A Contango-Constrained Model for Storable Commodities*

Diana R. Ribeiro† Stewart D. Hodges‡

October 13, 2004

Abstract

In this article we develop a model for the commodity price dynamics under the risk-neutral measure where the spot price switches between two distinct stochastic processes depending on whether or not inventory is being held. Specifically, whenever the drift of the spot price exceeds the cost of carrying inventory (interest rate plus storage costs) the inventory is being held. Conversely, whenever the drift of the spot price is less than the cost of carry, all the inventory is sold and the storage facility becomes empty. If inventory is being held, we assume that the spot price follows a geometric Brownian motion with drift equal to the cost of carrying inventory. Otherwise, the price follows a Ornstein-Uhlenbeck stochastic process. This model verifies arbitrage-free arguments since the commodity price process has a drift lees or equal to the cost of carry under the risk neutral measure. We illustrate and analyze the properties of the spot price and the forward curves implied by this model using numerical examples. The spot price sample paths and the corresponding forward curves are constructed by applying trinomial tree techniques. For comparison, we also provide the equivalent numerical examples for the single-factor model provided by Schwartz (1997), which correspond to the unconstrained version of the spot price process in this model.

*Diana Ribeiro thanks Fundacao para a Ciencia e Tecnologia, Portugal for the partial financial support provided for this project. We would like to thank Elizabeth Whalley for her helpful comments. All the errors remain our responsibility.

†Doctoral Researcher at the Warwick Business School, University of Warwick, Coventry, CV4 7AL, United Kingdom, phone +44 (0)24 76524465, fax: +44 (0)24 7652 7167 and email: D.Ribeiro@wbs.ac.uk

‡Director of the Financial Options Research Center (FORC), Warwick Business School, University of Warwick, Coventry CV4 7AL, United Kingdom, phone: +44 (0)24 76523606, fax: +44 (0)24 7652 4176 and email forcsh@wbs.warwick.ac.uk

1 Introduction

Over the past two decades energy markets have undergone significant changes. The market has evolved from a monopolistic stable pricing environment to a volatile pricing environment. This has created the need for new analytical tools to price accurately commodity related securities and projects. In particular, stochastic models for commodity prices play a central role in pricing these contracts. This has motivated the recent development of new models to replicate the dynamics of commodity spot prices and futures valuation.

Early models for commodity prices assume that commodity prices follow a geometric Brownian motion. This specification is inappropriate since it does not capture the mean reversion property of spot commodity prices. Accordingly, numerous authors have developed models that focus on replicating this property in the price dynamics. Mean reversion on commodity prices appears to be directly related to reaction to market events, which causes imbalances between demand and supply. Either a correction on the supply side, to match the demand side, or the actual dissipation of the events tends to cause the commodity prices to come back to their typical levels. Single-factor models that represent the spot price as a stochastic process of the Ornstein-Uhlenbeck type has been extensively used in the real options literature (e.g. Dixit and Pindyck (1994), Metcalf and Hasset (1995) and Epstein et al. (1998) and in the context of futures valuation (Schwartz (1997)). This representation, however, violates standard cash-and-carry no-arbitrage arguments and leads to severe disparities when pricing futures and options on futures contracts. Gibson and Schwartz (1990) and Schwartz (1997) introduce a two-factor model where the spot price and the convenience yield follow joint stochastic processes with constant correlation. Specifically, the spot price follows a geometric Brownian motion and the convenience yield follows a mean reverting stochastic process of the Ornstein-Uhlenbeck process and enters the drift of the spot price under the risk-neutral measure. Schwartz (1997) three-factor model, Miltersen and Schwartz (1998) and Hilliard and Reis (1998) add a third stochastic factor to the model to account for stochastic interest rates. Nevertheless, the inclusion of stochastic interest rates in the commodity price models does not have a significant impact in the pricing of commodity options and futures in practice. For this reason the two-factor class of models has been dominating the current literature and practice. However, modelling the convenience yield as a O-U process does not preclude negative values and therefore may create cash-and-carry arbitrage possibilities. Moreover, specifying the convenience yield as an exogenous "dividend yield" to the spot price seems illusive. In fact, the convenience yield is not an observable variable but it is derived from the relationship between spot and forward prices.

The empirical calibration of the two-factor models developed by Schwartz (1997), Nielsen and Schwartz (2004) and Ribeiro and Hodges (2004b) show that there is a very strong correlation between the spot price and the convenience yield dynamics, particularly in the oil market and industrial metals. This motivates the question to whether a two-factor model is essential to capture the commodity spot price movements or if an arbitrage-free single-factor model would explain much of the empirical behaviour. In this article, we explore the answer to this question by studying the properties of the simplest mean reverting model possible which satisfies the cash-and-carry arbitrage-free constrain. We develop a model for the commodity price dynamics under the risk-neutral measure where the spot price switches between two distinct stochastic processes depending on whether or not inventory is being held. Specifically, whenever the drift of the spot price exceeds the cost of carrying inventory (interest rate plus storage costs) the inventory is being held. Conversely, whenever the drift of the spot price is less than the cost of carry, all the inventory is sold and the storage facility becomes empty. If inventory is being held, we assume that the spot price follows a geometric Brownian motion (GBM) with drift equal to the cost of carrying inventory. Otherwise, the price follows a O-U stochastic process[1]. This model verifies arbitrage-free arguments since the commodity price process has a drift less or equal to the cost of carry under the risk neutral measure.

We illustrate and analyze the properties of the spot price and the forward curves implied by this model using numerical examples. The spot price sample paths and

[1]The unconstrained version of this model is the same as the single-factor model described by Schwartz (1997).

the corresponding forward curves are constructed by applying trinomial tree techniques. For comparison, we also provide the equivalent numerical examples for the single-factor model provided by Schwartz (1997), which correspond to the unconstrained version of the spot price process in this model. This comparison assists on understanding of how the introduction of storage in our model modifies the commodity forward curve and the spot price distribution generated by an unconstrained O-U stochastic process.

The remaining of this paper is organized as follows. Section 2 describes the model. Section 3 describes the numerical implementation of the model. Section 4 presents and analyses the numerical results. Section 5 concludes.

2 Model Definition

The model is specified under the risk-neutral measure. We assume that the risk-free interest rate, r, is constant and therefore forward and futures prices are equivalent. Thereafter, we may refer to one or another without distinction. We also assume that the storage cost is given by $c \times p_t$ per unit stored, where $0 < c < 1$ represents a proportion of the spot price.

The commodity spot price process switches between two distinct stochastic processes, depending on whether or not inventory is being held. In the absence of storage, the spot price follows a mean reverting process of the O-U type. Whenever the drift in the spot price, p_t, exceeds the cost of carrying inventory (interest rate, r, plus storage cost, c) the stock holders hold stock. Whenever inventory is being held, the commodity price has the drift equal to $(r+c)$. Therefore, we assume that the spot price process follows a standard GBM whenever the stock is being held. This rule leads to the existence of a single critical price, p^*, in our model. Accordingly, the inventory holders buy stock as soon as p_t falls below p^*. On the other hand, as soon as p_t rises above p^*, all the inventory is sold. In the latter case, the spot price switches back to the mean-reverting stochastic process.

Let p_t be the commodity spot price. The stochastic process for the spot price switches between the two following components:

$$dp_t = \alpha(m - \ln p_t)p_t dt + \sigma p_t dB_t, \qquad \text{if} \qquad p_t \geq p^* \tag{1}$$

$$dp_t = (r + c)p_t dt + \sigma p_t dB_t, \qquad \text{otherwise,} \tag{2}$$

where:

- p_t is the spot price;
- m is a constant;
- α is the speed of mean reversion, a constant;
- σ is the volatility of the spot price, a constant;
- B_t is a standard Wiener process;
- r is the constant risk-free rate;
- c is the storage cost, which is a constant proportion of the spot price;
- p^* is the critical spot price.

Following the description above, the critical value spot price, p^*, is given by:

$$(r + c) = \alpha(m - \ln(p^*)), \qquad \text{that is,} \tag{3}$$

$$\ln(p^*) = m - \frac{(r + c)}{\alpha} \tag{4}$$

Equation (1) represents an alternative format to the standard geometric O-U process adopted by Dixit and Pindyck (1994), Metcalf and Hasset (1995) and Epstein et al. (1998) in the context of real options. This alternative eases the numerical implementation of the O-U process. Possibly also due to this advantage, Schwartz (1997) also uses this format to represent the single-factor model of a mean-reverting commodity price. Accordingly, Schwartz single-factor model corresponds to the unconstrained version of our model, that is, in the absence of the spot price drift constraint.

Defining $x_t = \ln p_t$ and applying Ito's lemma[2], the log price follows the O-U stochastic process:

$$dx_t = \alpha(\overline{x} - x_t)dt + \sigma dB_t, \qquad \text{if} \qquad x_t \geq x^* \tag{5}$$
$$dx_t = \left((r+c) - \frac{1}{2}\sigma^2\right) dt + \sigma dB_t, \qquad \text{otherwise,} \tag{6}$$

where $x^* = \ln(p^*)$.

The relation between $\overline{x}$ and m is given by:

$$\overline{x} = m - \frac{\sigma^2}{2\alpha}. \tag{7}$$

The log price, x_t, corresponding to the O-U process given by equation (5) is normally distributed with mean and variance given by:

$$E[x_T|x_t] = \overline{x} + (x_t - \overline{x})e^{-\alpha(T-t)}; \tag{8}$$
$$V[x_T|x_t] = \frac{\sigma^2}{2\alpha}\left(1 - e^{-\alpha(T-t)}\right). \tag{9}$$

The branch of x_t corresponding to the GBM given by equation (6) is normally distributed with mean and variance given by:

$$E[x_T|x_t] = x_t + \left((r+c) - \frac{1}{2}\sigma^2\right)(T-t) \tag{10}$$
$$V[x_T|x_t] = \sigma^2(T-t) \tag{11}$$

Note, however, that the commodity price keeps switching between the two processes

[2]This derivation is described in Appendix A.

above and therefore none of the equations (8), (9), (10) and (11) describe the conditional moments of the resulting price distribution. However, these moments will be used in the numerical implementation of the commodity price lattice.

Under the risk-neutral measure, the futures price at time t for delivery at some future time $T > t$ is given by the expected spot price at time T conditional on the information available at t, that is:

$$F_{t,T} = E_t[p_T], \tag{12}$$

where E_t denotes the conditional expectation under the risk neutral measure given the information at time t. We apply this relationship to compute the forward curve for the commodity prices in the numerical implementation as described in the next section.

We compute the convenience yield from the relationship between the futures and the spot price of a commodity when the interest rate and the convenience yield are deterministic. Since the storage costs are expressed as a proportion c of the spot price, the convenience yield, δ is defined so that:

$$F_{t+\Delta t} = p_t e^{(r+c-\delta)\Delta t} \tag{13}$$

Based on this relationship, we calculate the annualized convenience yield for the time interval between t and $t + \Delta t$ by using pairs of adjacent maturities futures contracts according to the following formula:

$$\delta_{t,t+\Delta t} = (r + c) - \frac{1}{\Delta t} \ln\left(\frac{F_{t+\Delta t}}{F_t}\right) \tag{14}$$

Accordingly, the curve is in backwardation if the convenience yield, $\delta_{t,t+\Delta t}$ is greater than $(r+c)$; the curve is in contango if the convenience yield is less than $(r+c)$, that is, the forward curve increases with time to maturity. A negative convenience yield would imply the violation of the standard arbitrage-free condition for commodity prices:

$$F_{t+\Delta t} \leq p_t e^{(r+c)\Delta t}. \tag{15}$$

3 Numerical Implementation

We apply trinomial tree techniques to illustrate numerical examples of the forward curve and to analyze the properties implied by this model. In particular, we study the properties of the spot price distribution, the forward curve and the corresponding convenience yield. In order to understand the price dynamics generated by imposing a constraint on the drift of the spot prices in our model, we also compute a numerical examples of the unconstrained version of our model. As mentioned above, the unconstrained version of our model is the same as the Schwartz (1997) single-factor model.

Excess kurtosis and right skewness are two fundamental properties typically observed in storable commodity price distributions. In other words, commodity prices are characterized by long periods of stagnant prices interrupted by sharp upward prices. This asymmetry is a consequence of the inventory non-negativity constraint. Inventory can always be added to keep current spot prices from being too low. In other words, the existence of storage "cuts out" the left tail of the price distribution. However, stockholders are unable to respond to sudden demand/supply imbalances that lead to sudden upward rises in commodity prices. Hence, we are interested in verifying if our model generates right skewness and excess kurtosis in the commodity prices distribution.

We compute a single tree for the logarithm of the commodity prices paths, $x_t = \ln(p_t)$, which results from the combination of equations (5) and (6). That is, x_t follows a O-U process described by equation (5) if $x_t \geq x^*$ and follows the process given by equation (6) otherwise. We describe separately the local probability structure of the tree corresponding to each of the processes and then explain how to combine both procedures in order to obtain the final lattice for the model.

3.1 Lattice Description

A more detailed description of the trinomial lattice method for mean reverting stochastic process is presented in the Appendix B. Nevertheless, we repeat here the main steps of the procedure. The method used to implement a trinomial tree to represent the standard O-U process given by equation (5) is based on the technique described Hull and White (1993, 1994) and later revised by Clewlow and Strickland (1998) and Clewlow and Strickland (2000), which was originally designed to implement short interest rates that follow a mean-reverting arithmetic stochastic process.

The trinomial tree is constructed by using time steps of length Δt and x-steps of length Δx. At the end of each time step, x takes the value $x_0 + j\Delta x$, where j can be either positive or negative and x_0 is the initial value. (i, j) is defined as the node for which $t = i\Delta t$ and $x = j\Delta x$. As described in Appendix B, the trinomial branching process can take three different forms. It can take a normal branching process where we can move up by Δx, stay the same and move down by Δx; a branching process when $x_{i,j}$ is currently low and $x_{i,j}$ can stay the same, move up by Δx and move up by $2\Delta x$. Finally when $x_{i,j}$ is currently high, the price path can stay the same, move down by Δx and move down by $2\Delta x$[3]. In other words, the three nodes emanating from node (i, j) are $(i + 1, k + 1)$ - the "upper" node, (i+1, k) - the "middle" node and $(i + 1, k - 1)$ - the "lower" node. The value of k is chosen so that $x_{i+1,k}$ is as close as possible to the expected value of x, which by definition is given by $x_{i,j} + \mu_{i,j}$, where $\mu_{i,j} = \alpha(\overline{x} - (x_0 + j\Delta x))\Delta t$. For the normal branching process $k = j$, when $x_{i,j}$ is currently low $k = j + 1$ and when $x_{i,j}$ is currently high $k = j - 1$, respectively.

In order to obey stability and convergence conditions Hull and White (1990) suggest that a good relationship between Δt and the space step Δz is:

[3]See Figure 7 in Appendix A.

$$\Delta x = \sigma\sqrt{3\Delta t} \tag{16}$$

Define (i,j) as the node for which $t = i\Delta t$ and $x = j\Delta x$. Define $pr^u_{i,j}$, $pr^m_{i,j}$ and $pr^d_{i,j}$ as the probabilities of the highest, middle and lowest branches emanating from node (i,j). The probabilities are chosen to match the expected change and variance in x over the next interval Δt. The probabilities must also sum to unity. Accordingly, the resulting probabilities are given by:

$$\begin{aligned} pr^u_{i,j} &= \frac{\sigma^2\Delta t + \eta^2}{2\Delta x^2} + \frac{\eta}{2\Delta x} \\ pr^m_{i,j} &= 1 - \frac{\sigma^2\Delta t + \eta^2}{\Delta x^2} \\ pr^d_{i,j} &= \frac{\sigma^2\Delta t + \eta^2}{2\Delta x^2} - \frac{\eta}{2\Delta x} \end{aligned} \tag{17}$$

where $\eta = \mu_{i,j} + (j-k)\Delta x$ and $k = j-1, j$ and $j+1$, depending on the type of branching, as described above and $\mu_{i,j} = E[\Delta x|(i,j)] = \alpha(\overline{x} - x_{i,j})\Delta t$ is the conditional expectation of the discretized x_t process at node (i,j) and $\sigma^2_{i,j}\Delta t + (\mu_{i,j}\Delta t)^2 = E[\Delta x^2|(i,j)]$, where $\sigma_{i,j} = \sigma$ is from equation (5). Provided that Δx is within the range $\sigma\sqrt{3\Delta t}/2$ to $2\sigma\sqrt{\Delta t}$, the probabilities are always between 0 and 1 (Hull and White (1993)).

The implementation of a trinomial lattice for the GBM is simpler than it is in the case of a mean reversion stochastic process as described above. The procedure is standard (see, e.g. Clewlow and Strickland (1998)). Now, there are no branching decisions to be made, that is, we always have a "normal" branching process.

As before, we work in terms of $x_t = \ln(p_t)$. The trinomial tree is constructed by using time steps of length Δt and x-steps of length Δx. Again, at the end of each time-step, x takes the value $x_0 + j\Delta x$, where j can be either positive or negative and x_0 is the initial value. Each node of the tree is represented by (i,j), for which $t = i\Delta t$ and $x = j\Delta x$. At each node, $x_{i,j}$ can go up by Δx, stay the same or go down by Δx, with probabilities p^u, p^m and p^d, respectively. As before, we choose $\Delta x = \sigma\sqrt{3\Delta t}$. The probabilities are obtained by matching the mean and variance over the time interval Δt and requiring that the probabilities sum to one. Accordingly, we obtain:

$$
\begin{aligned}
p_u &= \frac{1}{2}\left(\frac{\sigma^2\Delta t + \mu^2}{\Delta x^2} + \frac{\mu}{\Delta x}\right), \\
p_m &= 1 - \frac{\sigma^2\Delta t + \mu^2}{\Delta x^2} \text{ and} \\
p_d &= \frac{1}{2}\left(\frac{\sigma^2\Delta t + \mu^2}{\Delta x^2} - \frac{\mu}{\Delta x}\right),
\end{aligned}
\tag{18}
$$

where $\left((r+c) - \frac{1}{2}\sigma^2\right)\Delta t = E[\Delta x|(i,j)] = \mu$ and $\sigma^2\Delta t + \left((r+c) - \frac{1}{2}\sigma^2\right)^2 \Delta t^2 = \sigma^2\Delta t + \mu^2\Delta t^2 = E[\Delta x^2|(i,j)]$. The value of the commodity price relative to the initial commodity price at node (i,j), which corresponds to the i^{th} time step and level j in the tree is $p_{i,j} = p_0 \exp\{j\Delta x\}$.

The complete tree is constructed according to space-time description above, that is, using time-steps Δt and x-steps of length $\Delta x = \sigma\sqrt{3\Delta t}$. Each node is represented by (i,j) for $i = 0, \ldots, N$ and $j = j_{min}, \ldots, j_{max}$, where N is the total number of time-steps and j_{min} and j_{max} are the minimum and maximum levels in the tree at each time step i. At every node, we test whether the commodity spot price, p_t, is greater than the threshold price, p^*, and choose the probabilities $p^u_{i,j}$, $p^m_{i,j}$ and $p^d_{i,j}$ accordingly. That is, the local probabilities are defined by (17) if $p_{i,j} > p^*$, which means that the price follows the stochastic process described by equation (5); otherwise the local probabilities are given by equation (18). Note that the branching decision between (a), (b) and (c) in Figure 7 in Appendix B is only necessary to be considered when p_t follows the mean-reverting process; otherwise, we have the "normal" branching process as described above.

4 Numerical Results

In this section we illustrate and analyze examples of forward curves generated by both our model and by the mean-reverting single-factor model described by Schwartz (1997). As mentioned before, the latter model represents the unconstrained version of our model. This enables us to study the effect of the introduction of storage through the imposition of an arbitrage-free constraint in our model. The lattices are implemented for a period of five years, that is, from $t = 0$ to $T = 5$

α	σ	$\overline{x}$	r	c
3	0.2	ln 45	0.05	0.1

Table 1: Parameters values used in the lattice computation.

using the parameter values displayed in Table 1 below.

Figure 1 illustrates a set of forward curves generated by our model, which corresponds to five different spot price values at time $t = 0$, that is, $p_0 = 25,\ 35,\ 45,\ 55$ and 65. This range of values is representative for the spot prices that are generated by this model for the parameters shown in Table 1. Figure 2 illustrates the forward curves produced by the unconstrained single-factor Schwartz model. These curves are computed for the same range of initial spot prices and the same parameters values.

Figures 1 and 2 show that, for high spot prices at time $t = 0$, the curves are in backwardation; conversely, for low spot prices at $t = 0$, the curves are in contango. In either case, the curves eventually move towards a long-run state independent forward price, F_∞. For the unconstrained version of our model, the long-run mean is equal to 44.85, which is equal to the long-run mean of the O-U process[4]. For our model, the long run mean is slightly lower and approximately equal to 42.3. However, when $p_0 = 25$ the forward curve implied by our model does not reach a long run steady price for the five year period considered. We will discuss this issue later below. The length of time each curve remains in contango/backwardation is directly proportional to the distance between the spot price at time zero and the long-run equilibrium price, $\overline{p}$.

Figures 3 and 4 illustrate the corresponding convenience yield curves. If the initial spot price, p_0, is significantly above the long-run mean, the corresponding convenience yield is also high. High spot prices signal tight demand and supply conditions in the market, which imply the possibility of a stockout. This in turn leads to high convenience yield values since there is a high benefit from holding inventory. As a result we observe backwardation. On the other hand, if the spot prices are low because there is low-demand/high supply in the market, the stockholder builds

[4]In the long-run, x_t follows a normal distribution with $E[x_t] = m - \frac{\sigma^2}{2\alpha} = 3.80$. Since $p_t = \exp(x_t)$, in the long-run $E[p_t] = \exp(m - \frac{\sigma^2}{4\alpha}) = 44.85$.

inventory in the expectation of a rise in the spot prices. These circumstances imply a low convenience yield and therefore we observe contango. Both these conditions are only temporary because the supply/demand conditions in the market tend to adjust and the forward curves move towards an equilibrium long-run forward price, F_∞.

The main differences between the forward curves implied by the unconstrained single-factor model by Schwartz (1997) and the constrained model presented here are observed when the market is in contango, while the forward curves are similar in backwardation. The degree of contango measured by the slope of the convenience yield is much greater in the unconstrained model, whereby we observe negative values in the convenience yield. This implies that the arbitrage-free condition given by equation (15) is violated. We also observe that the contango and backwardation phenomena observed in the one-factor model is symmetric, whereas it is asymmetric in our model . This is an immediate consequence of the nature of the arbitrage-free condition that restricts the spot price drift when stock is being held, which only affects the contango relationship in the forward curve.

We would expect that the futures return is equal to the cost of carrying inventory when the curve is in contango and it is equal to zero in the long run since the forward price becomes equal to a steady-state constant, $F_\infty = \overline{P}$. This property is consistent with the forward curve analysis presented by Routledge et al. (2000) and Ribeiro and Hodges (2004a). However, the convenience yield implied by the model presented here is initially equal to zero and starts to rise slowly afterwards until it becomes equal to the cost of carry. Moreover, if the initial spot price is very low compared with $\overline{P}$ (for example $p_0 = 25$), the convenience yield and therefore the forward curve do not reach the steady-state for the period of five years considered. In particular, the futures returns given by the slope of the forward curve, do not cover the cost of holding inventory.

Figures 5 and 6 illustrate the probability density functions for the spot prices sample paths at final time $T = 5$ when $p_0 = 45$[5] generated by our model and by the unconstrained mean reverting model, respectively. Table 2 shows correspond-

[5]Without loss of generality.

ing sample moments. As expected, the unconstrained model presents a perfectly centered distribution with a gaussian kurtosis since the log spot prices generated by this model have a steady-state normal distribution in the long-run. The introduction of the constraint in the spot prices drift skews the spot price distribution to the left and adds a significant amount of kurtosis. The latter property is desirable commodity prices typically exhibit high kurtosis. However, the left skewness is not a property that we usually observe in commodity markets. On the contrary, commodity prices are generally right skewed, since we frequently observe upward (not downward) price spikes in commodity markets. As mentioned previously, this behaviour results from the non-negativity constraint in the inventory.

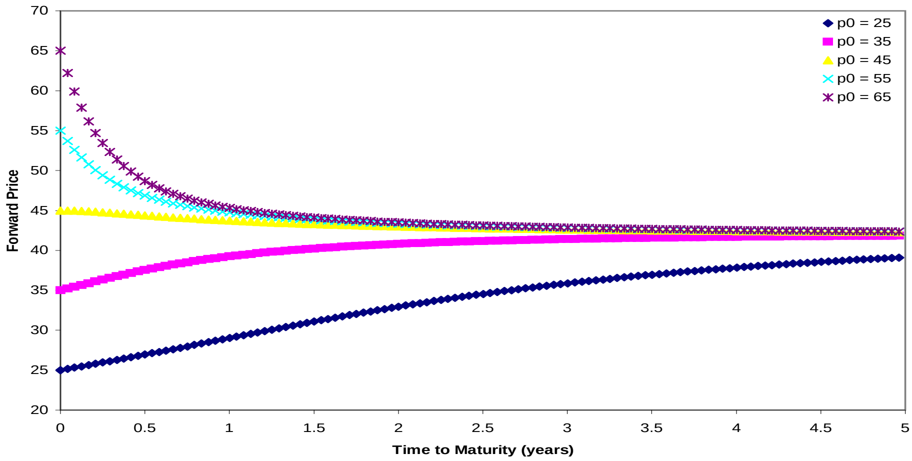

Figure 1: Forward curves generated by our model at different spot price levels for a 5-year period.

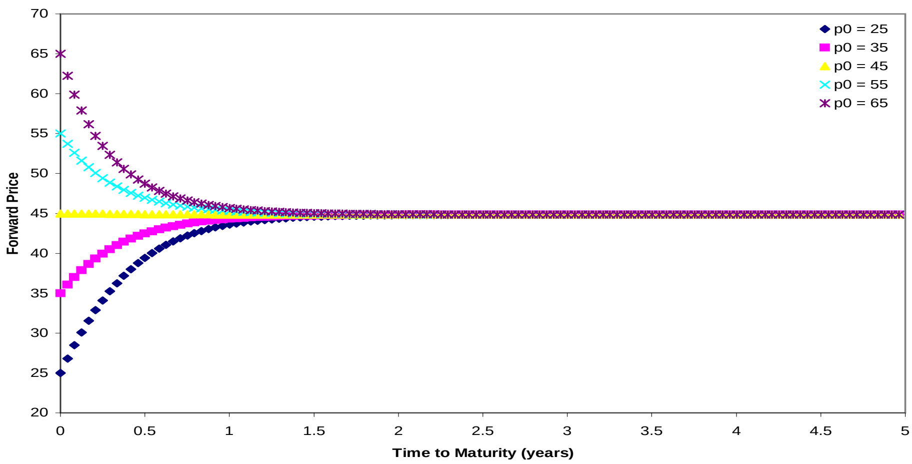

Figure 2: Forward curves generated by the one-factor mean-reverting model at different spot price levels for a 5-year period.

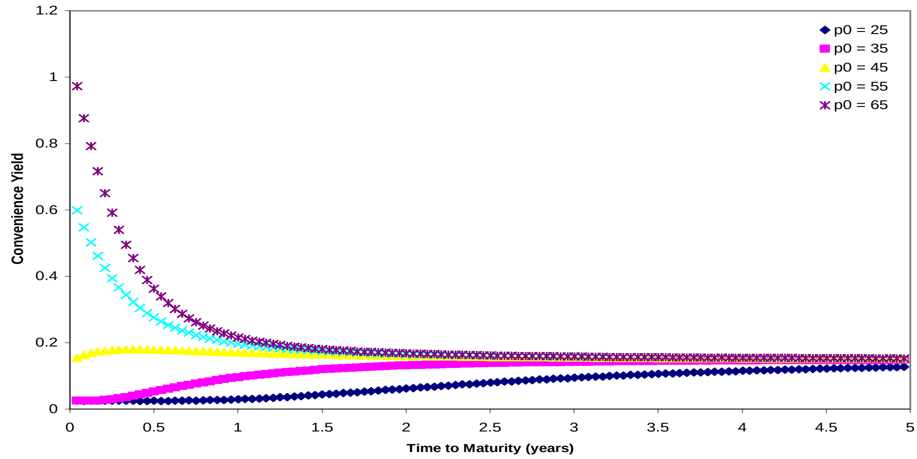

Figure 3: Convenience yield term structures generated by our model at different spot price levels for a 5-year period.

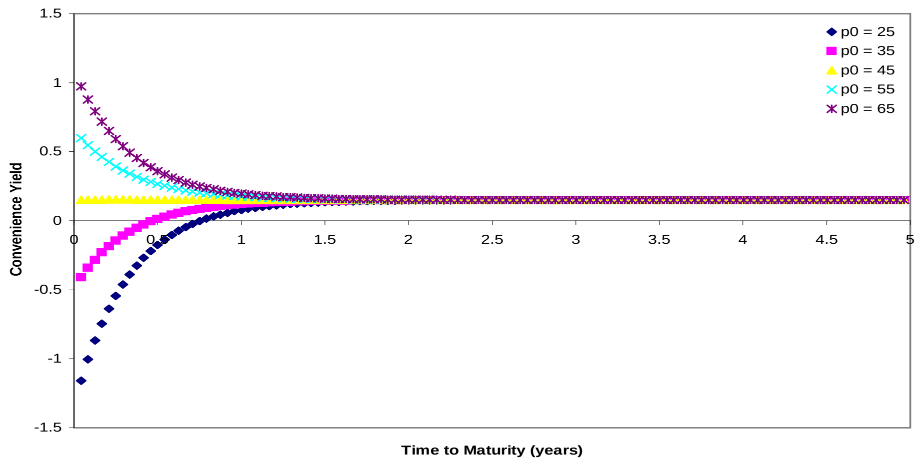

Figure 4: Convenience yield term structures generated by the one-factor mean-reverting model at different spot price levels for a 5-year period.

	Mean	*St.Dev.*	*Skewness*	*Kurtosis*
Our Model	3.73	0.15	-1.35	6.07
Unconstrained Model	3.80	0.08	0.00	3.00

Table 2: Sample moments at final time $T = 5$ for the log price sample paths x_T.

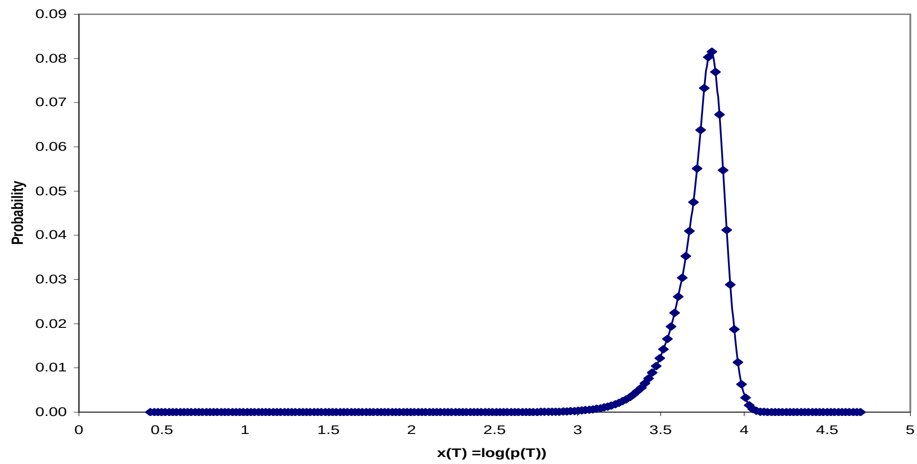

Figure 5: Probability density function for the spot prices sample paths at final time $T = 5$ for our model when the initial spot price is equal to 45.

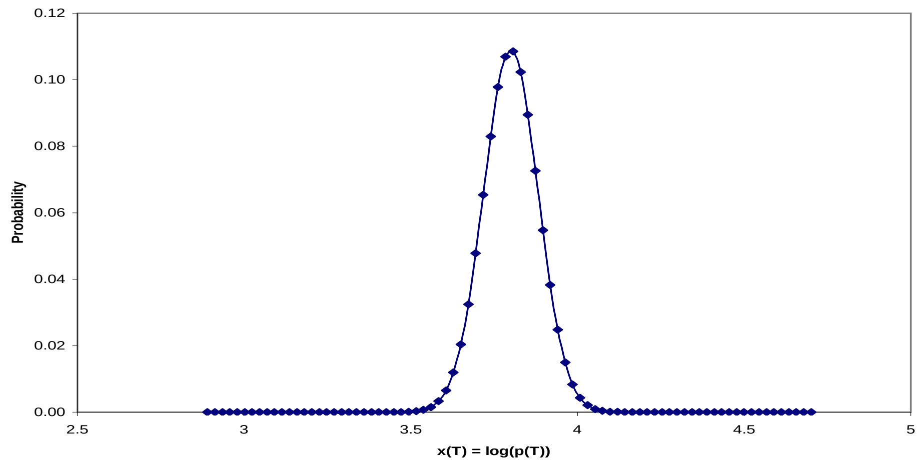

Figure 6: Probability density function for the spot prices sample paths at final time $T = 5$ for the unconstrained model when the initial spot price is equal to 45.

5 Conclusion

In this study, we have explored the properties of the simplest mean reverting model for commodity prices that satisfy the arbitrage-free contango constraint. Rather than modelling explicitly the convenience yield we impose a cash-and-carry arbitrage-free constraint in the O-U stochastic process that drives the spot prices. Our model generates a rich set of forward curve dynamics without violating arbitrage-free conditions. The general properties of the forward curves are consistent with the theory of storage - low initial spot prices generate contango and high initial spot prices generate backwardation. In addition, the degree of backwardation (contango) is increasing (decreasing) with the initial spot prices. In the long run, however, the forward move towards a long-run value, which is independent of the initial spot price. We ensure that the convenience yield is always positive and therefore this model does not violate arbitrage-free conditions. Additionally, our model implies that the spot price distribution presents excess kurtosis, which is a property observed in most commodity markets. From this perspective, our model is more appropriate than the standard ones such as Gibson and Schwartz (1990) and Schwartz (1997).

Nevertheless we recognize two main misspecifications in our model. First, the futures returns observed in contango are smaller than the cost of carrying inventory. This implies that if the initial spot price is very low when compared with the long-run mean price, the forward curve does not reach a steady-state equilibrium within a realistic length of time. This result diverges from the analysis of the forward curve provided by Routledge et al. (2000) and Ribeiro and Hodges (2004a). The other drawback is that the spot price distribution is left skewed, which is not a desirable property in commodity price distributions. Commodity prices are characterized by long periods of stagnant prices interrupted by sudden upward spikes. This means that the empirical commodity spot price distributions have excess kurtosis and are right skewed as a consequence of stockholders buying when prices are low but being unable to sell when they are high and the inventory is empty. Notwithstanding these two shortcomings, this model provides an exceptionally simple representation of commodity prices precluding arbitrage which is capable of capturing much of the essential behaviour.

A Derivation of the Process Followed by $x_t = \ln(p_t)$

Ito's Lemma tells us that is the process Y_t follows the diffusion $dY = \mu(\cdot)dt + \sigma(\cdot)dB_t$, where B_t represents a standard diffusion process, and $f(Y_t)$ is twice continuously differentiable, then:

$$df = \frac{\partial f}{\partial \mathrm{Y}_t} dY + \frac{1}{2} \frac{\partial^2 f}{\partial \mathrm{Y}_t^2} (dY)^2 \tag{19}$$

For our problem, $f(\cdot) = \ln(\cdot)$, and thus:

$$dx = \frac{1}{p_t} dp_t - \frac{1}{2p_t^2} (dp_t)^2 \tag{20}$$

Applying Ito's lemma to $x_t = \ln(p_t)$ where:

$$dp_t = \alpha(m - \ln(p_t))p_t dt + \sigma p_t dB_t, \tag{21}$$

which corresponds to equation (5.1) gives:

$$dx_t = \alpha(\overline{x} - x_t)dt + \sigma dB_t \tag{22}$$

where:

$$\overline{x} = m - \frac{1}{2}\sigma^2. \tag{23}$$

Similarly, applying Ito's lemma to the GBM described by equation:

$$dp_t = (r + c)p_t dt + \sigma p_t dB_t, \tag{24}$$

which corresponds to equation (5.2) gives:

$$dx_t = \left((r + c) - \frac{1}{2}\sigma^2 \right) dt + \sigma dB_t \tag{25}$$

Equations (22) and (25) correspond to equations (5.5) and (5.6).

B Lattice Model for the arithmetic Ornstein-Uhlenbeck Process

In this section, we describe a general and efficient procedure involving the use of a trinomial tree to implement the standard O-U process given by:

$$dz_t = \alpha(\overline{z} - z_t)dt + \sigma dB_t, \qquad t \geq 0 \tag{26}$$

where:

- α is the speed of mean reversion;
- $\overline{z}$ is the long-run mean, that is, the level to which z reverts as t goes to infinity;
- σ is the (constant) volatility;
- B_t is a standard Wiener process;

This procedure is based on the method described by Hull and White (1993, 1994) and later revised by Clewlow and Strickland (1998) and Clewlow and Strickland (2000). Originally, the method was designed to implement short interest rate that follow a mean-reverting arithmetic stochastic process. Examples of such interest rate models are the Vasicek or Hull-White model (Hull and White (1993)), the Ho-Lee model Ho and Lee (1986) and the Black-Karasinski model (Black and Karasinski (1991)).

The trinomial tree is constructed by using time steps of length Δt and z-steps of length Δz. At the end of each time step, z takes the value $z_0 + j\Delta z$, where j can be either positive or negative and z_0 is the initial value. (i, j) is defined as the node for which $t = i\Delta t$ and $z = j\Delta z$. The trinomial branching process can take any of the forms represented in Figure 7. The branching process (a) is a normal branching process where we can move up by Δz, stay the same and move down by Δz. Branching process (b) occurs when $z_{i,j}$ is currently low and $z_{i,j}$ can stay the same, move up by Δz and move up by $2\Delta z$. Branching (c) occurs when $z_{i,j}$ is currently high and can stay the same, move down by Δz and move down by $2\Delta z$. In other words, the three nodes emanating from node (i, j) are $(i + 1, k + 1)$ - the "upper" node, (i+1, k) - the "middle" node and $(i+1, k-1)$ - the "lower" node. The value of k is chosen so that $z_{i+1,k}$ is as close as possible to the expected value of z, which by definition is given by $z_{i,j} + \mu_{i,j}$, where $\mu_{i,j} = \alpha(\overline{z} + (z_0 + j\Delta z))\Delta t$. For the normal branching process $k = j$, for the branching processes illustrated by (b) and (c), $k = j+1$ and $k = j-1$, respectively.

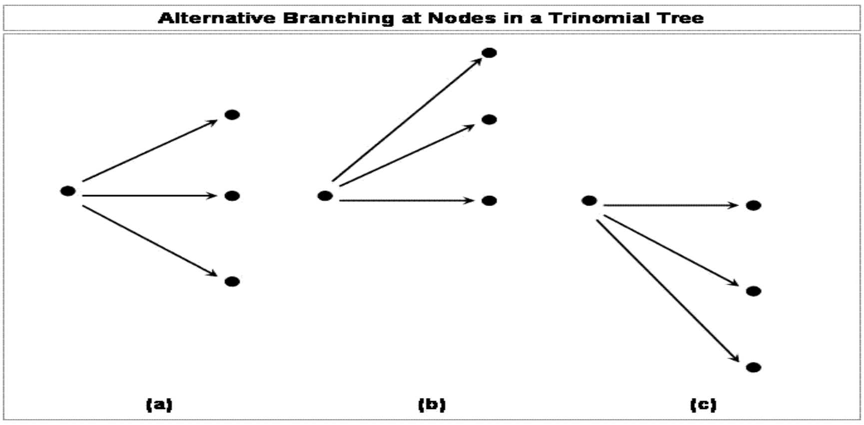

Figure 7: Alternative Branching at Nodes in a Trinomial Tree

Given the size of the time-step, Δt , Hull and White suggest that (see Hull and White (1990)), accordingly to stability and convergence conditions, a good relationship between Δt and the space step Δz is:

$$\Delta z = \sigma\sqrt{3\Delta t} \tag{27}$$

(i, j) is defined as the node for which $t = i\Delta t$ and $z = j\Delta z$. Define $pr^u_{i,j}$, $pr^m_{i,j}$ and $pr^d_{i,j}$ as the probabilities of the highest, middle and lowest branches emanating from node (i, j). The probabilities are chosen to match the expected change and variance in z over the next interval Δt. The probabilities must also sum to unity. Accordingly, the probabilities are determined by the three relationships:

$$\begin{aligned} pr^u_{i,j} + pr^m_{i,j} + pr^d_{i,j} &= 1 \\ (k+1-j)\Delta z pr^u_{i,j} + (k-j)\Delta z pr^m_{i,j} + (k-1-j)\Delta z pr^d_{i,j} &= \mu_{i,j} \\ ((k+1-j)\Delta z)^2 pr^u_{i,j} + ((k-j)\Delta z)^2 pr^m_{i,j} + (k-1-j)\Delta z)^2 pr^d_{i,j} &= \sigma^2_{i,j}\Delta t + (\mu_{i,j})^2 \end{aligned} \tag{28}$$

where $\mu_{i,j} = E[\Delta z|(i,j)] = \alpha(\overline{z} - z_{i,j})\Delta t$ is the conditional expectation of the discretized z_t process at node (i,j) and $\sigma_{i,j}^2 \Delta t + (\mu_{i,j}\Delta t)^2 = E[\Delta z^2|(i,j)]$, where $\sigma_{i,j} = \sigma$ form equation (26) above. The probabilities are given by:

$$
\begin{aligned}
pr_{i,j}^u &= \frac{\sigma^2\Delta t + \eta^2}{2\Delta z^2} + \frac{\eta}{2\Delta z} \\
pr_{i,j}^m &= 1 - \frac{\sigma^2\Delta t + \eta^2}{\Delta z^2} \\
pr_{i,j}^d &= \frac{\sigma^2\Delta t + \eta^2}{2\Delta z^2} - \frac{\eta}{2\Delta z}
\end{aligned}
\qquad (29)
$$

where $\eta = \mu_{i,j} + (j-k)\Delta z$ and $k = j-1, j$ and $j+1$, depending on the type of branching, as described above.

Provided that Δz is within the range $\sigma\sqrt{3\Delta t}/2$ to $2\sigma\sqrt{\Delta t}$, the probabilities are always between 0 and 1 (Hull and White (1993)).

References

Black, F. and Karasinski, P. (1991). Bond and option pricing when short rates are lognormal. *Financial Analysts Journal*, July-August:52–59.

Clewlow, L. and Strickland, C. (1998). *Implementing Derivatives Models*. John Wiley, London.

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

Dixit, A. K. and Pindyck, R. S. (1994). *Investment Under Uncertainty*. Princeton University Press, Princeton, New Jersey.

Epstein, N., Mayor, P., Schonbucher, A., Whalley, E., and Wilmott, P. (1998). The valuation of a firm advertising optimally. *The Quarterly Review of Economics and Finance*, 38(2):149–166.

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

Hilliard, J. and Reis, J. (1998). Valuation of commodity futures and options under stochastic convenience yields, interest rates, and jump diffusions on the spot. *Journal of Financial and Quantitative Analysis*, 33(1):61–86.

Ho, T. and Lee, S. (1986). Term structure movements and pricing of interest rate claims. *Journal of Finance*, 41:1011–1029.

Hull, J. and White, A. (1990). Valuing derivative secutiries using the explicit finite difference method. *Journal of Financial and Quantitative Analysis*, 25.

Hull, J. and White, A. (1993). One-factor interest rate models and the valuation of interest-rate derivative securities. *Journal of Financial and Quantitative Analysis*, 28(2):235–254.

Hull, J. and White, A. (1994). Numerical procedures for implementing term structure models i: Single-factor models. *Journal of Derivatives*, Fall:7–16.

Metcalf, G. E. and Hasset, K. A. (1995). Investment under alternative return assumptions comparing random walks and mean reversion. *Journal of Economic Dynamics and Control*, 19:1471–1488.

Miltersen, K. R. and Schwartz, E. S. (1998). Pricing of options on commodity futures with stochastic term structures of convenience yields and interest rates. *Journal of Financial and Quantitative Analysis*, 33(1):33–59.

Nielsen, M. J. and Schwartz, E. S. (2004). Theory of storage and the pricing of commodity claims. *Review of Derivatives Research*, 7(1):5–24.

Ribeiro, D. and Hodges, S. (2004a). Equilibrium storage model for commodity price dynamics: Competitive and monopolistic market. Working Paper, University Warwick.

Ribeiro, D. and Hodges, S. (2004b). A two-factor model for commodity prices and futures valuation. Working Paper, University Warwick.

Routledge, B., Seppi, D. J., and Spatt, C. (2000). Equilibrium forward curves for commodities. *The Journal of Finance*, 55(3):1297–1338.

Schwartz, E. S. (1997). The stochastic behaviour of commodity prices: Implications for valuation and hedging. *Journal of Finance*, 52(3):923–973.