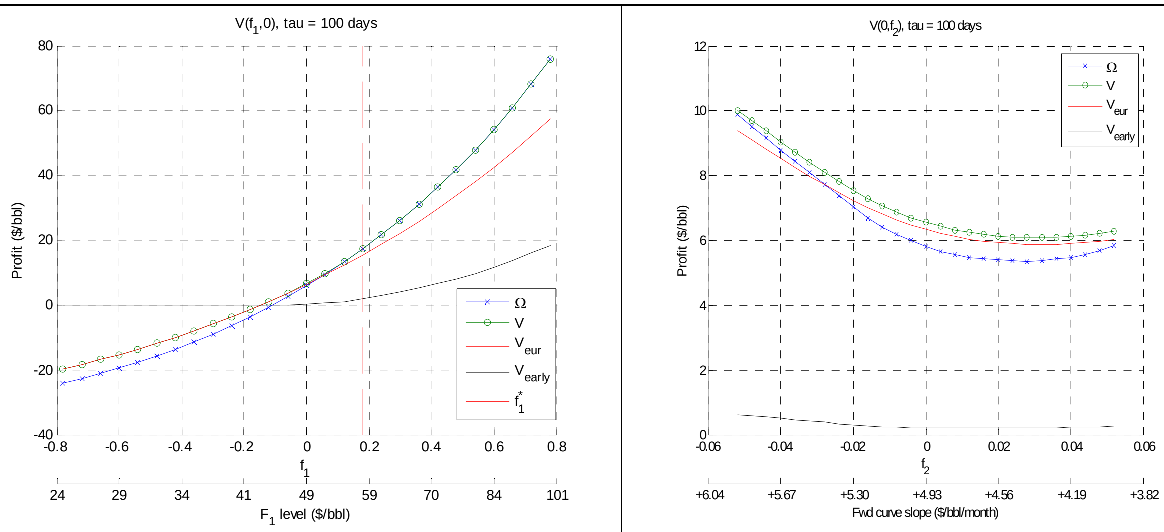

Figure 29. Cross-section of Ω and V at trade initiation as a function of f_1 (left) and f_2 (right)

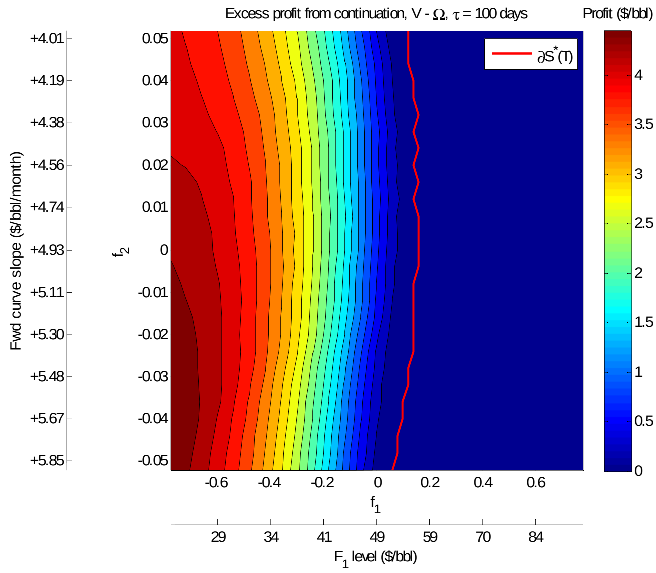

Figure 30. Expected excess profit from continuation V - Ω at trade initiation, in USD/bbl, and exercise boundary (red line)

In Figure 28 and Figure 29 we show the dependence of V and Ω on f_1 and f_2. The values corresponding to the initial forward curve are $V(0,0,0)$ and $\Omega(0,0,0)$, valued respectively at 6.55 $/bbl and 5.81 $/bbl. The other values correspond to a forward curve that has been shocked by the factor values f_1 and f_2. We can see that, predictably, a positive parallel shift ($f_1 > 0$) yields a higher expected profit. The slope with respect to the second factor is lower.

We can also see that V and Ω are the same at the maximal exposure time: this is the terminal condition that we impose. Furthermore, at trade initiation V is higher than Ω, and more so for low values of Ω. Thus there is value to keeping the options open. For negative values of Ω it is still possible to have positive values of V: there is a chance that prices will rebound enough to yield a profit during the trade period.

Figure 29 decomposes the value function V into two components: the European exercise value V_{eur} and the early exercise premium V_{early}. The European exercise value corresponds to the expected profit that would be earned if the cargo was held until the maximal exposure time T (100 days in this case), and then sold into the market. This value largely depends on the drifts of the factors. The mean-reverting model has a large impact in this respect. When the value of f_1 is negative, it is expected to increase, which pushes the expected value up compared to the arbitrage value. When the value of f_1 is positive, its expected value is lower, pushing the expected value down.

This is counteracted by the early-exercise premium, which is positive for high values of the first factor and for high absolute values of the second factor. This corresponds to situations where it is close to optimal to exercise.

The decision to exercise is made based on the difference between the expected profit V and the exercise profit Ω. We plot this difference in Figure 30. The darkest zone, where $V = \Omega$, is the exercise region. If (f_1, f_2) falls in this zone it is optimal to specify delivery of the cargo and harvest the profit Ω. Outside this region it is optimal to continue sailing and delay the decision about delivery time until later. The excess profit is seen to depend on the shape of the forward curve through the factor values f_1 and f_2. The excess profit is seen to be highest when the first factor is lowest: because it is mean-reverting, keeping the position open gives more upside exposure than downside exposure.

Dependence of expected profits on model parameters

The results above are presented for a two-factor model that has been calibrated on the crude oil market as detailed in Part 2. We have seen that the interpretation of the excess profits is linked to the model parameters α, σ and μ which determine the distribution of possible forward curves. It is therefore interesting to examine the dependence of the profits on the values of these parameters. We vary the parameters within reasonable ranges around the reference values that have been used before, and plot the dependence of V, V_{eur}, V_{early} and Ω on these parameters. The trade initiation date is December 18, 2008.

The results are presented in Figure 31. The dependence on the mean-reversion parameter is rather weak compared with the dependence on the other parameters. This can be explained by the fact that the maximal exposure time, 100 days, is rather short, and that the excess profit we are considering is taken at $f_1^0 = f_2^0 = 0$, such that the expected value of the factors is not affected by the mean-reversion parameter.

The dependence on the volatilities of the factors is very strong, with a doubling of σ_1 from 26% to 52% taking the excess expected profit from 0.68 \$/bbl to 2.19 \$/bbl. The effect is larger in the first factor because the magnitude of the excess profits coming from the first factor are much larger. But in relative terms, doubling σ_2 from 1.86% to 3.72% takes the excess profit attributable to the second factor, i.e. the difference between the profit for a σ_2 larger than zero and the profit for $\sigma_2 = 0$, from 7.78 c/bbl to 26.14 c/bbl, which is a significant increase.

The effect of the drift parameters μ_1 and μ_2 is to change the expectations about what the forward curve will look like in the future. In particular, a negative value for μ_1 means that the trader is taking a sharply negative view on the future level of prices. In that case it is more interesting to exercise early to take the profits given the current level of prices. A positive value for μ_1 is a positive view on levels and it will be preferable to wait to take advantage of rising prices. A negative value for μ_2 corresponds to a view of a sharper contango, which is beneficial to the trade, while more backwardation ($\mu_2 > 0$) is detrimental.

Figure 31. Dependence of expected profits on the model parameters

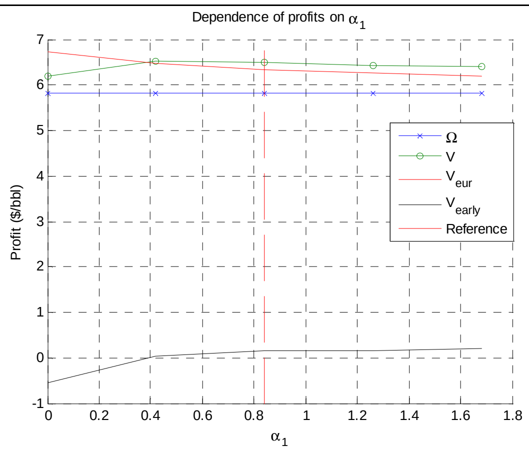

Dependence of profits on α_1. Reference is $\alpha_1 = 0.84$ yr^{-1}

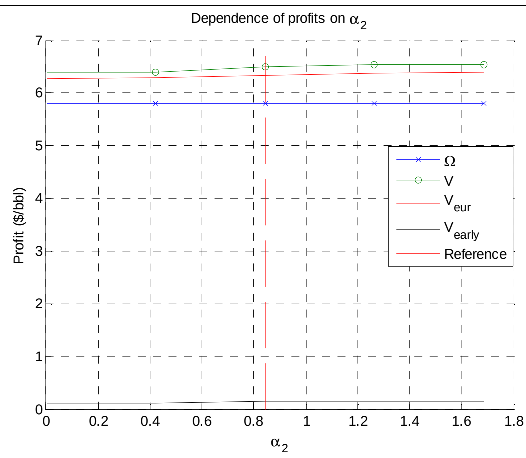

Dependence of profits on α_2. Reference is $\alpha_2 = 0.84$ yr^{-1}

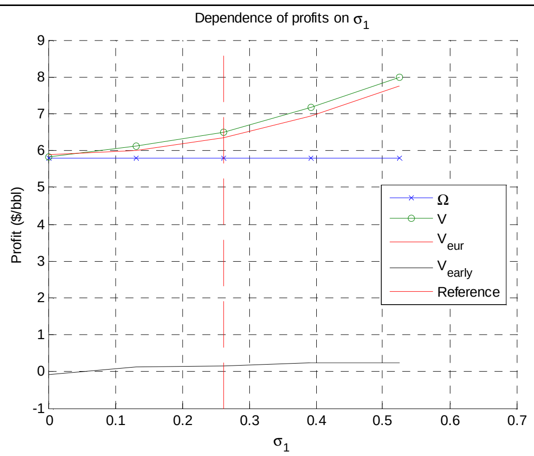

Dependence of profits on σ_1. Reference is $\sigma_1 = 26\%$ (annualized)

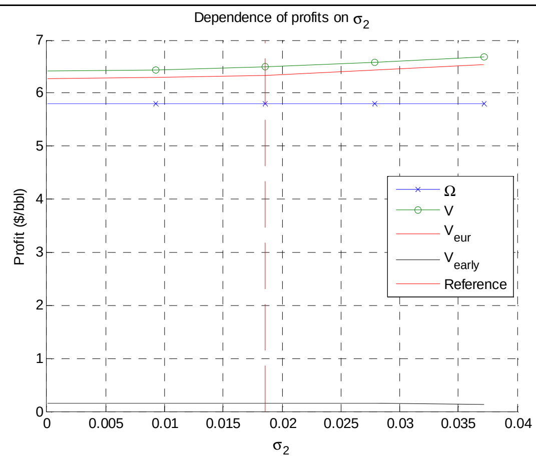

Dependence of profits on σ_2. Reference is $\sigma_2 = 1.86\%$ (annualized)

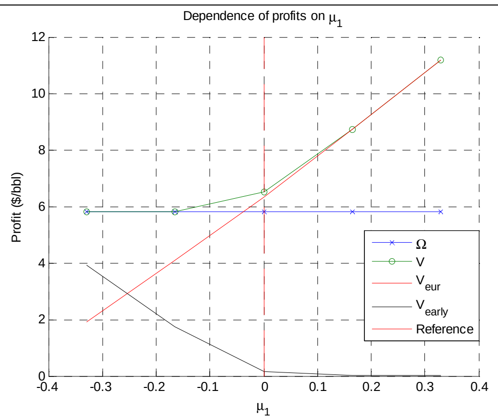

Dependence of profits on μ_1. Reference is $\mu_1 = 0$ yr^{-1}

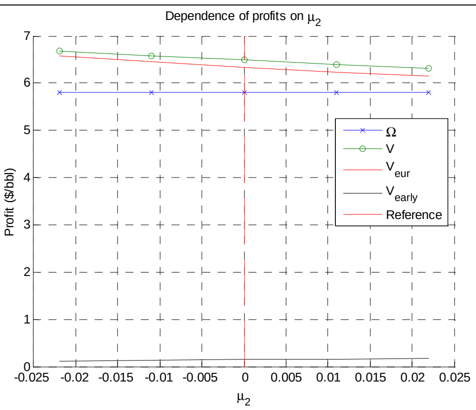

Dependence of profits on μ_2. Reference is $\mu_2 = 0$ yr^{-1}

Dependence of profits on the ship speed

In some circumstances, it can be beneficial for the trade to sail the ship slowly across the Atlantic in order to save on fuel costs. Intuitively, this will be especially useful when the forward curve is in a slight contango.

The speed will affect profits in three ways:

- A faster ship will be able to deliver its cargo earlier, which is important in a strong backwardation
- A faster ship will be chartered for less time, such that its total time charter cost will be lower
- A faster ship will consume more fuel. The fuel consumption function $FC(u)$ is approximately cubic in the speed u.

In Figure 32 we examine the variation of the profits with the speed of the ship u for a trade beginning on August 13, 2008 and April 28, 2009. We notice that the speed has a small influence on profits, of the order for 10 c/bbl for a speed varying from 8 to 17 knots. The speed is fixed during the voyage.

When the forward curve is in backwardation, there is incentive to deliver the cargo as soon as possible. A higher speed allows the trader to deliver the cargo earlier, but at the cost of higher fuel consumption. There is an optimal speed of around 13 knots yielding the best tradeoff. When the forward curve is in contango, the trade will involve some amount of floating storage at destination, such that fuel savings can be interesting. The excess profit $V - \Omega$, however, is not affected by the speed.

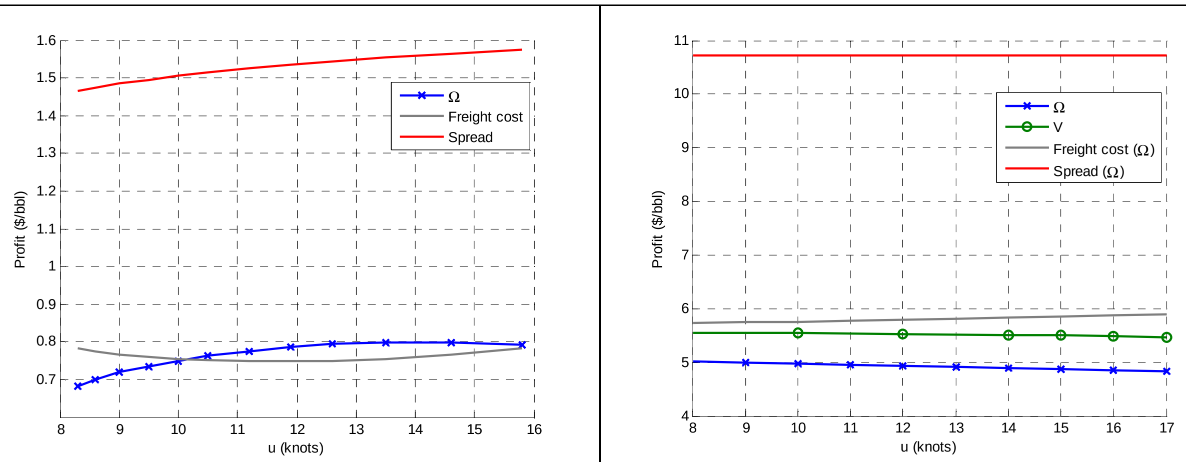

August 13, 2007 (backwardation)

April 28, 2009 (contango)

Figure 32. Variation of profits with vessel speed (fixed during the voyage) on two different dates, when the forward curve was in backwardation (left) and contango (right)

Time series of expected profit, risk and Sharpe ratio

For each week in the sample period, we perform the above calculations and derive:

- The arbitrage profit Ω, the expected profit V and the excess profit $V - \Omega$
- The expected risk, i.e. the standard deviation of W
- The a priori Sharpe ratio

We plot these values as a function of time.

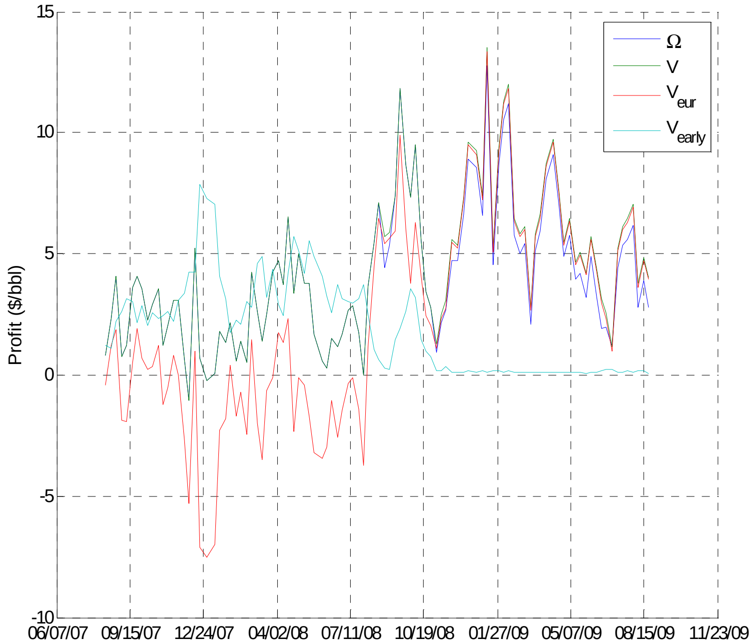

Figure 33. Arbitrage profit Ω and expected profit V, decomposed into V_{eur} and V_{early}, for weekly loading dates from August 2007 to August 2009

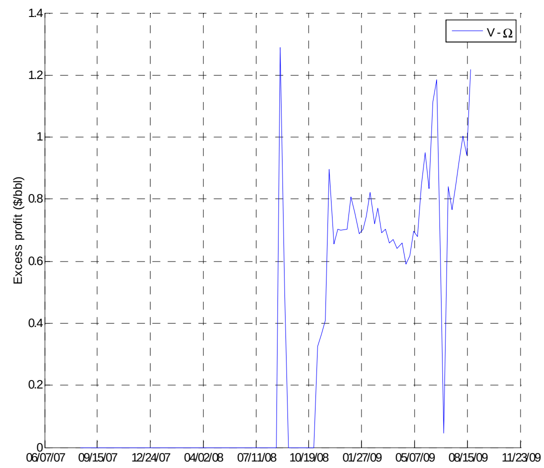

Figure 34. Expected excess profit V – Ω for different loading dates

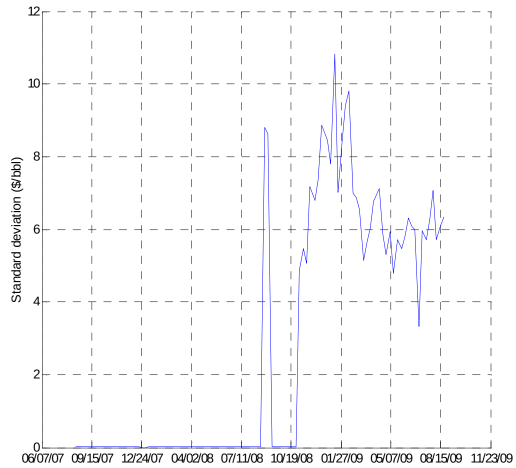

Figure 35. Expected standard deviation of realized profits. The zero values correspond to dates when exercise is immediate

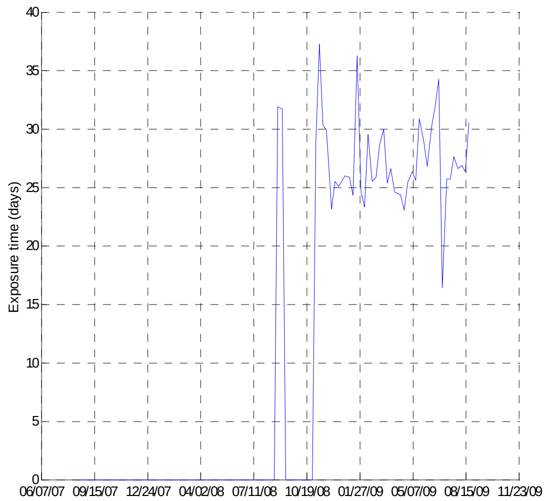

Figure 36. Expected exposure time $E[t^*]$ of the trade. The maximal exposure time is T = 100 days

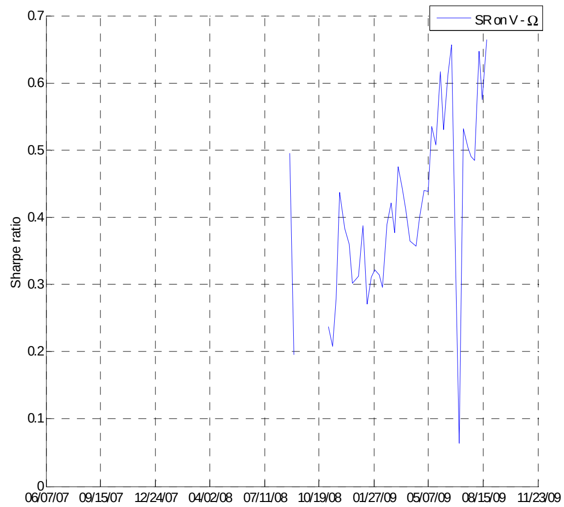

Figure 37. A priori annualized Sharpe ratio[1] of the excess profit V – Ω

[1] The average long-term Sharpe ratio of the S&P500 is about 0.4. The Sharpe ratios in Figure 37 are calculated on profits over the riskless arbitrage profit Ω. They are on the order of 6 when calculated over the risk-free rate.

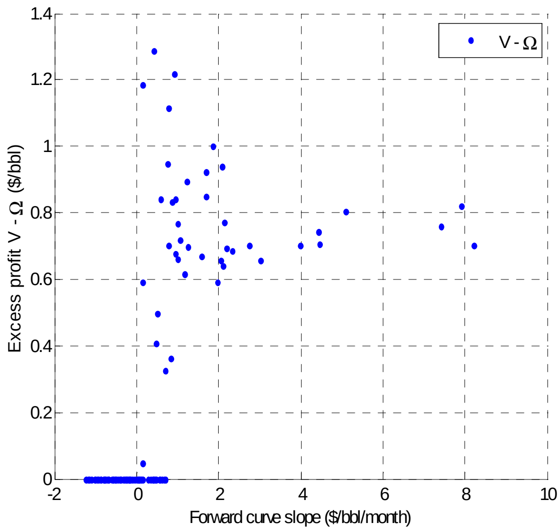

Figure 38. Excess expected profit V – Ω vs. Forward curve slope at delivery port

In Figure 33 we plot the arbitrage profit Ω and the expected profit V as a function of the trade initiation date. I.e. each week we examine the forward curve, spot price and shipping cost and determine what the arbitrage profit would be for a cargo loaded within the loading window τ_{load} (15 days), as well as the expected profit V from loading the cargo and executing the optimal trading strategy. We see that these profits are always at least as great as the arbitrage profits.

We have already studied the behavior of the time series of Ω, so we will concentrate here on the excess profit, V – Ω. We plot this expected excess profit as a function of time in Figure 34. Its value varies between 0 and 1 $/bbl, averaging 74 c/bbl in the period when the excess profit is positive.

We note that the period under consideration can be separated in two: from 2007 to mid-2008 the continuation value is zero, while after the market crash in 2008 the excess profit jumps to values around 75 c/bbl. The first period corresponds to a backwardated forward curve, while the second period corresponds to a period of strong contango and low freight rates following the crisis. In Figure 38 we show the relationship between the forward curve slope at the delivery port and the excess profit from continuation. Consistently with what was proved in Section 4.4, we find that a forward curve in backwardation or in slight contango yields a zero excess profit, while all the positive excess profits are associated with a forward curve in contango.

The standard deviation of the profits over the trade, presented in Figure 35, is significant, averaging 6.66 $/bbl in the period when the excess profit is positive. This is the risk associated with keeping the exposure to the forward curve open, and is accordingly zero when the cargo should be sold forward immediately, i.e. V = Ω. The expected time over which this exposure is held $E[t^*]$ is presented in Figure 36, and averages 27 days.

We note that the exposure time never reaches the maximal exposure time T that is set to 100 days here. Combining expected profit and risk we can calculate the annualized Sharpe Ratio associated with the strategy, which is presented in Figure 37. We consider this Sharpe ratio in excess of the riskless profit Ω. It averages 0.41 during the period. In Section 4.6 we examine the detail of these time series and attempt to explain the appearance of excess profits.

Realized profit and standard deviation

The functions $\Omega(t, f_1, f_2)$ and $V(t, f_1, f_2)$ define a physical trading strategy that can readily be put into practice. Given historical time series of the actual moves in the forward curve we can calculate the profits that would have been realized by following this strategy, and compare them to the expected profits and risks presented above.

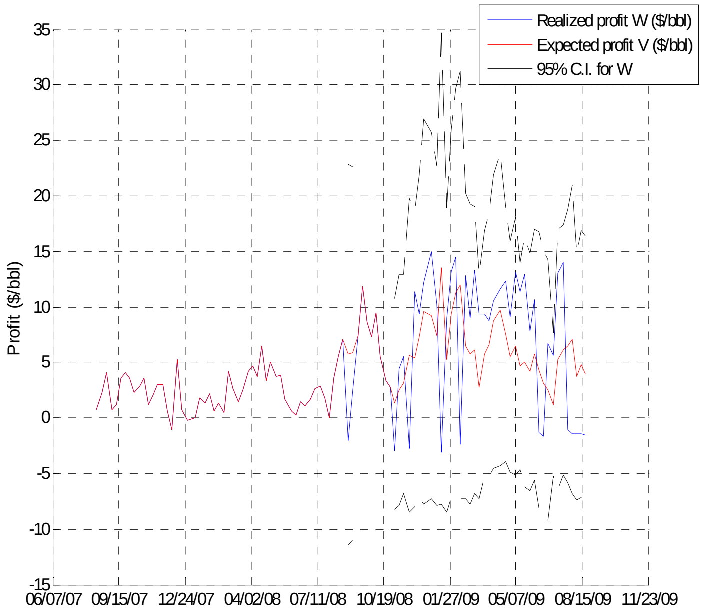

Statistics[1]	
E[W – V]	0.74 $/bbl
Std[W – V]	6.21 $/bbl
SR on W	5.95
SR on W – Ω	1.99

Figure 39. Expected profit V and realized profit W on different dates

[1] Statistics are not calculated on the entire period, but on the period when V > Ω

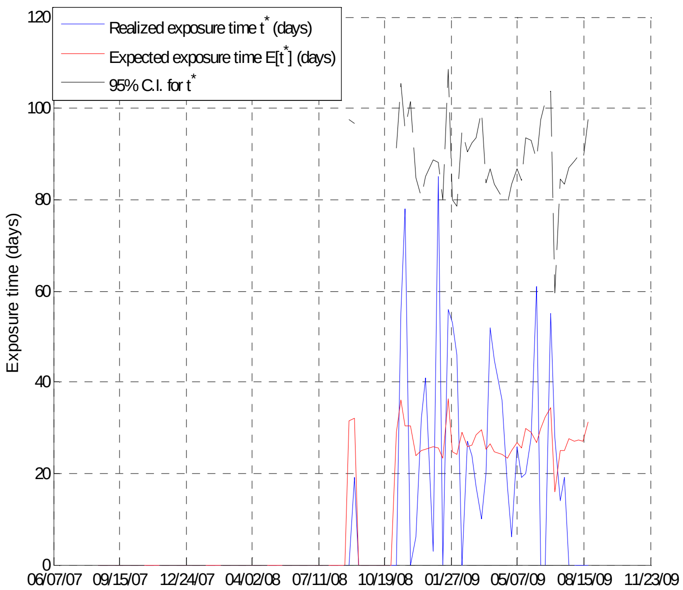

Statistics	
$E[t^*_{realized} - E[t^*]]$	-1.5 days
$\mathrm{Std}[t^*_{realized} - E[t^*]]$	15 days

Figure 40. Expected and realized exposure times t^*

On each week in the sample period, having calculated the functions $V_{t_0}(t, f_1, f_2)$ and $\Omega_{t_0}(t, f_1, f_2)$ with trade initiation date t_0, we execute the trading strategy defined by:

- if $V_{t_0}(0,0,0) > 0$, the trade is expected to be profitable, so initiate the trade by buying the cargo and chartering the vessel
- At each time from trade initiation $t > 0$, observe the forward curve $F(t,\tau)$ and calculate the factor values using the orthogonality condition in (2.11),

$$f_k(t) = \int_0^{\tau_{\max}} u_k(\tau) \log\left(\frac{F(t,\tau)}{F_0(\tau)}\right) d\tau \qquad (4.52)$$

- If $t < T$ (maximal exposure time, 100 days in this case), compare the exercise and continuation profits:
  - If $V_{t_0}(t, f_1(t), f_2(t)) > \Omega_{t_0}(t, f_1(t), f_2(t))$, then continue sailing at speed u
  - If $V_{t_0}(t, f_1(t), f_2(t)) = \Omega_{t_0}(t, f_1(t), f_2(t))$, it is optimal to exercise, so sell the cargo forward and collect $\Omega_{t_0}(t, f_1(t), f_2(t))$

- If $t = T$, sell the cargo forward and collect $\Omega_{t_0}(T, f_1(T), f_2(T))$

If the exercise time is t^* , the realized profit on this trade is then

$$W(t^*, f_1(t^*), f_2(t^*)) = \int_0^{t^*} -g(t)dt + \Omega(t^*, f_1(t^*), f_2(t^*)) \tag{4.53}$$

As can be seen from Figure 39, the realized profit is highly variable – but it stays within the bounds of the 95% confidence interval for *W* based on the expected risk calculated previously. The standard deviation of *W*–*V* calculated over the period when there are excess profits is 6.21 $/bbl, close to the average standard deviation seen in Figure 35. The exposure time, presented in Figure 40, varies widely

The annualized Sharpe ratio of the strategy over this period is 1.99 if calculated on the profits in excess of Ω, and 5.95 if considered in excess of the risk-free rate.

Realized profits and standard deviation with hedging

The significant standard deviation of the realized profits *W* versus the expected profits *V* comes from the exposure of the trade to the risk factors f_1 and f_2. Using the hedge ratios computed from the expected profit function *V* we can simulate what the realized profit is when the profit is delta-hedged with respect to the first or second factor.

At time *t* into the trade, assuming the cargo has not been sold, the value function has deltas δ_1 and δ_2 with respect to f_1 and f_2:

$$\delta_k = \frac{\partial V}{\partial f_k}(t, f_1(t), f_2(t)) \tag{4.54}$$

In order to eliminate the risk from factor 1, for example, we take a position $-\delta_1$ in the factor f_1. How to achieve this with the available futures contracts is explained in Section 2.9. The impact of this position on the evolution of the expected portfolio P&L $\tilde{U}$ is

$$\begin{aligned} d\tilde{U} &= \delta_1\sigma_1 dW_1 + \delta_2\sigma_2 dW_2 - \delta_1 df_1 \\ &= \delta_1(\alpha_1 f_1 - \mu_1)dt + \delta_2\sigma_2 dW_2 \end{aligned} \tag{4.55}$$

Hence the realized P&L at the end of the trade is

$$\tilde{U}(\tau^*, f_1, f_2) = V(0, f_1^0, f_2^0) + \int_0^{\tau^*} \delta_1(t, f_1(t), f_2(t))(\alpha_1 f_1 - \mu_1)dt + \int_0^{\tau^*} \delta_2(t, f_1(t), f_2(t))\sigma_2 dW_2(t) \tag{4.56}$$

The risk tied to the first factor has therefore been eliminated – but this also reduces the expected profit. The same approach can be applied to the second factor, hedging out tilts. A common practice is to hedge out the parallel shift factor, which is the major risk factor, and keep the exposure to tilts.

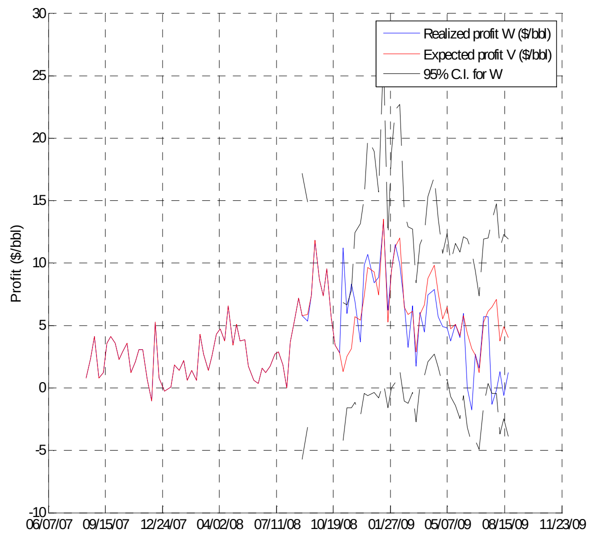

Statistics	
E[W – V]	-0.68 $/bbl
Std[W – V]	3.03 $/bbl
SR on W	6.48
SR on W – Ω	0.06

Figure 41. Expected profit V and realized profit W when hedging the first factor

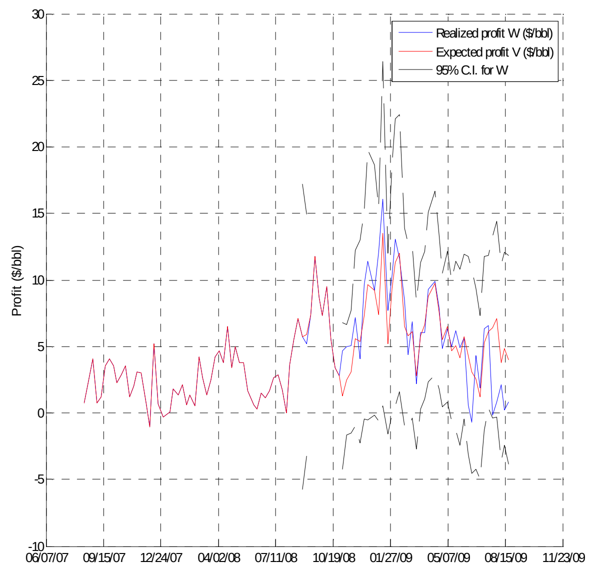

Statistics	
E[W – V]	0.01 $/bbl
Std[W – V]	2.42 $/bbl
SR on W	9.15
SR on W – Ω	1.12

Figure 42. Expected profit V and realized profit W when hedging the first and second factors

As can be seen in Figure 41 and Figure 42, the hedging does indeed diminish the risk of the strategy. The historical standard deviation of $W - V$ is

- 3.03 \$/bbl when hedging the first factor
- 2.42 \$/bbl when hedging the first and second factor

This should be compared to the unhedged standard deviation of 6.21 \$/bbl.

It is interesting to note that even hedging both factors does not render the strategy riskless, contrary to theory. There are two reasons for this:

- The delta-hedging is only daily and not continuous, and high-amplitude movements (jumps) in the factors will not be hedged perfectly
- The forward curve does not only move in shifts and tilts, and only those movements have been hedged out

7. Origins of excess profits

We have shown that in addition to significant arbitrage profits to be made on arbitraging crude oil between Europe and the United States, following an optimal storage and selling strategy could lead to significant excess profits. It is interesting to understand the origin of these profits in order to understand in what fundamental situations they might appear.

We will make a distinction in what follows between

- The origin of excess *expected* profits
- The origin of *realized* profits, i.e. when the trading strategy performs well

Origin of excess expected profits

We have established in Section 3.5 that the period August 2007 – August 2009 can be decomposed into two periods: August 2007 to October 2008, when the forward curve for crude oil was in backwardation and there were no expected excess profits, and October 2008 to August 2009, when the forward curve was in contango and there could be found excess profits in keeping exposure to the forward curve open.

We will concentrate on the second period here. We have already established that the forward curve (net of freight cost) being in contango is a necessary condition for the excess profit to be positive. We can gain more insight into the origins of excess profits by decomposing the excess profit as follows

$$\begin{aligned}EP_{\text{total}} &= V(0, f_1, f_2) - \Omega(0, f_1, f_2) = EP_{\text{drift}} + EP_{\text{convexity}} + EP_{\text{early}} \\ EP_{\text{drift}} &= \Omega(T, Ef_1(T), Ef_2(T)) - G(T) - \Omega(0, f_1, f_2) \\ EP_{\text{convexity}} &= V_{eur} - (\Omega(T, Ef_1(T), Ef_2(T)) - G(T)) \\ EP_{\text{early}} &= V_{early}\end{aligned} \tag{4.57}$$

When considering the initial expected profit, $f_1 = f_2 = 0$ such that $Ef_1(T) = Ef_2(T) = 0$ and

$$EP_{\text{drift}} = \Omega(T, 0, 0) - G(T) - \Omega(0, 0, 0) \tag{4.58}$$

This expected profit will generally be zero for a forward curve in contango. It can, however, be significant for non-zero factor values because of their mean-reverting property. The excess profit from convexity can be written as

$$EP_{\text{convexity}} = E\left[\Omega(T, f_1(T), f_2(T))\right] - \Omega(T, Ef_1(T), Ef_2(T)) \tag{4.59}$$

and captures the non-linearity of Ω. As for the early-exercise premium, it captures the possibility of selling the cargo before the date *T*.

We present the time series of the excess expected profits and its decomposition in Figure 43. We notice that the major part of the excess profit comes from the convexity, averaging 84% of the total excess profit. The convexity and early-exercise premia are rather regular.

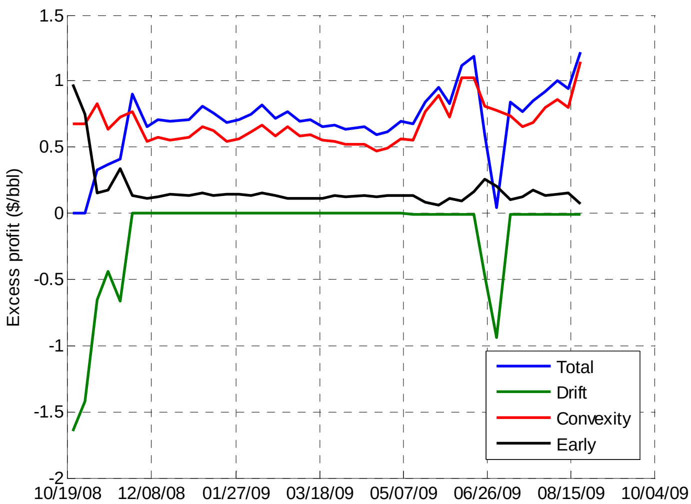

Figure 43. Decomposition of expected excess profits as a function of time

Based on these observations we can conclude that

- The existence of an excess profit is conditional on the forward curve net of shipping cost being in contango
- When the contango condition is satisfied, the expected excess profits are fairly stable.

Trade performance and origin of realized profits

We have identified in what situations excess profits are expected. However, in a trading situation it is important to know in what cases the trade will succeed and in which cases it will yield a loss, in order to understand the expected profits and risk manage the position.

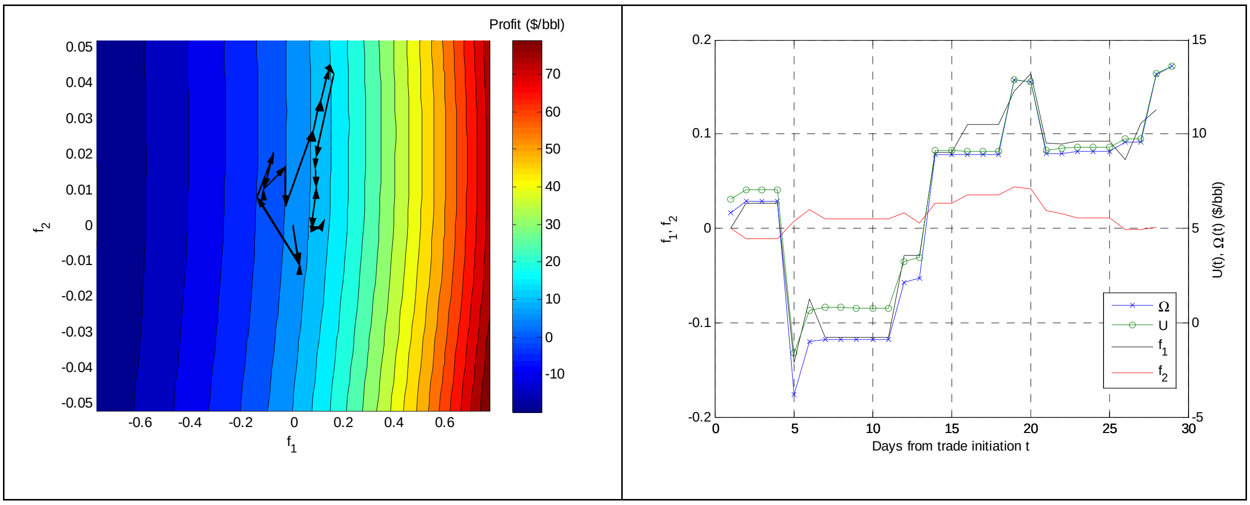

Figure 44. Evolution of f_1, f_2, expected P&L U and exercise profit Ω during the trade starting on December 18, 2008. Left, the path of ($f_1(t)$, $f_2(t)$) during the trade overlaid on the expected profit V. Right, these functions as a function of days from trade initiation. The delivery of the cargo is specified after 29 days.

In Figure 44 we present the evolutions of the factor values and the expected and exercise profits U and Ω during the physical trade initiated on December 18, 2008. In this particular case, the cargo is exercised after 29 days, when the expected profit and exercise profit are seen to converge. The realized profit W at the end of the trade is 13.6 $/bbl. We also present the evolutions of the factor values f_1 and f_2 on the same figure. As we have already seen, V has the strongest delta with respect to the first factor, and the realized profit is highly correlated with the value of f_1 during the trade. When hedging the first factor, however, the realized profits are more correlated with the second factor, as is seen in Figure 45.

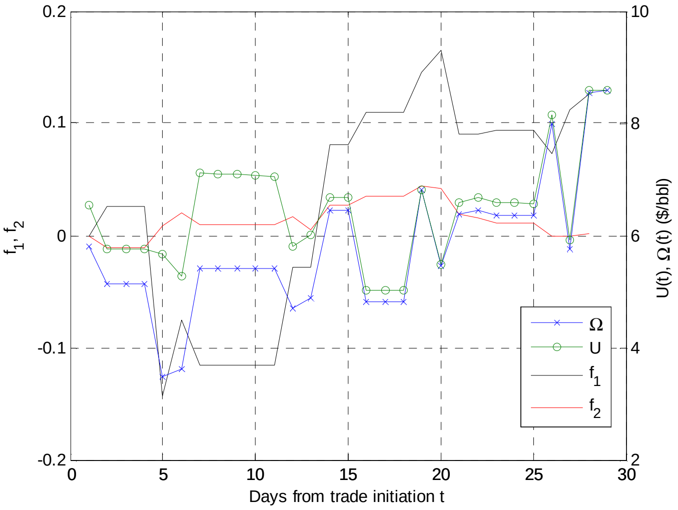

Figure 45. Evolution of the factor values and profits during the trade starting on December 18, 2008, when hedging the first factor

In order to assess how the trade will perform based on the evolutions of the two factors it is useful to recall the shape of the payoff function Ω as a function of both factors.

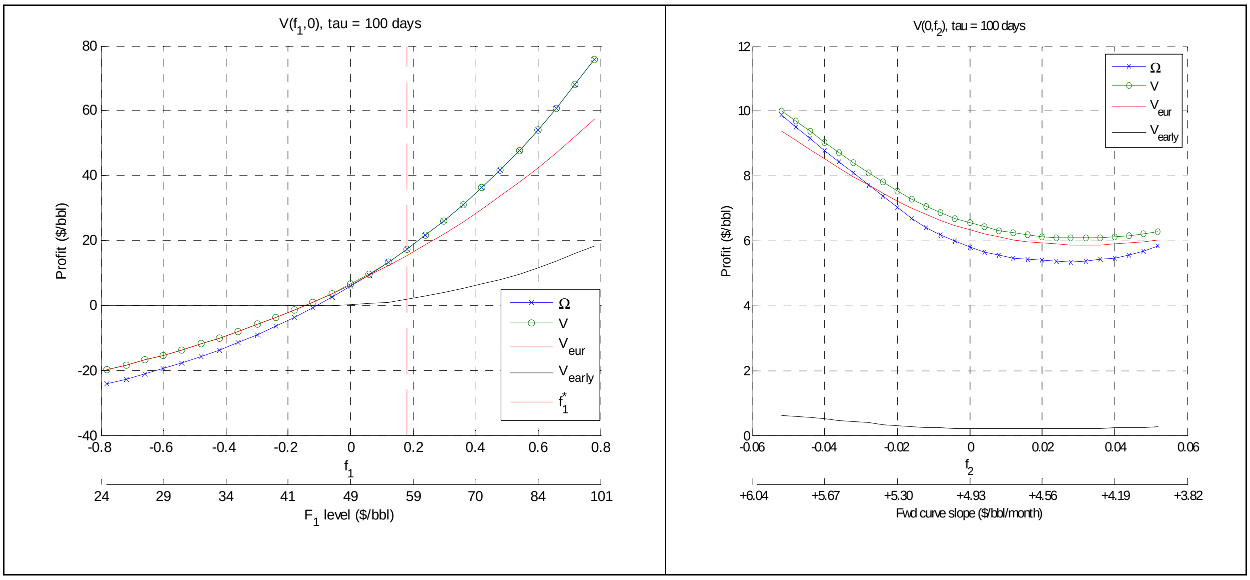

Figure 46. Cross-section of Ω and V at trade initiation as a function of f_1 (left) and f_2 (right)

As can be seen in Figure 46 the dependence of the trade has the following characteristics with respect to f_1 and f_2:

- It is directional with respect to f_1, similar to a forward exposure
- It is a volatility trade with respect to f_2 : the payoff is convex and has higher payoff for large movements of f_2 in either direction. This is closer to a straddle option.

Figure 47. Realized profits as a function of realized drifts and volatilities during the trade period, for a trade starting on December 18, 2008

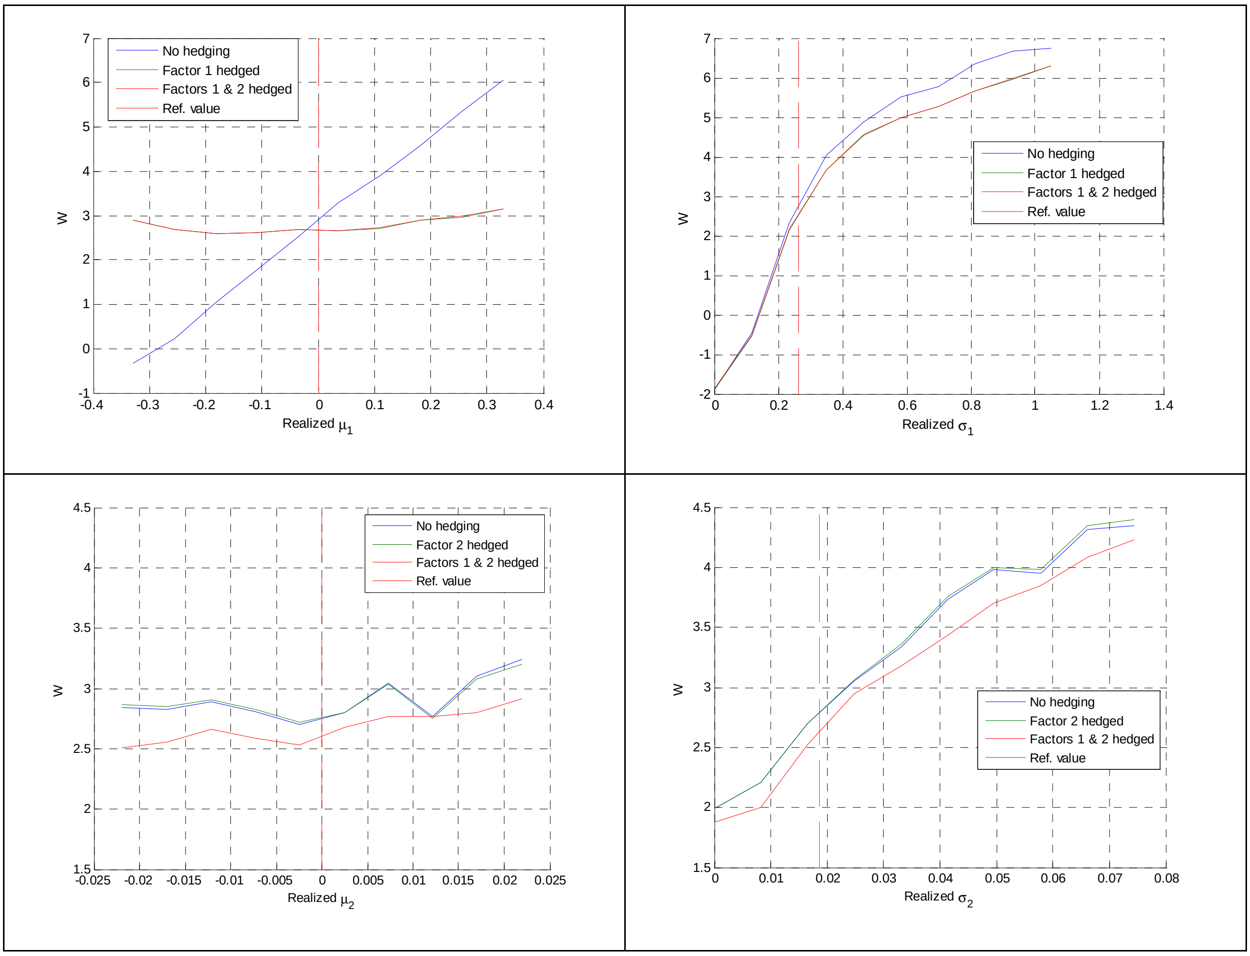

This intuition is confirmed by the results in Figure 47. In this figure we present the realized profits as functions of *realized* drifts and volatilities, imposed in a Monte Carlo simulation, different from the *a priori* drifts and volatilities used when valuing the floating storage opportunity. We can clearly see the directional nature of the position in f_1, with realized profits that are linear in the drift μ_1. These are also increasing in the volatility σ_1 because of the slight convexity of the payoff function. On the other hand, the realized profits are independent of the drift μ_2 of the second factor, but strongly related to its realized volatility.

8. General commodity trading problem

The problem we have been considering is limited to a single delivery port and a single tanker speed. The only choice that is left to the trader is time of delivery.

In general, an oil (or other commodity) cargo that has not yet been sold forward can be rerouted to a different port. The ship can also sail slower in order to save fuel. These optionalities make the cargo more valuable to a trader than what has been calculated previously. A decision model for optimal ship routing should take into account the forward prices at different potential delivery ports and open for the possibility of delaying the choice of delivery port to a later date.

For example, a cargo of Bonny Light crude oil loaded in Nigeria could potentially be delivered to Europe or the United States. Instead of choosing a delivery location immediately the oil trader could choose to route the ship northbound in the mid-Atlantic, and waiting to see if the spread evolves.

We will formulate the stochastic control equations governing how the ship should be routed to maximize profit. The notations are the same as previously, but we now introduce:

- A set of destination ports X_k^P at which the cargo can be delivered, each with a forward curve $F_k(\tau)$
- The speed of the ship *u* can be varied within bounds $[u_1, u_2]$, usually between 8 and 16 knots
- The instantaneous direction of the ship is the unit vector $\vec{d}$

Exercise profit

We define $\Omega(X, F_k(\tau))$ to be the profit that can be earned on the cargo if the ship is at location *X* and the forward curve in port *k* is $F_k(\tau)$, by committing to a specific delivery price sometime in the future and sailing to deliver at that time. In effect, the trader gives up the possibility of changing delivery port and time.

At exercise, one chooses a delivery port *k*, a time-to-delivery τ and a sailing speed *u*. For one choice of these parameters, the profit is

$$\begin{aligned} &\tau_{sail}(k,u) = d(X, X_k^P)/u \\ &\omega(k,\tau,u) = F_k(\tau) - S_0 - H\tau - B\cdot(FC(u)\tau_{sail}(k,u) + FC_a(\tau - \tau_{sail}(k,u))) - BH_k \quad (\tau \geq \tau_{sail}(k,u)) \end{aligned} \tag{4.60}$$

The exercise profit is obtained by maximizing $\omega(k,\tau,u)$ over all possible ports, speeds and times to delivery:

$$\Omega(X, F_k(\tau)) = \max_{\substack{k,u \\ \tau \geq \tau_{sail}(k,u)}} \omega(k,\tau,u) \tag{4.61}$$

Expected profit and optimal route

Let us define $V(X, F_k(\tau))$ as the expected profit from the cargo when the tanker is at the location X and the forward curves are given by $F_k(\tau)$.

Let $g(X,u)$ be the daily cost of sailing at speed u when the ship is at location X, i.e.

$$\begin{aligned} g(X,u) &= H + B\cdot FC(u) && \text{if the ship is sailing} \\ g(X_k^P,u) &= H + B\cdot FC_a && \text{if the ship is at anchor at port k} \end{aligned} \tag{4.62}$$

When the ship is at a location X, the trader has two choices:

- either "exercise" and sell the cargo forward, thereby earning the exercise profit $\Omega(X, F_k(\tau))$

- or choose to continue speculating during a time dt without exercising. If the ship is at sea he can choose the optimal speed u and direction $\vec{d}$ and the expected profit is:

$$V_C(X, F_k(\tau)) = \max_{u,\vec{d}} E\left[V(X + udt\vec{d}, F_k(\tau) + dF_k)\right] - g(X,u)dt \tag{4.63}$$

If the ship is in port (floating storage), the expected profit is

$$V_C(X_l^P, F_k(\tau)) = E\left[V(X, F_k(\tau) + dF_k)\right] - g(X_l^P, 0)dt \tag{4.64}$$

This gives the continuation value V_C. The forward curves are evolved during the time period dt using the two-factor model.

Hence the expected profit at location X given the forward curves $F_k(\tau)$ is

$$V(X, F_k(\tau)) = \max\left[\Omega(X, F_k(\tau)), \quad V_C(X, F_k(\tau))\right] \tag{4.65}$$

Hamilton-Jacobi-Bellman equation

When we assume that the underlying factors follow diffusions we can derive a continuous-time equation to evaluate V. This equation is known as the Hamilton-Jacobi-Bellman (HJB) equation for the stochastic control problem, see Morimoto (2010) and Chang (2004).

We assume that each of the forward curves $F_k(\tau)$ is governed by a two-factor model, such that

$$\log F_t^k(\tau) = \log F_0^k(t+\tau) + \psi_1^k(t,\tau) + \psi_2^k(t,\tau) + u_1^k(\tau) f_1^k(t) + u_2^k(\tau) f_2^k(t) \tag{4.66}$$

where:

$$
\begin{aligned}
df_j^k(t) &= -\alpha_j^k f_j^k(t)dt + \sigma_j^k dW_j^k(t), \quad j = 1,2 \\
d\psi_j^k(t,\tau) &= -\frac{1}{2}\left(\sigma_j^k u_j^k(t+\tau)\right)^2 dt
\end{aligned}
\tag{4.67}
$$

For the sake of simplicity we renumber the factors f_j^k as a sequence $(f_i)_{i=1,\ldots,M}$. While the two factors for a single forward curve are uncorrelated, factors for different forward curves will be correlated, such that in general

$$
df_k df_l = \sigma_k \sigma_l \rho_{kl} dt \tag{4.68}
$$

Consider a location *X*, time *t* and factor values f_i and assume that $V_C > \Omega$, i.e. we are in the continuation region, s.t. $V = V_C$. We develop *V* using Ito's formula:

$$
\begin{aligned}
V(X + \vec{d} \cdot dX, t + dt, f_1 + df_1, \ldots, f_M + df_M) &= V + \frac{\partial V}{\partial t} dt + \frac{\partial V}{\partial X} \cdot \vec{d} \cdot dX + \sum_{i=1}^{M} \frac{\partial V}{\partial f_i} df_i + \frac{1}{2} \sum_{i=1}^{M} \sum_{j=1}^{M} df_i \frac{\partial^2 V}{\partial f_i \partial f_j} df_j \\
&= V + \left[\frac{\partial V}{\partial t} + u \frac{\partial V}{\partial X} \cdot \vec{d} + \sum_{i=1}^{M} (\mu_i - \alpha_i f_i) \frac{\partial V}{\partial f_i} + \frac{1}{2} \sum_{i=1}^{M} \sum_{j=1}^{M} \rho_{ij} \sigma_i \sigma_j \frac{\partial^2 V}{\partial f_i \partial f_j} \right] dt + \sum_{i=1}^{M} \frac{\partial V}{\partial f_i} \sigma_i dW_i
\end{aligned}
\tag{4.69}
$$

Taking expectations in the definition of V_C and simplifying, we finally get the equation:

$$
0 = \max_{u,\vec{d}} \left(-g(X,u) + \frac{\partial V}{\partial t} + u \frac{\partial V}{\partial X} \cdot \vec{d} + \sum_{i=1}^{M} (\mu_i - \alpha_i f_i) \frac{\partial V}{\partial f_i} + \frac{1}{2} \sum_{i=1}^{M} \sum_{j=1}^{M} \rho_{ij} \sigma_i \sigma_j \frac{\partial^2 V}{\partial f_i \partial f_j} \right) \tag{4.70}
$$

This is a typical example of a Hamilton-Jacobi-Bellman equation. It is valid in the continuation region, i.e. $(f_1, \ldots, f_M) \in S^*(t)$. The boundary conditions are given by the smooth pasting condition on the free boundary $\partial S^*(t)$:

$$
\forall (f_1, \ldots, f_M) \in \partial S^*(t, X), \begin{cases} V(t, X, f_1, \ldots, f_M) = \Omega(t, X, f_1, \ldots, f_M) \\ \nabla V(t, X, f_1, \ldots, f_M) = \nabla \Omega(t, X, f_1, \ldots, f_M) \end{cases} \tag{4.71}
$$

The terminal condition is that at the maximal exposure time *T*, the cargo should be delivered:

$$
V(T, X, f_1, \ldots, f_M) = \Omega(T, X, f_1, \ldots, f_M) \tag{4.72}
$$

Solving the general problem

This problem can in principle be solved numerically by dynamic programming or finite differences. However, the potentially large number of state variables can make it challenging to solve using these methods. The preferred method for such a problem would be a least-squares Monte Carlo simulation as presented in Longstaff and Schwartz (2001). This requires work on finding appropriate basis functions to project the solution on.

5. CONCLUSIONS

1. Summary of results

In Part 2 we have developed a two-factor model and given evidence that it is sufficient for modeling the term structure of volatility and the correlation surface of a number of commodities. We prove that it is easily formulated as a model involving two independent and mean-reverting factors that represent the change in level and slope of the forward curve. We find that the first factor is the dominant factor and the majority of variance of forward prices comes from the first factor. However, other factors cannot be ignored as they will affect portfolios that are weighted differently.

We also show that the spot price process implied by this two-factor model is consistent with the Schwartz and Smith (2000) formulation with short-term and long-term shocks driving the spot price. Furthermore, we show that the shapes of forward curves consistent with the two-factor model are exponentials of the factors weighted by their factor loadings. This allows for a simple calibration of forward curves to the market model and an interpretation of the factor values in terms of mean level and initial slope of the curve.

The applicability of this model to a number of forward markets, as well as its simple analytical formulation, makes it useful in different valuation settings involving commodity prices. In Part 3 we address the pricing of Asian options written on commodity forwards. We show that by understanding the term structure of volatility correctly, as well as the effect of the averaging on the volatility of the payoff, Asian options can be priced approximately but analytically in a simple way. Comparing our theoretical prices to market prices, we find that it correctly reproduces the term structure of implied volatilities. The understanding of this should increase liquidity in the freight options market.

The understanding of volatility and its value also has a profound impact on valuation and operational decisions that involve commodities. In Part 4 we study the floating storage trade involving crude oil and tankers using the two-factor model. This trade can be viewed as the sum of a cross-Atlantic and temporal arbitrage trade – arbitraging crude oil between Europe and the United States and between now and the future – and of a storage trade where the trader can choose the optimal time to release the oil into the market. We show that while the arbitrage window has been open for most of the time during 2007-2009, the storage trade has only existed in the second half of this period.

The floating storage opportunity is associated with a forward curve in contango when netted of freight costs. When it is open, there is additional value involved in not selling the cargo immediately and taking advantage of the possibility of higher prices. The framework that we present allows us to evaluate the profits from such a strategy, the decision rules for running the trade, and its exposure to the two risk factors through hedge ratios. The excess value is understood as a combination of the drifts in the factors, of the payoff convexity and of an early exercise premium.

2. Suggestions for future research

As pertains to the market modeling, an essential improvement that is not performed in the present thesis would be to allow the model to be easily calibrated to market implied volatilities as well as the historical correlation surface. The crude oil market, for example, has a very liquid options market that can be used for such calibration.

The market modeling framework presented in this thesis can be applied to a number of problems related to commodity trading. It would be very interesting to see empirical results for the general commodity trading problem presented in Section 4.8, and understand what value is associated with the possibility of switching destinations. It can also be applied to Liquefied Natural Gas cargoes that are currently being rerouted from their long-term contract destinations in the United States to Europe or Asia.

Furthermore, this market-based routing problem should be integrated with the optimal weather routing problem developed in Avougleas and Sclavounos (2009). An underlying assumption in our formulation is that routes are deterministic and fuel consumption only depends on speed. In practice, ship routing and fuel consumption depends strongly on weather, and using forecasts and dynamic programming one can determine the optimal route to follow. Integrating this uncertainty with our model would give a much more precise evaluation of the commodity trade, especially when profits come from geographical spreads and not floating storage.

However, this general problem involves a large number of state variables and is difficult to solve using dynamic programming. Developing solution methods adapted to such a large-scale problem would greatly enhance its applicability. One promising method, applied for American options, is the Longstaff and Schwartz (2001) least squares Monte Carlo method. This would require finding suitable basis functions on which to project the solution.

In this thesis we view the shipping problem from the point of view of a physical oil trader who has the possibility of chartering a ship for one trade, before returning it to the market. Another direction would be to see the problem from the point of view of a shipowner or long-term charterer who can operate the ship continuously on several trades. In that case, the decision taken on one trade, such as storing oil, will have consequences for the next one. In some cases it might be more profitable to sell the oil, return to the loading port and take advantage of a better geographical spread. The same framework can be used, but the problem is of longer-term nature, of years rather than weeks.

6. APPENDIX

1. Traded volumes in commodity derivative markets

From ICE (2009) and CME (2009):

	Contract	Daily volume ('000 bbl)	Yearly volume ('000 bbl)
ICE	Brent Crude Futures	287,355	74,137,750
	Brent Crude Options	823	212,341
	WTI Crude Futures	179,820	46,393,671
	WTI Crude Options	70	18,200
NYMEX (Futures)	Crude oil physical	545,351	141,791,260
	Crude oil	4,521	1,175,460
	Miny WTI	13,369	1,737,970
	Brent Financial Futures	1,997	519,220
	Dubai Crude oil Calendar	3,665	952,900
NYMEX (Options)	WTI Calendar	4,198	1,091,480
	Brent Calendar Options	546	141,960
	Brent last day	74	19,240
	Crude oil 1mo	2,857	742,820
	Crude oil APO	12,713	3,305,380
	Crude oil physical	113,302	29,458,520
	Total	**1,170,661**	**301,698,172**

From Imarex (2009). One lot is 1000 metric tons.

Period	# trades	# lots
Dec '09	698	14 504
Nov '09	1 017	20 817
Oct '09	1 083	17 750
Sep '09	1 066	13 733
Aug '09	711	12 795
Jul '09	1 048	14 113
Jun '09	1 328	24 766
May '09	1 128	16 458
Apr '09	1 249	18 703
Mar '09	1 362	19 965
Feb '09	1 133	15 625
Jan '09	1 343	18 020
Total	**13 166**	**207 249**

2. Spot price process implied by the two-factor model

Using the forward curve process and the spot-forward relationship $S(t) = F(t,t)$, we get:

$$\frac{dF(t,T)}{F(t,T)} = \sigma_S e^{-\alpha(T-t)} dW_S + \sigma_L dW_L = \sigma_S(t,T) dW_S + \sigma_L(t,T) dW_L$$

$$\log S(t) = \log F(0,t) - \frac{1}{2}\int_0^t \sigma_S^2(s,t)ds + \int_0^t \sigma_S(s,t)dW_S(s) - \frac{1}{2}\int_0^t \sigma_L^2(s,t)ds + \int_0^t \sigma_L(s,t)dW_L(s)$$

Such that:

$$\begin{aligned} d\log S(t) = &\left[\frac{\partial \log F(0,t)}{\partial t} - \frac{1}{2}\sigma_S^2(t,t) - \int_0^t \sigma_S(s,t)\frac{\partial \sigma_S(s,t)}{\partial t}ds + \int_0^t \frac{\partial \sigma_S(s,t)}{\partial t} dW_S(s)\right. \\ &\left. - \frac{1}{2}\sigma_L^2(t,t) - \int_0^t \sigma_L(s,t)\frac{\partial \sigma_L(s,t)}{\partial t}ds + \int_0^t \frac{\partial \sigma_L(s,t)}{\partial t} dW_L(s)\right] dt \\ &+ \sigma_S(t,t)dW_S(t) + \sigma_L(t,t)dW_L(t) \end{aligned}$$

We have:

$$\frac{\partial \sigma_S(s,t)}{\partial t} = -\alpha\sigma_S(s,t), \qquad \frac{\partial \sigma_L(s,t)}{\partial t} = 0$$

Such that

$$\begin{aligned} \int_0^t \frac{\partial \sigma_S(s,t)}{\partial t} dW_S(s) &= -\alpha \int_0^t \sigma_S(s,t) dW_S(s) \\ &= -\alpha\left[\log S(t) - \log F(0,t) + \frac{1}{2}\int_0^t \sigma_S^2(s,t)ds + \frac{1}{2}\int_0^t \sigma_L^2(s,t)ds - \int_0^t \sigma_L(s,t)dW_2(s)\right] \end{aligned}$$

Let:

$$\mu(t) = \frac{1}{\alpha}\left[\frac{\partial \log F(0,t)}{\partial t} - \frac{1}{2}\sigma_S^2(t,t) - \frac{1}{2}\sigma_L^2 + \frac{\alpha}{2}\int_0^t (\sigma_S^2(s,t) - \sigma_L^2)ds + \alpha \log F(0,t) + \alpha\sigma_L(W_L(t) - W_L(0))\right]$$

$$\mu(t)=\frac{1}{\alpha}\left[\frac{\partial\log F(0,t)}{\partial t}-\frac{1}{2}(\sigma_S^2+\sigma_L^2)+\frac{\alpha}{2}\left(\frac{\sigma_S^2}{2\alpha}(1-e^{-2\alpha t})-\sigma_L^2 t\right)+\alpha\log F(0,t)+\alpha\sigma_L W_L(t)\right]$$
$$=\frac{1}{\alpha}\frac{\partial\log F(0,t)}{\partial t}+\log F(0,t)-\frac{\sigma_S^2}{4\alpha}(1+e^{-2\alpha t})-\frac{\sigma_L^2}{2\alpha}(1+\alpha t)+\sigma_L W_L(t)$$

Then:

$$d\log S(t)=\alpha(\mu(t)-\log S(t))dt+\sigma_S dW_S+\sigma_L dW_L$$
$$d\mu(t)=m(t)dt+\sigma_L dW_L$$

where:

$$m(t)=\frac{1}{\alpha}\frac{\partial^2\log F(0,t)}{\partial t^2}+\frac{\partial\log F(0,t)}{\partial t}+\frac{1}{2}\left(\sigma_S^2 e^{-2\alpha t}-\sigma_L^2\right)$$

3. Principal Components Analysis of the two-factor model

We want to find the functions u_k that are eigenvectors of the covariance matrix $\Sigma(\tau_1,\tau_2)$ with associated eigenvalues λ_k. For this we must choose some arbitrary maximal tenor *T*, and find eigenvalues λ_k and eigenvectors $u_k(\tau)$ satisfying:

$$\int_0^{\tau_{\max}}\Sigma(\tau_1,\tau_2)u_k(\tau_2)d\tau_2=\lambda_k u_k(\tau_1)$$
$$\int_0^{\tau_{\max}}u_k^2(\tau)d\tau=1$$
$$\int_0^{\tau_{\max}}u_k(\tau)u_l(\tau)d\tau=\delta_{kl}$$

Given the parametric form of $\Sigma(\tau_1,\tau_2)$ we find that

$$\int_0^{\tau_{\max}}\left[(\sigma_S e^{-\alpha\tau_1}+\rho\sigma_L)(\sigma_S e^{-\alpha\tau_2}+\rho\sigma_L)+(1-\rho^2)\sigma_L^2\right]u_k(\tau_2)d\tau_2=\lambda_k u_k(\tau_1)$$

Leaving out the *k* index and developing this equation we find that

$$\lambda u(\tau_1)=\left[\int_0^{\tau_{\max}}(\sigma_S^2 e^{-\alpha\tau_2}+\rho\sigma_S\sigma_L)u(\tau_2)d\tau_2\right]e^{-\alpha\tau_1}+\left[\int_0^{\tau_{\max}}(\rho\sigma_S\sigma_L e^{-\alpha\tau_2}+\sigma_L^2)u(\tau_2)d\tau_2\right]$$

Thus we see that $u(\tau)$ can be written in the form $u(\tau) = Ae^{-\alpha\tau} + B$ where A and B are constants. We replace this expression into the equation to find that:

$$\lambda\sigma(\tau) = \left[\int_0^{\tau_{\max}} (\sigma_S^2 e^{-\alpha\tau_2} + \rho\sigma_S\sigma_L)(Ae^{-\alpha\tau_2} + B)d\tau\right] e^{-\alpha\tau} + \left[\int_0^{\tau_{\max}} (\rho\sigma_S\sigma_L e^{-\alpha\tau_2} + \sigma_L^2)(Ae^{-\alpha\tau_2} + B)d\tau\right]$$

Equating the constant and exponential terms we get the matrix eigenvector equation:

$$\lambda\begin{bmatrix} A \\ B \end{bmatrix} = \begin{bmatrix} \int_0^{\tau_{\max}} (\sigma_S^2 e^{-\alpha\tau_2} + \rho\sigma_S\sigma_L)e^{-\alpha\tau_2} d\tau_2 & \int_0^{\tau_{\max}} (\sigma_S^2 e^{-\alpha\tau_2} + \rho\sigma_S\sigma_L) d\tau_2 \\ \int_0^{\tau_{\max}} (\rho\sigma_S\sigma_L e^{-\alpha\tau_2} + \sigma_L^2)e^{-\alpha\tau_2} d\tau_2 & \int_0^{\tau_{\max}} (\rho\sigma_S\sigma_L e^{-\alpha\tau_2} + \sigma_L^2) d\tau_2 \end{bmatrix} \begin{bmatrix} A \\ B \end{bmatrix}$$

This shows that λ is an eigenvalue and (A, B) an eigenvector of the two-dimensional matrix *M*. Thus there are only two distinct eigenfunctions $u(\tau)$ and eigenvalues λ - as expected for a two-factor model.

4. Evolution of the constant-maturity forward curve under the two-factor model

If we let $\sigma_k(t,T) = \sigma_k u_k(T-t) = \sigma_k(A_k e^{-\alpha_k(T-t)} + B_k)$, we have:

$$\log f(t,\tau) = \log F(t,t+\tau) = \log F(0,t+\tau) + \sum_{k=1}^{2} -\frac{1}{2}\int_0^t \sigma_k^2(s,t+\tau)ds + \int_0^t \sigma_k(s,t+\tau)dW_k(s)$$

Let:

$$s_k(t,\tau) = -\frac{1}{2}\int_0^t \sigma_k^2(s,t+\tau)ds + \int_0^t \sigma_k(s,t+\tau)dW_k(s)$$

Expand the stochastic component (dropping the index *k* for now):

$$\int_0^t \sigma(s,t+\tau)dW(s) = Ae^{-\alpha\tau}\int_0^t \sigma e^{-\alpha(t-s)}dW(s) + B\sigma\int_0^t dW(s)$$

$$= (Ae^{-\alpha\tau} + B)\sigma\int_0^t e^{-\alpha(t-s)}dW(s) + B\sigma\int_0^t (1 - e^{-\alpha(t-s)})dW(s)$$

Let:

$$f(t)=\sigma\int_0^t e^{-\alpha(t-s)}dW(s)$$

$$g(t)=B\sigma\int_0^t (1-e^{-\alpha(t-s)})dW(s)$$

$$\psi(t,\tau)=-\frac{1}{2}\int_0^t \sigma^2(s,t+\tau)ds$$

Then

$$s(t,\tau)=\psi(t,\tau)+g(t)+u(\tau)f(t)$$

Differentiate this

$$df(t)=-\alpha f(t)dt+\sigma dW(t)$$

$$dg(t)=B(\sigma dW+\alpha f(t)dt-\sigma dW)=B\alpha f(t)dt$$

$$d\psi(t,\tau)=\left[-\frac{1}{2}\sigma^2(t,t+\tau)-\int_0^t \sigma(s,t+\tau)\frac{\partial\sigma}{\partial T}(s,t+\tau)ds\right]dt=\mu(t,\tau)dt$$

We recognize that $f(t)$ is an Ornstein-Uhlenbeck process mean-reverting to 0, $g(t)$ is an integral of f and $\psi(t,\tau)$ is a deterministic drift.

We can calculate $\mu_k(t,\tau)$ explicitly

$$\begin{aligned}
\mu_k(t,\tau) &= -\frac{1}{2}\sigma_k^2(t,t+\tau)-\int_0^t \sigma(s,t+\tau)\frac{\partial\sigma}{\partial T}(s,t+\tau)ds \\
&= -\frac{1}{2}\sigma_k^2(A_k e^{-\alpha_k\tau}+B_k)^2+\sigma_k^2\alpha_k\int_0^t (A_k e^{-\alpha_k(t+\tau-s)}+B_k)A_k e^{-\alpha_k(t+\tau-s)}ds \\
&= -\frac{1}{2}\sigma_k^2(A_k e^{-\alpha_k\tau}+B_k)^2+\frac{1}{2}\sigma_k^2 A_k^2 e^{-2\alpha_k\tau}(1-e^{-2\alpha_k t})+\sigma_k^2 A_k B_k e^{-\alpha_k\tau}(1-e^{-\alpha_k t}) \\
&= -\frac{1}{2}\sigma_k^2 A_k^2 e^{-2\alpha_k(t+\tau)}-\sigma_k^2 A_k B_k e^{-\alpha_k(t+\tau)}-\frac{1}{2}\sigma_k^2 B_k^2 \\
\mu_k(t,\tau) &= -\frac{1}{2}\sigma_k^2(A_k e^{-\alpha_k(t+\tau)}+B_k)^2=-\frac{1}{2}\sigma_k^2 u_k(t+\tau)^2
\end{aligned}$$

Thus the constant-maturity futures price can be written as:

$$\log f(t,\tau)=\log F(0,t+\tau)+\psi_1(t,\tau)+\psi_2(t,\tau)+g_1(t)+g_2(t)+u_1(\tau)f_1(t)+u_2(\tau)f_2(t)$$

where:

$$df_k(t) = -\alpha_k f_k(t)dt + \sigma_k dW_k(t)$$
$$dg_k(t) = B_k \alpha_k f_k(t)dt$$
$$d\psi_k(t,\tau) = -\frac{1}{2}\sigma_k^2 u_k^2(t+\tau)dt$$

5. Impact of a third factor on the constant-maturity forward curve

We have, as in Appendix 4, that

$$\log f(t,\tau) = \log F(t,t+\tau) = \log F(0,t+\tau) + \sum_{k=1}^{3} -\frac{1}{2}\int_0^t \sigma_k^2(s,t+\tau)ds + \int_0^t \sigma_k(s,t+\tau)dW_k(s)$$

Let:

$$s_k(t,\tau) = -\frac{1}{2}\int_0^t \sigma_k^2(s,t+\tau)ds + \int_0^t \sigma_k(s,t+\tau)dW_k(s)$$

We consider only the third factor and will assume $k = 3$ in what follows. Let us consider first the stochastic part:

$$\int_0^t \sigma(s,t+\tau)dW(s) = \sigma\int_0^t (Ae^{-2\alpha(t+\tau-s)} + Be^{-\alpha(t+\tau-s)} + C)dW(s)$$
$$= u(\tau)\int_0^t \sigma e^{-2\alpha(t-s)}dW(s) + Be^{-\alpha\tau}\int_0^t \sigma e^{-\alpha(t-s)}(1-e^{-\alpha(t-s)})dW(s)$$
$$+ C\int_0^t \sigma(1-e^{-2\alpha(t-s)})dW(s)$$

Let:

$$f(t) = \int_0^t \sigma e^{-2\alpha(t-s)}dW(s), \quad g(t) = \int_0^t \sigma e^{-\alpha(t-s)}(1-e^{-\alpha(t-s)})dW(s), \quad h(t) = \int_0^t \sigma(1-e^{-2\alpha(t-s)})dW(s)$$

Then

$$df(t) = -2\alpha f(t) + \sigma dW(t) \qquad \text{(Ornstein-Uhlenbeck process)}$$
$$dg(t) = \alpha(f(t) - g(t))dt \qquad g(t) = \alpha\int_0^t e^{-\alpha(t-s)} f(s)ds$$
$$dh(t) = 2\alpha f(t)dt \qquad h(t) = 2\alpha\int_0^t f(s)ds$$

The process $f(t)$ is an Ornstein-Uhlenbeck process mean-reverting to zero with mean-reversion speed 2α and volatility σ. The processes $g(t)$ and $h(t)$ are stochastic drifts – integrals of $f(t)$ with different weights.

6. Black volatilities of the Average price contract

In-settlement

We consider a date $T_M \le t < T_{M+1}$.

$$\int_t^T \sigma_A^2(s,T)ds = \int_t^T \left[\left(\sigma_S \frac{1-e^{-\alpha(T-s)}}{\alpha c_M^{'}} + \rho \frac{T-s}{c_M^{'}} \sigma_L \right)^2 + (1-\rho^2) \left(\frac{T-s}{c_M^{'}} \right)^2 \sigma_L^2 \right] ds$$
$$= \int_0^{T-t} \left[\frac{\sigma_S^2}{\alpha^2 c_M^{'2}} (1-e^{-\alpha s})^2 + \frac{2\rho\sigma_S\sigma_L}{\alpha c_M^{'2}} (1-e^{-\alpha s}) s + \frac{s^2}{c_M^{'2}} \sigma_L^2 \right] ds$$

Let us calculate each of the terms separately:

$$\int_0^{T-t} \frac{\sigma_S^2}{\alpha^2 c_M^{'2}} (1-e^{-\alpha s})^2 ds = \frac{\sigma_S^2}{\alpha^2 c_M^{'2}} \left(T-t-\frac{2}{\alpha}\left(1-e^{-\alpha(T-t)}\right) + \frac{1-e^{-2\alpha(T-t)}}{2\alpha} \right)$$

$$\int_0^{T-t} \frac{2\rho\sigma_S\sigma_L}{\alpha c_M^{'2}} (1-e^{-\alpha s}) s ds = \frac{2\rho\sigma_S\sigma_L}{\alpha c_M^{'2}} \left[\left[\frac{s^2}{2} \right]_0^{T-t} - \left[s\frac{e^{-\alpha s}}{-\alpha} \right]_0^{T-t} + \left[\frac{e^{-\alpha s}}{\alpha^2} \right]_0^{T-t} \right]$$
$$= \frac{2\rho\sigma_S\sigma_L}{\alpha c_M^{'2}} \left[\frac{(T-t)^2}{2} + \frac{(T-t)e^{-\alpha(T-t)}}{\alpha} - \frac{1}{\alpha^2}\left(1-e^{-\alpha(T-t)}\right) \right]$$

$$\int_0^{T-t} \frac{s^2}{c_M^{'2}} \sigma_L^2 ds = \frac{1}{3} \frac{(T-t)^3}{c_M^{'2}} \sigma_L^2$$

Such that the square of the Black volatility is given by:

$$\sigma_{Black}(t,T)^2 = \frac{1}{T-t}\int_t^T \sigma_A^2(s,T)ds$$
$$= \frac{\sigma_S^2}{\alpha^2 c_M^{'2}}\left(1-\frac{2}{\alpha}\frac{1-e^{-\alpha(T-t)}}{T-t}+\frac{1}{2\alpha}\frac{1-e^{-2\alpha(T-t)}}{T-t}\right)+\cdots$$
$$\frac{2\rho\sigma_S\sigma_L}{\alpha c_M^{'2}}\left(\frac{T-t}{2}+\frac{e^{-\alpha(T-t)}}{\alpha}-\frac{1}{\alpha^2}\frac{1-e^{-\alpha(T-t)}}{T-t}\right)+\cdots$$
$$\frac{1}{3}\frac{(T-t)^2}{c_M^{'2}}\sigma_L^2$$

Let us consider the case when $\alpha c \ll 1$ and simplify this expression

$$\sigma_{Black}(t,T)^2 \approx \frac{1}{3}\frac{\sigma_S^2}{c_M^{'2}}(T-t)^2 + \frac{2\rho\sigma_S\sigma_L}{3c_M^{'2}}(T-t)^2 + \frac{1}{3}\frac{\sigma_L^2}{c_M^{'2}}(T-t)^2$$
$$\sigma_{Black}(t,T) \approx \frac{1}{\sqrt{3}}\frac{T-t}{c_M^{'}}\sqrt{\sigma_S^2+2\rho\sigma_S\sigma_L+\sigma_L^2}$$

and $c_M^{'} = T_N - T_M \approx T-t$ such that

$$\sigma_{Black}(t,T) \approx \frac{1}{\sqrt{3}}\sqrt{\sigma_S^2+2\rho\sigma_S\sigma_L+\sigma_L^2}$$

Pre-settlement

$$\int_t^T \sigma_A^2(s,T)ds = \int_t^{T_1}\sigma_A^2(s,T)ds + (T-T_1)\sigma_{Black}^2(T_1,T)$$

The second term is known from the calculations above. Let us calculate the first term.

$$\int_t^{T_1}\sigma_A^2(s,T)ds = \int_t^{T_1}\left[\left(\sigma_S e^{-\alpha(T-s)}\frac{e^{\alpha c}-1}{\alpha c}+\rho\sigma_L\right)^2+(1-\rho^2)\sigma_L^2\right]ds$$
$$= \int_t^{T_1}\sigma_S^2 e^{-2\alpha(T-s)}\left(\frac{e^{\alpha c}-1}{\alpha c}\right)^2 ds + \int_t^{T_1} 2\rho\sigma_S\sigma_L\frac{e^{\alpha c}-1}{\alpha c}e^{-\alpha(T-s)}ds + (T_1-t)\sigma_L^2$$
$$= \sigma_S^2\left(\frac{e^{\alpha c}-1}{\alpha c}\right)^2\frac{e^{-2\alpha(T-T_1)}-e^{-2\alpha(T-t)}}{2\alpha} + 2\rho\sigma_S\sigma_L\frac{e^{\alpha c}-1}{\alpha c}\frac{e^{-\alpha(T-T_1)}-e^{-\alpha(T-t)}}{\alpha} + (T_1-t)\sigma_L^2$$

If we assume that $\alpha c \ll 1$, and noticing that $T-T_1 = c$

$$\sigma_{Black}(t,T)^2 \approx \sigma_S^2 \frac{e^{-2\alpha c} - e^{-2\alpha(T-t)}}{2\alpha(T-t)} + 2\rho\sigma_S\sigma_L \frac{e^{-\alpha c} - e^{-\alpha(T-t)}}{\alpha(T-t)} + \left(1 - \frac{c}{T-t}\right)\sigma_L^2$$
$$+ \frac{c}{T-t}\frac{1}{3}(\sigma_S^2 + 2\rho\sigma_S\sigma_L + \sigma_L^2)$$
$$\sigma_{Black}(t,T)^2 \approx \sigma_S^2 \frac{c}{T-t}\left[\frac{1-e^{-2\alpha(T_1-t)}}{2\alpha c} + \frac{1}{3}\right] + 2\rho\sigma_S\sigma_L \frac{c}{T-t}\left[\frac{1-e^{-\alpha(T_1-t)}}{\alpha c} + \frac{1}{3}\right] + \sigma_L^2\left[1 - \frac{2}{3}\frac{c}{T-t}\right]$$

We check that when $t = T_1$ (i.e the contract enters settlement):

$$\sigma_{Black}(T_1,T)^2 \approx \frac{1}{3}\sigma_S^2 + \frac{2\rho\sigma_S\sigma_L}{3} + \frac{1}{3}\sigma_L^2$$

and when $T - t \to \infty$, $\sigma_{Black}(t,T) \to \sigma_2$

7. Semi-analytical solution to the optimal stopping problem

We begin by presenting the analysis in the simple case of one factor. The continuation region is then given by $S^*(t) = (-\infty, f^*(t))$. The equation satisfied by the value function V is

$$\frac{\partial V}{\partial t} + (\mu - \alpha f)\frac{\partial V}{\partial f} + \frac{1}{2}\sigma^2\frac{\partial^2 V}{\partial f^2} - g(t), \quad f < f^*(t)$$
$$V(t, f) = \Omega(t, f), \quad f \geq f^*(t)$$
$$V(T, f) = \Omega(T, f)$$

Let

$$\mathrm{L}V = (\mu - \alpha f)\frac{\partial V}{\partial f} + \frac{1}{2}\sigma^2\frac{\partial^2 V}{\partial f^2} - g(t)$$

such that

$$\frac{\partial V}{\partial t} + LV = \psi(t, f), \qquad \psi(t, f) = \begin{cases} 0 & f < f^*(t) \\ \frac{\partial \Omega}{\partial t} + L\Omega & f \geq f^*(t) \end{cases}$$

The transition density for the Ornstein-Uhlenbeck process is

$$p(f',\tau,f,t)=\frac{1}{\left(2\pi\sigma^2\frac{1-e^{-2\alpha(\tau-t)}}{2\alpha}\right)^{1/2}}\exp\left[-\frac{1}{2}\frac{\left(f'-\left(fe^{-\alpha(\tau-t)}+\frac{\mu}{\alpha}(1-e^{-\alpha(\tau-t)})\right)\right)^2}{\sigma^2\frac{1-e^{-2\alpha(\tau-t)}}{2\alpha}}\right]$$

This transition density satisfies the equation

$$\frac{\partial p}{\partial t}+(\mu-\alpha f)\frac{\partial p}{\partial f}+\frac{1}{2}\sigma^2\frac{\partial^2 p}{\partial f^2}=0$$

The value function can then be written as

$$\begin{aligned}V(t,f)&=\int_{-\infty}^{\infty}p(f',T;f,t)(\Omega(T,f')-G(t,T))df'+\int_t^T\int_{-\infty}^{\infty}p(f',s;f,t)(-\psi(s,f'))df'ds\\&=V_{eur}(t,f)+V_{early}(t,f)\end{aligned}$$

where $G(t,T)$ is the cost of shipping between times t and T:

$$G(t,T)=\int_t^T g(s)ds$$

Let us verify this result by differentiating the above formula. Consider first the European value $V_{eur}(t,f)$:

$$\begin{aligned}\frac{\partial V_{eur}}{\partial t}&=\int_{-\infty}^{\infty}\frac{\partial p}{\partial t}(f',T;f,t)(\Omega(T,f')-G(t,T))df'-\int_{-\infty}^{\infty}p(f',T;f,t)\frac{\partial G}{\partial t}df'\\&=-(\mu-\alpha f)\frac{\partial V_{eur}}{\partial f}-\frac{1}{2}\sigma^2\frac{\partial^2 V_{eur}}{\partial f^2}+g(t)\end{aligned}$$

and the early-exercise premium:

$$\begin{aligned}\frac{\partial V_{early}}{\partial t}&=\int_{-\infty}^{+\infty}p(f',t;f,t)\psi(t,f')df'+\int_t^T\int_{-\infty}^{+\infty}\frac{\partial p(f',s;f,t)}{\partial t}(-\psi(s,f'))df'ds\\&=\psi(t,f)+\int_t^T\int_{-\infty}^{+\infty}\left(-(\mu-\alpha f)\frac{\partial p}{\partial f}-\frac{1}{2}\sigma^2\frac{\partial^2 p}{\partial f^2}\right)(-\psi(s,f'))df'ds\\&=-(\mu-\alpha f)\frac{\partial V_{early}}{\partial f}-\frac{1}{2}\sigma^2\frac{\partial^2 V_{early}}{\partial f^2}+\psi(t,f)\end{aligned}$$

Furthermore, V_{eur} and V_{early} satisfy the terminal conditions

$$V_{eur}(T,f)=\int_{-\infty}^{\infty} p(f',T;f,T)(\Omega(T,f')-G(T,T))df'=\Omega(T,f)$$

$$V_{early}(T,f)=\int_{T}^{T}\int_{-\infty}^{\infty} p(f',s;f,T)(-\psi(s,f'))df'ds=0$$

such that $V=V_{eur}+V_{early}$ solves the equation.

The early exercise premium can be written in terms of the stopping boundary as:

$$V_{early}(t,f)=\int_{t}^{T}\int_{f^*(s)}^{\infty} p(f',s;f,t)(-\psi(s,f'))df'ds$$

This formulation gives a closed form expression of *V*. However, it involves the values of $f^*(s)$ for $t\le s\le T$. These are determined by the continuity condition

$$\begin{aligned}\Omega(t,f^*(t))&=V(t,f^*(t))\\ &=\int_{-\infty}^{\infty} p(f',T;f^*(t),t)(\Omega(T,f')-G(t,T))df'+\int_{t}^{T}\int_{f^*(s)}^{\infty} p(f',s,f^*(t),t)(-\psi(s,f'))df'ds\end{aligned}$$

The value function at maturity t and the stopping boundary $f^*(t)$ can be determined recursively as follows: discretize the dates as $t_0=0,t_1,...,t_N=T$. If $f^*(t_N),...,f^*(t_{k+1})$ have been calculated, let

$$\begin{aligned}F(f_k^*)=&\int_{-\infty}^{\infty} p(f',T,f_k^*,t)(\Omega(T,f')-G(t,T))df'+\sum_{j=k+1}^{N-1}\Delta t\int_{f^*(t_j)}^{\infty} p(f',t_j,f_k^*,t_k)(-\psi(t_j,f'))df'\\ &+\Delta t\int_{f_k^*}^{\infty} p(f',t_k;f_k^*,t_k)(-\psi(t_k,f'))df'\end{aligned}$$

Finding the stopping boundary $f^*(t_k)$ at time t_k involves finding, numerically,

$$f^*(t_k)=\min\left\{f_k^*,F(f_k^*)=\Omega(t_k,f_k^*)\right\}$$

Once this stopping boundary has been located the value function can be calculated for all *f* using

$$V(t,f)=\int_{-\infty}^{\infty} p(f',T,f,\tau)(\Omega(T,f')-G(t,T))df'+\sum_{j=k}^{N-1}\Delta t\int_{f^*(t_j)}^{\infty} p(f',t_j,f,t_k)(-\psi(t_j,f'))df'$$

We can extend this analysis to two factors, by using the fact that they are independent. The transition density function for the joint Ornstein-Uhlenbeck process is

$$p(f',\tau;f,t)=p_1(f_1',\tau;f_1,t)\,p_2(f_2',\tau;f_2,t)$$

where p_1 and p_2 are the transition densities for the one-dimensional Ornstein-Uhlenbeck processes. This two-dimensional transition density function solves the partial differential equation

$$\frac{\partial p}{\partial t}+(\mu_1-\alpha_1 f_1)\frac{\partial p}{\partial f_1}+\frac{1}{2}\sigma_1^2\frac{\partial^2 p}{\partial f_1^2}+(\mu_2-\alpha_2 f_2)\frac{\partial p}{\partial f_2}+\frac{1}{2}\sigma_2^2\frac{\partial^2 p}{\partial f_2^2}=0$$

The equation for the value function is

$$\frac{\partial V}{\partial t}+LV=\psi(t,f)$$

$$\psi(t,f)=\begin{cases}0 & f\in S^*(t)\\ \dfrac{\partial\Omega}{\partial t}+L\Omega & f\notin S^*(t)\end{cases}$$

where

$$\mathrm{L}V=(\mu_1-\alpha_1 f_1)\frac{\partial V}{\partial f_1}+(\mu_2-\alpha_2 f_2)\frac{\partial V}{\partial f_2}+\frac{1}{2}\sigma_1^2\frac{\partial^2 V}{\partial f_1^2}+\frac{1}{2}\sigma_2^2\frac{\partial^2 V}{\partial f_2^2}-g(t)$$

such that the value function can be written as

$$\begin{aligned}V(t,f)&=\int_{\mathbb{R}^2} p(f',T;f,t)(\Omega(T,f')-G(t,T))df'+\int_t^T\int_{\mathbb{R}^2\setminus S^*(s)} p(f',s;f,t)(-\psi(s,f'))df'ds\\&=V_{eur}(t,f)+V_{early}(t,f)\end{aligned}$$

In the two-dimensional case the continuation region $S^*(t)$ is defined as

$$S^*(t)=\{(f_1,f_2),V(t,f_1,f_2)>\Omega(t,f_1,f_2)\}$$

The boundary $\partial S^*(t)$ of this domain has to be determined for each date t. We write it as a function of the second factor

$$\partial S^*(t)=\{(f_1^*(t,f_2),f_2),f_2\in\mathbb{R}\}$$

such that the equation to be solved by $f_1^*(t,f_2)$ is

$$
\begin{aligned}
V(t, f_1^*(t, f_2), f_2) &= \int_{\mathbb{R}^2} p(f_1^{'} f_2^{'}, T; f_1^*(t, f_2), f_2, t)(\Omega(T, f') - G(t,T))df' \\
&\quad + \int_t^T \int_{-\infty}^{+\infty} \int_{f_1^*(s,f_2)}^{+\infty} p_1(f_1^{'}, f_2^{'}, s; f_1^*(t, f_2), f_2, t)(-\psi(s, f'))df'ds \\
&= \Omega(t, f_1^*(t, f_2), f_2)
\end{aligned}
$$

To find the function $f_1^*(t, f_2)$ we proceed recursively as in the one-factor case. Having determined $f_1^*(t_j, f_2)$ for $j > k$, we calculate

$$
\begin{aligned}
F(f_1^*, f_2) &= \int_{\mathbb{R}^2} p(f_1^{'}, f_2^{'}, T; f_1^*, f_2, t_k)(\Omega(T, f') - G(t_k, T))df_1^{'}df_2^{'} \\
&\quad + \sum_{j=k+1}^{N-1} \Delta t \int_{-\infty}^{+\infty} \int_{f_1^*(t_j, f_2^{'})}^{\infty} p(f_1^{'}, f_2^{'}, t_j; f_1^*, f_2, t_k)(-\psi(t_j, f'))df_1^{'}df_2^{'} \\
&\quad + \Delta\tau(-\psi(t_k, f_1^*, f_2))
\end{aligned}
$$

and we find the exercise boundary by varying f_1^*:

$$
f_1^*(t_k, f_2) = \min\left\{ f_1^*, F(f_1^*, f_2) = \Omega(t_k, f_1^*, f_2)\right\}
$$

Once this exercise boundary has been located the value function can be calculated for all *f* using

$$
V(t_k, f_1, f_2) = \int_{\mathbb{R}^2} p(f', T; f, t_k)(\Omega(T, f') - G(t_k, T))df' + \sum_{j=k}^{N-1} \Delta\tau \int_{-\infty}^{+\infty} \int_{f_1^*(t_j, f_2)}^{\infty} p(f', t_j; f, t_k)(-\psi(t_j, f'))df'
$$

8. Routes, cargoes and ships used in the floating storage trades

Crude oil Sullom Voe – LOOP

Route	Sullom Voe – LOOP	**Ship**[1]	
Distance *d*	4535 Nm	Type	Very Large Crude Carrier (VLCC)
Cargo	Brent	Cargo size	270 000 mt
Barrel factor	7.578 bbl/mt	DWT	300 000 mt
Loading port	Sullom Voe	Speed u	15 knots
Loading price S_0	Dated Brent 10-21 days	Fuel consumption sailing: FC(u) anchor: FC_a	 87.5 mt/day (laden) 74 mt/day (ballast) 85 mt/day (pumping) 15 mt/day (anchor)
Loading delay τ_{load}	15 days	Timecharter price H	VLCC average timecharter equivalent (Baltic Exchange)
IFO price B	Fuel Oil 3.5% CIF NWE (Platts)	**Delivery port**	LOOP
		Delivery price $F(t,\tau)$	LLS forward curve

Heating oil ARA – NYH

Route	ARA – NYH	**Ship**	
Distance *d*	3383 Nm	Type	Very Large Crude Carrier (VLCC)
Cargo	No. 2 fuel oil	Cargo size	270 000 mt
Barrel factor	312.63 gal/mt	DWT	300 000 mt
Loading port	Amsterdam-Rotterdam-Antwerp	Speed u	15 knots
Loading price S_0	ICE Gasoil front month price	Fuel consumption sailing: FC(u) anchor: FC_a	 87.5 mt/day (laden) 74 mt/day (ballast) 85 mt/day (pumping) 15 mt/day (anchor)
Loading delay τ_{load}	15 days	Timecharter price H	VLCC average timecharter equivalent (Baltic Exchange)
IFO price B	Fuel Oil 3.5% CIF NWE (Platts)	**Delivery port**	New York Harbor
		Delivery price $F(t,\tau)$	Nymex heating oil forward curve

[1] Corresponds to the modern double-hull VLCC from Clarksons (2009)

7. REFERENCES

Albanese, C. & Campolieti, G. (2006). *Advanced Derivatives Pricing and Risk Management.* Elsevier Academic Press.

Avougleas, K. & Sclavounos, P.D. (2009). Fuel efficient ship routing. *Massachusetts Institute of Technology*, Working paper.

Baker, M.P., Mayfield, E.S., & Parsons, J.E. (1998). Alternative models of uncertain commodity prices for use with modern asset pricing methods. *The Energy Journal, 19, 1,* 115-148

Basilevsky, A. (1994). *Statistical Factor Analysis and Related Methods.* John Wiley & Sons.

Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics, 3*, 167-179.

Brennan, M.J. & Schwartz, E.S. (1985). Evaluating Natural Resource Investments. *The Journal of Business*, Vol. 58. No. 2 (Apr. 1985), pp. 135-157

Carr, P., Ewald, C.-O., & Xiao, Y. (2008). On the qualitative effect of volatility and duration on prices of Asian options. *Finance Research Letters*, 5, pp. 162-171.

Central Intelligence Agency (2010). Country Comparison: oil production. *CIA World Factbook*, https://www.cia.gov/library/publications/the-world-factbook/. Accessed May 7, 2010.

Chang, F.-R. (2004). *Stochastic Optimization in Continuous Time.* Cambridge University Press.

Clarksons (2009). *Sources & Methods for the Shipping Intelligence Weekly*, May 2009.

Clewlow, L. & Strickland, C. (1998).*Implementing Derivative Models.* John Wiley & Sons.

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

CME Group (2009). *NYMEX/COMEX Exchange ADV Report – Monthly*, December 2009.

Dewynne, J.N. & Wilmott, P. (1995). A Note on Average Rate Options with Discrete Sampling. *SIAM Journal on Applied Mathematics*, Vol. 55, No. 1 (Feb. 1995), pp. 267-276

Dixit, A.K., & Pindyck, R.S. (1994). *Investment Under Uncertainty.* Princeton University Press.

Geman, H., & Yor, M. (1992) Quelques relations entre processus de bessel, options asiatiques, et fonctions confluences hypergeometriques. *Comptes Rendus De l'Académie Des Sciences De Paris, I*, 471-474.

Gibson, R. & Schwartz, E.S. (1990). Stochastic Convenience Yield and the Pricing of Oil Contingent Claims. *The Journal of Finance*, Vol. 53, No. 3 (Jul. 1990), pp. 959-976.

Haug, E.G. (2006). *The Complete Guide to Option Pricing Formulas.* McGraw-Hill, 2nd edition.

Haykov, J.M. (1993). A better control variate for pricing standard Asian options. *Journal of Financial Engineering*, Vol. 2, No. 3, pp. 207-216.

Heath, D., Jarrow, R.A. & Morton, A. (1992) Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Valuation. *Econometrica*, Vol. 60, No. 1, pp. 77-105.

InterContinental Exchange (2009). *ICE Report Center, MTD/YTD Volume for 2009.* https://www.theice.com/marketdata/reports/ReportCenter.shtml. Accessed May 7, 2010.

International Maritime Exchange (2010). *Tanker FFA Trading Statistics.* http://www.exchange.imarex.com/information-services/volume-statistics/tanker-ffa/. Accessed May 7, 2010.

Joy, C., Boyle, P.P. & Tan, K.S. (1996). Quasi-Monte Carlo Methods in Numerical Finance. *Management Science*, Vol. 42, No. 6 (June 1996). pp. 926-938.

Kemna, A.G.Z & Vorst, A.C.F. (1989). A pricing method for options based on average asset values. *Journal of Banking and Finance*, Vol. 14, pp. 113-129.

Koekebakker, S., & Ollmar, F. (2005). *Clarkson Securities Freight Option Model – Technical Document.* http://www.clarksonsecurities.com/files/pdf/clarkson_technical_doc.pdf.

Koekebakker, S., Ådland, R., & Sødal, S. (2007). Pricing freight rate options. *Transportation Research Part E, 43*, 535-548.

Levy, E. (1992). Pricing european average rate options. *Journal of International Money and Finance, 14*, 474-491.

Longstaff, F.A., Santa-Clara, P. & Schwartz, E.S. (2001) Throwing away a billion dollars: The Cost of Suboptimal Exercise Strategies in the Swaptions Market. *Journal of Financial Economics*, Vol. 62, No. 1 (Oct. 2001), pp. 39-66.

Longstaff, F.A. & Schwartz, E.S. (2001) Valuing American Options by Simulation: A Simple Least-Squares Approach. *The Review of Financial Studies*, Vol. 14, No. 1, pp. 113-147

Morimoto, H. (2010). *Stochastic control and mathematical modeling. Applications in Economics.* Cambridge University Press

Miltersen, K.R. & Schwartz, E.S. (1998). Pricing of Options on Commodity Futures with Stochastic Term Structures of Convenience Yields and Interest Rates. *The Journal of Financial and Quantitative Analysis*, Vol. 33, No. 1 (Mar., 1998), pp. 33- 59

Musiela, M. & Rutkowski, M. (2008). *Martingale Methods in Financial Modeling.* Springer.

Paddock, J.L., Siegel, D.R. and Smith, J.L. (1988), Option Valuation of Claims on Real Assets: The Case of Offshore Petroleum Leases. *The Quarterly Journal of Economics*, Vol. 103, No. 3 (Aug. 1988), pp 479-508.

Pilipovic, D. (2007). *Energy Risk. Valuing and Managing Energy Derivatives.* McGraw Hill.

Rebonato, R. (2002). *Modern pricing of interest rate derivatives: The LIBOR market model and beyond.* Princeton University Press.

Rogers, L.C.G & Shi, Z. (1995). The Value of an Asian Option. *Journal of Applied Probability*, Vol. 32, No. 4 (Dec. 1995), pp. 1077-1088.

Samuelson, P.A. (1965). Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review*, Vol. 6, No. 2 (Spring 1965), pp. 41-49

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

Sclavounos, P. D., & Ellefsen, P. E. (2009). Multi-factor model of correlated commodity forward curves for crude oil and shipping markets. *Massachusetts Institute of Technology. Working Paper. Center for Energy and Environmental Policy Research (CEEPR)*

Stopford, M. (2009). *Maritime Economics*. Routledge, 3rd edition.

Turnbull, S.M. & Wakeman, L.M. (1991). A Quick Algorithm for Pricing European Average Options. *The Journal of Financial and Quantitative Analysis*, Vol. 26, No. 3 (Sep. 1991), pp. 377 -389.

Yergin, D. (2008). *The Prize: The Epic Quest for Oil, Money & power.* Free Press, New edition.

Ådland, R. (2003). *The stochastic behavior of spot freight rates and the risk premium in bulk shipping*. PhD Thesis, Massachusetts Institute of Technology.